

Strain-minimising Stream Surfaces

Michael Bartoň^a, Jiří Kosinka^b, Victor M. Calo^a

^aNumerical Porous Media Center, King Abdullah University of Science and Technology, Thuwal 23955-6900, Kingdom of Saudi Arabia

^bComputer Laboratory, University of Cambridge, 15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

Abstract

We study the problem of finding strain-minimising stream surfaces in a divergence-free vector field. These surfaces are generated by motions of seed curves that propagate through the field in a strain minimising manner, i.e., they move without stretching or shrinking, preserving the length of their arbitrary arc. In general fields, such curves do not exist. However, the divergence-free constraint gives rise to these ‘strain-free’ curves that are locally arc-length preserving when infinitesimally propagated. Several families of strain-free curves are identified and used as initial guesses for stream surface generation. These surfaces are subsequently globally optimised to obtain the best strain-minimising stream surfaces in a given divergence-free vector field.

Our algorithm was tested on benchmark datasets, proving its applicability to incompressible fluid flow simulations, where our strain-minimising stream surfaces realistically reflect the flow of a flexible univariate object.

Keywords: Stream surface, divergence-free vector field, strain, flow visualization

1. Introduction

We investigate a special class of stream surfaces generated by seed curves that minimise a certain arc-length energy. With flow simulations of viscous fluids in mind, consider a seed curve as a sequence of infinitely many liquid drops. We seek to answer the following question: Given a time-independent vector field in 3D, is there a curve that propagates in time without changing the strain between any two of its neighbouring drops? In other words, does a curve that moves in the vector field without shrinking or stretching any of its arcs exist?

It is well known [Dav67] that divergence-free vector fields are volume preserving, i.e., the volume of any 3D object remains constant when propagated in time. In general, this is not true for lower-dimensional objects. Thus, a natural question arises: Are there lower-dimensional objects (surfaces, curves) that propagate in time in the same manner, i.e., preserving their lower-dimensional measures (area, length)? Or, stated differently, is there a region where the divergence-free vector field acts on an object more than in volume preserving manner, namely by preserving its 1D or 2D measure? This paper investigates this question in the curve case.

Our motivation is straightforward: if a sufficiently elastic univariate object (an elastic rod or drops of another fluid) is put into the flow along a special curve, the deformation that acts on it is bending only; it contains no ‘strain’ forces.

Problem formulation. Given a steady (time independent), divergence-free vector field \mathbf{v} in some domain $\Omega \subset \mathbb{R}^3$, ($\nabla \cdot \mathbf{v}(\mathbf{p}) = 0, \forall \mathbf{p} \in \Omega$), find a *stream surface* such that the *seed*

curve that defines it propagates in time along the surface so that its arc-length changes as little as possible.

We combine theoretical investigations and a practical algorithm for finding such stream surfaces. The main steps and contributions of our method are:

- We theoretically investigate families of seed curves based on certain strain-minimising energies (Section 3).
- These candidate seed curves are used to generate initial stream surfaces (Section 4.1).
- The initial stream surfaces are globally optimised and ranked according to their strain energies (Section 4.2).

Our implementation of the method is presented in Section 5. We have validated our theoretical results on several benchmark datasets and demonstrated the applicability of our method on numerous examples (Section 6 and the accompanying video). Possible extensions of our method are discussed in Section 7 and the paper is concluded in Section 8.

2. Related work

Stream surfaces, used as a tool for visualising characteristic features of vector fields, have been extensively studied in the visualisation literature; see [BP02, MSRT13, ELM*12, SEG*14], the survey paper [MLP*10] and the references cited therein. Classical methods [Hul92] are usually based on trial-and-error approaches: the user inserts seed curves (typically straight lines), stream surfaces are computed, and, if they do not capture desired features well, the initial seed curves are modified and the whole process is repeated. Since visualising vector fields by stream surfaces (compared to using streamlines) has become more popular [MLP*10], research in *automatic* stream surface

Email addresses: Michael.Barton@kaust.edu.sa (Michael Bartoň),
Jiri.Kosinka@cl.cam.ac.uk (Jiří Kosinka),
Victor.Calo@kaust.edu.sa (Victor M. Calo)




Figure 1: Strain minimisation in terms of differential geometry. At any point \mathbf{p} on the seed curve $\gamma(s) = \mathbf{S}(s, 0)$, the directional derivative \mathbf{S}_{st} of the vector \mathbf{u} along \mathbf{v} is perpendicular to \mathbf{u} ; see (4).

seeding has recently become very active. Our method fits in this modern family of automatic stream surface algorithms.

Divergence-free vector fields are used in many areas and applications such as incompressible fluid simulations [LZF10], smoke visualisation [vFWTS08], and are also a favourite modelling/deformation tool [vFTS06] due to their volume preserving property.

Another research area, which our work is connected to, relates to curve evolution [You10], where typically the curve and the property to preserve (e.g. arc-length) are given and the corresponding evolving vector fields are sought after. For example, [HPP05] seeks 3D smoothing flows that satisfy additional spatial constraints, and [BSK*13] uses the flow to evolve curves in 3D space while preserving their arc-length and curvature.

3. Strain-energy minimising curves

Let $\mathbf{v}(\mathbf{p})$ be a steady differentiable vector field defined over a domain $\Omega \subset \mathbb{R}^3$ and assume that \mathbf{v} is divergence free, i.e., $\text{div } \mathbf{v}(\mathbf{p}) = \langle \nabla, \mathbf{v}(\mathbf{p}) \rangle = 0, \forall \mathbf{p} \in \Omega$. Let \mathbf{J} be the Jacobian matrix of $\mathbf{v}(\mathbf{p})$, i.e., $\mathbf{J}_{ij} = \frac{\partial v_i}{\partial p_j}$ with $\mathbf{p} = (p_1, p_2, p_3)$.

Consider a regular curve $\gamma(s)$ parametrised by arc-length, $s \in [s_0, s_1]$. We regard γ as a *seed curve* that gives rise to a *stream surface* $\mathbf{S}(s, t)$ with normal $\mathbf{n}(s, t)$, i.e., $\langle \mathbf{v}(\mathbf{S}(s, t)), \mathbf{n}(s, t) \rangle = 0$ for all (s, t) in the surface domain $[s_0, s_1] \times [t_0, t_1]$, and $\mathbf{S}(s, 0) = \gamma(s)$, $0 \in [t_0, t_1]$; see Fig. 1. The partial derivatives of \mathbf{S} will be denoted $\mathbf{S}_s, \mathbf{S}_{st}$, etc.

