

This is a postprint version of the following published document:

Gomez, J. V., Alvarez, D., Garrido, S. & Moreno, L.
(2019). Fast Methods for Eikonal Equations: An
Experimental Survey. *IEEE Access*, 7, 39005–39029.

DOI: [10.1109/access.2019.2906782](https://doi.org/10.1109/access.2019.2906782)

© 2019, IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Fast Methods for Eikonal Equations: an Experimental Survey

Javier V. Gómez · David Álvarez ·
Santiago Garrido · Luis Moreno

Received: date / Accepted: date

Abstract The Fast Marching Method is a very popular algorithm to compute times-of-arrival maps (distances map measured in time units). Since their proposal in 1995, it has been applied to many different applications such as robotics, medical computer vision, fluid simulation, etc. Many alternatives have been proposed with two main objectives: to reduce its computational time and to improve its accuracy. In this paper, we collect the main approaches which improve the computational time of the standard Fast Marching Method, focusing on single-threaded methods and isotropic environments. 9 different methods are studied under a common mathematical framework and experimentally in representative environments: Fast Marching Method with binary heap, Fast Marching Method with Fibonacci Heap, Simplified Fast Marching Method, Untidy Fast Marching Method, Fast Iterative Method, Group Marching Method, Fast Sweeping Method, Lock Sweeping Method and Double Dynamic Queue Method.

Keywords Fast Marching Method · Fast Sweeping Method · Fast Iterative Method · Eikonal Equation

1 Introduction

The Fast Marching Method and its derivatives (Fast Methods in short) have been extensively applied since they were firstly proposed in 1995 [35] as a

This work was supported by the Spanish Ministry of Science and Innovation under project DPI2010-17772 and by the Comunidad Autonoma de Madrid and Structural Funds of the EU under project S2013/MIT-2748.

RoboticsLab. Carlos III University of Madrid Avda. de la Universidad 30, 28911, Leganes, Madrid, Spain.
Tel.: +34-916248812
E-mail: jvgomez,dasanche,sgarrido,moreno@ing.uc3m.es

solution to isotropic control problems using first-order semi-Lagrangian discretizations on Cartesian grids. Their main field of application are robotics [15, 11, 23] and computer vision [12], mainly medical image segmentation [4, 2]. However, it has proven to be useful in many other applications such as tomography [26] or seismology [39].

The first approach was proposed by Tsitsiklis [35] but the most popular solution was given few months later by Sethian [29] using first-order upwind-finite differences in the context of isotropic front propagation. Differences and similarities between both works can be found in [33].

Fast Methods were originally proposed to simulate a wavefront propagation through a regular discretization of the space. However, many different approaches have been proposed, extending these methods to other discretizations and formulations. For a more detailed history of Fast Marching methods, we refer the interested readers to [7]. Despite the vast amount of Fast Methods literature, there is a lack of in-depth comparison and benchmarking.

In this work 9 sequential (mono-thread), isotropic, grid-based Fast Methods are detailed in the following sections: Fast Marching Method (FMM), Fibonacci-Heap FMM (FMMFib), Simplified FMM (SFMM), Untidy FMM (UFMM), Group Marching Method (GMM), Fast Iterative Method (FIM), Fast Sweeping Method (FSM), Lock Sweeping Method (LSM) and Double Dynamic Queue Method (DDQM). All these algorithms provide exactly the same solution except for UFMM and FIM, which have bounded errors. However, the question of which one is the best for which applications is still open. For example, [16] compares only FMM and FSM in spite of the fact that GMM and UFMM were already published. This survey [20] mentions most of the algorithms but only compares FMM and SFMM. A more recent work compares FMM, FSM and FIM in 2D [6]. However, FIM is parallelized and implemented in CUDA providing a biased comparison. Fig. 1 schematically shows the comparisons among algorithms carried out in the literature. Methods such as UFMM have been barely compared to their counterparts while others as FMM and FSM are compared in many works despite the fact that it is well known when each perform better: FSM is faster in simple environments with constant velocity. Also, results from one work cannot be directly extrapolated to other works since the performance of these methods highly depend on their implementation.

Statement of contributions: Three main contributions are included in this paper: 1) Based on previous works, a common formulation and notation is given for all the algorithms. This way, it is possible to easily understand their working principles and mathematical formulation. More detailed formulations are available in the literature but we decided to focus from a practical perspective. 2) A recent survey of the work on designing sequential Fast Methods during the last years. We are not taking into account parallel or high-accuracy approaches as these fields are dense enough to fill another survey. 3) Extensive and systemic comparison among the mentioned methods with experiments designed taking into account their applications and results previously reported.

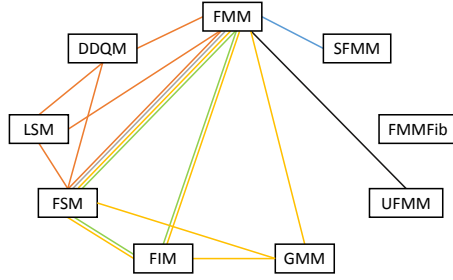


Fig. 1 Comparisons among algorithms. Colors refer to different works: orange [3], gray [16], yellow [19], green [6], black [37], and blue [20].

It is important to note that we have not taken into account the three methods proposed in [7] (Fast Marching-Sweeping Method, Heap Cell Method and Fast Heap Cell Method) for 3 different reasons: 1) their error is not mathematically bounded, 2) the error highly depends on the environment, and 3) they assume that the velocities are approximately the same in regions of arbitrary size, which is a strong assumption.

This paper is organized as follows: next section introduces the common notation and formulation to be used throughout the document. Following, Fast Methods are detailed, classified attending to their algorithm family. Thus, Section 3 includes Fast Marching-like algorithms, Section 4 Fast Sweeping-based, and Section 5 contains other Fast Methods. The benchmark and its results are included in section Section 6, followed by a discussion in section Section 7. Finally, Section 8 outlines the conclusions and proposed future works.

2 Problem Formulation

Fast Methods are built to solve the nonlinear boundary value problem¹. That is, given a domain Ω and a function $F : \Omega \rightarrow \mathbb{R}_+$ which represents the local speed of the motion, drive a system from a starting set $\mathcal{X}_s \subset \Omega$ to a goal set $\mathcal{X}_g \subset \partial\Omega$ through the fastest possible path. The Eikonal equation computes the minimum time-of-arrival function $T(\mathbf{x})$ as follows:

$$\begin{aligned} |\nabla T(\mathbf{x})| F(\mathbf{x}) &= 1, \text{ on } \Omega \subset \mathbb{R}^N \\ T(\mathbf{x}) &= 0, \mathbf{x} \text{ in } \mathcal{X}_s \end{aligned} \quad (1)$$

Once solved, $T(\mathbf{x})$ represents a distances (times-of-arrival) field containing the time it takes to go from any point \mathbf{x} to the closest point in \mathcal{X}_s following the velocities on $F(\mathbf{x})$.

We assume, without loss of generality, that the domain is a unit hypercube of N dimensions: $\Omega = [0, 1]^N$. The domain is represented with a rectangular Cartesian grid $\mathcal{X} \subset \mathbb{R}^N$, containing the discretizations of the functions $F(\mathbf{x})$ and $T(\mathbf{x})$, \mathcal{F} and \mathcal{T} respectively. We refer to grid points $\mathbf{x}_{ij} = (x_i, y_i)$, $\mathbf{x}_{ij} \in \mathcal{X}$

¹ This problem formulation closely follows [30]

as the point $\mathbf{x} = (x, y)$ in the space corresponding to a cell (i, j) of the grid (for the 2D case). For notation simplicity, we will denote $T_{ij} = T(\mathbf{x}_{ij}) \approx T(\mathbf{x})$, $T_{ij} \in \mathcal{T}$, that is, T_{ij} represents an approximation to the real value of the function $T(\mathbf{x})$. Analogously, $F_{ij} = F(\mathbf{x}_{ij}) \approx F(\mathbf{x})$, $F_{ij} \in \mathcal{F}$. For a general grid of N -dimensions, we will refer to cells by their index (or key) i as \mathbf{x}_i , since a flat representation is more efficient for such datastructure. We will denote the set of Von-Neumann (4-connectivity in 2D) neighbors of grid point \mathbf{x}_{ij} as $\mathcal{N}(\mathbf{x}_{ij})$.

2.1 n-Dimensional Discrete Eikonal Equation

In this section the most common first-order discretization of the Eikonal equation is detailed. There exist many other first-order and higher-order approaches on grids, meshes and manifolds [22, 32, 1, 24].

We are deriving the discrete Eikonal equation in 2D for better understanding. The most common first-order discretization is given in [25], which uses an upwind-difference scheme to approximate partial derivatives of $T(\mathbf{x})$ ($D_{ij}^{\pm x}$ represents the one-sided partial difference operator in direction $\pm x$):

$$\begin{aligned} T_x(\mathbf{x}) &\approx D_{ij}^{\pm x} T = \frac{T_{i\pm 1, j} - T_{ij}}{\pm \Delta_x} \\ T_y(\mathbf{x}) &\approx D_{ij}^{\pm y} T = \frac{T_{i, j\pm 1} - T_{ij}}{\pm \Delta_y} \end{aligned} \quad (2)$$

$$\left\{ \begin{array}{l} \max(D_{ij}^{-x} T, 0)^2 + \min(D_{ij}^{+x} T, 0)^2 + \\ \max(D_{ij}^{-y} T, 0)^2 + \min(D_{ij}^{+y} T, 0)^2 \end{array} \right\} = \frac{1}{F_{ij}^2} \quad (3)$$

Simpler but less accurate solution to Eq. (3) is proposed in [31]:

$$\left\{ \begin{array}{l} \max(D_{ij}^{-x} T, -D_{ij}^{+x} T, 0)^2 + \\ \max(D_{ij}^{-y} T, -D_{ij}^{+y} T, 0)^2 \end{array} \right\} = \frac{1}{F_{ij}^2} \quad (4)$$

and Δx and Δy are the grid spacing in the x and y directions. Substituting Eq. (2) in Eq. (4) and letting

$$\begin{aligned} T &= T_{i,j} \\ T_x &= \min(T_{i-1,j}, T_{i+1,j}) \\ T_y &= \min(T_{i,j-1}, T_{i,j+1}) \end{aligned} \quad (5)$$

we can rewrite the Eikonal Equation, for a discrete 2D space as:

$$\max\left(\frac{T - T_x}{\Delta_x}, 0\right)^2 + \max\left(\frac{T - T_y}{\Delta_y}, 0\right)^2 = \frac{1}{F_{ij}^2} \quad (6)$$

Since we are assuming that the speed of the front is positive ($F > 0$), T must be greater than T_x and T_y whenever the front wave has not already passed over the coordinates (i, j) . Therefore, it is safe to simplify Eq. (6) to:

$$\left(\frac{T - T_x}{\Delta x}\right)^2 + \left(\frac{T - T_y}{\Delta y}\right)^2 = \frac{1}{F_{ij}^2} \quad (7)$$

Eq. (7) is a regular quadratic equation of the form $aT^2 + bT + c = 0$, where:

$$\begin{aligned} a &= \Delta_x^2 + \Delta_y^2 \\ b &= -2(\Delta_y^2 T_x + \Delta_x^2 T_y) \\ c &= \Delta_y^2 T_x^2 + \Delta_x^2 T_y^2 - \frac{\Delta_x^2 \Delta_y^2}{F_{ij}^2} \end{aligned} \quad (8)$$

In order to simplify the notation for the n-dimensional case, let us assume that the grid is composed by cubic cells, that is, $\Delta_x = \Delta_y = \Delta_z = \dots = h$. Let us denote T_d as the generalization of T_x or T_y for dimension d , up to N dimensions. We also denote as F the propagation velocity for point with coordinates (i, j, k, \dots) . Therefore, the discretization of the Eikonal is a quadratic equation with parameters:

$$\begin{aligned} a &= N \\ b &= -2 \sum_{d=1}^N T_d \\ c &= \left(\sum_{d=1}^N T_d^2 \right) - \frac{h^2}{F^2} \end{aligned} \quad (9)$$

2.2 Solving the nD discrete Eikonal equation

Wavefront propagation follows causality. That is, in order to reach a point with higher time of arrival, it should firstly travel through neighbors of such point with smaller values. The opposite would imply a jump in time continuity and therefore the solutions would be erroneous.

