

Partial Optimal Transport for a Constant-Volume Lagrangian Mesh with Free Boundaries

Bruno Lévy

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

June 3, 2021

Abstract

This article introduces a representation of dynamic meshes, adapted to some numerical simulations that require controlling the volume of objects with free boundaries, such as incompressible fluid simulation, some astrophysical simulations at cosmological scale, and shape/topology optimization. The algorithm decomposes the simulated object into a set of convex cells called a Laguerre diagram, parameterized by the position of N points in 3D and N additional parameters that control the volumes of the cells. These parameters are found as the (unique) solution of a convex optimization problem – semi-discrete Monge-Ampère equation – stemming from optimal transport theory. In this article, this setting is extended to objects with free boundaries and arbitrary topology, evolving in a domain of arbitrary shape, by solving a partial optimal transport problem. The resulting Lagrangian scheme makes it possible to accurately control the volume of the object, while precisely tracking interfaces, interactions, collisions, and topology changes.

1 Introduction

Some numerical simulations require to control the volume of an object while allowing it to change shape and topology. For instance, in incompressible fluid simulation (see e.g. [1]), the volume of fluid is conserved, while the shape of the fluid can considerably vary throughout the simulation, and can change topology (split and merge). In some astrophysic simulations [2, 3, 4], the Universe considered at a cosmological scale can be modeled as a "fluid"¹. With the simplifying assumption that this fluid obeys the incompressible Euler equation, one can reconstruct the full dynamics of the Universe from the knowledge of the density fluid at current time [2, 3, 4]. To name another example, in shape and

¹The "atoms" of this "fluid" correspond to galaxy clusters !

topology optimization (see e.g. [5]), one wants to find the shape of a given volume with maximum resistance (minimum compliance). Again, in this example, the considered shape can have an arbitrary topology, that can change during computation. There are several difficulties in the three applications mentioned above: (1) choosing a representation that can account for changes of topology, (2) controlling the volume, (3) tracking the interfaces and the changes of topology, and (4) tracking the interactions and the collisions between the simulated object and the boundary of the domain.

Clearly, to account for changes of topology, it is possible to represent the object as a density supported by a Eulerian grid (see for instance the *homogenization method* for shape optimization [6]) and enforce the constraint of volume conservation using Lagrange multipliers or similar techniques, then track iso-surfaces in this grid. However, we think it is interesting to experiment with an alternative Lagrangian representation that directly represents the simulated object and its boundary. The proposed alternative representation has the following properties:

- The new representation is a Lagrangian mesh that continuously depends on a set of N points in 3D space;
- each cell of the mesh has a prescribed volume for any position of the N points;
- changes of topology, interfaces, interactions and collisions are accurately tracked.

The approach builds on recent advances in numerical optimal transport, that resulted in Lagrangian schemes for fluid simulation [7] or early Universe reconstruction [4]. In the works cited above, the object fills the entire simulation domain. In the rest of this article, after summarizing the underlying semi-discrete optimal transport method (Section 2), I show how the mathematical setting can be extended it to objects with free boundaries, by solving a *partial* optimal transport problem (Section 3). Then I detail the numerical algorithm that can solve this partial optimal transport problem (Section 4). Finally, I demonstrate some applications of the method to free-surface fluid simulation (Section 5).

2 Volume control through optimal transport

In this section, I summarize the existing semi-discrete optimal transport framework with the objective of giving an intuition on the aspects that are important to control the volume of the cells in a Lagrangian mesh ². Then, in the next section, I explain how to extend this framework to simulated objects with free boundaries.

²The reader is referred to [8, 9, 10, 11, 12, 13] for an extensive and general introduction on optimal transport.

We suppose for now that the simulated object entirely fills a simulation domain Ω , that can be the 2D unit square or the 3D unit cube³. The goal is now to find a partition of Ω into a set of N cells V_i with the following properties:

- the set of cells $(V_i)_{i=1}^N$ depends on a set of parameters $(\mathbf{x}_i)_{i=1}^N$. These parameters can be interpreted as 2D (resp. 3D) points (hence a *Lagrangian* representation);
- the area (resp. volume) of each V_i 's is controlled: $|V_i| = \nu_i$ for some prescribed ν_i 's such that $\sum_{i=1}^N \nu_i = |\Omega| = 1$;
- the cells V_i *continuously* depend on the \mathbf{x}_i 's;

The Monge-Ampère equation

Such a Lagrangian mesh parameterization can be obtained through a specific (semi-discrete) version of the Monge-Ampère (MA) equation. In the general (continuous) setting, the MA equation may be thought of as seeking for an application $T : \Omega \rightarrow \Omega$ with controlled Jacobian, by solving for a potential $\psi : \Omega \rightarrow \mathbb{R}$:

$$\begin{cases} \det(\text{Hess}(\psi)(\mathbf{x})) = \nu(\mathbf{x}) & \forall \mathbf{x} \in \Omega \\ \text{s.t. } \psi^{cc} = \psi & \text{where } \psi^c(\mathbf{x}) = \inf_{\mathbf{y} \in \Omega} [\|\mathbf{x} - \mathbf{y}\|^2 - \psi(\mathbf{y})] \end{cases} \quad (1)$$

where $\nu : \Omega \rightarrow \mathbb{R}^+$ is a square-integrable density⁴. In the constraint (second line), ψ^c corresponds to the Legendre-Fenchel transform of ψ , and the condition that applying it twice to ψ does not change ψ means that ψ is convex (because the graph of ψ^{cc} corresponds to the convex hull of the graph of ψ). From the solution of Eq. 1, one can deduce an *optimal transport map* $T : \mathbf{x} \mapsto \mathbf{x} - \frac{1}{2} \nabla \psi^c(\mathbf{x})$ that minimizes a "transport cost" while satisfying mass conservation [14]:

$$\begin{cases} \inf_T [\int_{\Omega} \|T(\mathbf{x}) - \mathbf{x}\|^2 d\mathbf{x}] & \text{s. t.} \\ \int_{T^{-1}(B)} d\mathbf{x} = \int_B \nu(\mathbf{x}) d\mathbf{x} & \forall \text{ measurable set } B \subset \Omega. \end{cases} \quad (2)$$

With the viewpoint of optimal transport, ψ can be considered as the Lagrange multiplier of the volume conservation constraint (second line). The optimal transport map $T = \text{Id} - \frac{1}{2} \nabla \psi^c$ is deduced from the gradient of the Legendre transform of the solution ψ of the Monge-Ampère equation (1). The left-hand side of the Monge-Ampère equation corresponds to the Jacobian of T . Hence, the Monge-Ampère equation is a mean of controlling the Jacobian of an application. The Kantorovich dual of (a relaxed version of) the optimization problem in Eq. 2 writes:

$$\begin{cases} \sup_{\psi} [K(\psi) = \int_{\Omega} \psi^c(\mathbf{x}) d\mathbf{x} + \int_{\Omega} \psi(\mathbf{x}) \nu(\mathbf{x}) d\mathbf{x}] \\ \text{s. t. } \psi^{cc} = \psi \end{cases} \quad (3)$$

³We will consider later in this article the case where Ω is an arbitrary simplicial set.

⁴In its general form, the Monge-Ampère equation also has a density μ on the left-hand side, but in the context of this article, we consider a uniform density $\mu = 1$.




Figure 1: Left: Example of Laguerre diagram. Center and Right: unlike Voronoi cells, Laguerre cells do not necessarily contain the points they are associated with, depending on both the distribution of the points (center) and the values of the ψ_i 's (right).

The Kantorovich dual is smooth and convex. These properties are interesting because they can be used to establish the existence and uniqueness of the optimal transport map and the associated Lagrange multiplier ψ . They can be also exploited to design a numerical solution mechanism. In the next subsection, I present the so-called *semi-discrete* setting, where the target density ν is replaced with a discrete probability measure, supported by a pointset $(\mathbf{x}_i)_{i=1}^N$.

Semi-discrete Optimal Transport

The equivalence between the Kantorovich dual problem (Eq. 3) and the optimal transport problem (Eq. 2), established in Brenier's polar factorization theorem [14] characterizes mathematical objects (probability measures) that can be less regular than the densities involved in the Monge-Ampère equation (Eq. 1). For instance, the density ν can be replaced with a weighted sum of Dirac masses $\nu = \sum_{i=1}^N \nu_i \delta_{\mathbf{x}_i}$. A subset B of Ω is measured as follows by a so-defined ν :

$$\int_B \nu(\mathbf{x}) d\mathbf{x} = \sum_{j|\mathbf{x}_j \in B} \nu_j.$$

In this setting, the potential function ψ is parameterized by a vector $(\psi_i)_{i=1}^N$. The functional K of the dual problem (Eq. 3) becomes a function $K : \mathbb{R}^N \rightarrow \mathbb{R}$ that depends on the vector $(\psi_i)_{i=1}^N$:

$$\begin{cases} K(\psi) &= \int_{\Omega} \psi^c(\mathbf{x}) d\mathbf{x} & + \int_{\Omega} \psi(\mathbf{x}) \nu(\mathbf{x}) d\mathbf{x} \\ &= \int_{\Omega} \psi^c(\mathbf{x}) d\mathbf{x} & + \sum_i \psi_i \nu_i \\ &= \int_{\Omega} \inf_i [\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i] d\mathbf{x} & + \sum_i \psi_i \nu_i. \end{cases} \quad (4)$$

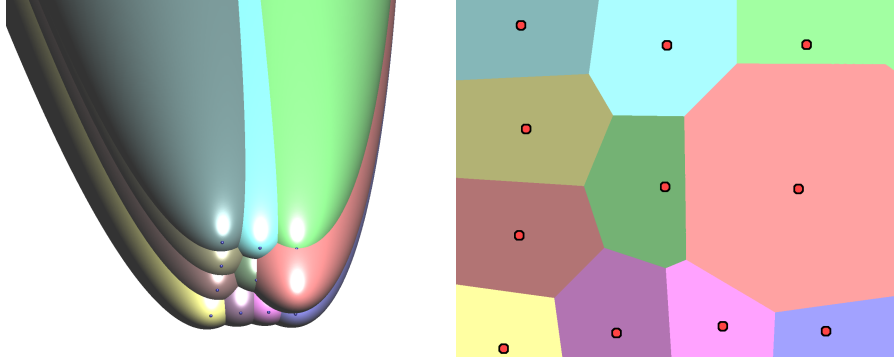


Figure 2: Laguerre diagrams considered from an alternative viewpoint. Left: the Laguerre diagram is the minimization diagram of the family of functions $f_i(\mathbf{x}) = \|\mathbf{x}, \mathbf{x}_i\|^2 - \psi_i$. Right: looking at these graphs from above, one sees a set of convex polygonal cells.