In general, curves which maintain their arc-length constant (i.e., equal to $s_1 - s_0$ independently of t) when deformed by \mathbf{v} do not exist. However, as we show below, it is possible to find curves which approximate this property to first or even second order. To make this concept precise, we formulate the following

Lemma 3.1. *The Taylor expansion of the arc length of γ with respect to \mathbf{v} at $t = 0$ is given by*

$$\int_{s_0}^{s_1} \|\mathbf{S}_s(s, t)\| ds = (s_1 - s_0) + c_1 t + c_2 t^2 + \mathcal{O}(t^3) \quad (1)$$




Figure 2: Computing first-order vectors on the boundary of Ω . Self-conjugate vectors associated with \mathbf{J} (see (4)) form a quadratic cone C . If restricted to the tangent space of Ω at \mathbf{p} , at most two first-order vectors \mathbf{d}_1 and \mathbf{d}_2 exist.

with

$$\begin{aligned} c_1 &= \int_{s_0}^{s_1} \langle \mathbf{S}_s, \mathbf{S}_{st} \rangle ds|_{t=0}, \\ c_2 &= \int_{s_0}^{s_1} \langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle + \langle \mathbf{S}_s, \mathbf{S}_{stt} \rangle - \langle \mathbf{S}_s, \mathbf{S}_{st} \rangle^2 ds|_{t=0}. \end{aligned} \quad (2)$$

A straightforward proof can be found in Appendix A.

Our aim is to identify curve(s) γ in Ω for which c_1 (and also c_2 , if possible) vanishes. These will subsequently be used to identify strain-minimising stream surfaces in an optimisation procedure.

3.1. First-order strain energy

We have $\mathbf{S}_t = \mathbf{v}$ and we denote $\mathbf{u} := \dot{\gamma}(s) = \mathbf{S}_s(s, 0)$. Then

$$\langle \mathbf{S}_s, \mathbf{S}_s \rangle|_{t=0} \equiv \langle \mathbf{u}, \mathbf{u} \rangle \equiv 1. \quad (3)$$

Differentiation with respect to t yields (see Fig. 1)

$$\langle \mathbf{S}_s, \mathbf{S}_{st} \rangle|_{t=0} \equiv \mathbf{u} \mathbf{J} \mathbf{u}^T \equiv 0. \quad (4)$$

This observation leads us to

Definition 3.1. The *first-order strain energy* of γ is given by

$$E_1(\gamma) = \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} \langle \mathbf{S}_s, \mathbf{S}_{st} \rangle^2 ds|_{t=0}. \quad (5)$$

A curve γ on which E_1 vanishes will be called a *first-order curve*, and their collection denoted

$$\Gamma_1 = \{\gamma | \gamma \subset \Omega, E_1(\gamma) = 0\}.$$

Moreover, a vector \mathbf{d} such that $\mathbf{d} \mathbf{J} \mathbf{d}^T = 0$, i.e., a self-conjugate vector associated with \mathbf{J} , will be called a *first-order vector*.

Note that $E_1(\gamma) = 0$ implies $c_1 = 0$ and thus the first order energy is well defined. In other words, E_1 measures the change of the magnitude of $\dot{\gamma}$ along γ for an infinitesimal increment of t . $E_1(\gamma) = 0$ implies the deformation given by \mathbf{v} that acts on γ preserves, up to *first order*, the magnitude of its tangent vector $\dot{\gamma}$ and hence is locally arc-length preserving.




Figure 3: Initialisation of first-order curves. (a) A divergence-free vector field $\mathbf{v} = [x + y^2 + 2z^3, 10x^3 + 2y, 2x^2y - 3z]$ within a unit domain Ω . (b) Two first-order vector fields (yellow) on the boundary quad given by $y = 0$. Integration gives rise to first-order boundary curves (γ_1 and γ_2 in red). (c) The stream surface obtained by propagating γ_1 in \mathbf{v} . Observe that the preservation of arc length is guaranteed only locally, i.e., close to the boundary, up to first order via (4).

In order to find seed curves that belong to Γ_1 with respect to \mathbf{v} , we seek first-order vectors. To this end, we denote *strain rate* $\mathbf{J}^+ := (\mathbf{J} + \mathbf{J}^T)/2$ and *vorticity* $\mathbf{J}^- := (\mathbf{J} - \mathbf{J}^T)/2$ as the symmetric and antisymmetric parts of the Jacobian matrix of \mathbf{v} . A vector field \mathbf{v} for which \mathbf{J}^+ is regular will be called *non-degenerate*.

Lemma 3.2. *Let $\mathbf{p} \in \Omega$ and \mathbf{v} be divergence free and non-degenerate in a neighbourhood of \mathbf{p} . Then all first-order vectors \mathbf{d} form a quadratic cone with apex at \mathbf{p} .*

Proof. Since $\mathbf{u}\mathbf{J}^-\mathbf{u}^T = 0$ for any \mathbf{u} , we obtain the condition $\mathbf{d}\mathbf{J}^+\mathbf{d}^T = 0$. By definition, $\text{div } \mathbf{v}(\mathbf{p}) = 0$ is equivalent to $\text{tr}(\mathbf{J}) = 0$ at \mathbf{p} and thus $\text{tr}(\mathbf{J}^+) = 0$. From the spectral theorem it follows that \mathbf{J}^+ has three real eigenvalues such that $\lambda_1 + \lambda_2 + \lambda_3 = 0$. This in turn implies that the signature of \mathbf{J}^+ is either $(+, +, -)$ or $(+, -, -)$ by non-degeneracy of \mathbf{v} . Consequently, all solutions of $\mathbf{d}\mathbf{J}\mathbf{d}^T = 0$ form a quadratic cone. \square

Remark 1. In the special case (of measure zero) when \mathbf{J}^+ is singular, the space of first-order vectors \mathbf{d} at \mathbf{p} that solve $\mathbf{d}\mathbf{J}\mathbf{d}^T = 0$ is either given by two planes intersecting in a line incident with \mathbf{p} (the signature of \mathbf{J}^+ is $(+, -, 0)$) or any vector is a first-order vector (the signature of \mathbf{J}^+ is $(0, 0, 0)$). Consequently, first-order vectors exist at any point $\mathbf{p} \in \Omega$ for divergence-free vector fields.

From Lemma 3.2 and Remark 1 it follows that there exist infinitely many first-order curves γ passing through every point in Ω . They can be obtained by integrating first-order vectors \mathbf{d} , which, however, form a multi-valued field. Thus, the set Γ_1 of these curves is too large to be practical. We therefore explore three conditions that select special classes of first-order curves from Γ_1 :

1. minimise a certain second-order strain energy;
2. restrict Γ_1 to curves on the boundary only, i.e., $\gamma \subset \partial\Omega$;
3. constrain the variation of $\mathbf{d} = \dot{\gamma}$ along γ .

We now address each of these strategies in detail.