The proposed Eikonal solution (quadratic equation with parameters of Eq. (9)) does not guarantee the causality of the resulting distance map, as F and h can have arbitrary values. Therefore, before accepting a solution as valid its causality has to be checked. For instance, in 2D the Eikonal is solved as:

$$T = \frac{T_x + T_y}{2} + \frac{1}{2} \sqrt{\frac{2h^2}{F^2} - (T_x - T_y)^2} \quad (10)$$

called the *two-sided* update, as both parents T_x and T_y are taken into account. The solution is only accepted if $T \geq \max(T_x, T_y)$. The *upwind condition* [7] shows that:

$$T \geq \max(T_x, T_y) \iff |T_x - T_y| \leq \frac{h}{F} \quad (11)$$

If this condition fails, the *one-sided update* is applied instead:

$$T = \min(T_x, T_y) + \frac{h}{F} \quad (12)$$

This is a top-down approach: the parents are iteratively discarded until a causal solution is found. To generalize Eq. (11) is complex. Therefore, we choose to use a bottom-up approach: Eq. (12) is solved and parents are iteratively included until the time of the next parent is higher than the current solution: $T_k > T$. The procedure is detailed in Algorithms 1 and 2. The `MINTDIM()` function returns the minimum time of the neighbors in a given dimension (left and right for $dim = 1$, bottom and top for $dim = 2$, etc.). Our experiments found this approach more robust for 3 or more dimensions with negligible impact on the computational performance.

Algorithm 1 Solve Eikonal Equation

```

1: procedure SOLVEEIKONAL( $\mathbf{x}_i, \mathcal{T}, \mathcal{F}$ )
2:    $a \leftarrow N$ 
3:   for  $dim = 1 : N$  do
4:      $min_T \leftarrow \text{MINTDIM}(dim)$ 
5:     if  $min_T \neq \infty$  and  $min_T < T_i$  then
6:        $T_{\text{values}}.\text{push}(min_T)$ 
7:     else
8:        $a \leftarrow a - 1$ 
9:   if  $a = 0$  then ▷ FSM can cause this situation.
10:    return  $\infty$ 
11:    $\mathcal{T}_{\text{values}} \leftarrow \text{SORT}(T_{\text{values}})$ 
12:   for  $dim = 1 : a$  do
13:      $\tilde{T}_i \leftarrow \text{SOLVENDIMS}(\mathbf{x}_i, dim, T_{\text{values}}, \mathcal{F})$ 
14:     if  $dim = a$  or  $\tilde{T}_i < \mathcal{T}_{\text{values}, dim+1}$  then
15:       break
16:   return  $\tilde{T}_i$ 

```

Algorithm 2 Solve Eikonal for n dimensions

```

1: procedure SOLVENDIMS( $\mathbf{x}_i, dim, T_{\text{values}}, \mathcal{F}$ )
2:   if  $dim = 1$  then
3:     return  $T_{\text{values}, 1} + \frac{h}{F_1}$ 
4:    $sumT \leftarrow \sum_{i=1}^{dim} T_{\text{values}, i}$ 
5:    $sumT^2 \leftarrow \sum_{i=1}^{dim} T_{\text{values}, i}^2$ 
6:    $a \leftarrow dim$ 
7:    $b \leftarrow -2sumT$ 
8:    $c \leftarrow sumT^2 - \frac{h^2}{F_1}$ 
9:    $q \leftarrow b^2 - 4ac$ 
10:  if  $q < 0$  then ▷ Non-causal solution
11:    return  $\infty$ 
12:  else
13:    return  $\frac{-b + \text{sqrt}(q)}{2a}$ 

```

3 Fast Marching Methods

The Fast Marching Method (FMM) [29] is the most common Eikonal solver. It can be classified as a label-setting, Dijkstra-like algorithm [10]. It uses a first-order upwind-finite difference scheme to simulate an isotropic front propagation. The main difference with Dijkstra's algorithm is the operation carried out on every node. Dijkstra's algorithm is designed to work on graphs. Therefore, the value for every node \mathbf{x}_i only depends on one parent \mathbf{x}_j , following the Bellman's optimality principle [5]:

$$T_i = \min_{\mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)} (c_{ij} + T_j) \quad (13)$$

In other words, a node \mathbf{x}_i is connected to the parent \mathbf{x}_j in its neighborhood $\mathcal{N}(\mathbf{x}_i)$ which minimizes (or maximizes) the function value (in this case T_i) composed by the value of T_j plus the addition of the cost of *traveling* from \mathbf{x}_j to \mathbf{x}_i , represented as c_{ij} .

The FMM follows Bellman's optimality principle but the value for every node is computed following first-order upwind discretization of the Eikonal equation detailed in Section 2. This discretization takes into account the spatial representation (i.e. a rectangular grid) and the value of all the causal upwind neighbors. Thus, the times-of-arrival field computed by FMM is more accurate than Dijkstra's.

The algorithm divides the cells in three different sets: 1) **Frozen**: those cells which value is computed and cannot change, 2) **Unknown**: cells with no value assigned, to be evaluated, and 3) **Narrow band** (or just **Narrow**): frontier between **Frozen** and **Unknown** containing those cells with a value assigned that can be improved. These sets are mutually exclusive, that is, a cell cannot belong to more than one of them at the same time. The implementation of the **Narrow** set is a critical aspect of FMM. A more detailed discussion will be carried out in Section 3.1.

The procedure is detailed in Algorithm 3. Initially, all points² in the grid belong to the **Unknown** set with infinite arrival time. The initial points (wave sources) are assigned a value 0 and introduced in **Frozen** (lines 2-7). Then, the main FMM loop starts by choosing the element with minimum arrival time from **Narrow** (line 9). All its non-**Frozen** neighbors are evaluated: for each of them the Eikonal is solved and the new arrival time value is kept if it is improved. In case the cell is in **Unknown**, it is transferred to **Narrow** (lines 10-16). Finally, the previously chosen point from **Narrow** is transferred to **Frozen** (lines 17 and 18) and a new iteration starts until the **Narrow** set is empty. The arrival times map \mathcal{T} is returned as the result of the procedure.

² From now on, we will indistinctly use point, cell or node to refer to each element of the grid.

Algorithm 3 Fast Marching Method

```

1: procedure FMM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s$ )
  Initialization:
2:    $\text{Unknown} \leftarrow \mathcal{X}, \text{Narrow} \leftarrow \emptyset, \text{Frozen} \leftarrow \emptyset$ 
3:    $T_i \leftarrow \infty \forall \mathbf{x}_i \in \mathcal{X}$ 
4:   for  $\mathbf{x}_i \in \mathcal{X}_s$  do
5:      $T_i \leftarrow 0$ 
6:      $\text{Unknown} \leftarrow \text{Unknown} \setminus \{\mathbf{x}_i\}$ 
7:      $\text{Narrow} \leftarrow \text{Narrow} \cup \{\mathbf{x}_i\}$ 

  Propagation:
8:   while  $\text{Narrow} \neq \emptyset$  do
9:      $\mathbf{x}_{\min} \leftarrow \arg \min_{\mathbf{x}_i \in \text{Narrow}} \{T_i\}$  ▷ Narrow top operation.
10:    for  $\mathbf{x}_i \in (\mathcal{N}(\mathbf{x}_{\min}) \cap \mathcal{X} \setminus \text{Frozen})$  do ▷ For all neighbors not in Frozen.
11:       $\tilde{T}_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_i, \mathcal{T}, \mathcal{F})$ 
12:      if  $\tilde{T}_i < T_i$  then
13:         $T_i \leftarrow \tilde{T}_i$  ▷ Narrow increase operation if  $\mathbf{x}_i \in \text{Narrow}$ .
14:        if  $\mathbf{x}_i \in \text{Unknown}$  then ▷ Narrow push operation.
15:           $\text{Narrow} \leftarrow \text{Narrow} \cup \{\mathbf{x}_i\}$ 
16:           $\text{Unknown} \leftarrow \text{Unknown} \setminus \{\mathbf{x}_i\}$ 
17:         $\text{Narrow} \leftarrow \text{Narrow} \setminus \{\mathbf{x}_{\min}\}$  ▷ Narrow pop operation: add to Frozen.
18:         $\text{Frozen} \leftarrow \text{Frozen} \cup \{\mathbf{x}_{\min}\}$ 
19:  return  $\mathcal{T}$ 

```

3.1 Binary and Fibonacci Heaps

FMM requires the implementation of the **Narrow** set to have four different operations: 1) Push: to insert a new element to the set, 2) Increase: to reorder an element already existing in the set which value has been improved, 3) Top: retrieve the element with minimum value, and 4) Pop: remove the element with minimum value. As stated before, this is the most critical aspect of the FMM implementation. The most efficient way to implement **Narrow** is by using a min-heap data structure. A heap is an ordered tree in which every parent is ordered with respect to its children. In a min-heap, the minimum value is at the root of the tree and the children have higher values. This is satisfied for any parent node of the tree.

Among all the existing heaps, FMM is usually implemented with a binary heap [17]. However, the Fibonacci Heap [13] has a better amortized time for Increase and Push operations. However, it has additional computational overhead with respect other heaps. For relatively small grids, where the narrow band is composed by few elements and the performance is still far from its asymptotic behavior, the binary heap performs better. Table 1 summarizes the time complexities for these heaps³ (the priority queue will be detailed in Section 3.2). Note that n for the methods complexity is the number of cells in the map, as the worst case is to have all the cells in the heap.

Each cell is pushed and popped at most once in the heap. For each loop, the top of **Narrow** is accessed ($\mathcal{O}(1)$), the Eikonal is solved for at most 2^N neighbors

³ <http://bigocheatsheet.com/>

Table 1 Summary of amortized time complexities for common heaps used in FMM (n is the number of elements in the heap).

	Push	Increase	Top	Pop
Fibonacci	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(\log n)$
Binary	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n)$	$\mathcal{O}(1)$	$\mathcal{O}(\log n)$
Priority Queue	$\mathcal{O}(\log n)$	–	$\mathcal{O}(1)$	$\mathcal{O}(\log n)$

($\mathcal{O}(1)$ for a given N), these cells are pushed or increased ($\mathcal{O}(\log n)$ in the worst case), and finally the top cell is popped ($\mathcal{O}(\log n)$). Therefore each loop is at most $\mathcal{O}(\log n)$. Since this loop is executed at most n times, the total FMM complexity is $\mathcal{O}(n \log n)$, where n represents the total number of cells of the grid in the worst case scenario.

3.2 Simplified Fast Marching Method

The Simplified Fast Marching Method (SFMM) [20] is a relatively unknown variation of the standard FMM but with an impressive performance. SFMM is a reduced version of FMM, where **Narrow**, implemented as a simple priority queue, can contain different instances of the same cell with different values. Additionally, it can happen that the same cell belongs to **Narrow** and **Frozen** at the same time. The simplification occurs since no Increase operation is required. Every time a cell has an updated value, it is pushed to the priority queue. Once it is popped and inserted in **Frozen**, the remaining instances in the queue are simply ignored.

The advantage is that all the increase operations are substituted by push operations. Although both have the same computational complexity, the constant for push is much lower (increase requires removal and Push operations). Note that the computational complexity is maintained, $\mathcal{O}(n \log n)$.

3.3 Untidy Fast Marching Method

The Untidy Fast Marching Method (UFMM) [37, 27] follows exactly the same procedure as FMM. However, a special heap structure is used which reduces the overall computational complexity to $\mathcal{O}(n)$: the *untidy* priority queue.