The second line is obtained by taking into account the discrete definition of ν , and the third one by replacing ψ^c by its definition (in Equation 1). The integral in the first term of $K(\cdot)$ can be rearranged:

$$K(\psi) = \sum_i \int_{Lag_i^\psi} (\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i) d\mathbf{x} + \sum_i \psi_i \nu_i \quad (5)$$

where the partition of Ω into the sets Lag_i^ψ , called a *Laguerre diagram*, is defined as follows:

Definition 1. Given the unit square Ω (resp. unit cube), a set $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ of points in Ω , a vector $\psi = (\psi_1, \dots, \psi_n) \in \mathbb{R}^N$, the Laguerre diagram is the partition of Ω into the N regions Lag_i^ψ defined by:

$$Lag_i^\psi = \{ \mathbf{x} \in \Omega \mid \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i \leq \|\mathbf{x} - \mathbf{x}_j\|^2 - \psi_j \quad \forall j \neq i \}$$

Each region Lag_i^ψ is referred to as a *Laguerre cell*. Laguerre diagram can be considered as a generalization of Voronoi diagram (they are equivalent in the specific case where all the ψ_i 's have the same value). Laguerre diagrams have been extensively studied and characterized in computational geometry [15] (in this context, if the cost corresponds to the squared distance, like in our case, Laguerre diagrams are called *power diagrams*). Two examples of Laguerre diagrams are shown in Figure 1. It is worth mentioning that depending on the ψ_i 's, a Laguerre cell *does not necessarily contain the point* it is associated with (unlike Voronoi cells). For instance, it is easy to check that one can translate the entire diagram by an arbitrary vector \mathbf{u} , without changing the \mathbf{x}_i 's, by setting $\psi_i = \sqrt{\max_j (\mathbf{u} \cdot \mathbf{x}_j) - \mathbf{u} \cdot \mathbf{x}_i}$, hence the cells can be arbitrarily far away from the \mathbf{x}_i 's.

The Laguerre diagram can also be considered as the minimization diagram of the family of functions $f_i(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i$. The graphs of these functions (paraboloids) for a 2D diagram is pictured in Figure 2-Left. The coefficient ψ_i 'shifts' the associated paraboloid along the Z axis. Seen from above (Figure 2-Right), the function graphs appear as convex polygonal cells. The common boundary between two cells corresponds to the projection of a parabola included in a plane orthogonal to the picture, hence a straight segment.

Still considering Figure 2-Left, each coefficient ψ_i corresponds to the amount of translation of the associated paraboloid along the Z axis. As one can imagine, lowering a paraboloid (that is, *increasing* ψ_i , because of the "-" sign) increases the size of the associated Laguerre cell. Conversely, raising it decreases the size of the cell⁵. One can also see that the Laguerre diagram does not change when adding a constant to all ψ_i 's (the entire diagram is shifted along the Z axis and the image does not change when viewed from above).

Now the question is: given a vector of prescribed volumes $\nu_i \geq 0$, such that $\sum_i \nu_i = |\Omega| = 1$, is it possible to find the values of ψ_i such that $|Lag_i^\psi| = \nu_i$ for all i ? In other words, is it possible to tune the heights of the paraboloids in such a way that the areas of the cells seen from above match the prescribed areas? The (positive) answer is given by the following theorem [16] (it is also a direct consequence of Brenier's more general polar factorization theorem [14] considered in the specific semi-discrete setting). We summarize below the main argument of the proof in [16]:

Theorem 1. *Given a set of n points $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ in Ω , a set of positive prescribed volumes $(\nu_1, \nu_2, \dots, \nu_n)$ such that $\sum_i \nu_i = 1$, there exists a unique (up to a translation) set of scalars $(\psi_1, \psi_2, \dots, \psi_n)$ such that each Laguerre cell Lag_i^ψ has the prescribed volume V_i .*

proof (summarized, see [16, 9] for a complete proof) Consider the Kantorovich dual function $K(\psi) : \mathbb{R}^N \rightarrow \mathbb{R}$ defined by:

$$\begin{aligned} K(\psi) &= \sum_i \int_{Lag_i^\psi} (\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i) d\mathbf{x} + \sum_i \psi_i \nu_i \\ &= \int_{\Omega} \inf_i [\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i] d\mathbf{x} + \sum_i \psi_i \nu_i \end{aligned}$$

The function K has the following properties:

1. The function K is concave;
2. the function K is differentiable up to the second order;
3. the components of the gradients of K are given by:
$$\partial K / \partial \psi_i = \nu_i - |Lag_i^\psi|.$$

⁵A cell can even completely disappear if the corresponding paraboloid is shifted above all the other graphs.

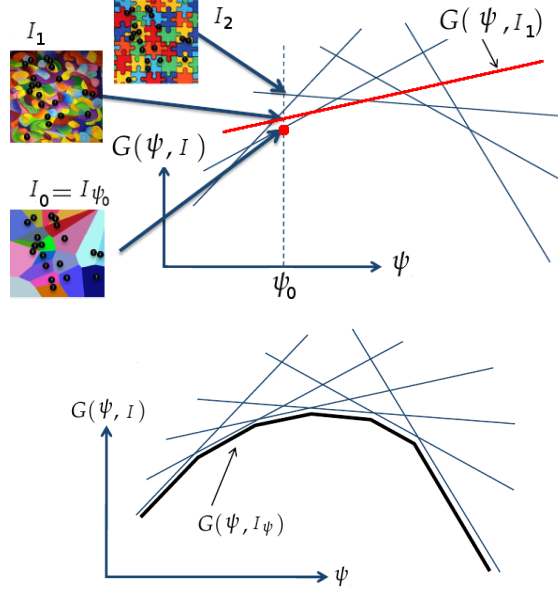


Figure 3: Convexity of K : the graph of K is the lower envelope of a family of affine functions $G(\psi, I)$, parameterized by the Lagrange multipliers ψ and an index map I , that maps each point of Ω to an index in $[1 \dots N]$ (symbolized by colors on the left).

We explain below how to obtain the concavity of K (the two other properties can be obtained by direct calculation. The second-order differentiability of K with respect to ψ is studied in [17, 18], and its differentiability with respect to \mathbf{x} in [19], as well as the expression of the second-order derivatives.

To show the concavity of K , we introduce the following functional G , parameterized by ψ and by an *index map* $I : \Omega \rightarrow [1, \dots, N]$ that associates one of the points $\mathbf{x}_{I(\mathbf{x})}$ to each point $\mathbf{x} \in \Omega$.

$$G(\psi, I) = \int_{\mathbb{R}^3} (\|\mathbf{x} - \mathbf{x}_{I(\mathbf{x})}\|^2 - \psi_{I(\mathbf{x})}) d\mathbf{x}.$$

Clearly, it is easy to verify that for a fixed index map I , G *affinely* depends on ψ . Among all the possible index maps I , let us consider now I_ψ that maps each point \mathbf{x} to the index of the Laguerre cell it belongs to:

$$I_\psi(\mathbf{x}) = i \quad \forall \mathbf{x} \in \text{Lag}_i^\psi$$

We also have (from the definition of Lag_i^ψ):

$$I_\psi(\mathbf{x}) = \arg \min_i [\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i]$$

which implies that for a given ψ_0 , among all possible index maps I , I_{ψ_0} minimizes $G(\psi_0, I)$. Thus, the graph of $G(\psi, I_\psi)$ is the lower envelope of a family of affine functions, hence $\psi \mapsto G(\psi, I_\psi)$ is a concave function (see Figure 3). Since $G(\psi, I_\psi)$ corresponds to the first term of K , and since the second term of K is linear in ψ , K is also concave. K has a unique maximizer ψ^* , where its gradient vanishes, and the expression of the derivatives indicates that the volumes of the Laguerre cells match the prescribed volumes. ■

It is easy to verify that the optimal transport map $T_{\psi^*} = \text{Id} - \frac{1}{2}\nabla\psi^c$ maps a point \mathbf{x} of Ω to the point $\mathbf{x}_{I_{\psi^*}}(\mathbf{x})$:

$$\begin{aligned} T_{\psi^*}(\mathbf{x}) &= \mathbf{x} - \frac{1}{2}\nabla(\inf_i[\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i]) \\ &= \mathbf{x} - \frac{1}{2}\nabla(\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i) \quad \text{where } i = I_{\psi^*}(\mathbf{x}) \\ &= \mathbf{x} - \frac{1}{2}(2\mathbf{x} - 2\mathbf{x}_i) \\ &= \mathbf{x}_i \end{aligned}$$

In other words, for all $\mathbf{x} \in \text{Lag}_i^{\psi^*}$, $T(\mathbf{x}) = \mathbf{x}_i$. Thus, the Laguerre cell Lag_i^{ψ} corresponds of the set of points \mathbf{x} of Ω mapped to \mathbf{x}_i through the optimal transport map. The theorem implies that the following three statements are equivalent:

1. The application T_{ψ^*} is the solution to the optimal transport problem (Eq 2);
2. the ψ_i^* 's are such that $|\text{Lag}_i^{\psi^*}| = \nu_i$ for all $1 \leq i \leq n$;
3. the vector $(\psi_i^*)_{i=1}^N$ is the unique (up to a translation) maximizer of the Kantorovich dual K .

In our context, our initial motivation (controlling the volumes of a Lagrangian mesh) is satisfied by the second statement. The optimal transport map T_{ψ^*} in the first one can be considered as a "by product" that we are not going to use directly. In the (2D) context depicted in Figure 1, the theorem means that by shifting correctly the paraboloids along the Z axis, one can make the areas of the Voronoi cells match prescribed areas. In addition, the third statement and the second-order differentiability of K [17, 19] can be exploited to design a Newton algorithm [18]. It means that from any set of points $(\mathbf{x}_i)_{i=1}^N$, it is possible to compute a partition of Ω into a set of cells $V_i = \text{Lag}_i^{\psi}$ of prescribed volumes, that is $|V_i| = \nu_i$ for a set of positive ν_i 's such that $\sum_i \nu_i = |\Omega|$. About the convexity constraint $\psi^{cc} = \psi$ in Eq. 3, in the semi-discrete setting, it is equivalent to the absence of empty Laguerre cells in the diagram ($\text{Lag}_i^{\psi} \neq \emptyset \quad \forall 1 \leq i \leq N$), as can be shown by direct computation (see [9] for details).

Before detailing the numerical solution mechanism that computes the ψ_i 's, I shall explain how the geometric setting can be extended to objects with free boundaries.

