Journal of Symbolic Computation 78 (2017) 3–40

Contents lists available at ScienceDirect

Journal of Symbolic Computation

www.elsevier.com/locate/jsc

Certified computation of planar Morse–Smale complexes

Amit Chattopadhyay [a,b], Gert Vegter [b], Chee K. Yap [c]

[a] International Institute of Information Technology, Bangalore, India
[b] Johann Bernoulli Institute of Mathematics and Computer Science, University of Groningen, Groningen, The Netherlands
[c] Courant Institute of Mathematical Sciences, New York University, New York, USA

ARTICLE INFO

Article history:
Received 29 January 2015
Accepted 22 September 2015
Available online 29 March 2016

Keywords:
Morse–Smale complex
Certified computation
Interval arithmetic

ABSTRACT

The Morse–Smale complex is an important tool for global topological analysis in various problems of computational geometry and topology. Algorithms for Morse–Smale complexes have been presented in case of piecewise linear manifolds (Edelsbrunner et al., 2003a). However, previous research in this field is incomplete in the case of smooth functions. In the current paper we address the following question: Given an arbitrarily complex Morse–Smale system on a planar domain, is it possible to compute its certified (topologically correct) Morse–Smale complex? Towards this, we develop an algorithm using interval arithmetic to compute certified critical points and separatrices forming the Morse–Smale complexes of smooth functions on bounded planar domain. Our algorithm can also compute geometrically close Morse–Smale complexes.

1. Introduction

Geometrical shapes occurring in the real world are often extremely complex. To analyze them, one associates a sufficiently smooth scalar field with the shape, e.g., a density function or a function interpolating gray values. Using this function, topological and geometrical information about the shape may be extracted, e.g., by computing its *Morse–Smale complex*. The cells of this complex are maximal

E-mail addresses: A.Chattopadhyay@iiitb.ac.in (A. Chattopadhyay), G.Vegter@rug.nl (G. Vegter), yap@cs.nyu.edu (C.K. Yap).

http://dx.doi.org/10.1016/j.jsc.2016.03.006
0747-7171/

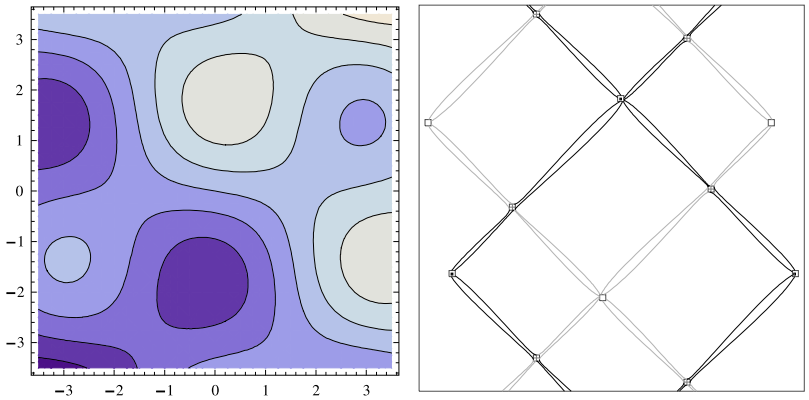

(a) Contour plot (left) and Morse-Smale complex (right) of $h(x, y) = \cos x \ \sin y + 0.2\,(x + y)$ inside box $[-3.5, 3.5] \times [-3.5, 3.5]$. CPU-time: 11 seconds.

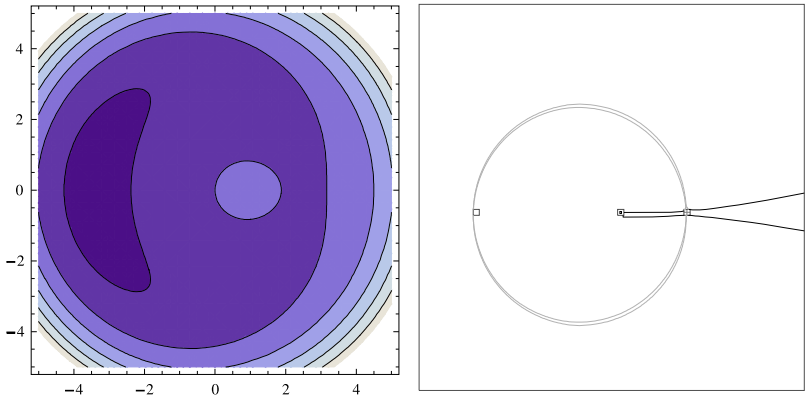

(b) Contour plot (left) and Morse-Smale Complex (right) of $h(x, y) = 10x - \frac{13}{2}(x^2 + y^2) + \frac{1}{3}\,(x^2 + y^2)^2$ inside box $[-5, 5] \times [-5, 5]$. CPU-time: 0.5 seconds.