This untidy priority queue is closer to a look-up table than to a tree. It assumes that the \mathcal{F} values are bounded, hence the \mathcal{T} values are also bounded. The untidy queue is a circular array which divides the maximum range of \mathcal{T} into a set of k consecutive buckets. Each bucket contains an unordered list of cells with similar T_i value. The threshold values for each bucket evolve with the algorithm, trying to maintain a uniform distribution of the elements in **Narrow** among the buckets.

Since the index of the corresponding bucket can be analytically computed, Push is $\mathcal{O}(1)$ as well as Top and Pop. The Increase operation is, in average, $\mathcal{O}(1)$ as long as $\#_{buckets} < \mathcal{O}(n)$. Therefore, the total UFMM complexity is

Algorithm 4 Simplified Fast Marching Method

```

1: procedure SFMM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s$ )
  Initialization as FMM (in Algorithm 3)

  Propagation:
2:   while Narrow  $\neq \emptyset$  do
3:      $\mathbf{x}_{\min} \leftarrow \arg \min_{\mathbf{x}_i \in \text{Narrow}} \{T_i\}$  ▷ Narrow top operation.
4:     if  $\mathbf{x}_{\min} \in \text{Frozen}$  then
5:       Narrow  $\leftarrow \text{Narrow} \setminus \{\mathbf{x}_{\min}\}$ 
6:     else
7:       for  $\mathbf{x}_i \in (\mathcal{N}(\mathbf{x}_{\min}) \cap \mathcal{X} \setminus \text{Frozen})$  do ▷ For all neighbors not in Frozen.
8:          $\tilde{T}_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_i, \mathcal{T}, \mathcal{F})$ 
9:         if  $\tilde{T}_i < T_i$  then ▷ Update arrival time.
10:           $T_i \leftarrow \tilde{T}_i$ 
11:          Narrow  $\leftarrow \text{Narrow} \cup \{\mathbf{x}_i\}$  ▷ Narrow push operation.
12:          if  $\mathbf{x}_i \in \text{Unknown}$  then
13:            Unknown  $\leftarrow \text{Unknown} \cup \{\mathbf{x}_i\}$ 
14:          Narrow  $\leftarrow \text{Narrow} \setminus \{\mathbf{x}_{\min}\}$  ▷ Narrow pop operation.
15:          Frozen  $\leftarrow \text{Frozen} \cup \{\mathbf{x}_{\min}\}$ 
16:   return  $\mathcal{T}$ 

```

$\mathcal{O}(n)$. However, since elements within a bucket are not sorted (FIFO strategy is applied in each bucket), errors are being introduced in the final result. In fact, it is shown that the accumulated additional error is bounded by $\mathcal{O}(h)$, which is the same order of magnitude as in the original FMM.

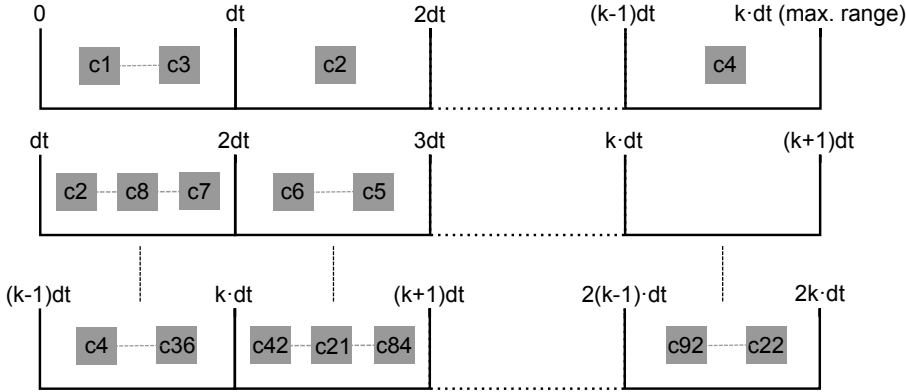


Fig. 2 Untidy priority queue representation. Top: first iteration, the four neighbors of the initial point are pushed. Middle: the first bucket became empty, so the circular array advances one position. Cell $c2$ is first evaluated because it was the first pushed in the bucket. Bottom: after a few iterations, a complete loop on the queue is about to be completed.

4 Fast Sweeping Methods

The Fast Sweeping Method (FSM) [34,40] is an iterative algorithm which computes the times-of-arrival map by successively *sweeping* (traversal) the whole grid following a specific order. FSM performs Gauss-Seidel iterations in alternating directions. These directions are chosen so that all the possible characteristic curves of the solution to the Eikonal are divided into the possible quadrants (or octants in 3D) of the environment. For instance, a bi-dimensional grid has 4 possible Gauss-Seidel iterations (the combinations of traversing x and y dimensions forwards and backwards): are North-East, North-West, South-East and South-West, as shown in Fig. 3.

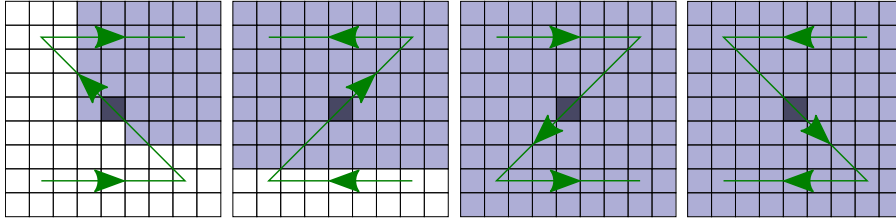


Fig. 3 FSM sweep directions in 2D represented with arrows. The darkest cell is the initial point and the shaded cells are those analyzed by the current sweep (time improved or maintained).

The FSM is a simple algorithm: it performs sweeps until no value is improved. In each sweep, the Eikonal equation is solved for every cell. However, to generalize this algorithm to N -dimensions is complex. Up to our knowledge, there are only 2D and 3D implementations. However, in Algorithm 5 we introduce an N -dimensional version. We will denote the sweeping directions as a binary array **SweepDirs** with elements 1 or -1 , with 1 (-1) meaning forwards (backwards) traversal in that dimension. This array is initialized to 1 (North-East in the 2D case or North-East-Top in 3D) and the grid is initialized as in FMM (lines 2-5). The main loop updates **SweepDirs** and a sweep is performed in the new direction (lines 8-9).

The GETSWEEPDIRS() procedure (see Algorithm 6) is in charge of generating the appropriate Gauss-Seidel iteration directions. If a 3D **SweepDirs** = $[1, 1, 1]$ vector is given, the following sequence will be generated:

$$\begin{array}{lll}
 1 - [-1, -1, -1] & 4 - [1, 1, -1] & 7 - [-1, 1, 1] \\
 2 - [1, -1, -1] & 5 - [-1, -1, 1] & 8 - [1, 1, 1] \\
 3 - [-1, 1, -1] & 6 - [1, -1, 1] &
 \end{array} \quad (14)$$

Note that this sequence creates a sweep pattern which is not exactly the same as detailed in the literature, but it is equally valid as the same directions are visited and the same number of sweeps are required to cover the whole grid.

Finally, the `SWEEP()` procedure (see Algorithm 7) recursively generates the Gauss-Seidel iterations following the traversal directions specified by the corresponding value of `SweepDirs` (Line 4). Once the most inner loop is reached, the corresponding cell is evaluated and its value updated if necessary (lines 7-11).

Algorithm 5 Fast Sweeping Method

```

1: procedure FSM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s$ )
   Initialization.
2:   SweepDirs  $\leftarrow [1, \dots, 1]$  ▷ Initialize sweeping directions.
3:    $T_i \leftarrow \infty \forall \mathbf{x}_i \in \mathcal{X}$ 
4:   for  $\mathbf{x}_i \in \mathcal{X}_s$  do
5:      $T_i \leftarrow 0$ 

   Propagation:
6:   stop  $\leftarrow \text{False}$ 
7:   while stop  $\neq \text{True}$  do
8:     SweepDirs  $\leftarrow \text{GETSWEEPDIRS}(\mathcal{X}, \text{SweepDirs})$ 
9:     stop  $\leftarrow \text{SWEEP}(\mathcal{X}, \mathcal{T}, \mathcal{F}, \text{SweepDirs}, N)$ 
10:  return  $\mathcal{T}$ 

```

Algorithm 6 Sweep directions algorithm

```

1: procedure GETSWEEPDIRS( $\mathcal{X}, \text{SweepDirs}$ )
2:   for  $i = 1 : N$  do
3:     SweepDirs $i$   $\leftarrow \text{SweepDirs}_i + 2$ 
4:     if SweepDirs $i$   $\leq 1$  then
5:       break ▷ Finish For loop.
6:     else
7:       SweepDirs $i$   $\leftarrow -1$ 
8:   return SweepDirs

```

Algorithm 7 Recursive sweeping algorithm

```

1: procedure SWEEP( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \text{SweepDirs}, n$ )
2:   stop  $\leftarrow \text{True}$ 
3:   if  $n > 1$  then
4:     for  $i \in \mathcal{X}_n$  following SweepDirs $n$  do
5:       stop  $\leftarrow \text{SWEEP}(\mathcal{X}, \mathcal{T}, \mathcal{F}, \text{SweepDirs}, n - 1)$ 
6:   else
7:     for  $i \in \mathcal{X}_1$  following SweepDirs $1$  do
8:        $\tilde{T}_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_i, \mathcal{T}, \mathcal{F})$  ▷  $\mathbf{x}_i$  is the corresponding cell.
9:       if  $\tilde{T}_i < T_i$  then
10:         $T_i \leftarrow \tilde{T}_i$ 
11:       stop  $\leftarrow \text{False}$ 
12:  return stop

```

The FSM carries out as many grid traversals as necessary until the value T_i for every cell has converged. Since no ordering is implied, the evaluation of each cell is $\mathcal{O}(1)$. As there are n cells, the total computational complexity of FSM is $\mathcal{O}(n)$. However, note that the constants highly depend on the velocity function $F(\mathbf{x})$. In the case of an empty environment with constant $F(\mathbf{x})$, only 2^N sweeps will be required as the characteristic directions are straight lines. However, for environment with obstacles or complex velocities functions, where the characteristics directions change frequently, the number of sweeps required can be much higher and therefore FSM will take longer to return a solution. Note as well that the \mathcal{T} returned by FSM is exactly the same as all the FMM-like algorithms (except UFMM).

4.1 Lock Sweeping Methods

The Lock Sweeping Method (LSM) [3] is a natural improvement over FSM. The FSM might spend computation time recomputing T_i even if none of the neighbors of \mathbf{x}_i has improved their value since the last sweep. LSM labels a cell as *unlocked* if any of its neighbors has changed and thus its value can be improved. Otherwise, the cell is labeled as *locked* and it will be skipped.

The LSM procedure is detailed in Algorithm 8. It is basically the same as FSM but with the addition of tracking if a point is locked (**Frozen**) or unlocked (**Narrow**). Analogously, the LOCKSWEEP() (see Algorithm 9) procedure is similar to SWEEP() with two differences: 1) if a point is not unlocked it is skipped (see Line 8), and 2) neighbors of \mathbf{x}_i are unlocked if the new value T_i is better than their current value.