3 Free boundaries through partial optimal transport

Let us now consider a different setting, where the simulated object does not fill the volume Ω entirely. That is, the sum of the prescribed volumes ν_i is smaller than the volume of Ω . This problem is referred to as a *partial optimal transport* problem. I show how partial optimal transport can be considered as a particular instance of optimal transport. To do so, we still consider the optimal transport problem (Eq. 2) and its Kantorovich dual (Eq. 3), with the difference that this time transport is done towards a set of "objects" (pointsets) \mathcal{O}_i :

$$\begin{aligned} & \inf_I \left[\int_{\Omega} d^2(\mathbf{x}, \mathcal{O}_{I(\mathbf{x})}) d\mathbf{x} \right] \\ & \text{subject to } |I^{-1}(\mathcal{O}_i)| = \nu_i \quad \forall i \end{aligned} \tag{6}$$

where:

- the index map $I : \Omega \rightarrow [1..N]$ assigns an object \mathcal{O}_i to each point \mathbf{x} of Ω ;
- each \mathcal{O}_i is a set of points of Ω (that can contain either a single point, or an integer number of points, or be a continuous subset of Ω);
- each \mathcal{O}_i is supposed to receive a prescribed quantity of matter ν_i (constraint);
- the distance between a point \mathbf{x} and an object \mathcal{O} is defined by $d(\mathbf{x}, \mathcal{O}) = \inf_{\mathbf{y} \in \mathcal{O}} [\|\mathbf{x} - \mathbf{y}\|]$;
- $I^{-1}(\mathcal{O}_i) = \{\mathbf{x} \mid I(\mathbf{x}) = i\}$;
- Given an index map I , the corresponding map $T_I : \Omega \rightarrow \Omega$ is defined by $T_I(\mathbf{x}) = \arg \inf_{\mathbf{y} \in \mathcal{O}_{I(\mathbf{x})}} \|\mathbf{x} - \mathbf{y}\|^2$.

Consider now the dual problem:

$$\begin{aligned} & \sup_{\psi} \left[K(\psi) = \int_{\Omega} \inf_i [d^2(\mathbf{x}, \mathcal{O}_i) - \psi_i] d\mathbf{x} + \sum_i \psi_i \nu_i \right] . \\ & \text{subject to } \psi^{cc} = \psi \end{aligned}$$

It keeps the structure of our initial Kantorovich dual (Eq. 3,4), and the argument for the convexity, smoothness, existence and uniqueness of the solution still hold, but now we have a more general setting that we can use to account for free boundaries: we consider now a set of $N + 1$ objects $\mathcal{O}_{i=0}^N$:

- The first object \mathcal{O}_0 is a set of M points $\{\mathbf{y}_1, \dots, \mathbf{y}_M\}$ uniformly distributed in Ω ;
- each \mathcal{O}_i for $1 \leq i \leq N$ is the singleton $\{\mathbf{x}_i\}$;
- the object \mathcal{O}_0 is associated with $\nu_0 = |\Omega| - \sum_{i=1}^N \nu_i$, that is all the volume of Ω not affected to the points $\{\mathbf{x}_i\}$.




Figure 4: Left: approximation of partial optimal transport using a discretization of the background ("ghost cells"). Right: exact solution to partial optimal transport, where the background is considered as a continuum. Each cell V_i^ψ is the intersection between the laguerre cell Lag_i^ψ and the ball centered on \mathbf{x}_i of radius $\sqrt{\psi_i}$.

This configuration is depicted in Figure 4-Left: the cells of the fluid are displayed in blue. Each of them is a convex polygon. The background points $\mathbf{y}_{i=1}^M$ fill the entire domain Ω , including the part occupied by the fluid. All the \mathbf{y}_i points share the same Lagrange multiplier ψ_0 . This may be interpreted as the fact that it takes no cost for background cells to exchange matter (air). For a part occupied by the fluid, each cell Lag_i^ψ that correspond to the fluid has a Lagrange multiplier ψ_i larger than ψ_0 , then the fluid cell "shadows" the background cells. On the boundary of the fluid, the presence of the background cells limit the size of the fluid cells. The external boundary of the fluid is formed by straight segment shared by a fluid cell and a background cell (closeup in Figure 4-Left).

A similar technique was used in [20], where it is proposed to insert points \mathbf{y}_i in the vicinity of the boundary of the fluid (in the cited article, the corresponding cells are called "ghost cells"). The main advantage it that it allows directly reusing a discrete optimal transport implementation to solve transport problems with free boundaries. However, the "ghost cells" technique has two drawbacks:

- first, it is difficult to know *in advance* where to insert the ghost cells to make sure the boundary of the fluid is accurately represented, that is, there is no easy way of predetermining which ghost cells influence the fluid;
- second, the ghost cells have a significant computational cost: computing a Laguerre diagram costs $O(N \sqrt[d]{N})$ (where $d \in \{2, 3\}$ is the dimension), and the Newton algorithm converges in $O(N \log(N))$ iterations (empirical results in [4]). The "ghost cells" techniques computes a diagram with $N + M$ vertices (instead of N), where M needs to be sufficiently large to accurately capture the free boundary.




Figure 5: Left: The minimization diagram of the family of functions $f_i(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i$ augmented with the graph of the function $f(\mathbf{x}) = 0$ (horizontal plane). Right: seen from below, the obtained cells V_i^ψ are the intersections between the Laguerre cells Lag_i^ψ and the disks of radii $\sqrt{\psi_i}$ centered on the \mathbf{x}_i 's (or the empty set if ψ_i is negative).

To overcome these limitations, I propose an alternative technique: let us now imagine that the number of background points tends to infinity. More simply put, the background object \mathcal{O}_0 becomes the entire Ω domain (and it is still associated with $\nu_0 = |\Omega| - \sum_{i=1}^N \nu_i$). Let us also remember that the vector $\psi_{i=0}^N$ is independent on a translation, thus, w.l.o.g. we can choose $\psi_0 = 0$. Let us now consider a cell V_i^ψ associated with a fluid particle:

$$\begin{aligned}
 V_i^\psi &= \{ \mathbf{x} \mid d^2(\mathbf{x}, \mathcal{O}_i) - \psi_i < d^2(\mathbf{x}, \mathcal{O}_j) - \psi_j \quad \forall j \neq i \quad \} \\
 &= \{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i < \|\mathbf{x} - \mathbf{x}_j\|^2 - \psi_j \quad \forall 1 \leq j \neq i \leq N \\
 &\quad \text{and} \quad \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i < d^2(\mathbf{x}, \mathcal{O}_0) - \psi_0 \quad \} \\
 &= \{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i < \|\mathbf{x} - \mathbf{x}_j\|^2 - \psi_j \quad \forall 1 \leq j \neq i \leq N \\
 &\quad \text{and} \quad \|\mathbf{x} - \mathbf{x}_i\|^2 < \psi_i \quad \}.
 \end{aligned}$$

The last line is obtained by remembering that $\psi_0 = 0$, and by noticing that $d^2(\mathbf{x}, \mathcal{O}_0) = 0$ for all \mathbf{x} in Ω since $\mathcal{O}_0 = \Omega$. In other words, V_i^ψ is the intersection between the Laguerre cell Lag_i^ψ and a disk of radius $\sqrt{\psi_i}$ centered on \mathbf{x}_i , as shown in Figure 4-Right⁶. From this observation, not only "ghost cells" are no longer needed, but also we can much more accurately compute the boundary of the fluid, by computing the intersection between the Laguerre cells and a set of disks. Put differently, what we compute is the limit case where the number of ghost cells M tends to infinity. Not only the result will be more precise, but the overall computational cost will be significantly reduced, since the number of vertices in the Laguerre diagram is not increased.

⁶or the empty set if ψ_i is negative. Note that a negative ψ_i cannot happen in our context, since it would contradict the convexity of ψ ($\psi^{cc} = \psi \Leftrightarrow V_i^\psi \neq \emptyset \forall 1 \leq i \leq N$), guaranteed by all the iterations of the KMT Newton algorithm [18] (more on this below).

As shown in Figure 5, to gain more intuition about this setting, it is also possible to take the "minimization diagram" point of view: the distance to $\mathcal{O}_0 = \Omega$ is equal to zero on Ω , and $\psi_0 = 0$, thus the graph of the associated function f_0 is the horizontal plane $Z = 0$. The other functions f_i are shifted paraboloids just like before. Now, looking at the diagram from above, one will see the tip of the paraboloids intersected by the horizontal planes (disks). The paraboloids can also intersect each other, forming polygonal cells with straight edges (that correspond to projected parabolas), just like in the previous configuration.

4 Numerical Solution Mechanism

I shall now explain how to design a numerical solution mechanism. The associated algorithm takes the following inputs and produces the following outputs:

- Input:**
- The domain Ω (can be $[0, 1]^d$ or a simplicial mesh)
 - a set of N points $\mathbf{x}_i \in \Omega$
 - prescribed volumes $(\nu_i)_{i=1}^N$ such that $\sum_i \nu_i \leq |\Omega|$
- Output:**
- the (unique) vector $\psi^* \in \mathbb{R}^N$ that maximizes $K(\cdot)$
 - the cells $(V_i^{\psi^*})_{i=1}^N$ such that $|V_i^{\psi^*}| = \nu_i$

The core of the algorithm solves the following optimization problem:

$$\sup_{\psi} \left[K(\psi) = \sum_{i=1}^N \int_{V_i^{\psi}} (\|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i) d\mathbf{x} + \sum_{i=1}^N \nu_i \psi_i \right] \quad (7)$$

where $V_i^{\psi} = \text{Lag}_i^{\psi} \cap \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_i\|^2 \leq \psi_i\}$

subject to $V_i^{\psi} \neq \emptyset \quad \forall 1 \leq i \leq N$

The KMT algorithm due to Kitagawa, Merigot and Thibert [18] is a Newton algorithm that produces a series of iterates $\psi^{(k)}$ that provably converges to the solution of Eq. 7. Each iterate satisfies the convexity constraint $\psi^{cc} = \psi$ (that is, $V_i^{\psi} \neq \emptyset \quad \forall k, \forall 1 \leq i \leq N$, see also [9]). Adapting the KMT algorithm to our specific context has numerical and geometrical aspects detailed below.