Fig. 1. Contour plots of Morse–Smale functions, and their Morse–Smale complexes.

connected sets consisting of orthogonal trajectories of the contour lines—curves of steepest ascent—with the same critical point of the function as origin and the same critical point as destination. The leftmost plots in Figs. 1(a) and 1(b) illustrate the level sets of such a density function h, and the rightmost pictures the Morse–Smale complex of h as computed by the algorithm in this paper. This complex reveals the global topology of the shape. Recently, the Morse–Smale complex has been successfully applied in different areas like molecular shape analysis, image analysis, data and vector field simplification, visualization and detection of voids and clusters in galaxy distributions (Cazals et al., 2003; Gyulassy et al., 2008).

1.1. Problem statement

A Morse function $h : \mathbb{R}^2 \to \mathbb{R}$ is a real-valued function with *non-degenerate critical points* (i.e., critical points with non-singular Hessian matrix). Non-degenerate critical points are *isolated*, and are either maxima, or minima, or saddle points. They correspond to singular points of the *gradient vector field* ∇h of h, of type sink, source or saddle, respectively. Regular integral curves of the gradient vector field ∇h are orthogonal trajectories of the regular level curves of h. We are interested in the configuration of integral curves of the gradient vector field. An *unstable (stable) separatrix* of a saddle point is the set of all regular points whose forward (backward) integral curve emerges from the

saddle point. Section 2 contains a more precise definition. A non-degenerate saddle point has two stable and two unstable separatrices. A *Morse–Smale function* is a Morse function whose stable and unstable separatrices are disjoint. In particular, the unstable separatrices flow into a sink, and separate the unstable regions of two sources. Similarly, the stable separatrices emerge from a source, and separate the stable regions of two sinks. The corresponding gradient vector field is called a *Morse–Smale system* (MS-system). The *Morse–Smale complex* (MS-complex for short) is a complex consisting of all singularities, separatrices and the open regions forming their complement, of the *MS-system*. In other words, a cell of the MS-complex is a maximal connected subset of the domain of h consisting of points whose integral curves have the same origin and destination. See also Edelsbrunner et al. (2003a, 2003b), Palis and de Melo (1982) and Section 2. The MS-complex describes the global structure of a Morse–Smale function.

Existing algorithms for MS-complexes (Edelsbrunner et al., 2003a, 2003b) compute the complex of a piecewise linear function on a piecewise linear manifold, or, in other words, of a discrete gradient-like vector field. When h is an analytic function, we cannot use these algorithms without first creating a piecewise linear approximation $\widetilde{h}$. However, the MS-complex of $\widetilde{h}$ is not guaranteed to be combinatorially equivalent to the MS-complex of the smooth vector field. The topological correctness depends on how close the approximation $\widetilde{h}$ is to h. Here "topological correctness" of the computed MS-complex $\widetilde{M}$ means that there is a homeomorphism f of the domain that induces a homeomorphism of each cell $\widetilde{c} \in \widetilde{M}$ to a cell $c \in M$ where M is the true MS-complex, and, moreover, this induced map $\widetilde{c} \mapsto c$ is an isomorphism of $\widetilde{M}$ and M. An isomorphism of two MS-complexes preserves the types of cells and their incidence relations. We can also require f to be an *ε-homeomorphism* for some specified $\varepsilon > 0$, i.e., the distance of a point and its f-image does not exceed ε. As far as we know this problem has never been rigorously studied. Therefore, the main problem of this paper is *to compute a piecewise-linear complex that is ε-homeomorphic to the MS-complex of a smooth Morse–Smale function h*. In short, we seek an exact computation in the sense of the Exact Geometric Computation (EGC) paradigm (Li et al., 2004). Note that it is unclear whether many fundamental problems from analysis are exactly computable in the EGC sense. In particular, the current state-of-the-art in EGC does not (yet) provide a good approach for coping with degenerate situations, and, in fact, this paradigm needs to be extended to incorporate degeneracies. Therefore, we have to *assume* that the gradients we start out with are Morse–Smale systems. However, *generic* gradients are Morse–Smale systems (Palis and de Melo, 1982), so the presence of degenerate singularities and of saddle–saddle connections is exceptional. Note that in restricted contexts, like the class of polynomial functions, absence of degenerate critical points (the first, and local, Morse–Smale condition) can be detected. However, even (most) polynomial gradient systems cannot be integrated explicitly, so absence of saddle–saddle connections (the second, and global, Morse–Smale condition) cannot be detected with current approaches. Detecting such connections even in a restricted context remains a challenging open problem.