Algorithm 8 Lock Sweeping Method

```

1: procedure LSM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s$ )
   Initialization.
2:   Frozen  $\leftarrow \mathcal{X}$ , Narrow  $\leftarrow \emptyset$ 
3:    $T_i \leftarrow \infty \forall \mathbf{x}_i \in \mathcal{X}$ 
4:   SweepDirs  $\leftarrow [1, \dots, 1]$  ▷ Initialize sweeping directions.
5:   for  $\mathbf{x}_i \in \mathcal{X}_s$  do
6:      $T_i \leftarrow 0$ 
7:     for  $\mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)$  do ▷ Unlocking neighbors of starting cells.
8:       Frozen  $\leftarrow \text{Frozen} \setminus \{\mathbf{x}_j\}$ 
9:       Narrow  $\leftarrow \text{Narrow} \cup \{\mathbf{x}_j\}$ 

   Propagation:
10:  stop  $\leftarrow \text{False}$ 
11:  while stop  $\neq \text{True}$  do
12:    SweepDirs  $\leftarrow \text{GETSWEEPDIRS}(\mathcal{X}, \text{SweepDirs})$ 
13:    stop  $\leftarrow \text{LOCKSWEEP}(\mathcal{X}, \mathcal{T}, \mathcal{F}, \text{SweepDirs}, N)$ 
14:  return  $\mathcal{T}$ 

```

Note that the asymptotic computational complexity of FSM is kept, $\mathcal{O}(n)$. The number of required sweeps is also maintained. However, in practice it

Algorithm 9 Recursive sweeping algorithm

```

1: procedure LOCKSWEEP( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \text{SweepDirs}, n$ )
2:   stop  $\leftarrow$  True
3:   if  $n > 1$  then
4:     for  $i \in \mathcal{X}_n$  following  $\text{SweepDirs}_n$  do
5:       stop  $\leftarrow$  LOCKSWEEP( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \text{SweepDirs}, n - 1$ )
6:   else
7:     for  $i \in \mathcal{X}_1$  following  $\text{SweepDirs}_1$  do
8:       if  $\mathbf{x}_i \in \text{Narrow}$  then
9:          $T_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_i, \mathcal{T}, \mathcal{F})$   $\triangleright \mathbf{x}_i$  is the corresponding cell.
10:        if  $\tilde{T}_i < T_i$  then
11:           $T_i \leftarrow \tilde{T}_i$ 
12:          stop  $\leftarrow$  False
13:          for  $\mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)$  do
14:            if  $T_i < T_j$  then  $\triangleright$  Add improvable neighbors to Narrow.
15:               $\text{Frozen} \leftarrow \text{Frozen} \setminus \{\mathbf{x}_j\}$ 
16:               $\text{Narrow} \leftarrow \text{Narrow} \cup \{\mathbf{x}_j\}$ 
17:           $\text{Narrow} \leftarrow \text{Narrow} \setminus \{\mathbf{x}_i\}$   $\triangleright$  Add  $\mathbf{x}_i$  to Frozen.
18:           $\text{Frozen} \leftarrow \text{Frozen} \cup \{\mathbf{x}_i\}$ 
19:   return stop

```

turns out that most of the cells are locked during a sweep. Therefore, the computation time saved is important.

5 Other Fast Methods

5.1 Group Marching Method

The Group Marching Method (GMM) [21] is an FMM-based Eikonal solver which solves a group of grid points in **Narrow** at once, instead of sorting them in a heap structure.

Consider a front propagating. At a given time, **Narrow** will be composed by the set of cells belonging to the wavefront. GMM selects a group G out of **Narrow** composed by the global minimum and the local minima in **Narrow**. Then, every neighboring cell to G is evaluated and added to **Narrow**. These points in G have to be chosen carefully so that causality is not violated since GMM does not sort the **Narrow** set. For that, GMM selects those points following Eq. (15):

$$G = \{\mathbf{x}_i \in \text{Narrow} : T_i \leq \min(T_{\text{Narrow}}) + \delta_\tau\} \quad (15)$$

where

$$\delta_\tau = \frac{1}{\max(\mathcal{F})} \quad (16)$$

In the original GMM work [21], $\delta_\tau = \frac{h}{\max(\mathcal{F})\sqrt{N}}$. However, we have chosen Eq. (16) as referred in [19]. Although this second formula is not mathematically

proven, the results for the original δ_τ are much worse than FMM in most of the cases, reaching an order of magnitude of difference.

If the time difference between two adjacent cells is larger than δ_τ , their values will barely affect each other since the wavefront propagation direction is more perpendicular than parallel to the line segment formed by both cells. However, the downwind points (those to be evaluated in future iterations) can be affected by both adjacent cells. Therefore, points in G are evaluated twice to avoid instabilities.

GMM is detailed in Algorithm 10. Its initialization is FMM-like. Note that δ_τ depends on the maximum velocity value in the grid. The main loop updates the threshold T_m every iteration. Firstly, it carries out a reverse traversal through the selected points, computing and updating their value (lines 17-21). Then, lines 22-31 perform a forward traversal with the same operations as the reverse traversal but updating the **Narrow** and **Frozen** sets in the same way as FMM.

Note that GMM returns the same solution as FMM. GMM evaluates twice every node before inserting it in **Frozen** while FMM only evaluates it once. However, GMM does not require any sorting. Therefore, GMM is an $\mathcal{O}(n)$ iterative algorithm that converges in only 2 iterations (traversals). The value of δ_τ can be modified: higher δ_τ would require more iterations to converge. However, smaller δ_τ will require also 2 traversals but the group G will be composed by less cells. As GMM authors point out, GMM can be interpreted as an intermediary point between FMM ($\delta_\tau = 0$) and a purely iterative method [28] ($\delta_\tau = \infty$).

Additionally, a generalized N -dimensional implementation is straightforward.

5.2 Dynamic Double Queue Method

The Dynamic Double Queue Method (DDQM) [3] is inspired in the LSM but resembles to GMM. DDQM is conceptually simple. **Narrow** is divided into two non-sorted FIFO queues: one with cells to be evaluated sooner and the other one with cells to be evaluated later. Every iteration takes an element from the first queue and evaluates it. If the time is improved, the neighboring cells with higher time are unlocked and added to the first or second queue depending on the value of the cell updated. Once the first queue is empty, queues are swapped and the algorithm continues. The purpose is to achieve a pseudo-ordering of the cells, so that cells with lower value are evaluated first.

Since queues are not sorted, it could require to solve many times the same cell until its value converges. DDQM dynamically computes the threshold value depending on the number of points inserted in each queue, trying to reach an equilibrium. The original paper includes an important analysis about this threshold update. Initially, the step value of the threshold is increased every iteration and is computed as:

Algorithm 10 Group Marching Method

```

1: procedure GMM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s$ )
  Initialization:
2:    $\text{Unknown} \leftarrow \mathcal{X}, \text{Narrow} \leftarrow \emptyset, \text{Frozen} \leftarrow \emptyset$ 
3:    $T_i \leftarrow \infty \forall \mathbf{x}_i \in \mathcal{X}$ 
4:    $\delta_\tau \leftarrow \frac{1}{\max(\mathcal{F})}$ 
5:   for  $\mathbf{x}_i \in \mathcal{X}_s$  do
6:      $T_i \leftarrow 0$ 
7:      $\text{Unknown} \leftarrow \text{Unknown} \setminus \{\mathbf{x}_i\}$ 
8:      $\text{Frozen} \leftarrow \text{Frozen} \cup \{\mathbf{x}_i\}$ 
9:     for  $\mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)$  do ▷ Adding neighbors of starting points to Narrow.
10:       $T_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_j, \mathcal{T}, \mathcal{F})$ 
11:      if  $T_i < T_m$  then
12:         $T_m \leftarrow T_i$ 
13:       $\text{Unknown} \leftarrow \text{Unknown} \setminus \{\mathbf{x}_i\}$ 
14:       $\text{Narrow} \leftarrow \text{Narrow} \cup \{\mathbf{x}_i\}$ 

  Propagation:
15:   while  $\text{Narrow} \neq \emptyset$  do
16:      $T_m \leftarrow T_m + \delta_\tau$ 
17:     for  $\mathbf{x}_i \in (\text{Narrow} \leq T_m)$  REVERSE do ▷ Reverse traversal.
18:       for  $\mathbf{x}_j \in (\mathcal{N}(\mathbf{x}_i) \cap \mathcal{X} \setminus \text{Frozen})$  do
19:          $\tilde{T}_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_j, \mathcal{T}, \mathcal{F})$ 
20:         if  $\tilde{T}_i < T_i$  then
21:            $T_i \leftarrow \tilde{T}_i$ 
22:       for  $\mathbf{x}_j \in (\text{Narrow} \leq T_m)$  FORWARD do ▷ Forward traversal.
23:         for  $\mathbf{x}_i \in (\mathcal{N}(\mathbf{x}_j) \cap \mathcal{X} \setminus \text{Frozen})$  do
24:            $\tilde{T}_j \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_i, \mathcal{T}, \mathcal{F})$ 
25:           if  $\tilde{T}_j < T_j$  then
26:              $T_j \leftarrow \tilde{T}_j$ 
27:         if  $\mathbf{x}_i \in \text{Unknown}$  then
28:            $\text{Unknown} \leftarrow \text{Unknown} \setminus \{\mathbf{x}_i\}$ 
29:            $\text{Narrow} \leftarrow \text{Narrow} \cup \{\mathbf{x}_i\}$ 
30:        $\text{Narrow} \leftarrow \text{Narrow} \setminus \{\mathbf{x}_i\}$ 
31:        $\text{Frozen} \leftarrow \text{Frozen} \cup \{\mathbf{x}_i\}$ 
32:   return  $\mathcal{T}$ 

```

$$step = \frac{1.5hn}{\sum_i F_i} \quad (17)$$

where n is the total number of cells in the grid. Originally, this step was proposed as $step = \frac{1.5n}{h \sum_i \frac{1}{F_i}}$. However, step should have time units and this

expression $[t^{-1}]$ (probably an error due to the ambiguity of using speed F or slowness $f = \frac{1}{F}$). Therefore, we propose as alternative Eq. (17).

Every time the first queue is empty, `UPDATESTEP()` (see Algorithm 11) is called, with the value of the current $step$, c_1 , and c_{total} which are the number of cells inserted in the first queue and the total number of cells, correspondingly. $step$ is modified so that the number of cells inserted in the first queue are between 65% and 75%. This is a conservative approach, since the closer

this percentage is to 50% the faster DDQM is. However, the penalization for percentages lower than 50% is much important than for higher percentages.

Note that the step is increased by a factor 1.5 but decreased by a factor of 2. This makes *step* to converge to a value instead of overshooting around the optimal value. Dividing by a larger number causes the first queue to become empty earlier. Thus, next iteration will finish faster and a better *step* value can be computed.

Algorithm 11 DDQM Threshold Increase

```

1: procedure UPDATESTEP(step,  $c_1$ ,  $c_{total}$ )
2:    $m \leftarrow 0.65$ 
3:    $M \leftarrow 0.75$ 
4:    $Perc \leftarrow 1$ 
5:   if  $c_1 > 0$  then
6:      $Perc \leftarrow \frac{c_1}{c_{total}}$ 
7:   if  $Perc \leq m$  then
8:      $step \leftarrow step * 1.5$ 
9:   else if  $Perc \geq M$  then
10:     $step \leftarrow \frac{step}{2}$ 
11:  return step

```

DDQM is detailed in Algorithm 12. As in LSM, points are locked (**Frozen**) or unlocked (**Narrow**). Initialization sets all points as frozen except the neighbors of the start points, which are added to the first queue (lines 2-14). While first queue is not empty, its front element is extracted and evaluated (lines 16-18). If its value is improved, all its locked neighbors with higher value are unlocked and added to its corresponding queue.

In the original work, three methods were proposed: 1) single-queue (SQ), and therefore simpler algorithm, 2) two-queue static (TQS), where the *step* is not updated, and 3) two-queue dynamic (which we call DDQM). SQ and TQS slightly improve DDQM in some experiments, but when DDQM improves SQ and TQS (for instance environments with noticeable speed changes) the difference can reach one order of magnitude. Therefore, we decided to include DDQM instead of SQ and TQS since it has shown a more adaptive behaviour. In any case, any of these methods return the same solution as FMM.

In the worst case, the whole grid is contained in both queues and traversed many times during the propagation. However, since queue insertion and deletion are $\mathcal{O}(1)$ operations, the overall complexity is $\mathcal{O}(n)$. Note that SWAP() can be efficiently implemented in $\mathcal{O}(1)$ as a circular binary index, or updating references (or pointers). There is not need for a real swap operation.