4.1 Numerical aspects

The KMT algorithm follows the classical structure of a Newton algorithm:

- (1) : $\psi \leftarrow [0 \dots 0]$
- (2) : Loop
- (3) : Compute the cells $(V_i)_{i=1}^N$
- (4) : Compute the gradient $\nabla K(\psi)$
- (5) : If $\|\nabla K(\psi)\|_\infty < \epsilon_K$ then Exit loop
- (6) : Compute the Hessian matrix $\nabla^2 K(\psi)$
- (7) : Solve for $\mathbf{p} \in \mathbb{R}^n$ in $\nabla^2 K(\psi)\mathbf{p} = -\nabla K(\psi)$
- (8) : Find the descent parameter α
- (9) : $\psi \leftarrow \psi + \alpha\mathbf{p}$
- (10) : End loop

The algorithm above needs to evaluate the gradient and Hessian matrix of $K(\cdot)$ at each main loop iteration. The coefficients of the gradient and Hessian matrix can be deduced from the cells $(V_i^\psi)_{i=1}^N$ that are computed at step (3). The algorithm that computes the cells is detailed later in the next subsection on the geometric aspects. The components $\partial K / \partial \psi_i$ of the gradient $\nabla K(\cdot)$ are given by the following expression [18, 9]:

$$\frac{\partial K}{\partial \psi_i} = \nu_i - |V_i^\psi|. \quad (8)$$

In other words, each component of the gradient corresponds to the prescribed volume ν_i associated with a point \mathbf{x}_i minus the volume of the cell V_i^ψ . For the vector ψ^* that maximizes $K(\cdot)$, all components of the gradient vanish, which means that each cell $V_i^{\psi^*}$ has exactly the prescribed volume ν_i .

This expression of the gradient leads also to a natural stopping criterion (line 5), the largest component of the gradient corresponds to the maximum volume error. We stop the algorithm as soon as it is smaller than a prescribed ϵ_K (typically one percent of ν_i).

We now consider the Hessian matrix computed at step (6). Adapting the formulas in [18, 9] to our context, the coefficients of the Hessian are given by:

$$\begin{aligned} \frac{\partial^2 K}{\partial \psi_i \partial \psi_j} &= \frac{1}{2} \frac{|V_{ij}^\psi|}{\|\mathbf{x}_j - \mathbf{x}_i\|} \quad \text{if } i \neq j \\ \frac{\partial^2 K}{\partial \psi_i^2} &= - \left(\sum_{j \neq i} \frac{\partial^2 K}{\partial \psi_i \partial \psi_j} \right) - \frac{1}{2} \frac{|V_{i0}^\psi|}{\sqrt{\psi_i}} \end{aligned} \quad (9)$$

where $|V_{ij}^\psi| = |\partial V_i^\psi \cap \partial V_j^\psi|$ denotes the area of the intersection between the border of the cell V_i^ψ and the border of the cell V_j^ψ , that is, a polygonal facet in the Laguerre diagram clipped by the two balls $B(\mathbf{x}_i, \sqrt{\psi_i})$ and $B(\mathbf{x}_j, \sqrt{\psi_j})$. The

term $|V_{i0}^\psi|$ corresponds to the free surface area in ∂V_i^ψ , that is the portion of the border of V_i^ψ included in the sphere $S(\mathbf{x}_i, \sqrt{\psi_i})$, that touches the unoccupied portion of Ω .

Step (7) of the algorithm computes the Newton step vector \mathbf{p} , by solving a linear system. Except the boundary term in $\partial^2 K / \partial \psi_i^2$, this linear system is identical as the one solved in [4]: it corresponds to a Poisson equation discretized with finite elements. The matrix is sparse, with a non-zero coefficient at coefficient (i, j) if and only if the cells V_i^ψ and V_j^ψ touch each other. The additional term $(1/2) \left(|V_{i0}^\psi| / \sqrt{\psi_i} \right)$ in $\partial^2 K / \partial \psi_i^2$ does not change the sparsity pattern of the Hessian: since the Lagrange multiplier associated with ψ_0 is fixed and equal to 0, there is no partial derivative with respect to it except the diagonal term. The same linear solver as in [4] can be used (Jacobi-preconditioned conjugate gradient [21] with sparse matrix stored in CRS format, with optional GPU acceleration), as well as the same stopping criterion $\|H\mathbf{p} - \mathbf{g}\| / \|\mathbf{g}\| \leq 10^{-3}$.

Once the step vector \mathbf{p} is computed, we need to find a good descent parameter α (step (8)). In the KMT algorithm [18], provably convergent, the descent parameter α is determined as follows:

- (1): $\alpha \leftarrow 1$
- (2): Loop
- (3): If $\inf_i |V_i^{\psi + \alpha \mathbf{p}}| > a_0$
- (4): and $\|\nabla K(\psi + \alpha \mathbf{p})\| \leq (1 - \alpha/2) \|\nabla K(\psi)\|$
- (5): then Exit loop
- (6): $\alpha \leftarrow \alpha/2$
- (7): Compute the cells $(V_i^{\psi + \alpha \mathbf{p}})_{i=1}^N$
- (8): End loop

where $a_0 = \frac{1}{2} \min \left(\inf_i |V_i^{\psi(0)}|, \inf_i (\nu_i) \right)$.

The KMT algorithm iteratively halves the descent parameter α until two criteria are met: the volume of the smallest cell needs to be larger than a threshold a_0 (line 3), and the norm of the gradient needs to decrease sufficiently (line 4). The threshold a_0 for the minimum cell volume corresponds to (half) the minimum cell volume for $\psi = 0$ (also called Voronoi diagram) and minimum prescribed area ν_i .

Equipped with the KMT algorithm above, we can now compute the descent parameter α , by plugging the algorithm above into line (8) of the Newton algorithm at the beginning of this section.

4.2 Geometrical aspects

To evaluate the components of the gradient and coefficients of the Hessian of $K(\cdot)$, the KMT algorithm needs to construct the cells $V_i^\psi = \text{Lag}_i^\psi \cap B(\mathbf{x}_i, \sqrt{\psi_i})$,




Figure 6: The Bowyer-Watson algorithm computes the Laguerre diagram (cell boundaries displayed in white and grey). It uses internally a dual representation (triangulation in black). One also needs to compute the intersection with Ω (blue square in this example).

that is, the intersection between Laguerre cells and balls centered on the \mathbf{x}_i 's for each $\psi^{(k)}$ iterate. To do so, we first compute the Laguerre diagram, using the classical Bowyer-Watson algorithm [22, 23], see Figure 6. A readily-available open-source implementation is available in our GEOGRAM library [24] (and also in the CGAL library [25] for programmers who have a taste for C++ templates). The details of our highly optimized implementation in GEOGRAM are given in [4].

Given the points \mathbf{x}_i and the vector ψ , the Bowyer-Watson algorithm computes the Laguerre diagram. Internally, it uses the dual triangulation, formed by the black triangles in Figure 6 (in 3D, they become tetrahedra). Each vertex of the Laguerre diagram is deduced from one of the triangles (resp. tetrahedra), and the Laguerre cells are obtained by traversing the triangles (tetrahedra) incident to a given \mathbf{x}_i . Note that the Bowyer-Watson algorithm computes the Laguerre diagram in the entire \mathbb{R}^d domain, with infinite Laguerre cells (grey straight lines in the Figure). To obtain the V_i cells, one needs to compute intersections with the domain Ω (and also with the balls centered on the \mathbf{x}_i 's):

$$V_i^\psi = \text{Lag}_i^\psi \cap \Omega \cap B(\mathbf{x}_i, \sqrt{\psi_i})$$

In 2D, constructing the cells V_i^ψ means computing intersections between convex polygons and disks, which does not represent too much difficulty (see Figure 4). However, it becomes a challenging task in 3D: as can be seen, we need to compute a large number of intersections, and 3D mesh intersection is known to be a delicate operation, subject to numerical precision problems: the intersection algorithm depend on a small set of functions, called *geometric predicates*, that take as input some points, and that returns a discrete value. For instance, such a geometric predicate can tell whether a point \mathbf{p}_1 is above or below another point \mathbf{p}_2 . These geometric predicates are often polynomials in the coordinates of the points. The limited precision of floating point numbers can have catastrophic consequences, for instance when the algorithm estimates that a certain point \mathbf{p}_1 is above another point \mathbf{p}_2 , and later inconsistently estimates that \mathbf{p}_2 is above \mathbf{p}_1 . Such inconsistent behavior can result in invalid combinatorics (see e.g. [26]).

It is important to stress that we are going to compute a *huge* number of intersections (billions), hence these numerical issues that correspond to what may me initially thought of as very unlikely corner cases occur in fact *thousands times* in a typical simulation. The goal now is to find a strategy to overcome this robustness issue while keeping reasonable performance. The strategy is to express the problem in terms of a small number of "atomic" operations, that can be used as solid foundations to build the rest of the algorithm. These "atomic" operations are simple and can be robustly implemented:

- (1) the algorithm will solely manipulate *convex polytopes* and their intersections. A convex polytope C is defined as the intersection of N_v

half-spaces Π_i^+ :

$$C = \bigcap_{i=1}^{N_v} \Pi_i^+ = \bigcap_{i=1}^{N_v} \{\mathbf{x} \mid a_i x + b_i y + c_i z + d_i \leq 0\}$$

- (2) to avoid inconsistencies, all the used predicates will be computed in arbitrary precision [26], with arithmetic filters to speed-up computations in the easy cases [27].

A consequence of (1) is that the balls $B(\mathbf{x}_i, \sqrt{\psi_i})$ will be approximated by convex polytopes:

$$B(\mathbf{x}, R) \simeq \hat{B}(\mathbf{x}, R) = \bigcap_{i=1}^{N_u} \{\mathbf{y} \mid \|\mathbf{x} - \mathbf{y}\|^2 \leq \|\mathbf{x} + 2\mathbf{u}_i - \mathbf{y}\|^2\}$$

where $(\mathbf{u}_i)_{i=1}^{N_u}$ is a predetermined set of unit vectors \mathbf{u} that uniformly sample the unit sphere. While this strategy shares with the "ghost cells" approach the fact that it approximates the balls as convex polytopes, it differs from it in two important aspects:

- first, computations are *local* to a given V_i^ψ cell, with a small number of vertices (typically 50 to 200, as compared to the total number of points N), whereas the "ghost cells" strategy needs inserting all additional vertices into the *global* Laguerre diagram, with $N + M$ vertices;
- second, each $\hat{B}(\mathbf{x}, R)$ can be instanced from a *precomputed* unit sphere template $\hat{B}(\mathbf{0}, 1)$ without needing to compute any intersection, by scaling and translating the unit sphere template (no combinatorics needs to be computed).

This means a large number N_u (typically 100 to 200) of unit directions \mathbf{u}_i can be used without significantly increasing the computation time: the "partial optimal transport" algorithm computes a solution that is equivalent to using $N_u \times N$ "ghost cells", without the associated computational cost. Not only it significantly reduces the complexity, but also it makes it possible to use a large value for N_u , that gives a precise approximation of the balls $B(\mathbf{x}_i, \sqrt{\psi_i})$. Alternatively, one can also use a data structure for convex polytopes with spherical faces [28], but this comes with longer execution times, and robustness is more difficult to ensure.