3.2. Second-order strain energy

Similarly to the first-order strain energy E_1 , we now define its second-order counterpart. Differentiating (4) with respect to t gives

$$\langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle|_{t=0} + \langle \mathbf{S}_s, \mathbf{S}_{st} \rangle|_{t=0} \equiv \mathbf{u}\mathbf{K}\mathbf{u}^T \equiv 0, \quad (6)$$




Figure 4: Second-order curves. Left: At an interior point (red) of Ω , second-order vectors, generically and if they exist, correspond to the intersection of two quadratic cones. Right: Integrating second-order vectors gives γ , an integral curve whose tangent vectors $\dot{\gamma}$ solve (4) and (6). The stream surface emanating from γ for the same vector field as in Fig. 3 is shown.

where $\mathbf{K} = \mathbf{J}\mathbf{J}^T + \mathbf{J}^2$. This leads us to

Definition 3.2. The *second-order strain energy* of γ is given by

$$E_2(\gamma) = \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} (\langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle + \langle \mathbf{S}_{st}, \mathbf{S}_s \rangle)^2 ds|_{t=0}. \quad (7)$$

A curve γ on which both E_1 and E_2 vanish will be called a *second-order curve*, and their collection denoted

$$\Gamma_2 = \{\gamma \mid \gamma \subset \Omega, E_1(\gamma) = E_2(\gamma) = 0\}.$$

Moreover, a first-order vector \mathbf{d} such that $\mathbf{d}\mathbf{K}\mathbf{d}^T = 0$ will be called a *second-order vector*.

Note that $E_1(\gamma) = E_2(\gamma) = 0$ implies that $c_1 = c_2 = 0$ in Lemma 3.1 and thus the second-order energy is well defined: second-order curves are characterised by the property that their deformation given by \mathbf{v} that acts on γ preserves the magnitude of the tangent vector up to *second order* at $t = 0$; cf. Lemma 3.1.

Whereas the existence of first-order vectors \mathbf{d} is guaranteed at every \mathbf{p} (see Fig. 2 and Lemma 3.2), a non-zero solution of (6) need not exist. The generic cases with respect to $\mathbf{d}\mathbf{K}\mathbf{d}^T = 0$ are categorised by the signature of \mathbf{K}^+ : $(+, +, +)$ or $(-, -, -)$ yield no non-zero solution; $(+, +, -)$ or $(+, -, -)$ give, as in the first-order case, a quadratic cone. In the former case, there are no second-order vectors at \mathbf{p} and alternatives must be sought;

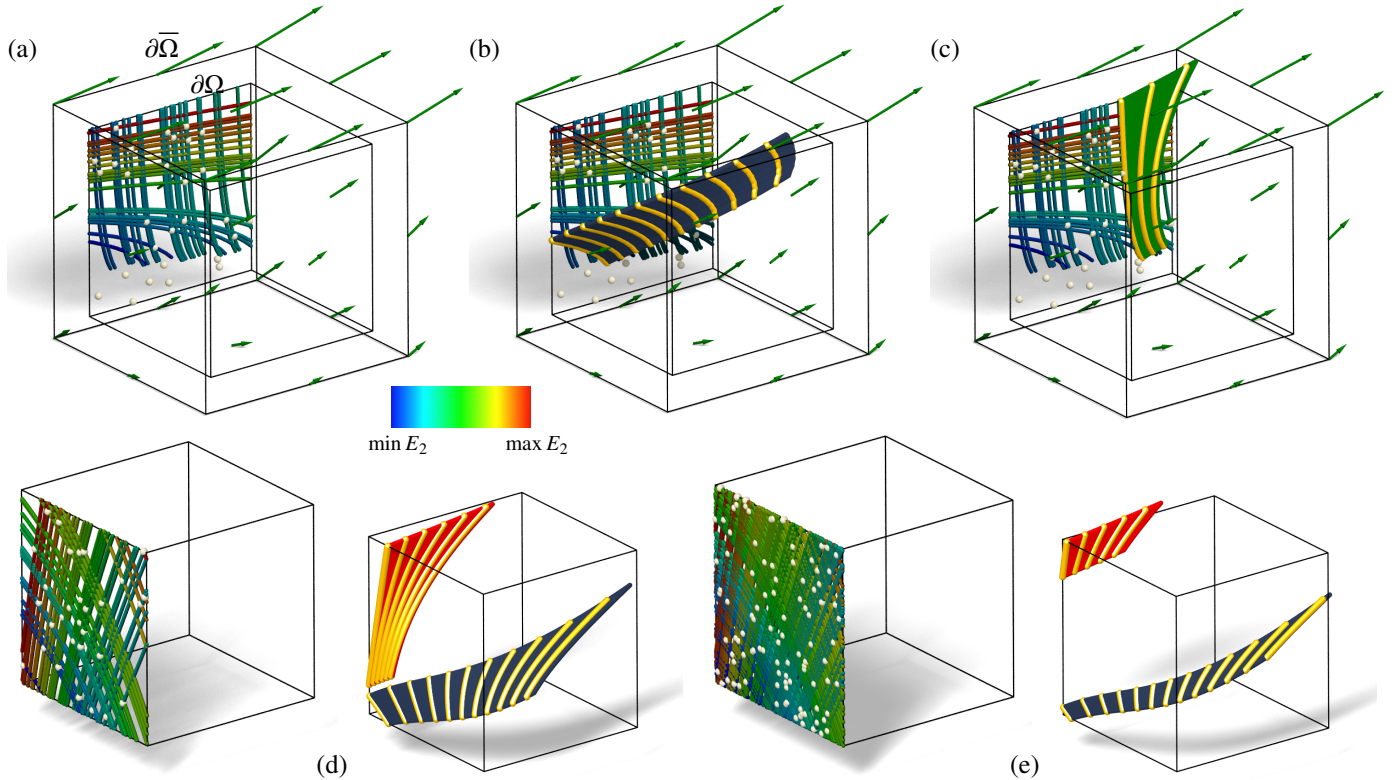


Figure 5: Exploring the space of first-order boundary curves. (a) At sampled boundary points, the curves are computed and colour-coded according to the second-order energy E_2 ; see (7). (b)–(c) The stream surface obtained from the curves corresponding to the best and the median value of E_2 , respectively. Seed curve integration is terminated when the outer offset of $\partial\Omega$, $\partial\bar{\Omega}$, is reached. (d)–(e) The analogous situation for a different part of $\partial\Omega$. Left: First-order boundary curves arisen from $n = 30$ and 100 random samples, respectively. Right: The stream surfaces obtained by integrating the vector field starting from the best (blue) and the worst (red) first-order boundary curves according to E_2 .

see Sections 3.3 and 3.4. The latter case leads to the intersection of two quadratic cones. The situation is shown in Fig. 4, left. Up to four second-order vectors can be found by solving a quartic equation or, more geometrically, by reducing the problem via a cubic equation to the intersection of a quadratic cone with two planes; see Appendix B. Non-generic cases are, for the sake of brevity, not considered.