1.2. Our contribution

We present an algorithm for computing such a certified approximation of the MS-complex of a given smooth Morse–Smale function on the plane, as illustrated in Figs. 1(a)–(b). In particular, the algorithm produces:

- (arbitrarily small) isolated *certified boxes* each containing a unique saddle, source or sink;
- *certified initial and terminal intervals* (on the boundary of saddleboxes), each of which is guaranteed to contain a unique point corresponding to a stable or unstable separatrix;
- disjoint *certified funnels (strips)* around each separatrix, each of which contains exactly one separatrix and can be as close to the separatrix as desired.

Note. The current article is an extensive elaboration of our previous paper (Chattopadhyay et al., 2012) by incorporating all the theoretical algorithmic details necessary to establish our method of certified Morse–Smale complex computation. The aim in Chattopadhyay et al. (2012) was more on providing an water-tight algorithm; however, the scope of showing all the theoretical details was limited. We

complete that analysis part in the current extensive version. In Section 3, under certified critical-box computation, we provide the details of the relevant lemmas those were missing in Chattopadhyay et al. (2012). In Section 4, we establish rigorous theoretical foundations for refining the saddle-, source- and sink boxes that are used in computing the initial and terminating intervals of the stable and unstable separatrices. All the theoretical results in this section are new additions to the current version. The final method section (Section 5) for the computation of disjoint certified funnels (strips) is now restructured into three subsections, each completes the relevant theoretical and algorithmic analysis.

1.3. Overview

Section 2 starts with a brief review of Morse–Smale systems, their singular points and their invariant manifolds. We also recall the basics of Interval Arithmetic, the computational context which provides us with the necessary certified methods. The construction of the Morse–Smale complex of a gradient system ∇h consists of two main steps: constructing disjoint certified boxes for its singular points, and constructing disjoint certified strips (funnels) enclosing its separatrices. Singular points of the gradient system are computed by solving the system of equations $h_x(x, y) = 0$, $h_y(x, y) = 0$. This is a special instance of the more general problem of solving a generic system of two equations $f(x, y) = 0$, $g(x, y) = 0$. Generic means that the Jacobi matrix at any solution is non-singular, or, geometrically speaking, that the two curves $f = 0$ and $g = 0$ intersect transversally. In our context, this genericity condition reduces to the fact that at singular points of the gradient ∇h the Hessian matrix is non-singular. Section 3 presents a method to compute disjoint isolating boxes for all solutions of such generic systems of two equations in two unknowns. This method yields disjoint isolating boxes for the singular points of the gradient system. In Section 4 these boxes are refined further. Saddle-boxes are augmented with four disjoint intervals in their boundary, one for each intersection of the boundary with the stable and unstable separatrices of the enclosed saddle point. We also show that these intervals can be made arbitrarily small, which is crucial in the second stage of the algorithm. Sink- and source-boxes are refined by computing boxes—not necessarily axis-aligned—around the sink or source on the boundaries of which the gradient system is transversal (pointing into the sink-box and out of the source-box). This implies that all integral curves reaching (emerging from) such a refined sink-box (source-box) lie inside this box beyond (before) the point of intersection.

Section 5 describes the second stage of the algorithm, in which isolating strips (funnels) for the stable and unstable separatrices are constructed. The boundary curves of funnels enclosing an unstable separatrix are polylines with initial point on a saddle box and terminal point on a sink box. The gradient vector field is transversally pointing inward at each point of these polylines. The initial points of the polylines are connected by the unstable interval through which the separatrix leaves its saddle box, and, hence, enters the funnel. The terminal points of these polylines lie on the boundary of the same sink-box. See also Fig. 17. Given this direction of the gradient system on the boundary of the funnel, the unstable separatrix enters the sink-box and tends to the enclosed sink, which is its ω-limit. Although the width of the funnel may grow exponentially in the distance from the saddle-box, this growth is controlled. We exploit the computable (although very conservative) upper bound on this growth rate to obtain funnels that isolate separatrices from each other, and, hence, form a good approximation of the Morse–Smale complex together with the source- and sink-boxes. These upper bounds are also used to prove that the algorithm, which may need several subdivision steps, terminates.