The domain Ω can be either a convex polytope, or a simplicial mesh $\Omega = \bigcup_{i=1}^{N_T} T_i$. In both cases the intersections between Ω , Lag_i^ψ and $\hat{B}(\mathbf{x}_i, \sqrt{\psi_i})$ involved in the definition of V_i^ψ can be computed as intersections between convex polytopes.

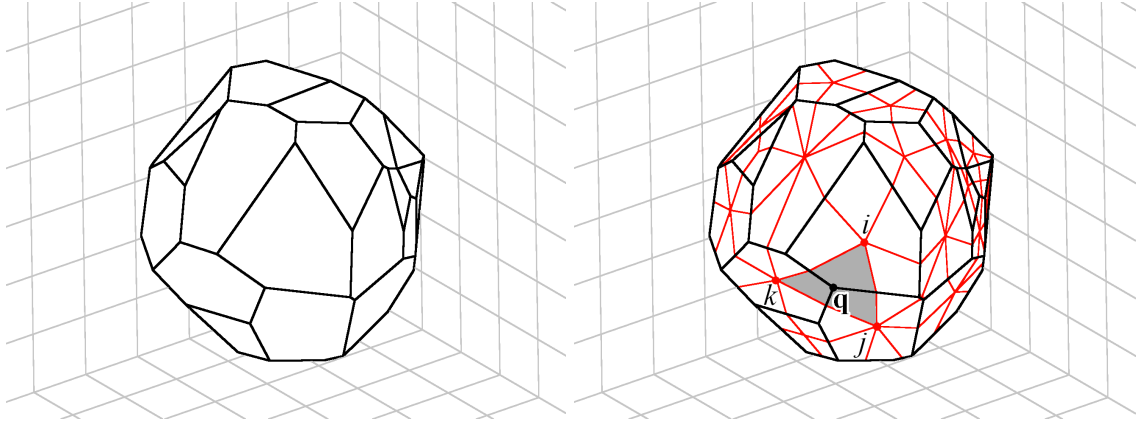


Figure 7: Representation of convex polytopes.

4.2.1 Convex polytopes: H and V representation

The "atomic" operation computes the intersection between a convex polytope C and a half-space Π^+ :

$$C \leftarrow C \cap \Pi^+.$$

In the end, to estimate the gradient and Hessian of $K(\cdot)$, we will need to compute the volume of the cells and the areas of the cell borders. To do so, we need to convert from the initial representation of the cell, that is, as an intersection of half-spaces" (H-representation), $(\Pi_i^+)_{i=1}^{N_v}$ into an explicit representation of all the vertices and facets of the cell (V-representation). As shown in Figure 7, if the planes are in generic position⁷, each vertex (black) is shared by three planes Π_i, Π_j, Π_k . It is then natural to represent the polytope in dual form (see Figure 7), by a triangulation $(T_l = \{i_l, j_l, k_l\})_{l=1}^{N_t}$ shown in red in the Figure (the same representation is exploited by the Boywer-Watson algorithm that we use to compute the Laguerre diagram). The coordinates q_x, q_y, q_z of the vertex \mathbf{q}_l associated with triangle T_l are the solution of:

$$\begin{bmatrix} a_i & b_i & c_i \\ a_j & b_j & c_j \\ a_k & b_k & c_k \end{bmatrix} \begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix} = - \begin{bmatrix} d_i \\ d_j \\ d_k \end{bmatrix}$$

Using Cramer's formula, the coordinates q_x, q_y, q_z are given by:

$$\begin{bmatrix} q_x \\ q_y \\ q_z \end{bmatrix} = - \frac{1}{\begin{vmatrix} a_i & b_i & c_i \\ a_j & b_j & c_j \\ a_k & b_k & c_k \end{vmatrix}} \begin{bmatrix} \begin{vmatrix} d_i & b_i & c_i \\ d_j & b_j & c_j \\ d_k & b_k & c_k \end{vmatrix} \\ \begin{vmatrix} a_i & d_i & c_i \\ a_j & d_j & c_j \\ a_k & d_k & c_k \end{vmatrix} \\ \begin{vmatrix} a_i & b_i & d_i \\ a_j & b_j & d_j \\ a_k & b_k & d_k \end{vmatrix} \end{bmatrix} \quad (10)$$

⁷we will discuss degenerate cases further on.




Figure 8: Clipping a convex polytope by a half-space.

(the i -th component of the solution corresponds to the determinant of the system with its i -th column replaced with the r.h.s divided by the determinant of the system, see [15, 29] for similar computations).

The algorithm to compute the intersection between a cell C and a half-space $\Pi^+(a, b, c, d)$ works in three phases (see Figure 8):

- Input:** – A convex polytope C and a half-space Π^+
Result: – replaces C with $C \cap \Pi^+$
(1) : classify the vertices
(2) : discard the triangles
(3) : triangulate the hole

The first phase determines the triangles that correspond to vertices \mathbf{q}_l that are on the negative side of Π^+ (shown in black in Figure 8-Left), that is, $aq_x + bq_y + cq_z + d < 0$. This can be rewritten as:

$$\begin{vmatrix} a_i & b_i & c_i \\ a_j & b_j & c_j \\ a_k & b_k & c_k \end{vmatrix} \times \begin{vmatrix} a_i & b_i & c_i & d_i \\ a_j & b_j & c_j & d_j \\ a_k & b_k & c_k & d_k \\ a & b & c & d \end{vmatrix} \leq 0$$

by replacing the coordinates of \mathbf{q} by their expression (one obtains an expression of the determinant above, developed w.r.t. to its last row). Note that the vertices of the triangles can be consistently oriented in such a way that the sign of the first term (the 3×3 determinant) is always positive. In the end, to classify a vertex, one only needs to compute the sign of the 4×4 determinant $|\Pi_i, \Pi_j, \Pi_k, \Pi|$ formed by the coefficients of the planes. This is the only predicate⁸ used by the algorithm. Note also that during the computation of the intersections, only the equations of the Π_i planes are used, the \mathbf{q} points do

⁸besides `orient3d` and `in_weighted_sphere` used by the Bowyer-Watson algorithm that constructs the Laguerre diagram.

not need to be explicitly computed. To compute the exact sign of these determinants, I use expansion arithmetics [26], and arithmetic filters [27] to quickly determine the sign in the easy cases. The implementation in [30] generates the code that computes the predicate from its formula. This strategy ensures that all combinatoric decisions taken by the algorithm will remain consistent. The points \mathbf{q} are computed once all clipping operations are computed, right before evaluating the volumes and areas involved in the gradient and Hessian of $K(\cdot)$.

The second phase discards the triangles (i, j, k) such that $|\Pi_i, \Pi_j, \Pi_k, \Pi| < 0$ (shown in grey in Figure 8-Center). This leaves a hole in the triangulation (the border of the hole is shown in blue). As often done in implementations of the Bowyer-Watson algorithm, the discarded triangles are kept in a linked list, so that they can be reused in the subsequent steps;

The third phase constructs a new triangle for each edge on the border of the hole created at the previous phase (see Figure 8-Right).

The algorithm has the same structure as the Bowyer-Watson algorithm, with the difference that the `in_sphere` predicate, that determines which triangle to discard, is replaced with `sign` $|\Pi_i, \Pi_j, \Pi_k, \Pi|$. Another difference is the way the Bowyer-Watson determines the list of triangles to be discarded. In our case we test all the triangles. We could instead adapt the Bowyer-Watson strategy, that starts from a random vertex, and walks along the edges of the dual graph until it finds a triangle to be discarded, but since the number of triangles per polytope remains small in our case (typically < 100), we did not observe a speedup using this technique.

4.2.2 Computing the cells V_i^ψ

I shall now explain how to use the algorithm of the previous paragraph to compute the cells. Each cell $V_i^\psi = Lag_i^\psi \cap \Omega \cap \hat{B}(\mathbf{x}_i, \sqrt{\psi_i})$ is the intersection between three objects: the Laguerre cell Lag_i^ψ , the domain Ω , represented by a simplicial mesh, that is, a set of tetrahedra, and the (approximated) ball $\hat{B}(\mathbf{x}_i, \sqrt{\psi_i})$. Note that the domain Ω is not necessarily convex, but since it is decomposed into simplices, we only need to compute intersections between convex objects. The algorithm that computes a cell V_i^ψ is detailed below:

- Input:**
- a simplicial mesh Ω
 - the set of points $(\mathbf{x}_i)_{i=1}^N$ and the vector $(\psi_i)_{i=1}^N$
 - the index i of a cell
- Output:**
- A set of convex polytopes $\mathbf{C} = \{C_k\}$ such that $V_i^\psi = \bigcup C_k$

- (1) : $\mathcal{T} \leftarrow \{t \subset \Omega \mid B(\mathbf{x}_i, \sqrt{\psi_i}) \cap t \neq \emptyset\}$
- (2) : If $\exists t \in \mathcal{T} \mid \partial t \cap \partial\Omega \neq \emptyset$
- (3) : For each $t \in \mathcal{T}$
- (4) : $C \leftarrow t$
- (5) : $C \leftarrow C \cap \text{Lag}_i^\psi$
- (6) : $C \leftarrow \hat{B}(\mathbf{x}_i, \sqrt{\psi_i}) \cap C$
- (7) : $\mathbf{C} \leftarrow \mathbf{C} \cup \{C\}$
- (8) : End for
- (9) : Else
- (10) : $C \leftarrow \text{Lag}_i^\psi$
- (11) : $C \leftarrow \hat{B}(\mathbf{x}_i, \sqrt{\psi_i}) \cap C$
- (12) : $\mathbf{C} \leftarrow \mathbf{C} \cup \{C\}$
- (13) : End if

In line (1), we first determine the set of simplices \mathcal{T} that have a non-empty intersection with the ball $B(\mathbf{x}_i, \sqrt{\psi_i})$. It is done using an axis-aligned bounding box tree (AABB tree), see e.g [31].

Then we distinguish two different cases: **first case:** if one of the intersected tetrahedra touches the boundary $\partial\Omega$ of the domain, then V_i^ψ is not necessarily convex, but can be easily decomposed into convex objects: to do so, we compute the intersections $\text{Lag}_i^\psi \cap \hat{B}(\mathbf{x}_i, \sqrt{\psi_i}) \cap t$ for each t in \mathcal{T} . **Second case:** the ball $B(\mathbf{x}_i, \sqrt{\psi_i})$ is included in Ω (no tetrahedron in \mathcal{T} touches $\partial\Omega$), then $V_i^\psi = \text{Lag}_i^\psi \cap B(\mathbf{x}_i, \sqrt{\psi_i})$ is convex and can be directly computed (lines (9) to (13)).