Therefore, testing whether a real non-zero solution of the system given by (4) and (6) exists is a cheap closed-form operation. This allows us to quickly explore Ω for regions where second-order vectors exist and, in the positive case, to integrate them to obtain curves in Γ_2 . An example of a second-order curve is shown in Fig. 4, right. Such curves then form input for our optimisation algorithm (Section 4.2).

If no second-order curves exist in Ω or if the existing ones are not satisfactory for visualisation purposes, we employ our second strategy: boundary curves.

3.3. First-order boundary curves

Our second strategy is to restrict the set Γ_1 to curves on the boundary of Ω . This is a reasonable restriction from the point of view of visualisation: boundary seed curves capture the behaviour of a given flow as it enters/exits the domain; see Fig. 3. When the point \mathbf{p} lies on the boundary $\partial\Omega$ of Ω and \mathbf{v} is non-degenerate, we see that there exist at most two first-order vectors \mathbf{d}_1 and \mathbf{d}_2 in the tangent space of $\partial\Omega$ at \mathbf{p} ; see Fig. 2.

First-order boundary curves form a subset of Γ_1 , which we denote

$$\Gamma_1^\partial = \{\gamma \mid \gamma \subset \partial\Omega, E_1(\gamma) = 0\}. \quad (8)$$

The first-order strain energy E_1 was defined to measure the change of the magnitude of the unit tangent vector $\dot{\gamma}$ of γ , when being instantaneously moved by the vector field \mathbf{v} . Every curve in Γ_1^∂ is a first-order curve, but is not, in general, a second-order curve. Nevertheless, E_2 can be used as a ranking criterion to determine good candidate seed curves among those in Γ_1^∂ ; see Fig. 5.

Alternative ranking energies. In the case when the strain energy is not the main objective, one may consider alternative components of a general ranking energy E :

$$E(\gamma) = w_1 E_{in} + w_2 E_{ortho} + w_3 E_{para} + w_4 E_{rigid} + E_2, \quad (9)$$

where

$$\begin{aligned} E_{in}(\gamma) &= \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} 1 - \langle \mathbf{S}_t, \mathbf{m} \rangle^2 ds|_{t=0}, \\ E_{ortho}(\gamma) &= \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} \langle \mathbf{S}_t, \mathbf{S}_s \rangle^2 ds|_{t=0}, \\ E_{para}(\gamma) &= \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} \langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle^2 ds|_{t=0}, \\ E_{rigid}(\gamma) &= \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} \langle \bar{\mathbf{c}} + \mathbf{c} \times \gamma(s), \mathbf{S}_t \rangle^2 ds|_{t=0}. \end{aligned} \quad (10)$$

The particular components of E have the following meaning: E_{in} votes for curves that start moving to the interior part of Ω ,

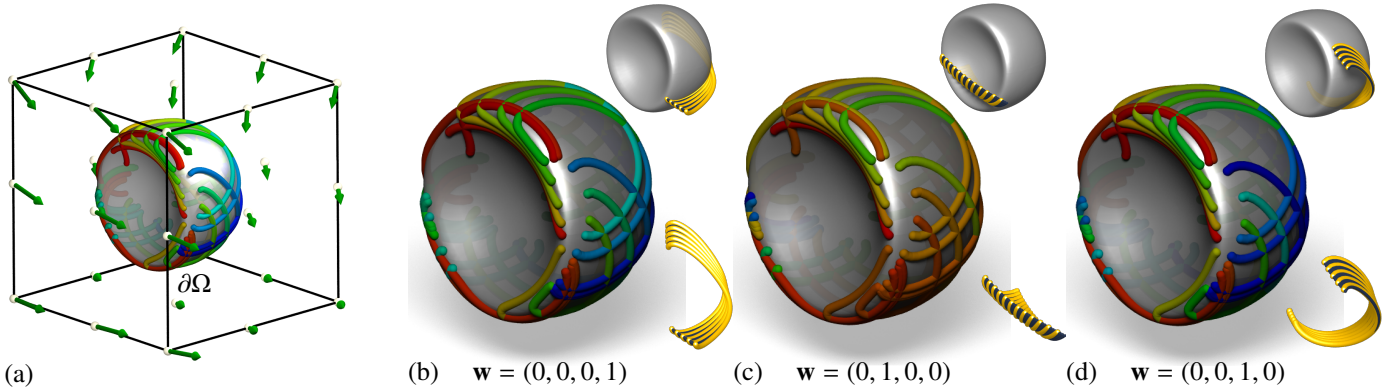


Figure 6: A family of first-order boundary curves from Γ_1^∂ , see (8), on the smooth surface $\partial\Omega$. These curves are subsequently ranked by E given in (9) considering several objectives: (b) the curves are additionally required to move as rigidly as possible, (c) as orthogonally as possible, (d) as parallel as possible. This is achieved by setting the appropriate weight vector \mathbf{w} in (9). The best seed curve and the corresponding stream surface are shown in each situation.

\mathbf{m} being the inward normal of $\partial\Omega$; E_{ortho} selects curves that start moving orthogonally to the flow, i.e., the tangent vectors $\mathbf{S}_s = \dot{\gamma}$ are (in the least square sense) as orthogonal as possible to the velocity vectors $\mathbf{S}_t = \mathbf{v}$; similarly E_{para} finds curves that propagate in a parallel fashion; and E_{rigid} favours curves that move as rigidly as possible, (c, \bar{c}) being the instantaneous motion; see [PW01, BSK* 13].

The behaviour of the ranking energy E depending on weights $\mathbf{w} = (w_1, w_2, w_3, w_4)$ when applied on a curved domain Ω is shown in Fig. 6.

However, there exist rare scenarios with no first-order boundary curves, i.e., the set Γ_1^∂ is empty, or the existing ones are insufficient for a particular application. If that is the case, we turn to our third strategy, which is always guaranteed to produce first-order curves.

3.4. First-order interior curves

In the rare situation when there are no second-order curves ($\Gamma_2 = \emptyset$) and no first-order boundary curves ($\Gamma_1^\partial = \emptyset$), we identify first-order curves that are, given an initial point and first-order direction, curvature minimising among those in Γ_1 . The benefit here is twofold: such curves are always guaranteed to exist, and they are particularly well suited for visualisation since they are as straight as possible. These curves are given by point-wise minimisation of $\|\dot{\gamma}(s)\|$ subject to $\dot{\gamma}(s)\mathbf{J}(\gamma(s))\dot{\gamma}(s)^\top = 0$; cf. (4).