5.3 Fast Iterative Method

The Fast Iterative Method (FIM) [19] is based on the iterative method proposed by [28] but inspired in FMM. It also resembles to DDQM (concretely to

Algorithm 12 Double Dynamic Queue Method

```

1: procedure DDQM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s$ )
  Initialization:
2:    $\text{Frozen} \leftarrow \mathcal{X}, \text{Narrow} \leftarrow \emptyset$ 
3:    $\mathcal{Q}_1 \leftarrow \emptyset, \mathcal{Q}_2 \leftarrow \emptyset$ 
4:    $c_1 \leftarrow 0$  ▷ Counters
5:    $c_{\text{total}} \leftarrow 0$ 
6:    $step = \frac{1.5hn}{\sum_i F_i}$  ▷  $n$  is the total number of cells.
7:    $th \leftarrow step$ 
8:    $T_i \leftarrow \infty \forall x_i \in \mathcal{X}$ 
9:   for  $x_i \in \mathcal{X}_s$  do
10:     $T_i \leftarrow 0$ 
11:    for  $x_j \in \mathcal{N}(x_i)$  do
12:       $\mathcal{Q}_1 \leftarrow \mathcal{Q}_1 \cup \{x_j\}$ 
13:       $\text{Unknown} \leftarrow \text{Unknown} \setminus \{x_i\}$ 
14:       $\text{Narrow} \leftarrow \text{Narrow} \cup \{x_i\}$ 

  Propagation:
15:   while  $\mathcal{Q}_1 \neq \emptyset$  or  $\mathcal{Q}_2 \neq \emptyset$  do
16:     while  $\mathcal{Q}_1 \neq \emptyset$  do
17:        $x_i \leftarrow \mathcal{Q}_1.\text{FRONT}()$  ▷ Extracts the front element.
18:        $\tilde{T}_i \leftarrow \text{SOLVEEIKONAL}(x_i, \mathcal{T}, \mathcal{F})$ 
19:       if  $\tilde{T}_i < T_i$  then
20:          $T_i \leftarrow \tilde{T}_i$ 
21:         for  $x_j \in (\mathcal{N}(x_i) \cap \text{Frozen})$  do
22:           if  $T_i < T_j$  then ▷ Add improvable neighbors to corresponding queue.
23:              $\text{Frozen} \leftarrow \text{Frozen} \setminus \{x_j\}$ 
24:              $\text{Narrow} \leftarrow \text{Narrow} \cup \{x_j\}$ 
25:              $c_{\text{total}} \leftarrow c_{\text{total}} + 1$ 
26:             if  $T_i \leq th$  then
27:                $\mathcal{Q}_1 \leftarrow \mathcal{Q}_1 \cup \{x_j\}$ 
28:                $c_1 \leftarrow c_1 + 1$ 
29:             else
30:                $\mathcal{Q}_2 \leftarrow (\mathcal{Q}_2 \cup \{x_j\})$ 
31:            $\text{Narrow} \leftarrow \text{Narrow} \setminus \{x_i\}$ 
32:            $\text{Frozen} \leftarrow \text{Frozen} \cup \{x_i\}$ 
33:        $step \leftarrow \text{UPDATESTEP}(step, c_1, c_{\text{total}})$ 
34:        $\text{SWAP}(\mathcal{Q}_1, \mathcal{Q}_2)$ 
35:        $c_1 \leftarrow 0$ 
36:        $c_{\text{total}} \leftarrow 0$ 
37:        $th \leftarrow th + step$ 
38:   return  $\mathcal{T}$ 

```

its single queue variant). It iteratively evaluates every point in **Narrow** until it converges. Once a node has converged its neighbors are inserted into **Narrow** and the process continues. *Narrow* is implemented as a non-sorted list. The algorithm requires a convergence parameter ϵ : if T_i is improved less than ϵ , it is considered converged. FIM has been also proposed for triangulated surfaces [14].

FIM is designed to be efficient for parallel computing, since all the elements in **Narrow** can be evaluated simultaneously. However, we are focusing

on its sequential implementation in order to have a fair comparison with other methods.

Algorithm 13 details FIM. Its initialization is the same as FMM. Then, for each element in **Narrow**, its value is updated (lines 11-12). If the value difference is less than ϵ , the neighbors are evaluated and added to **Narrow** in case their value is improved (lines 13-19). Since **Narrow** is a list, the new elements should be inserted just before the point being currently evaluated, \mathbf{x}_i . Finally, this point is removed from **Narrow** and labeled as **Frozen** (lines 20 and 21).

Algorithm 13 Fast Iterative Method

```

1: procedure FIM( $\mathcal{X}, \mathcal{T}, \mathcal{F}, \mathcal{X}_s, \epsilon$ )
  Initialization:
2:   Frozen  $\leftarrow \mathcal{X}$ , Narrow  $\leftarrow \emptyset$ 
3:    $T_i \leftarrow \infty \forall \mathbf{x}_i \in \mathcal{X}$ 
4:   for  $\mathbf{x}_i \in \mathcal{X}_s$  do
5:      $T_i \leftarrow 0$ 
6:     for  $\mathbf{x}_j \in (\mathcal{N}(\mathbf{x}_i) \cap \text{Unknown})$  do
7:       Frozen  $\leftarrow \text{Frozen} \setminus \{\mathbf{x}_i\}$ 
8:       Narrow  $\leftarrow \text{Narrow} \cup \{\mathbf{x}_i\}$ 

  Propagation:
9:   while Narrow  $\neq \emptyset$  do
10:    for  $\mathbf{x}_i \in \text{Narrow}$  do
11:       $\tilde{T}_i \leftarrow T_i$ 
12:       $T_i \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_i, \mathcal{T}, \mathcal{F})$ 
13:      if  $|T_i - \tilde{T}_i| < \epsilon$  then
14:        for  $\mathbf{x}_j \in (\mathcal{N}(\mathbf{x}_i) \cap \text{Frozen})$  do
15:           $\tilde{T}_j \leftarrow \text{SOLVEEIKONAL}(\mathbf{x}_j, \mathcal{T}, \mathcal{F})$ 
16:          if  $\tilde{T}_j < T_j$  then
17:             $T_j \leftarrow \tilde{T}_j$ 
18:            Frozen  $\leftarrow \text{Frozen} \setminus \{\mathbf{x}_j\}$ 
19:            Narrow  $\leftarrow \text{Narrow} \cup \{\mathbf{x}_j\}$  ▷ Insert in the list just before  $\mathbf{x}_i$ 
20:          Narrow  $\leftarrow \text{Narrow} \setminus \{\mathbf{x}_i\}$ 
21:          Frozen  $\leftarrow \text{Frozen} \cup \{\mathbf{x}_i\}$ 
22:   return  $\mathcal{T}$ 

```

A node can be added several times to **Narrow** during FIM execution, since every time an upwind (parent) neighbor is updated, the node can improve its value. In the worst case, **Narrow** contains the whole grid and the loop would go through all the points several times. Operations on the list are $\mathcal{O}(1)$. Therefore, the overall computational complexity of FIM is $\mathcal{O}(n)$.

For a small enough ϵ (depending on the environment), FIM will return the same solution as FMM. However, it can be speed up allowing small errors bounded by ϵ .

6 Experimental Comparison

6.1 Experimental setup

In order to give an impartial and meaningful comparison, all the algorithms have been implemented from scratch, in C++11 using the Boost.Heap library⁴. An automatic benchmarking application has been also created so that the experiments are carried out and evaluated in the most systematic possible way.

This implementation is focused on time performance and compiled using G++ 4.9.2 with optimizations flag `-Ofast`. However, no special optimizations have been included. All algorithms use the same primitive functions for grid and cell computations. The times reported correspond to an Ubuntu 14.04 64 bits computer running in a Dual Core 3.3 GHz with 4Gb of RAM. However, all experiments were carried out in one core. Only propagation times are taking into account. The computation time used in the initialization has been omitted since it can be done offline, besides it is similar for all algorithms and represents a little percentage of the total computation time.

Although the algorithms are deterministic (their result only depends on the grid, which is not modified), the results shown are the mean of 10 runs for every algorithm, so that the deviation of the results is practically 0.

For UFMM, the default is a maximum range of \mathcal{T} of 2 units and 1000 buckets (the checkerboard experiment required different parameters, see Section 6.1.4). The ϵ parameter for FIM is set to 0 (actually 10^{-47} to provide robust 64bit double comparison).

Although error analysis is not in the scope of this paper, these can be compared among the existing papers since they are implementation-independent. UFMM errors are reported in those experiments with non-constant velocity. Usually, L_1 and L_∞ norms of the error are reported. Most of the works compute norm L_1 as:

$$|\mathcal{T}|_1 = \sum_{\mathbf{x}_i \in \mathcal{X}} |T_i| \quad (18)$$

where \mathcal{X} is treated as a regular vector. However, following [7], we treat the numerical solutions as elements of L_p spaces (generalization of the p -norm to vector spaces), where L_1 norm is defined as an integral over the function. The result is a norm closely related to its physical meaning and independent of the cell size. L_1 is numerically integrated over the domain and therefore computed as (assuming cubic cells and grids):

$$|\mathcal{T}|_1 = \sum_{\mathbf{x}_i \in \mathcal{X}} |T_i h^N| = h^N \sum_{\mathbf{x}_i \in \mathcal{X}} |T_i| \quad (19)$$

Four different experiments have been carried out, which represent characteristic cases. By combining these problems it is possible to get close to any

⁴ The source code is available at <https://github.com/jvgomez/fastmarching>

situation. These experiments were chosen attending also to the most common situations tested in the literature.

6.1.1 Empty map

This experiment is designed to show the performance of the methods in the most basic situation, where most of the algorithms perform best. An empty map with constant velocity represents the simplest possible case for the Fast Methods. In fact, analytical methods could be implemented by computing the euclidean distance from every point to the initial point. However, it is interesting because it shows the performance of the algorithms on open spaces which, in a real application, can be part of large environments.

The same environment is divided into a different number of cells to study how the algorithms behave as the number of cell increases. Composed by an empty 2D, 3D and 4D hyper-cubical environment of size $[0, 1]^N$, with $N = 2, 3, 4$. Constant velocity $F_i = 1$ on Ω . The wavefront starts at the center of the grid.

The number of cells was chosen so that an experiment has the same (or as close as possible) number of cells in all dimensions. For instance, a 50x50 2D grid has 2500 cells. Therefore, the equivalent 3D grid is 14x14x14 (2744) and in 4D is 7x7x7x7 (2401). This way, it is possible to also analyze the performance of the algorithms for a different number of dimensions. Thus, we have chosen the following number of cells for each dimension for 2D grid:

$$2D : \{50, 100, 200, 400, 800, 1000, 1500, 2000, 2500, 3000, 4000\}$$

Consequently, the 3D and 4D cells are:

$$3D : \{14, 22, 34, 54, 86, 100, 131, 159, 184, 208, 252\}$$

$$4D : \{7, 10, 14, 20, 28, 32, 39, 45, 450, 55, 63\}$$

6.1.2 Alternating barriers

In this case, we want to analyze how the algorithms behave with obstacles ($F_i = 0$) in a constant velocity environment ($F_i = 1$). The obstacles cause the characteristics to change.

The experiment contains a 2D environment of constant size $[0, 1] \times [0, 2]$ discretized in a 1000x2000 grid. A variable number of alternating barriers are equally distributed along the longest dimension. The number of barriers goes from 0 to 9. Examples are shown in Fig. 4. Analogously, in 3D a $[0, 1] \times [0, 1] \times [0, 2]$ environment represented by a 100x100x200 is chosen, with equally-distributed alternating barriers (from 0 to 9) along the z axis. The wavefront starts in all cases close to a corner of the map.

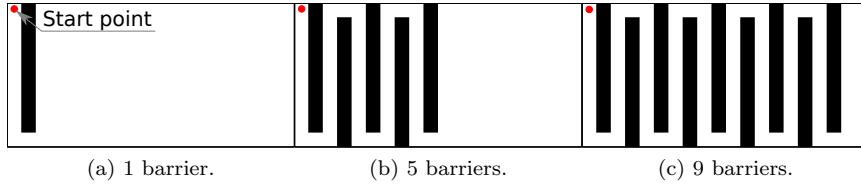


Fig. 4 2D alternating barriers environments.