In lines (6) and (11), we compute the intersection between a convex polytope and the approximated ball $\hat{B}(\mathbf{x}_i, \sqrt{\psi_i})$. If C 's bounding sphere centered on \mathbf{x}_i has a radius smaller than $\sqrt{\psi_i}$ (that is, $\max_{\mathbf{q} \in C} \{\|\mathbf{x}_i - \mathbf{q}\| < \sqrt{\psi_i}\}$), then C is entirely contained in the ball and one can skip the intersection computation. Else, the approximated ball $\hat{B}(\mathbf{x}_i, \sqrt{\psi_i})$ is instanced (copied from the unit ball, translated and scaled), and then clipped by all the Π_i 's from the H-representation of C (in this order, because $\hat{B}(\mathbf{x}_i, \sqrt{\psi_i})$ has a larger number of facets than C in general).

In line (4), the convex polytope C is initialized from a tetrahedron of Ω . Combinatorics are initialized from a fixed tetrahedron template, and the plane equation of each facet $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$ is given by $ax + by + cz + d = 0$, where $(a, b, c) = (\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$ and $d = -(a, b, c) \cdot \mathbf{p}_1$.

In line (10), C is initialized from a Laguerre cell. Combinatorics are copied from the Laguerre diagram, also stored in dual representation (tetrahedra). The triangles of C correspond to the facets of the tetrahedra incident to vertex i opposite to i . The plane equations $\Pi_{i,j}$ are given by: $\mathbf{x} \in \Pi_{i,j} \Leftrightarrow \|\mathbf{x} - \mathbf{x}_i\|^2 - \psi_i = \|\mathbf{x} - \mathbf{x}_j\|^2 - \psi_j$ or: $ax + by + cz + d = 0$ where $(a, b, c) = 2(\mathbf{x}_i - \mathbf{x}_j)$ and $d = \mathbf{x}_j^2 - \mathbf{x}_i^2 - \psi_j + \psi_i$.

Up to now we have supposed that all planes were in generic position (that is, each vertex of the polytope is shared by exactly three planes). To handle the specific configurations where more than three planes meet at a single vertex,

one would think of using the symbolic perturbations, as in the "simulation of simplicity" technique [29]. In our case, the symbolic perturbation is trivial: one can use the predicate $|\Pi_i, \Pi_j, \Pi_k, \Pi| \leq 0$ as is (provided that it is implemented with exact arithmetics, see [30] and references herein). It means that each time a vertex \mathbf{q} is exactly located on the boundary of a clipping half-space Π^+ the associated triangle will be discarded. To evaluate the efficiency of this technique, degenerate configurations are tested in the next section.

It is not necessary to store the set of convex polytopes \mathbf{C} : one can process them in a "streaming" manner, when assembling the gradient and Hessian of $K(\cdot)$. Then, the algorithm stores a single C at any time, and has only a small memory overhead besides the storage taken by the Laguerre diagram (stored as a 3D triangulation). Moreover, the algorithm is very simple to parallelize (then it stores one convex polytope per thread). When assembling the gradient and Hessian, it just needs a synchronization primitive (spinlock) to ensure that two threads do not modify the same entry at the same time.

5 Numerical experiments

I shall now test the different components of the algorithm, first the lowest-level one (convex polytopes clipping), then the partial optimal transport, and finally demonstrate the algorithm used to implement a basic free-surface fluid simulator in 3D.

5.1 Testing convex polytope clipping

The robustness of the convex polytope clipping algorithm is tested using some highly degenerate configurations that are created on purpose. In Figure 9-Left, the algorithm computes the intersection between 10000 (top) and 30000 (bottom) half-spaces tangent to a sphere with random normal vector. The same sequence of half-spaces is fed to the algorithm 5 times, which creates degenerate configurations with half-spaces that exactly correspond to existing facets. On the right, a cube is clipped by 30 (top) and 500 (bottom) random half-spaces tangent to a cone. This creates a highly degenerate vertex, common to all facets. Again, the sequence of half-spaces is fed to the algorithm 5 times. In both configurations, the implemented exact predicates robustly handles all the degeneracies. Execution time is smaller than 1 second for both tests.

5.2 Testing partial optimal transport

We now test partial optimal transport. In the first test, shown in Figure 10, we compute partial optimal transport between a cube and a set of 100 points randomly distributed in the lower half of a cube. The cells are constrained to occupy 10%, 50% and 90% of the cube respectively. The balls are approximated by $N_u = 162$ half-spaces.

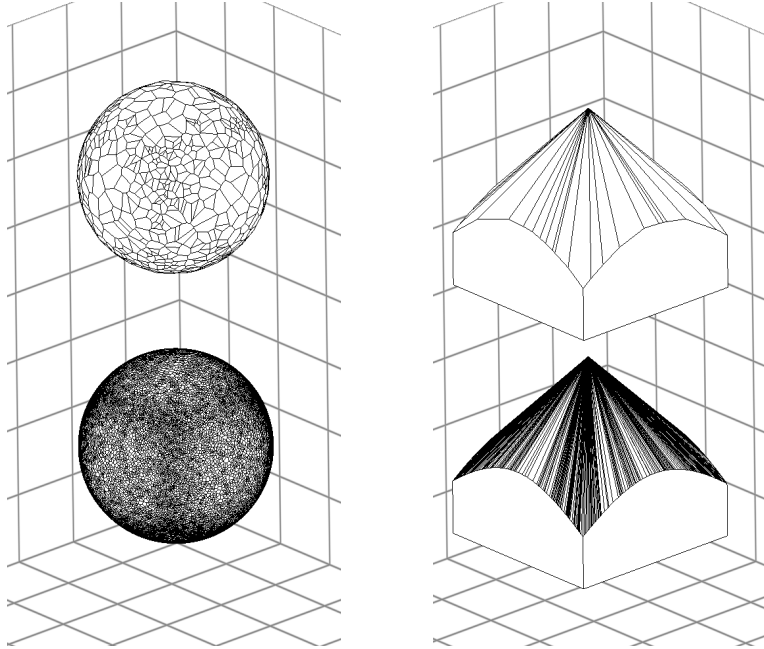


Figure 9: Testing the robustness of the convex polytope clipping algorithm.




Figure 10: Testing partial optimal transport in a cube, with 100 random points in the lower half of the cube. From left to right: fluid volume = 10%, 50% and 90%.




Figure 11: Testing partial optimal transport in a tet mesh, with 100 random points in the lower half of the domain. Left: fluid volume = 50%, Right: 75%.

fraction	nb iter	max err	avg err	time
10%	4	0.19%	0.006%	0.28s
50%	5	0.54%	0.1%	0.39s
90%	7	0.008%	0.001%	0.58s

The table above reports the number of Newton iteration (second column) spent to reach convergence, that is, when the largest difference between requested and actual cell volume among all cells (third column) is smaller than $\epsilon_K = 1\%$. The average volume difference is also reported (fourth column), as well as total execution time (fourth column). As can be seen, when the fraction of the volume occupied by the cells is larger, the algorithm takes a larger number of iterations to converge. This is because the interactions between the cells introduce mode dependencies in $K(\cdot)$.

We now test partial optimal transport in a domain Ω of arbitrary shape, here a genus 3 torus represented by a tessellated mesh (see Figure 11). The tetrahedral mesh has 56K elements. The points are distributed in the bottom half of Ω . Statistics and timings are reported in the table below.

fraction	nb iter	max err	avg err	time
10%	3	0.0032%	0.0031%	1.9s
50%	4	0.9%	0.2%	4s
75%	5	0.45%	0.12%	4s
90%	6	0.02%	0.003%	5.88s

5.3 Testing the Lagrangian mesh representation: a toy free-boundary fluid simulator

I shall now demonstrate how the algorithm can be used to implement a Lagrangian mesh with controlled volume that can change topology, with a very simple incompressible Navier-Stokes simulator for fluids with free boundaries. In the context of this article, this "toy" simulator is meant to demonstrate that our Lagrangian mesh can reproduce some of the typical behaviors of a fluid

while accurately tracking collisions and changes of topology. A more extensive quantitative evaluation of the simulator and the calibration of its parameters will be the subject of another article.

The simulator is obtained by adding viscosity and surface tension terms to the Gallouet-Merigot scheme [32], that simulates an incompressible Euler fluid. Inspired by Brenier and Benamou's point of view on incompressible fluids [14, 33], the Gallouet-Merigot scheme softly projects the motion onto the incompressibility constraint by adding a "spring" force, that moves each point \mathbf{x}_i towards the centroid of its cell V_i . This additional "pressure" force makes the motion of the fluid tangent to the manifold of incompressible motions, and the computed fluid motion provably converges to the solution of an incompressible Euler fluid, that is, a geodesic on the manifold of incompressible motions (see [34, 35] on geodesic flows and the Euler-Arnold equation). The reader is referred to [32] for all the details.

The fluid is modeled as a set of V_i cells, parameterized by the \mathbf{x}_i points. The volume of each V_i cell is controlled and remains constant throughout the whole simulation. All the cells have the same mass m . Each cell is subjected to the following forces:

$$\begin{aligned}
F_p(\mathbf{x}_i) &= \frac{1}{\epsilon_p^2}(\mathbf{x}_i - \mathbf{g}_i) && \text{where } \mathbf{g}_i = \frac{1}{|V_i|} \int_{V_i} \mathbf{x} dx \\
F_g(\mathbf{x}_i) &= -mg\mathbf{z} \\
F_v(\mathbf{x}_i) &= \mu \hat{\Delta} \mathbf{v}(\mathbf{x}_i) &= \mu \sum_{j \in N_i} \frac{1}{2\|\mathbf{x}_j - \mathbf{x}_i\|} |V_{ij}| (\mathbf{v}_j - \mathbf{v}_i) \\
F_t(\mathbf{x}_i) &= \gamma \hat{\Delta} \mathbf{x}(\mathbf{x}_i) &= \gamma \sum_{j \in N_i} \frac{1}{2\|\mathbf{x}_j - \mathbf{x}_i\|} |V_{ij}| (\mathbf{x}_j - \mathbf{x}_i).
\end{aligned}$$

Following [32], the "pressure" F_p may be thought of as a "spring" that connects each point \mathbf{x}_i to the centroid \mathbf{g}_i of the associated cell V_i with a stiffness ϵ_p . The additional forces are viscosity and surface tension: viscosity F_v is proportional to the Laplacian of the velocity. Surface tension F_t is derived in a volumetric manner, from the mutual attraction in the fluid. Inside the fluid, the net result is zero. On the boundary, the net result is a force that pulls towards the interior of the fluid. Both viscosity and surface tension use the \mathbb{P}_1 Laplacian:

$$\hat{\Delta} f(\mathbf{x}) = \sum_{j \in N_i} \frac{1}{2\|\mathbf{x}_j - \mathbf{x}_i\|} |V_{ij}| (f(\mathbf{x}_j) - f(\mathbf{x}_i))$$

where $|V_{ij}| = |\partial V_i \cap \partial V_j|$ denotes the area of the facet common to the cells V_i and V_j , and where N_i denotes the set of cells touching V_i . In the expressions of F_v and F_t , it is computed component-wise.