Discretisation and implementation details are discussed in Section 5. Before all that, we proceed to the exploration of optimal stream surface generation.

4. Stream surfaces and optimisation

We start by defining strain minimising stream surfaces, which are then optimised with respect to a certain strain energy.

4.1. Strain minimising stream surfaces

Second-order curves, or first-order curves if the former ones do not exist, are good initial guesses for seed curves, called *candidate seed curves*, but only *locally*. Our goal is to find stream

surfaces for which their generating seed curves propagate *globally* in a certain strain-minimising manner. This is formalised in

Definition 4.1. The *strain energy* of a stream surface \mathbf{S} is given by

$$E_{\mathbf{S}} = \frac{1}{\text{area}(\mathbf{S})} \int_{t_0}^{t_1} E_1(\mathbf{S}(s, t)) dt. \quad (11)$$

A stream surface that minimises this energy will be called *strain minimising*.

In other words, $E_{\mathbf{S}}$ measures the strain given by the deformation of the seed curve of \mathbf{S} through the field by accumulating the strain energies of all its timelines. This energy is, as in the case with seed curves, used as a ranking criterion for stream surfaces.

4.2. Optimisation of stream surfaces

Candidate seed curves generate stream surfaces with relatively low strain energy. However, their effect is still only *local*, whilst we are interested in minimising the strain energy *globally*. This is achieved by employing an optimisation procedure; see Fig. 8 for a schematic overview.

Candidate curves are used for initialising an optimisation cycle, which works as follows; see Fig. 7. A candidate seed curve γ_0 at $t = 0$ is integrated to form its stream surface \mathbf{S}_0 , which is subsequently optimised with respect to its strain energy (11). The optimised surface \mathbf{S}_0^{opt} , however, is not, in general, a stream surface any more. Therefore, its timelines are back-integrated to the initial time $t = 0$ to form a set of space curves \mathcal{S}_1 . These curves are used to compute an updated γ_1 by least square fitting, which is forward integrated to build \mathbf{S}_1 and so on. The algorithm continues updating $\mathcal{S}_i \leftarrow \mathcal{S}_{i+1}$ until the surface strain energy $E_{\mathbf{S}_i}$ stops being improved, or when the maximum number of iterations, set to ten if not stated otherwise, is reached. The particular steps of the algorithm are explained, including implementation details, in the next section.

5. Implementation

Having explored the smooth setting of seed curves and stream surfaces, we now discretise these notions and discuss the implementation of our algorithm; see Fig. 8 for an overview.

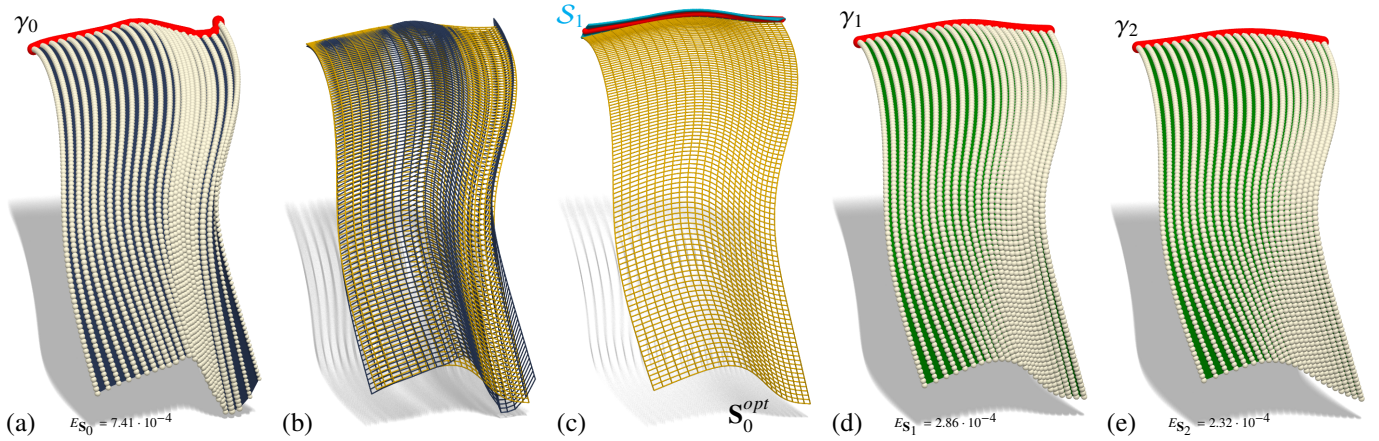


Figure 7: Stream surface optimisation. (a) The seed curve γ_0 (red) is integrated, resulting in the initial stream surface \mathbf{S}_0 (blue). The radii of the spheres correspond to arc-length change. (b) The stream surface is discretised into a quad mesh \mathbf{Q} and optimised. (c) Since the optimised surface \mathbf{S}_0^{opt} is not a stream surface any more, its timelines are back-integrated to $t = 0$, forming a set \mathbf{S}_1 from which the original seed curve is improved to γ_1 (red). (d) The improved seed curve gives rise to an updated stream surface \mathbf{S}_1 . (e) The stream surface \mathbf{S}_1 after another iteration of the optimisation cycle.

Finding candidate curves. Our approach is based on (adaptive) sampling. In the case of first-order boundary curves, $\partial\Omega$ is explored; see examples in Figs. 5 and 6. Otherwise, the interior of Ω is sampled and first- and second-order curves are computed. These are ranked by (11) and the regions with low values of E_S are sampled with higher density. This sampling is repeated recursively, if not stated otherwise, three times and the best 5% are taken as candidate curves.

Seed curve computation proceeds as follows. Starting with a sampled seed point, the admissible tangent field is computed (see Sections 3.2, 3.3, and 3.4 for particular cases) and integrated in all admissible directions. The step-size is, by default, set as 1% of the domain's diameter $\text{diam}(\Omega)$, but can be refined if the sampling is not sufficient; see Fig 10. Only curves of length greater or equal to $\text{diam}(\Omega)/10$ are considered as candidates. The integration terminates if there is no admissible direction to continue in, the maximum length set to $\text{diam}(\Omega)$ is achieved, or, for non-boundary curves, the boundary is reached. Integration of first-order interior curves can be found in Appendix C.

Integrating the vector field. Various techniques for numerical *streamline* integration have been studied [NHM97]. Higher-order methods with an adaptive stepsize were shown to be the preferable choice when considering accuracy versus speed trade-offs. In our case, however, since the back-integrated timelines of the optimised stream surface are taken as input for seed curve optimisation, our prime focus is on accuracy. Therefore, we employed the classical fourth order Runge-Kutta method with constant stepsize.