6.1.3 Random velocities

In order to test the performance of the algorithms with random velocities, or noisy images as in the case of medical computer vision. This experiment creates a 2D, 3D and 4D environment of size $[0, 1]^N$ with $N = 2, 3, 4$ discretized in a 2000×2000 grid in 2D, $159 \times 159 \times 159$ in 3D and $45 \times 45 \times 45 \times 45$ in 4D. These discretizations are chosen so that it is possible to compare directly with the empty map problem for the corresponding grid sizes. The wavefront starts in the center of the grid.

Additionally, the maximum velocity is increased from 10 to 100 (in steps of 10 units) to analyze how the algorithms behave with increasing velocity changes. 2D examples are shown in Fig. 5.

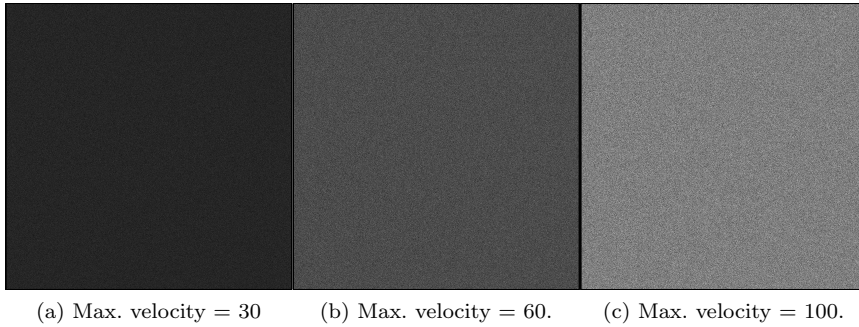


Fig. 5 2D random velocities environments. Lighter color means faster wave propagation.

6.1.4 Checkerboard

The random velocities experiment already tested changes in velocities. However, those are high-frequency changes because it is unlikely to have 2 adjacent cells with the same velocities. In this experiment low-frequency changes are studied. The same environment and discretization as in random velocities is now divided like a checkerboard, alternating minimum and maximum velocities. Analogously, the maximum velocity is increased from 10 to 100, while the

minimum velocity is always 1. There are 10 checkerboard divisions on each dimension. The wavefront starts in the center of the grid. 2D examples are shown in Fig. 5.

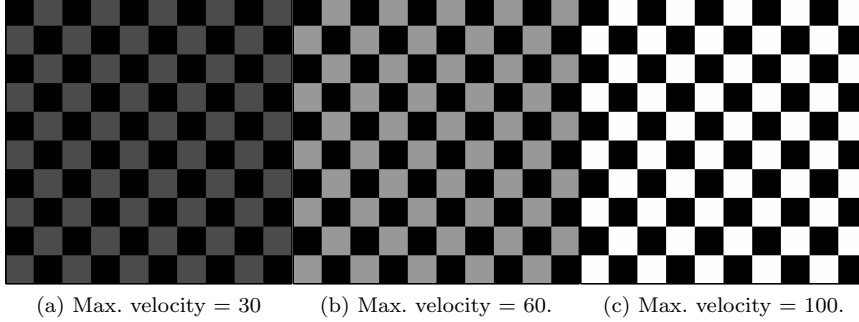


Fig. 6 2D checkerboard environments. Lighter color means faster wave propagation.

In this case, UFMM in 3D and 4D performed very poorly with the default parameters. Our previous experiments show that the best parameters for UFMM are approximately 1000 buckets with a maximum range of 0.01 in 3D. In the 4D case, 20000 buckets and maximum range of 0.025 are used.

6.2 Results

6.2.1 Empty map

An example of the times-of-arrival field computed by FMM is shown in Fig. 7. Note that all algorithms provide the same exact solution in this case. The higher the resolution the better the accuracy.

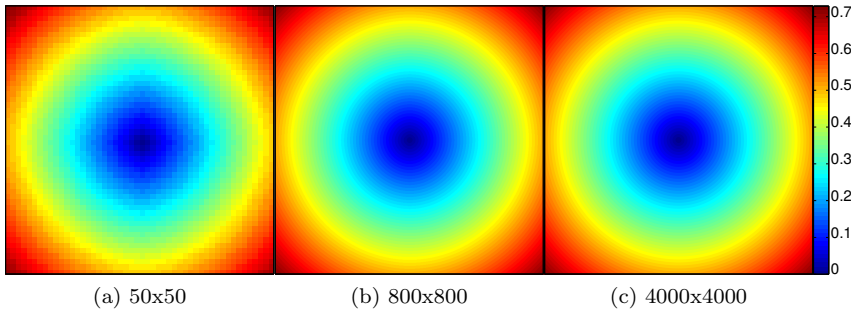


Fig. 7 Example of the resulting times-of-arrival maps applying FMM to the empty environment in 2D.

The results for the empty map experiment are shown in Fig. 8 for 2D, Fig. 9 for 3D, and Fig. 10 for 4D. In all cases 2 plots are included: raw computation times for each algorithm, and time ratios computed as:

$$ratio = \frac{\text{Alg. Time}}{\text{FMM Time}} \quad (20)$$

In all cases, both LSM and DDQM are the fastest algorithms. DDQM tends to perform better in smaller grids, specially in 4D. This was an expected result since this is the case in which they perform less sweeps (queue recomputations in DDQM). Besides, FSM is slower than UFMM in all cases and than FIM in 3D and 4D. As velocities are constant, UFMM provides the same solution as other methods,

GMM slightly improves FMM and FMMFib in 2D. However in higher dimensions this improvement becomes larger, but it always performs worse than UFMM, FIM, DDQM and LSM. In a previous comparison between GMM and FMM [19], GMM was about 50% faster in all cases than FMM. In this results GMM is at most 40% better. We attribute this difference to implementation, as the heaps for FMM and FMMFib are highly optimized. FIM is always faster as it only needs one iteration trough the narrow band, while GMM always performs two.

SFMM results are of special interest since it is a minor modification of FMM but, however, highly outperforms FMM in all cases, and even FIM in 2D and small 3D grids. In 4D FIM becomes faster. Besides UFMM outperforms FSM.

As expected, FMMFib is worse than FMM for almost all sizes. However, when dimensions increase FMMFib quickly outperforms FMM as the number of elements in the narrow band increases exponentially with the number of dimensions, therefore the better amortized times of Fibonacci heap become useful.

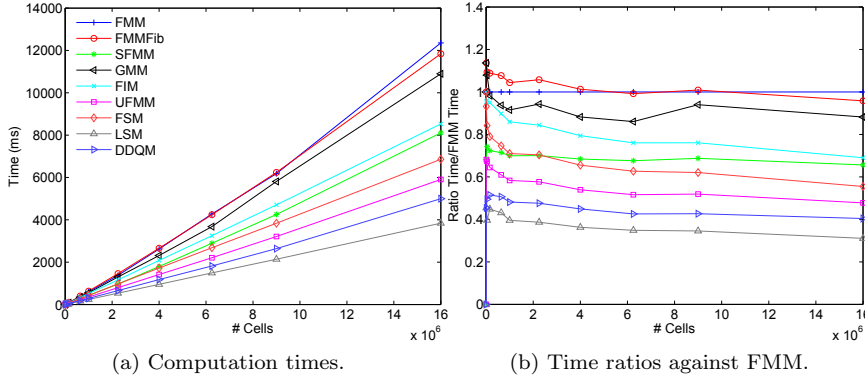


Fig. 8 Computation times and ratios for the empty map environment in 2D.

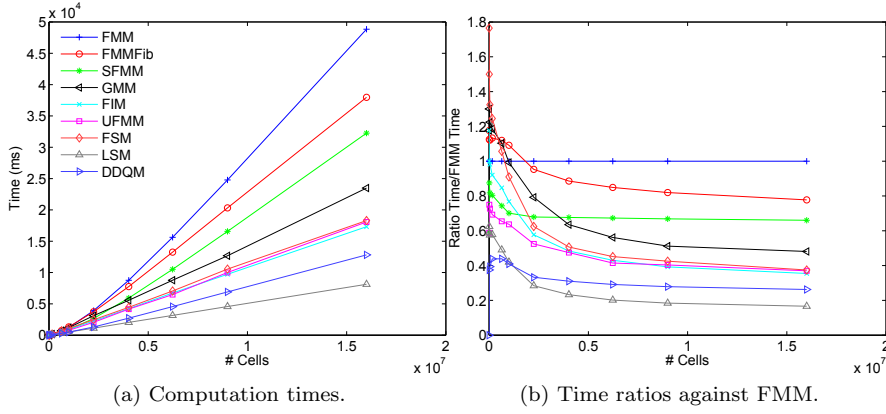


Fig. 9 Computation times and ratios for the empty map experiment in 3D.

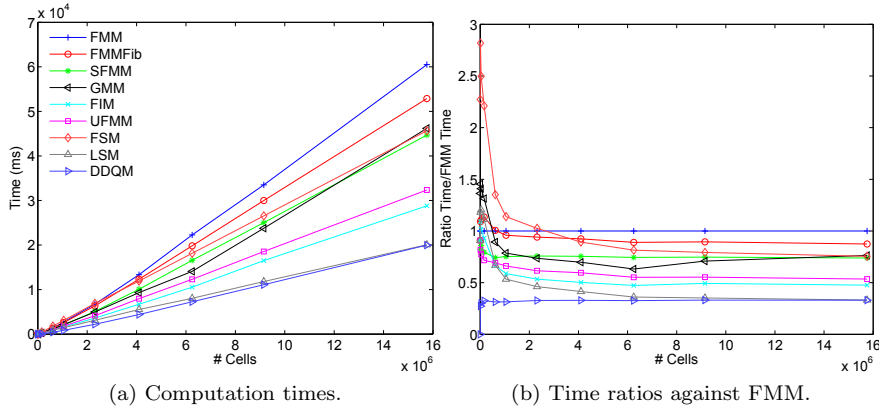


Fig. 10 Computation times and ratios for the empty map experiment in 4D.

6.2.2 Alternating barriers

An example of FMM results in some of the alternating barriers environment is shown in Fig. 11, while performance results (times and ratios) for 2D and 3D are shown in Fig. 12 and Fig. 13 correspondingly.

In this case, results are very similar to the empty map experiment. However, as the number of barrier increases and the environment becomes more complex, FSM and LSM decrease their performance. However, LSM is still faster than some algorithms in most cases.

DDQM also suffers from map complexity, however it affects much less and only in 3D. Changes in characteristic directions require more sweeps to compute the final times-of-arrival map. Therefore, FSM and LSM become worse as the number of barriers increases, despite the fact that the more barriers the

less cells are to evaluate. This is the reason why all other algorithms tend to lower times as the number of barriers increases.

UFMM provides again the same solution as other methods. It is the second fastest algorithm behind DDQM, closely followed by SFMM and FIM.

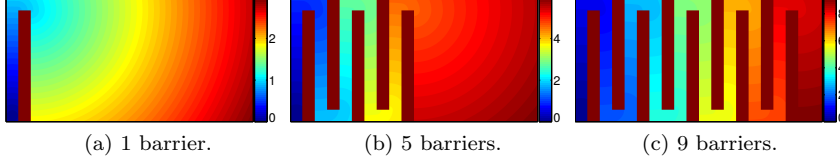


Fig. 11 Example of the resulting times-of-arrival maps applying FMM to some of the alternating barriers environment in 2D.