The simulation results shown below use a semi-implicit time-stepping, with implicit integration for the viscosity, as follows:

$$\begin{aligned}
\mathbf{x}^{(k+1)} &= \mathbf{x}^{(k)} + \delta t \mathbf{v}^{(k)} \\
\mathbf{v}^{(k+1)} &= \mathbf{v}^{(k)} + \frac{\delta t}{m} \left(F_p + F_g + F_t + \mu \hat{\Delta} \mathbf{v}^{(k+1)} \right) \\
\text{or: } \left(\text{Id} - \mu \frac{\delta t}{m} \hat{\Delta} \right) \mathbf{v}^{(k+1)} &= \mathbf{v}^{(k)} + \frac{\delta t}{m} (F_p + F_g + F_t) \\
\text{or: } \left(\hat{\Delta} - \frac{m}{\mu \delta t} \text{Id} \right) \mathbf{v}^{(k+1)} &= -\frac{1}{\mu} \left(\frac{m}{\delta t} \mathbf{v}^{(k)} + F_p + F_g + F_t \right).
\end{aligned}$$

The linear system is solved for each component x, y, z separately. Note that the matrix of this linear system corresponds to $\hat{\Delta}$ plus a diagonal term. To assemble it, the same code that constructs the Hessian of $K()$ can be reused.

Sub-stepping is used whenever the CFL condition is violated, that is whenever a displacement $\delta t \mathbf{v}$ is such that more than one cell is crossed, that is whenever there exists an i such that $\delta t |\mathbf{v}_i| > \sqrt{\psi_i}$.

The method is tested in four different configurations (see videos in the webpage <https://members.loria.fr/blevy/papers/POT>):

- a simulated crown splash, shown in Figure 12, computed with 500000 cells. The simulated mesh captures all the changes of topology, and preserves the fine details of the motion throughout the simulation;
- a fluid in a zigzag domain (Figure 13), demonstrating the interactions between the fluid and the boundary of the domain;
- the same zigzag domain (Figure 14) with a higher surface tension;
- a simulation of the experimental results in [36] shown in Figure 15, where a droplet slides along a slope and bounces over obstacles. An example of a simulation mesh is shown in Figure 16, as well as a cross-section, revealing the internal structure of the mesh, and the effect of surface tension that clusters the points \mathbf{x}_i near the surface (due to the volume constraint, the associated cells are more elongated than the others). Some frames of the full animation are shown in Figure 17. The fluid-boundary interactions and the typical oscillations of the droplet are reproduced in the simulation.

The table below indicates the parameters for the four simulations:

	g	m	μ	γ	N	comput. time
Splash	10	3	0.001	1.5	500000	98s
Zigzag 1	1	3	0.001	0.5	25000	12s
Zigzag 2	1	3	0.001	2	5000	2.8s
Hurdles	0.3	3	0.001	3	5000	3.5

The table reports the parameters for gravity (g), viscosity (μ), surface tension (γ), the number of cells (N) and the computation time for one timestep. The domain Ω is normalized in the $[0, 1]^3$ box. For fine-scale tension-dominated simulations (zigzag and hurdles), the simulation space is globally scaled-down by

applying a scaling factor to g . In the four simulations, the threshold for Newton iterations ϵ_K is set to 1%. During the whole simulation, each cell volume $|V_i|$ remains within a 1% tolerance from the specified volume. Note that unlike some methods that use an Eulerian mesh, volume errors do not accumulate, the 1% bound is met by all cells at all timesteps. The parameter ϵ_p for the "pressure" is set to 0.004, and the timestep is set to 0.004. The spheres are approximated using $N_u = 162$ half-spaces.

I shall now report timing breakdown, measured with the crown splash example (Intel Core i9-9880H, 2.3 GHz, multithreaded implementation):

phase	time
Laguerre (Bowyer-Watson)	9.7%
Linear solve	2.2%
Evaluate gradient	32.8%
Evaluate and assemble Hessian	41.5%
Euler update and implicit viscosity lin. solve	13.8%

The phases that take the largest amount of time are evaluating the gradient and Hessian. It is because these phases involve all the geometric computations (query the Axis-Aligned Bounding Box tree, compute intersections between convex polytopes and compute volumes and areas).

Concerning the geometric predicate used by the convex polytope clipping algorithm (the sign of a 4x4 determinant), during 100 timesteps of the crown splash example, it was called 15752996161 times (15.7 billions), and the arithmetic filter triggered the arbitrary precision mode 70034 times (that is, 0.0004 %). Clearly, degenerate configurations represent a tiny proportion, but given the huge number of computed intersections, they appear in each simulation. Our robust intersection algorithm properly handles them, with negligible overhead, thanks to the arithmetic filters (see also [30] and references herein).

6 Conclusions and future works

The experimental results tend to confirm that semi-discrete partial optimal transport can be used to implement a Lagrangian scheme for free-surface fluid simulation. The precise analysis, comparison to state of the art codes and calibration of the parameters will be the topic of another article. Note also that the convergence to the Merigot-Gallouet scheme to the solution of the incompressible Euler equation is guaranteed [32], but it is no longer the case for Navier-Stokes, with the added viscosity and surface tension terms.

I used here a simple semi-implicit time stepping, it would be interesting to implement a symplectic scheme: in the end this will result in a simulator with both precise conservation of volume and precise conservation of the Hamiltonian. However, this will require a careful analysis for sub-stepping.




Figure 12: Testing surface tension and topology changes: crown splash




Figure 13: Free-surface flow and interactions with boundaries: moderate surface tension.

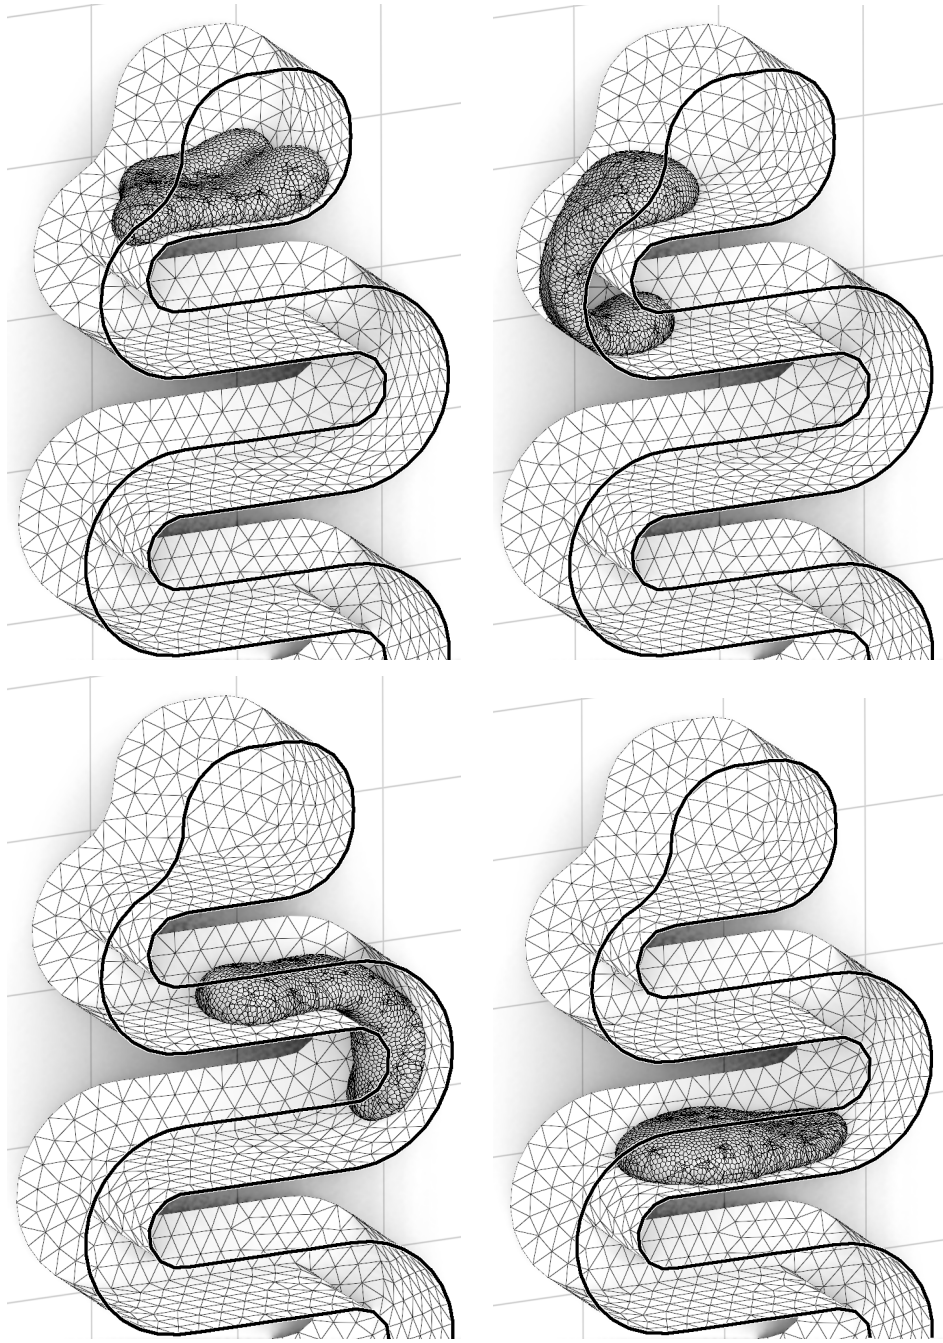


Figure 14: Free-surface flow and interactions with boundaries: higher surface tension.




Figure 15: Droplet hurdles race, courtesy H el ene de Maleprade, Rachid Bendimerad, Christophe Clanet and David Qu er e.




Figure 16: Simulated droplet hurdles race: Mesh details and cross-section.




Figure 17: Simulated droplet hurdles race

Adaptivity is another aspects that would be interesting to study, that is, creating and deleting cells dynamically, during the simulation, using larger cells in zones where the flow is simple, and smaller cells near the boundary of the fluid, to better capture the fluid shape and to more accurately represent surface tension. To even better adapt the geometry of the flow, anisotropic cells could be used, defined by a tensor attached to each point, and cells that are intersections between Laguerre cells and ellipsoids.

Performance can be improved, using a GPU implementation of the clipped Laguerre diagram [37, 38, 39]. This will require to find a way of making the representation more compact (convex polytopes with 200 faces do not fit well in GPU caches). Direct representation of curved polytopes [28] may be an option.