Optimising the stream surface. Given a stream surface \mathbf{S} , we want to reduce its strain energy (11). Sampling m points in the s -direction (seed curve) and n points in the t -direction (time), \mathbf{S} is discretised into a quad mesh \mathbf{Q} having $m \times n$ vertices $\mathbf{q}_{i,j}^0$. Denoting the vertices of the optimised surface by $\mathbf{q}_{i,j}$ and setting $\mathbf{e}_{i,j} := \mathbf{q}_{i+1,j+1} - \mathbf{q}_{i+1,j} + \mathbf{q}_{i,j} - \mathbf{q}_{i,j+1}$, the desired improvement is formulated as a non-linear least squares optimisation with the

objective function

$$F(\mathbf{Q}) = F_{strain}(\mathbf{Q}) + \mu_1 F_{fair}(\mathbf{Q}) + \mu_2 F_{prox}(\mathbf{Q}) \quad \text{with} \quad (12)$$

$$\begin{aligned} F_{strain}(\mathbf{Q}) &= \sum_{i,j} \langle \mathbf{q}_{i+1,j} - \mathbf{q}_{i,j}, \mathbf{e}_{i,j} \rangle^2 \\ F_{fair}(\mathbf{Q}) &= \sum_{i,j} \|\mathbf{q}_{i+1,j} - 2\mathbf{q}_{i,j} + \mathbf{q}_{i-1,j}\|^2 + \\ &\quad + \|\mathbf{q}_{i,j+1} - 2\mathbf{q}_{i,j} + \mathbf{q}_{i,j-1}\|^2, \\ F_{prox}(\mathbf{Q}) &= \sum_{i,j} \|\mathbf{q}_{i,j} - \mathbf{q}_{i,j}^0\|^2, \end{aligned} \quad (13)$$

where F_{strain} reflects the strain minimising condition (cf. (11)), F_{fair} is a fairness term, and F_{prox} is a proximity term, a regulariser that forces the vertices of the optimised mesh not to deviate much from the input. The optimisation problem is solved using the Gauss-Newton method for all the examples in the paper and the accompanying video. Experimentally, the weights were set to $\mu_1 = 0.1$, $\mu_2 = 0.02$. This results in an optimised surface \mathbf{S}^{opt} .

Improving the seed curves. Having obtained the optimised surface \mathbf{S}^{opt} , its timelines $\mathbf{S}^{opt}(s, t_j) = \gamma_j$, $j = 1, \dots, n$ are back-integrated to the initial time instant $t = 0$, giving the set \mathbf{S}_1 of space curves $\gamma_j^{t=0}$, see Fig. 7(c). If \mathbf{S}^{opt} was an exact stream surface, all $\gamma_j^{t=0}$ would coincide. We improve the seed curve γ of \mathbf{S} by replacing it by the least squares approximation [FHK02] of $\gamma_j^{t=0}$, $j = 1, \dots, n$. Note that we have the information about the correspondence in the s -direction, i.e., for a fixed i , all the points $\mathbf{q}_{i,j}$, $j = 1, \dots, n$, need to correspond to a single point $\gamma(s_i)$. This fact simplifies the problem to point-wise averaging. Having numerically back-integrated \mathbf{S}^{opt} in a point-wise fashion resulting in $\mathbf{q}_{i,j}^{t=0}$, the seed curve update is achieved by setting $\gamma(s_i) = \frac{1}{n} \sum_j \mathbf{q}_{i,j}^{t=0}$.

Curve and surface trimming. So far, we have not discussed how to set the intervals $[s_0, s_1]$ and $[t_0, t_1]$. In practise, both of these have to be finite. In the s (seed curve) direction, the tangent vector field is integrated while admissible directions exist, or until the boundary is reached. In the t (time) direction, the integration is terminated if the timeline reaches the boundary,




Figure 8: Algorithm overview.

or, as shown in Fig. 9, when the maximum number of timesteps is reached.

6. Numerical results

The algorithm was tested on several benchmark datasets. The example in Fig. 11 is a direct numerical Navier-Stokes simulation by [CSBI05] that is publicly available [Int]. We used a uniformly resampled version, which has been provided by Tino Weinkauf and used in [vFwTS08]. The example is based on the last time instant of the unsteady flow.




Figure 9: Rayleigh-Bénard heat convection. The dataset used here is given by one time instant of an unsteady vector flow obtained from [NaS]. (a) Several best strain minimising stream surfaces detected by our algorithm (see Fig. 8) are shown. A total of 434 second order curves were used for initialisation (top red); the optimisation parameters were set to $\mu_1 = 0.1, \mu_2 = 0.02$; see (12). The vector field integration was terminated by (b–d) reaching the boundary and (e) by exceeding the upper bound on the number of timesteps.




Figure 11: The ‘square cylinder’ of [CSBI05]. (a) Incompressible flow moving from left to right in a box, encountering a rigid obstacle (yellow box). The best strain-minimising surfaces for seed curves of different lengths are shown. The length is demanded to be at least 5% (b) and 10% (c) of the domain’s diameter.

The simulation of a flow coming from a Francis turbine is shown in Fig. 10. The original vector field is unsteady, so we again used only one time instant. Note that the best strain-minimising stream surfaces are generated by seed curves that move in only one part of the turbine’s body, whilst the low ranked surfaces are those where the middle part of the body forces their timelines to split apart; see Fig. 10(g). This observation suggests that E_S could be used to detect this type of splitting in a flow.

Fig. 9 shows a time instant of Rayleigh-Bénard heat convection, where the fluid is heated at the bottom and cooled at the top of the boundary, resulting in a circulatory motion. This vector field is not incompressible. Nevertheless, as seen in Fig. 9(a), the second-order curves still exist in this field. Recall that the guarantee of existence of first-order curves applies to divergence-free vector fields only (Remark 1), while second-order curves are not guaranteed to exist.

Fig. 12 shows a comparison between first- and the second-order curves when used as seed curves in the vector field from Fig. 9. As expected, second-order curves offer more favourable results and give rise to initial stream surfaces with nearly van-




Figure 10: Turbine data of [BP02]. (a) One time instant of the unsteady vector flow generated by the Francis turbine is shown. (b) A total of 317 uniformly sampled points in the domain gave rise to 480 second-order seed curves. The vector field is forward and backward integrated and the curves are ranked and colour-coded according to the strain energy E_S of the associated stream surfaces; see (11). (c) Regions with curves with low values of E_S are further explored with higher sampling density. (d) The best-ranked initial stream surface and (e) its optimised counterpart; $E_S = 6.32 \cdot 10^{-6}$. (f) The second best-ranked stream surface with $E_S = 9.28 \cdot 10^{-6}$. (g) An example of a stream surface with a high strain energy that is split into two parts by the flow. Note that the length of the timelines grows significantly, resulting in exceeding $E_S = 9.28 \cdot 10^2$.