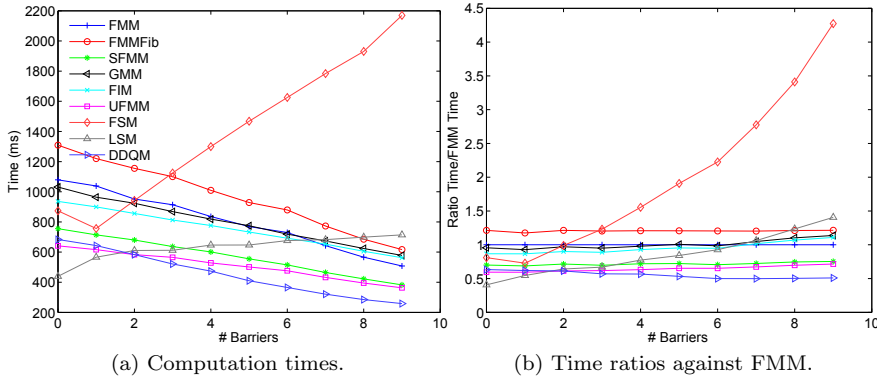


Fig. 12 Computation times and ratios for the alternating barriers experiment in 2D.

6.2.3 Random velocities

The output of the FMM for the random velocities map is apparently close to the one of the empty map (Fig. 14), but with the wavefronts slightly distorted because of the velocity changes.

Although the results on the times-of-arrival map are slightly different from the results obtained in the empty map, the performance of the algorithms are highly modified. 2D, 3D and 4D results are shown respectively in Fig. 15, Fig. 16, and Fig. 17. In this case, raw computation times are shown together with a zoomed view of the fastest algorithms to make the analysis easier. Note that all methods become slower with non-constant velocities. The reason is that the narrow band tends to have more elements in this cases.

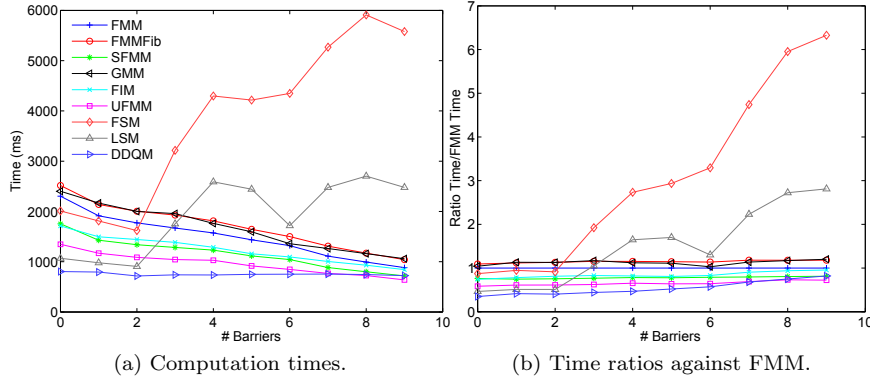


Fig. 13 Computation times and ratios for the alternating barriers experiment in 3D.

Some algorithms become unstable with non-uniform velocities: FSM, LSM and DDQM. DDQM is able to maintain the fastest time for slight velocity changes. But when these are sharper the double-queue threshold becomes unstable. However, for a high number of dimensions this effect vanishes (its computation time is barely modified with the number of dimensions) and DDQM becomes the fastest algorithms together with FIM and GMM.

SFMM provides one of the best performances across all dimensions. And FIM becomes relatively faster in 3D and 4D. However, in most of the cases GMM is the fastest algorithm. FIM requires now multiple iterations to converge to the solution while GMM guarantees convergence with only 2 iterations.

Finally, UFMM requires special attention as it does not return the same solution than the other Fast Methods. Its performance is highly affected by the number of dimensions. The main reason is the election of the parameters: they were experimentally chosen to optimize 2D performance. However, these parameters are no longer useful in other number of dimensions. We consider UFMM parameter tuning to be a complex task.

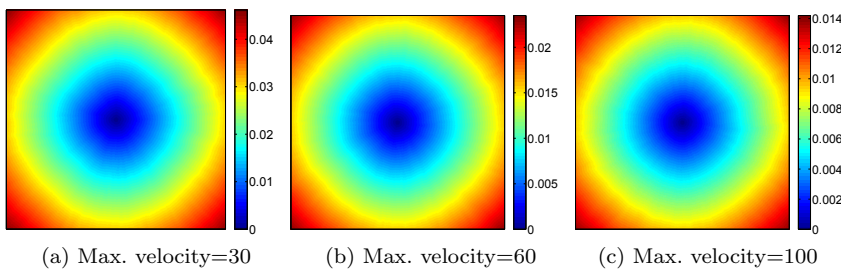


Fig. 14 Example of the resulting times-of-arrival maps applying FMM to the random velocities environment in 2D.

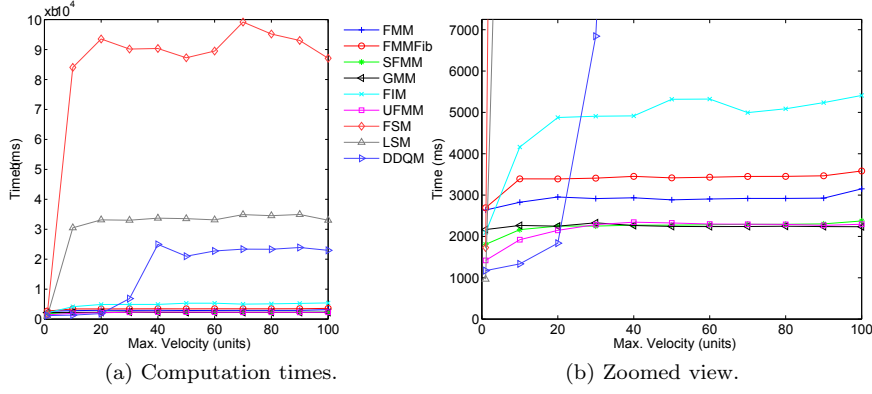


Fig. 15 Computation times and ratios for the random velocities experiment in 2D.

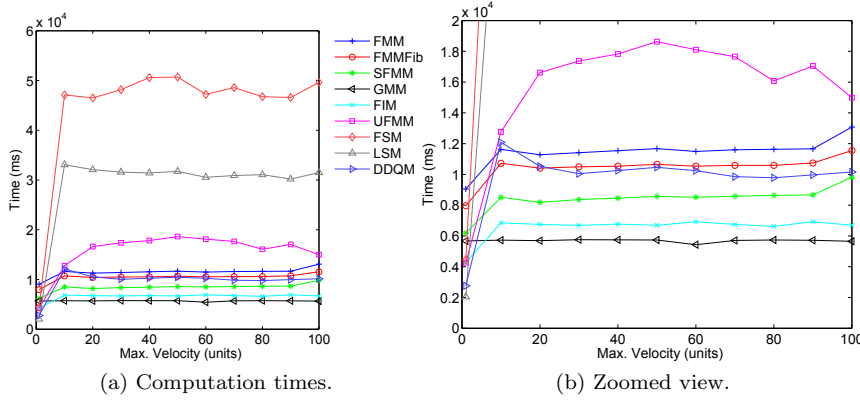


Fig. 16 Computation times and ratios for the random velocities experiment in 3D.

Table 2 summarizes the largest errors for this experiment. As the number of dimension is increased, the error decreases exponentially while the computation time increases exponentially. Therefore, by properly tuning parameters for 3D and 4D better times could be achieved while keeping a negligible error in most practical cases.

Table 2 Largest L_1 and L_∞ errors for UFMM in the random velocities experiment.

	2D	3D	4D
L_1	10^{-3}	$1.3 \cdot 10^{-10}$	$6.9 \cdot 10^{-12}$
L_∞	$4.8 \cdot 10^{-3}$	10^{-6}	10^{-7}

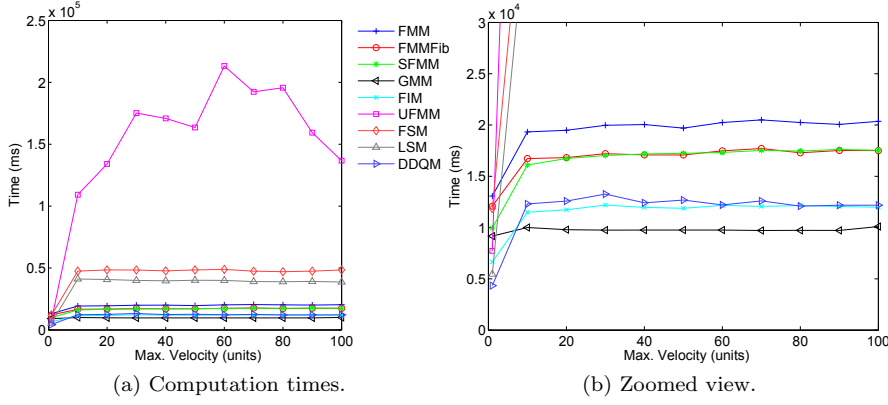


Fig. 17 Computation times and ratios for the random velocities experiment in 4D.

6.2.4 Checkerboard

Finally, different times-of-arrival maps returned by FMM applied to the checkerboard map are shown in Fig. 18. Numerical results of the computation times are included in Fig. 19 for 2D, Fig. 20 for 3D, and Fig. 21 for 4D.

The results are relatively close to the random velocities experiment. However, the differences for FSM, LSM and DDQM are much smaller. In fact, DDQM presents a poor performance in 2D, but in 3D and 4D it becomes the fastest algorithm for higher velocity modifications.

GMM and FIM have very similar results for 3D and 4D. Since the environment is structured, FIM does not require too many iterations to compute the final value.

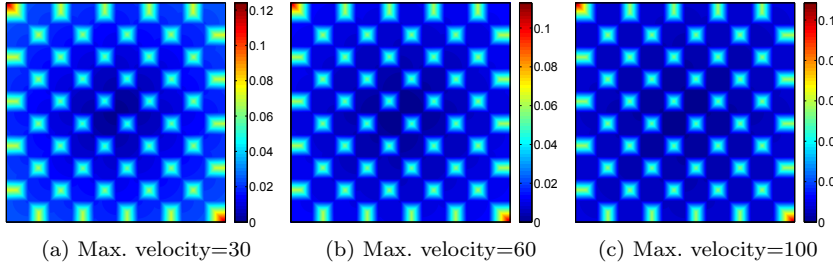


Fig. 18 Example of the resulting times-of-arrival maps applying FMM to the checkerboard environment in 2D.

UFMM errors are shown in Table 3. In this case, UFMM becomes worse with the number of dimensions but it is among the fastest algorithms in all cases.

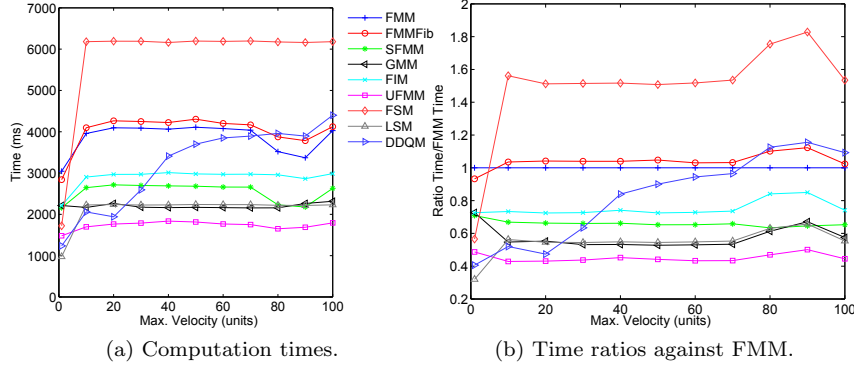


Fig. 19 Computation times for the checkerboard experiment in 2D.