Finally, it is important to stress the fact that the so-defined free surface Lagrangian mesh *smoothly* depends on the parameters \mathbf{x}_i . Namely, the function K is differentiable up to the second order in both ψ and \mathbf{x} , and the second-order derivatives have simple expressions (see [19]). As a consequence, this representation can be used to optimize an objective function that depends on a shape of prescribed volume, where the shape is parameterized by the \mathbf{x}_i 's. Then, the objective function will *smoothly* depend on the \mathbf{x}_i 's, making it possible to optimize it with efficient numerical methods (Newton, LBFGS, ...).

In future works, I plan to explore different applications, in astrophysics (extensions of [4]), and also in shape and topology optimization that may benefit from the differentiability of the representation. It is also interesting to notice that the current trend of numerical methods called "artificial intelligence" mostly rely on representations that are differentiable with respect to the parameters. I think the representation described here can be used to represent latent spaces of shapes, and efficiently fit them to databases of 3D objects.

Acknowledgements

I wish to thank Quentin Mériçot, Yann Brenier and Jean-David Benamou for many discussions, H el ene de Maleprade for sharing the photo of her amazing "droplet hurdles race" experiment (with Rachid Bendimerad, Christophe Clanet and David Qu er e), and Kevin Mattheus Moerman for indicating me H el ene's work on Twitter (and challenging me to simulate it !).

References

- [1] P. Frey, D. Kazerani, T. Ta, An adaptive numerical scheme for solving incompressible two-phase and free-surface flows, Int J. Numer. Methods in Fluids, 2018 (2018).
- [2] Y. Brenier, U. Frisch, M. H enon, G. Loeper, S. Matarrese, R. Mohayaei, A. Sobolevskii, Reconstruction of the early universe as a convex optimization problem, Mon. Not. R. Astron. Soc. 346 (501) (2003).

- [3] U. Frisch, S. Matarrese, R. Mohayaee, A. Sobolevskii, A reconstruction of the initial conditions of the universe by optimal mass transportation, *Nature* 417 (260) (2002).
- [4] B. Lévy, R. Mohayaee, S. von Hausegger, A fast semi-discrete optimal transport algorithm for a unique reconstruction of the early Universe, working paper or preprint (Dec. 2020).
URL <https://hal.archives-ouvertes.fr/hal-03081581>
- [5] P. Geoffroy-Donders, G. Allaire, O. Pantz, 3-d topology optimization of modulated and oriented periodic microstructures by the homogenization method, *Journal of Computational Physics* 401 (2020) 108994.
URL <https://hal.archives-ouvertes.fr/hal-01939201>
- [6] G. Allaire, Shape optimization by the homogenization method / Gregoire Allaire., *Applied mathematical sciences* ; 146, Springer, New York, 2002.
- [7] Q. Mérigot, J. Meyron, B. Thibert, An algorithm for optimal transport between a simplex soup and a point cloud, *CoRR* abs/1707.01337 (2017).
arXiv:1707.01337.
URL <http://arxiv.org/abs/1707.01337>
- [8] Q. Merigot, B. Thibert, Optimal transport: discretization and algorithms (2020). arXiv:2003.00855.
- [9] B. Lévy, E. L. Schwindt, Notions of optimal transport theory and how to implement them on a computer, *Computers & Graphics* 72 (2018) 135–148.
doi:10.1016/j.cag.2018.01.009.
URL <https://doi.org/10.1016/j.cag.2018.01.009>
- [10] G. Peyré, M. Cuturi, Computational optimal transport, Preprint (2017).
URL <https://optimaltransport.github.io/>
- [11] F. Santambrogio, Optimal transport for applied mathematicians, Vol. 87 of *Progress in Nonlinear Differential Equations and their Applications*, Birkhäuser/Springer, Cham, 2015, calculus of variations, PDEs, and modeling. doi:10.1007/978-3-319-20828-2.
URL <http://dx.doi.org/10.1007/978-3-319-20828-2>
- [12] C. Villani, Optimal transport : old and new, *Grundlehren der mathematischen Wissenschaften*, Springer, Berlin, 2009.
URL <http://opac.inria.fr/record=b1129524>
- [13] C. Villani, Topics in optimal transportation, *Graduate studies in mathematics*, American Mathematical Society, Providence (R.I.), 2003.
URL <http://opac.inria.fr/record=b1122739>
- [14] Y. Brenier, Polar factorization and monotone rearrangement of vector-valued functions, *Communications on Pure and Applied Mathematics* 44 (1991) 375–417.

- [15] F. Aurenhammer, Power diagrams: Properties, algorithms and applications, *SIAM J. Comput.* 16 (1) (1987) 78–96.
- [16] F. Aurenhammer, F. Hoffmann, B. Aronov, Minkowski-type theorems and least-squares partitioning, in: *Symposium on Computational Geometry*, 1992, pp. 350–357.
- [17] B. Lévy, A numerical algorithm for L_2 semi-discrete optimal transport in 3d, *ESAIM M2AN (Mathematical Modeling and Analysis)* (2015).
- [18] J. Kitagawa, Q. Mérigot, B. Thibert, A newton algorithm for semi-discrete optimal transport, *CoRR* abs/1603.05579 (2016). arXiv:1603.05579.
URL <http://arxiv.org/abs/1603.05579>
- [19] F. De Gournay, J. Kahn, L. Lebrat, Differentiation and regularity of semi-discrete optimal transport with respect to the parameters of the discrete measure, *Numerische Mathematik* 141 (2) (2019) 429–453.
URL <https://hal.archives-ouvertes.fr/hal-01721681>
- [20] F. de Goes, C. Wallez, J. Huang, D. Pavlov, M. Desbrun, Power particles: an incompressible fluid solver based on power diagrams, *ACM Trans. Graph.* 34 (4) (2015) 50:1–50:11. doi:10.1145/2766901.
URL <https://doi.org/10.1145/2766901>
- [21] M. R. Hestenes, E. Stiefel, Methods of conjugate gradients for solving linear systems, *J Res NIST* 49 (6) (1952) 409–436. doi:10.6028/jres.049.044.
URL <http://dx.doi.org/10.6028/jres.049.044>
- [22] A. Bowyer, Computing dirichlet tessellations, *Comput. J.* 24 (2) (1981) 162–166. doi:10.1093/comjnl/24.2.162.
URL <http://dx.doi.org/10.1093/comjnl/24.2.162>
- [23] D. Watson, Computing the n-dimensional delaunay tessellation with application to voronoi polytopes, *Comput. J.* 24 (2) (1981) 167–172.
- [24] Inria, project ALICE-PIXEL, Geogram: a programming library of geometric algorithms, <http://alice.loria.fr/software/geogram/doc/html/index.html> (2014-2021).
- [25] J. Boissonnat, O. Devillers, S. Pion, M. Teillaud, M. Yvinec, Triangulations in CGAL, *Comput. Geom.* 22 (1-3) (2002) 5–19. doi:10.1016/S0925-7721(01)00054-2.
URL [http://dx.doi.org/10.1016/S0925-7721\(01\)00054-2](http://dx.doi.org/10.1016/S0925-7721(01)00054-2)
- [26] J. R. Shewchuk, Robust adaptive floating-point geometric predicates, in: *Proceedings of the Twelfth Annual Symposium on Computational Geometry*, Philadelphia, PA, USA, May 24-26, 1996, 1996, pp. 141–150. doi:10.1145/237218.237337.
URL <http://doi.acm.org/10.1145/237218.237337>

- [27] A. Meyer, S. Pion, FPG: A code generator for fast and certified geometric predicates, in: Real Numbers and Computers, Santiago de Compostela, Spain, 2008, pp. 47–60.
URL <https://hal.inria.fr/inria-00344297>
- [28] H. Leclerc, Q. Mérigot, F. Santambrogio, F. Stra, Lagrangian discretization of crowd motion and linear diffusion, *SIAM J. Numer. Anal.* 58 (4) (2020) 2093–2118. doi:10.1137/19M1274201.
URL <https://doi.org/10.1137/19M1274201>
- [29] H. Edelsbrunner, E. P. Mücke, Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms, *ACM Trans. Graph.* 9 (1) (1990) 66–104. doi:10.1145/77635.77639.
URL <http://doi.acm.org/10.1145/77635.77639>
- [30] B. Lévy, Robustness and efficiency of geometric programs: The predicate construction kit (PCK), *Computer-Aided Design* 72 (2016) 3–12. doi:10.1016/j.cad.2015.10.004.
URL <http://dx.doi.org/10.1016/j.cad.2015.10.004>
- [31] P. Terdiman, Opcode 3d collision detection library, <http://www.codercorner.com/Opcode.htm> (2005).
- [32] T. O. Gallouët, Q. Mérigot, A lagrangian scheme à la Brenier for the incompressible euler equations, *Foundations of Computational Mathematics* (May 2017). doi:10.1007/s10208-017-9355-y.
URL <https://doi.org/10.1007/s10208-017-9355-y>
- [33] J. Benamou, Y. Brenier, A computational fluid mechanics solution to the monge-kantorovich mass transfer problem, *Numerische Mathematik* 84 (3) (2000) 375–393. doi:10.1007/s002110050002.
URL <https://doi.org/10.1007/s002110050002>
- [34] T. Tao, The euler-arnold equation, <http://terrytao.wordpress.com/2010/06/07/the-euler-arnold-equation/> (2010).
- [35] R. Abraham, J. Marsden, *Foundations of Mechanics*, Behamin-Cummings, London, 1978.
- [36] H. de Maleprade, R. Bendimerad, C. Clanet, D. Quéré, Droplet hurdles race, *Applied Physics Letters* 118 (171601) (2021).
URL <https://doi.org/10.1063/5.0043908>
- [37] D. Yan, W. Wang, B. Lévy, Y. Liu, Efficient computation of clipped voronoi diagram for mesh generation, *Computer-Aided Design* 45 (4) (2013) 843–852. doi:10.1016/j.cad.2011.09.004.
URL <http://dx.doi.org/10.1016/j.cad.2011.09.004>

- [38] N. Ray, D. Sokolov, S. Lefebvre, B. Lévy, Meshless voronoi on the GPU, *ACM Trans. Graph.* 37 (6) (2018) 265:1–265:12. doi:10.1145/3272127.3275092. URL <https://doi.org/10.1145/3272127.3275092>
- [39] J. Basselin, L. Alonso, N. Ray, D. Sokolov, S. Lefebvre, B. Lévy, Restricted Power Diagrams on the GPU, in: *Eurographics'2021*, Vienne, Austria, 2021. URL <https://hal.archives-ouvertes.fr/hal-03169575>