Table 1: A summary of our results. In the ‘Seed curves’ columns, $\#p$ is the number of sampled points, $\#\gamma$ is the number of (first- or second-order) seed curves emanating from them, and E_S^{ini} is the strain energy (11) of the best candidate before optimisation. In the ‘Optimisation’ columns, $\#S$ is the number of surfaces processed and E_S^{opt} is the energy of the best optimised stream surface. The last column lists total computation times in seconds.

| Fig. | Seed curves | | | Optimisation | | Time (sec.) |
|--------|-------------|------------|--------------|--------------|--------------|-------------|
| | $\#p$ | $\#\gamma$ | E_S^{ini} | $\#S$ | E_S^{opt} | |
| 5(a-c) | 30 | 41 | $6.04e^{-2}$ | n/a | n/a | 2 |
| 10 | 240 | 492 | $3.17e^{-4}$ | 54 | $6.32e^{-6}$ | 385 |
| 9 | 216 | 434 | $1.24e^{-5}$ | 76 | $2.44e^{-7}$ | 513 |
| 11(b) | 216 | 239 | $6.24e^{-4}$ | 32 | $8.17e^{-6}$ | 147 |
| 11(c) | 216 | 184 | $9.51e^{-4}$ | 21 | $1.04e^{-5}$ | 135 |

ishing strain energy E_S .

Statistics concerning the number of sampled points and seed curves, optimisation parameters, resulting energies E_S , and computation times are listed in Table 1. The timings differ depending on concrete parameter settings, ranging from a few seconds to several minutes. The most expensive part is the exploration of families of first- and/or second-order curves, with computing the energy E_S in particular, which requires integration of all the candidate seed curves. The reported timings were obtained on a machine running Windows with a double-core CPU (2.67 GHz) and 24GB RAM. Currently, only a single-core CPU implementation is available. However, the algorithm is well suited for parallelisation (e.g. curve/surface energies can be computed independently) on the CPU or even GPU.

7. Discussion, limitations and future work

We now discuss extensions, limitations and avenues for future research.

Rigid body flow. In our implementation, we have considered only generic cases, i.e., when $\det(\mathbf{J}) \neq 0$ and $\det(\mathbf{K}) \neq 0$. If a singular case was detected while integrating a seed curve, the integration was terminated. In the special case of rigid body flows, every curve is strain free and the problem becomes ill-posed. On the other hand, the singular cases can be easily detected.

Unstructured grids. All the vector fields tested in this paper were known at vertices of very fine structured grids (or analytically). This brings certain simplifications, e.g. when estimating the vector field outside the grid-points. Our implementation, if needed, could be easily extended to accommodate unstructured volumetric meshes as well.

Fields with non-vanishing divergence. We have focused on divergence-free vector fields since this guarantees the existence of strain minimising curves (of first order). Nevertheless, as shown in Fig. 9, such seed curves may exist even in general vector fields and explorations in this direction seem promising.

Unsteady vector flows. For the sake of simplicity, only steady vector fields were considered. The generalisation to the unsteady case is straightforward. This would extend the dimensionality of the space of seed curves from three to four.

Non-local first-order curves. One could consider ‘non-infinitesimally’ arc-length minimising/vanishing seed curves, i.e., curves that preserve their arc-length after a *non-zero* time increment Δt . The problem of finding the intersection of two quadratic cones, as encountered in Section 3.2, would turn into a sphere-ellipsoid intersection problem (induced by the linear map between the tangent spaces at $t = 0$ and $t = \Delta t$), which is




Figure 12: A comparison of first- and second-order curves. The best initial stream surface according to E_S given by first-order (top) and second-order (bottom) curves. (a) The initial first-order curve is constrained to lie on $\partial\Omega$ (top), whereas the second-order curve γ lies inside Ω (bottom). (b–e) Visualisation of the strain energy E_S along the stream surfaces is facilitated by spheres at uniformly sampled points along the seed curve. The spheres’ radii reflect the change of the magnitude of the tangent vectors of the time curves. Observe that second-order curves serve as perfect initial guesses for optimisation. All the spheres are of almost equal radii, which corresponds to E_S being close to zero, whilst the radius varies along the first order curve as is shown in the close-up of top (e). See also the accompanying video.

more demanding. Also, it is not obvious how to set an appropriate Δt .

Area-preserving surfaces. A natural generalisation of our method points towards area-preserving surfaces, i.e., surfaces that propagate in time while preserving the area of any of their sub-patches. As the family of first-order curves is large, it is reasonable to expect that such surfaces exist.

Stability. As a stream surface is determined by its seed curve, the seeding strategy is a crucial ingredient. We have tested random and uniform seeding schemes. Whereas the uniform sampling was applied in most of the examples, followed by adaptive sampling in the neighborhoods with low energy values, in Fig. 5 the test with random sampling of the boundary points and consequently boundary seeding curves is shown to validate the stability of the algorithm.

8. Conclusion

Combining theoretical and numerical techniques, we have attacked the difficult problem (with no known close-form solution) of finding curves that propagate in an arc-length preserving manner in divergence-free vector fields and thus give rise to strain-minimising stream surfaces. Our method successfully solves this problem.

Our contributions are both theoretical and practical. On the one hand, we have advanced the theory of strain-minimising curves and surfaces in divergence-free vector fields. We have shown that such objects always exist. On the other hand, we have illustrated by several examples and the accompanying

video that our theoretical results lead to immediate applications such as vector field exploration and visualisation.

Acknowledgments. We would like to thank Mike Schulze for providing the Turbine [BP02] and the ‘Square cylinder’ [vFWTS08] datasets, and Holger Theisel for his suggestions and comments. The second author were supported by EPSRC through Grant EP/H030115/1.

Appendix A

We now prove Lemma 3.1. By definition, $c_1 = C_1(0)$, where $C_1(t) = \frac{\partial}{\partial t} \int_{s_0}^{s_1} \|\mathbf{S}_s(s, t)\| ds$. We have that

$$\begin{aligned} C_1(t) &= \int_{s_0}^{s_1} \frac{\partial}{\partial t} \|\mathbf{S}_s(s, t)\| ds = \int_{s_0}^{s_1} \frac{\partial}{\partial t} \sqrt{\langle \mathbf{S}_s, \mathbf{S}_s \rangle} ds \\ &= \int_{s_0}^{s_1} \frac{\langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle}{\sqrt{\langle \mathbf{S}_s, \mathbf{S}_s \rangle}} ds. \end{aligned}$$