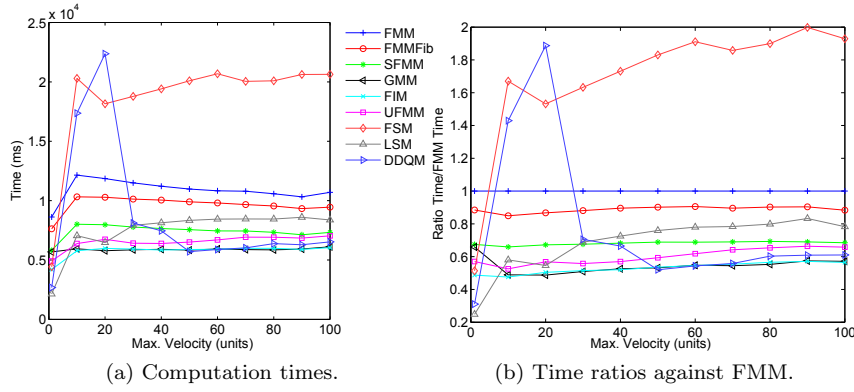


Fig. 20 Computation times for the checkerboard experiment in 2D.

The differences between this experiments and random velocities is that the environment presents a well-defined structure, and locally acts as a constant velocity environment, as in empty map experiment.

Table 3 Largest L_1 and L_∞ errors for UFMM in the checkerboard experiment.

	2D	3D	4D
L_1	$1.7 \cdot 10^{-7}$	$1.2 \cdot 10^{-9}$	$1.9 \cdot 10^{-10}$
L_∞	$2.5 \cdot 10^{-6}$	$5 \cdot 10^{-7}$	10^{-6}

7 Discussion

With the four experiments designed, the main characteristics of the Fast Methods have been shown. Any other environment can be thought as a combina-

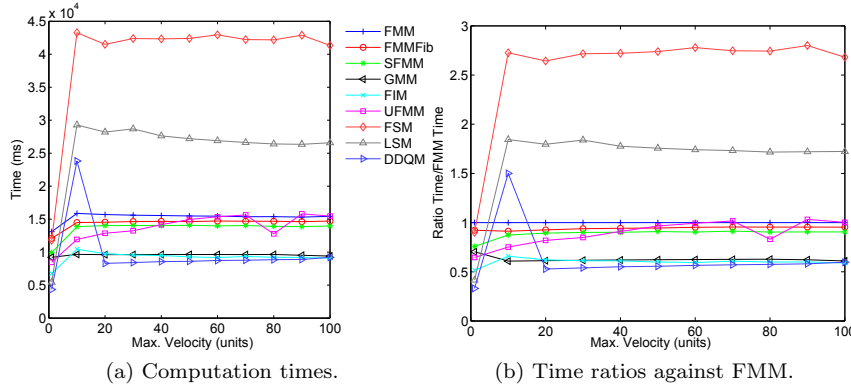


Fig. 21 Computation times for the checkerboard experiment in 4D.

tion of free space with obstacles and high-frequency or low-frequency velocity changes of different magnitude. Note also that the grid sizes covered by the experiments vary from extremely small to extremely big grids. In practice, it is hard to find applications requiring more than 16 million cells.

FIM results can be speeded up for non-constant velocity problems if larger errors are allowed. UFMM can be probably improved as well. However, our experience is that the configuration of its parameters is complex and requires a deep knowledge of the environment to be applied on.

Several conclusions can be extracted from the conducted experiments: 1) There is no practical reasons to use FMM or FMMFib as SFMM is faster in all cases with the same behaviour as its counterparts. 2) If a sweep-based method is required, LSM should be always chosen, as it greatly outperforms FSM. 3) In problems with constant velocity DDQM should be chosen, as it has shown the best performance for the empty map and the alternating barriers environments. 4) For variable velocities, but simple scenarios, GMM is the algorithm to choose. 5) UFMM is hard to tune and the result has errors. Also, it has been outperformed in most of the cases by DDQM in constant velocity scenarios, or by SFMM or FIM in experiments with variable velocities. 6) There is not a clear winner for complex scenarios with variable velocity. UFMM can perform well in all cases if tuned properly. Otherwise, SFMM is a safe choice, specially in cases where there is not too much information about the environment.

If a goal point is selected, cost-to-go heuristics could be applied [36], and thus enormously affect the results. Heuristics for FMM, FMMFib and SFMM are straightforward. They would improve the results in most of the cases. They could also be applied to UFMM. However, it is not clear if they can be applied to other Fast Methods. It is also of interest the solution of anisotropic problems [8], solved only by FMM-based methods.

8 Conclusions

Along this paper we have introduced the main Fast Methods in a common mathematical framework, focusing on a practical point of view.

With this work we aimed at closing the discussion about which method should be used in which case, as this is the first exhaustive comparison of the main Fast Methods (up to our knowledge).

The code is publicly available as well as the automatic benchmark programs. This code has been deeply tested and it can serve as a base for future algorithm design, as it provides all the tools required to easily implement and compare novel Fast Methods.

The future work focuses in 3 different aspects: develop the analogous work for parallel Fast Methods [9], study the application of these methods to anisotropic problems and also to the new Fast Marching-based solutions focused on path planning applications [18,38]. Finally, the combination of UFMM and SFMM seems straightforward and it would presumably outperform both algorithms.

Acknowledgements Authors want to gratefully acknowledge the contribution of Pablo Gely Muñoz to the GMM, FIM and UFMM implementations, and to Adam Chacon for the interesting discussions and suggestions towards improving the work.

References

1. Ahmed, S., Bak, S., McLaughlin, J., Renzi, D.: A third order accurate Fast Marching Method for the Eikonal equation in two dimensions. *SIAM Journal on Scientific Computing* **33**(5), 2402–2420 (2011)
2. Al Zaben, N., Madusanka, N., Al Shdefat, A., Choi, H.: Comparison of active contour and fast marching methods of hippocampus segmentation. In: 6th International Conference on Information and Communication Systems, pp. 106–110 (2015)
3. Bak, S., McLaughlin, J., Renzi, D.: Some improvements for the fast sweeping method. *SIAM Journal on Scientific Computing* **32**(5), 2853–2874 (2010)
4. Basu, S., Racoceanu, D.: Reconstructing neuronal morphology from microscopy stacks using fast marching. In: IEEE International Conference on Image Processing, pp. 3597–3601 (2014)
5. Bellman, R.: *Dynamic Programming*. Princeton University Press, Princeton, NJ (1957)
6. Capozzoli, A., Curcio, C., Lisenò, A., Savarese, S.: A comparison of fast marching, fast sweeping and fast iterative methods for the solution of the eikonal equation. In: 21st Telecommunications Forum, pp. 685–688 (2013)
7. Chacon, A.: *Eikonal equations: New two-scale algorithms and error analysis*. Ph.D. thesis, Cornell University (2014)
8. Clement, P., Pailhas, Y., Patron, P., Petillot, Y., Evans, J., Lane, D.: Path planning for autonomous underwater vehicles. *IEEE Transactions on Robotics* **23**(2), 331–341 (2007)
9. Detrixhe, M., Gibou, F., Min, C.: A parallel fast sweeping method for the eikonal equation. *Journal of Computational Physics* **237**, 46–55 (2013)
10. Dijkstra, E.: A note on two problems in connection with graphs. *Numerische Mathematik* **1**, 269–271 (1959)
11. Do, Q., Mita, S., Yoneda, K.: Narrow passage path planning using fast marching method and support vector machine. In: IEEE Intelligent Vehicles Symposium Proceedings, pp. 630–635 (2014)

12. Forcadel, N., Le Guyader, C., Gout, C.: Generalized fast marching method: applications to image segmentation. *Numerical Algorithms* **48**(1–3), 189–211 (2008)
13. Fredman, M., Tarjan, R.: Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the Association for Computing Machinery* **34**(3), 596–615 (1987)
14. Fu, Z., Jeong, W., Pan, Y., Whitaker, R.: A fast iterative method for solving the eikonal equation on triangulated surfaces. *SIAM Journal on Scientific Computing* **33**(5), 2468–2488 (2011)
15. Gómez, J.V.: Advanced Applications of the Fast Marching Square Planning Method. Master’s thesis, Carlos III University (2012)
16. Gremaud, P., Kuster, C.: Computational study of fast methods for the eikonal equation. *SIAM Journal on Scientific Computing* **27**(6), 1803–1816 (2006)
17. Jain, A.K., Hong, L., Pankanti, S.: Random insertion into a priority queue structure. Tech. rep., Stanford Univeristy Reports (1974)
18. Janson, L., Schmerling, E., Clark, A., Pavone, M.: Fast marching trees: a fast marching sampling-based method for optimal motion planning in many dimensions. *International Journal of Robotics Research* **Submitted** (2014)
19. Jeong, W., Whitaker, R.: A fast iterative method for eikonal equations. *SIAM Journal on Scientific Computing* **30**(5), 2512–2534 (2008)
20. Jones, M., Baerentzen, J., Sramek, M.: 3D distance fields: a survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics* **12**(7), 581–599 (2006)
21. Kim, S.: An $o(n)$ level set method for eikonal equations. *SIAM Journal on Scientific Computing* **22**(6), 2178–2193 (2001)
22. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. *Proceesings of the National Academy of Sciences* **95**(15), 8431–8435 (1998)
23. Liu, Y., Bucknall, R.: Path planning algorithm for unmanned surface vehicle formations in a practical maritime environment. *Ocean Engineering* **97**, 126–144 (2015)
24. Luo, S.: High-order factorizations and high-order schemes for point-source Eikonal equations. *SIAM Journal on Numerical Analysis* **52**(1), 23–44 (2014)
25. Osher, S., Sethian, J.A.: Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics* **79**(1), 12–49 (1988)
26. Qu, X., Liu, S., Wang, F.: A new ray tracing technique for crosshole radar traveltime tomography based on multistencils fast marching method and the steepest descend method. In: 15th International Conference on Ground Penetrating Radar, pp. 503–508 (2014)
27. Rasch, C., Satzger, T.: Remarks on the $o(n)$ implementation of the fast marching method (2008)
28. Rouy, E., Tourin, A.: A viscosity solutions approach to shape-from shading. *SIAM Journal on Numerical Analysis* **29**, 867–884 (1992)
29. Sethian, J.A.: A fast marching level set method for monotonically advancing fronts. *Proceesings of the National Academy of Sciences* **93**(4), 1591–1595 (1996)
30. Sethian, J.A.: Fast Marching Methods. *SIAM Review* **41**(2), 199–235 (1999)
31. Sethian, J.A.: Level Set Methods and Fast Marching Methods. Cambridge University Press (1999)
32. Sethian, J.A., Vladimirsky, A.: Fast Methods for the Eikonal and related Hamilton-Jacobi equations on unstructured meshes. *Proceesings of the National Academy of Sciences* **97**(11), 5699–5703 (2000)
33. Sethian, J.A., Vladimirsky, A.: Ordered upwind methods for static Hamilton-Jacobi equations: theory and algorithms. *SIAM Journal on Numerical Analysis* **41**(1), 325–363 (2003)
34. Tsai, Y., Cheng, L., Osher, S., Zhao, H.: Fast sweeping algorithms for a class of hamilton-jacobi equations. *SIAM Journal of Numerical Analysis* **41**(2), 659–672 (2003)
35. Tsitsiklis, J.N.: Fast Marching Methods. *IEEE Transactions on Automatic Control* **40**, 1528–1538 (1995)
36. Valero-Gómez, A., Gómez, J.V., Garrido, S., Moreno, L.: The Path to Efficiency: Fast Marching Method for Safer, More Efficient Mobile Robot Trajectories. *IEEE Robotics and Automation Magazine* **20**(4) (2013)

37. Yatziv, L., Bartesaghi, A., Sapiro, G.: $O(n)$ implementation of the fast marching algorithm. *Journal of Computational Physics* **212**, 393–399 (2005)
38. Yershov, D., Otte, M., Frazzoli, E.: Planning for optimal feedback control in the volume of free space. *ArXiv* (2015)
39. Zhang, X., Bording, R.: Fast marching method seismic traveltimes with reconfigurable field programmable gate arrays. *Canadian Journal of Exploration Geophysics* **36**(1), 60–68 (2011)
40. Zhao, H.: Fast sweeping method for eikonal equations. *Mathematics of Computation* **74**, 603–627 (2005)