Consequently, since $\sqrt{\langle \mathbf{S}_s, \mathbf{S}_s \rangle}|_{t=0} \equiv 1$, we can conclude that $c_1 = \int_{s_0}^{s_1} \langle \mathbf{S}_s, \mathbf{S}_{st} \rangle ds|_{t=0}$. Proceeding similarly in the case of c_2 , we have

$$\begin{aligned} C_2(t) &= \frac{\partial^2}{\partial t^2} \int_{s_0}^{s_1} \|\mathbf{S}_s(s, t)\| ds = \int_{s_0}^{s_1} \frac{\partial^2}{\partial t^2} \sqrt{\langle \mathbf{S}_s, \mathbf{S}_s \rangle} ds \\ &= \int_{s_0}^{s_1} \frac{\langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle + \langle \mathbf{S}_s, \mathbf{S}_{stt} \rangle}{\sqrt{\langle \mathbf{S}_s, \mathbf{S}_s \rangle}} - \frac{\langle \mathbf{S}_s, \mathbf{S}_{st} \rangle^2}{\sqrt{\langle \mathbf{S}_s, \mathbf{S}_s \rangle^3}} ds. \end{aligned}$$

Therefore,

$$c_2 = C_2(0) = \int_{s_0}^{s_1} \langle \mathbf{S}_{st}, \mathbf{S}_{st} \rangle + \langle \mathbf{S}_s, \mathbf{S}_{stt} \rangle - \langle \mathbf{S}_s, \mathbf{S}_{st} \rangle^2 ds|_{t=0}$$

as claimed. \square

Appendix B

We now present a geometric approach to the problem of computing the intersection vectors of two quadratic cones given by $\mathbf{dJd}^\top = 0$ and $\mathbf{dKd}^\top = 0$ with apexes at \mathbf{p} . The two cones belong to a pencil of cones given by $\mathbf{d}(\mathbf{J} + \lambda\mathbf{K})\mathbf{d}^\top = 0$ parametrised by λ . We identify the singular quadric in this pencil by setting $\det(\mathbf{J} + \lambda\mathbf{K}) = 0$, which, generically, leads to a cubic equation in λ . Thus, there exists at least one real λ_0 which determines a singular quadric in the family. If the intersection of the two cones is real and non-trivial (i.e., not \mathbf{p} or a cone), the real part of the singular quadric given by $\mathbf{J} + \lambda_0\mathbf{K}$ is a pair of planes (possibly coincident) or a straight line, and is incident with \mathbf{p} . Finally, intersecting this line or the planes with either of the input cones is a simple quadratic problem which gives the sought-after second-order vector(s), provided that they exist.

Appendix C

The integration procedure for first-order interior curves is shown in Fig. 13. Denote h the step-size (set to $\text{diam}(\Omega)/10$ by default). Given a point $\mathbf{p}(t_0) \in \Omega$ and an initial first-order vector $\mathbf{d}(t_0)$, the first-order vector at $t_1 = t_0 + h$ is obtained by projecting $\mathbf{d}(t_0)$ translated to $\mathbf{p}(t_1)$ onto the quadratic cone given by $\mathbf{dJ}(t_1)\mathbf{d}^\top = 0$.




Figure 13: Definition and computation of first-order interior curves. A first-order vector \mathbf{d} at $t = t_0$ (yellow) is projected (green) onto the first-order cone at $t_1 = t_0 + h$.

References

- [BP02] BAUER D., PEIKERT R.: Vortex tracking in scale-space. In *Proceedings of the Symposium on Data Visualisation* (2002), Eurographics Association, pp. 233–242.
- [BSK*13] BARTOŇ M., SHI L., KILIAN M., WALLNER J., POTTMANN H.: Circular arc snakes and kinematic surface generation. *Computer Graphics Forum* 32, 2 (2013), 1–10.
- [CSBI05] CAMARRI S., SALVETTI M.-V., BUFFONI M., IOLLO A.: Simulation of the three-dimensional flow around a square cylinder between parallel walls at moderate Reynolds numbers. In *XVII Congresso di Meccanica Teor. ed Applic.* (2005).
- [Dav67] DAVIS H.: *Introduction to vector field analysis*. Allyn and Bacon, Inc., Boston, 1967.
- [ELM*12] EDMUNDS M., LARAMEE R., MALKI R., MASTERS I., CROFT T., CHEN G., ZHANG E.: Automatic stream surface seeding: A feature centered approach. *Comp. Graph. Forum* 31 (2012), 1095–1104.
- [FHK02] FARIN G., HOSCHEK J., KIM M.-S.: *Handbook of Computer Aided Geometric Design*. Elsevier, Amsterdam, 2002.
- [HPP05] HILDEBRANDT K., POLTHIER K., PREUSS E.: Evolution of 3D curves under strict spatial constraints. In *CAD/Graphics* (2005), IEEE.
- [Hul92] HULTQUIST J. P. M.: Constructing stream surfaces in steady 3D vector fields. *IEEE Transactions on Visualization and Computer Graphics* (1992), 171–178.
- [Int] International CFD Database. URL: <http://cfd.cineca.it/>.
- [LZF10] LENTINE M., ZHENG W., FEDKIW R.: A novel algorithm for incompressible flow using only a coarse grid projection. *ACM SIGGRAPH '10*, pp. 114:1–114:9.
- [MLP*10] MCLOUGHLIN T., LARAMEE R. S., PEIKERT R., POST F. H., CHEN M.: Over Two Decades of Integration-Based, Geometric Flow Visualization. *Computer Graphics Forum* 29, 6 (2010), 1807–1829.
- [MSRT13] MARTINEZ ESTURO J., SCHULZE M., RÖSSL C., THEISEL H.: Global selection of stream surfaces. *Computer Graphics Forum* 32, 2 (2013), 113–122.
- [NaS] NaSt3DGP – A Parallel 3D Flow Solver. URL: <http://wissrech.iam.uni-bonn.de/research/projects/NaSt3DGP/index.htm>.
- [NHM97] NIELSON G., HAGEN H., MULLER H.: *Scientific Visualization*. IEEE Computer Society, 1997.
- [PW01] POTTMANN H., WALLNER J.: *Computational Line Geometry*. Springer, Heidelberg, 2001.
- [SEG*14] SCHULZE M., ESTURO J. M., GÜNTHER T., RÖSSL C., SEIDEL H.-P., WEINKAUF T., THEISEL H.: Sets of globally optimal stream surfaces for flow visualization. *Computer Graphics Forum* 33, 3 (2014), 1–10.
- [VFTS06] VON FUNCK W., THEISEL H., SEIDEL H.-P.: Vector field based shape deformations. *ACM SIGGRAPH '06*, pp. 1118–1125.
- [vFWTS08] VON FUNCK W., WEINKAUF T., THEISEL H., SEIDEL H.-P.: Smoke surfaces: An interactive flow visualization technique inspired by real-world flow experiments. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1396–1403.
- [You10] YOUNES L.: *Shapes and Diffeomorphisms*. Springer, Heidelberg, 2010.