We have implemented this algorithm, using Interval Arithmetic. Section 6 presents sample output of our algorithm. Appendix A contains guaranteed error bounds for the Euler method for solving ordinary differential equations, and Appendix B sketches a method for narrowing the separatrix intervals in the boundaries of the saddle boxes.

1.4. Related work

Milnor (1968) provides a basic set-up for Morse theory. The survey paper (Biasotti et al., 2008), focusing on geometrical-topological properties of real functions, gives an excellent overview of recent works on MS-complexes. Originally, Morse theory was developed for smooth functions on smooth

manifolds. Banchoff (1970) introduced the equivalent definition of critical points on polyhedral surfaces. Many of the recent developments on MS-complexes are based on this definition. A completely different discrete version of Morse theory is provided by Forman (1998).

Different methods for computation. In the literature there are two different methods for computing the Morse–Smale complexes: boundary based approaches and region based approaches. Boundary based methods compute boundaries of the cells of the MS-complex, i.e., the integral curves connecting a saddle to a source, or a saddle to a sink (Takahashi et al., 1995; Bajaj and Schikore, 1998; Edelsbrunner et al., 2003a). On the other hand, watershed algorithms for image segmentation are considered as region based approaches (Meyer, 1994). Edelsbrunner et al. (2003a) compute the Morse–Smale complex of piecewise linear manifolds using a paradigm called *Simulation of Differentiability*. In higher dimensions they give an algorithm for computing Morse Smale complexes of piecewise linear 3-manifolds (Edelsbrunner et al., 2003b).

Morse–Smale complexes have also been applied in shape analysis and data simplification. Computing MS-complexes is strongly related to vector field visualization (Helman and Hesselink, 1991). In a similar context, designing vector fields on surfaces has been studied for many graphics applications (Zhang et al., 2006). Cazals et al. (2003) applied discrete Morse theory to molecular shape analysis.

This paper contributes to the emerging area of Exact Numerical Algorithms for geometric problems (Yap, 2009). Recent algorithms of this genre (e.g., Plantinga and Vegter, 2007; Lin and Yap, 2011) are numerical subdivision algorithms based on interval function evaluation and sign evaluation.

2. Preliminaries

In this section we briefly review the necessary mathematical background on Morse functions, Morse–Smale systems, their singular points and invariant manifolds. We also recall the basics of our computational model and Interval Arithmetic that are necessary for developing our certified computation algorithm.

2.1. Mathematical background

Morse functions. A function $h : \mathcal{D} \subset \mathbb{R}^2 \to \mathbb{R}$ is called a *Morse function* if all its critical points are non-degenerate. The *Morse lemma* (Milnor, 1968) states that near a non-degenerate critical point a it is possible to choose local coordinates x, y in which h is expressed as $h(x, y) = h(a) \pm x^2 \pm y^2$. Existence of these local coordinates implies that non-degenerate critical points are isolated. The number of minus signs is called the *index* $i_h(a)$ of h at a. Thus a two variable Morse function has three types of non-degenerate critical points: minima (index 0), saddles (index 1) and maxima (index 2).

Integral curves. An integral curve $\mathbf{x} : I \subset \mathbb{R} \to \mathcal{D}$ passing through a point p_0 on $\mathcal{D}$ is a unique maximal curve satisfying: $\dot{\mathbf{x}}(t) = \nabla h(\mathbf{x}(t))$, $\mathbf{x}(0) = p_0$, for all t in the *interval* I. Integral curves corresponding to the gradient vector field of a smooth function $h : \mathcal{D} \to \mathbb{R}$ have the following properties:

1. Two integral curves are either disjoint or same.
2. The integral curves cover all the points of $\mathcal{D}$.
3. The integral curves of the gradient vector field of h form a partition of $\mathcal{D}$.
4. The integral curve $\mathbf{x}(t)$ through a critical point p_0 of h is the constant curve $\mathbf{x}(t) = p_0$.
5. The integral curve $\mathbf{x}(t)$ through a regular point p of h is injective, and if $\lim_{t\to\infty} \mathbf{x}(t)$ or $\lim_{t\to-\infty} \mathbf{x}(t)$ exists, it is a critical point of h. This implies integral curves corresponding to gradient vector field are never closed curves.
6. The function h is strictly increasing along the integral curve of a regular point of h.
7. Integral curves are perpendicular to the regular level sets of h.

Stable and unstable manifolds. Consider the integral curve $\mathbf{x}(t)$ passing through a point p. If the limit $\lim_{t\to\infty} \mathbf{x}(t)$ exists, it is called the ω-limit of p and is denoted by $\omega(p)$. Similarly, $\lim_{t\to-\infty} \mathbf{x}(t)$ is called

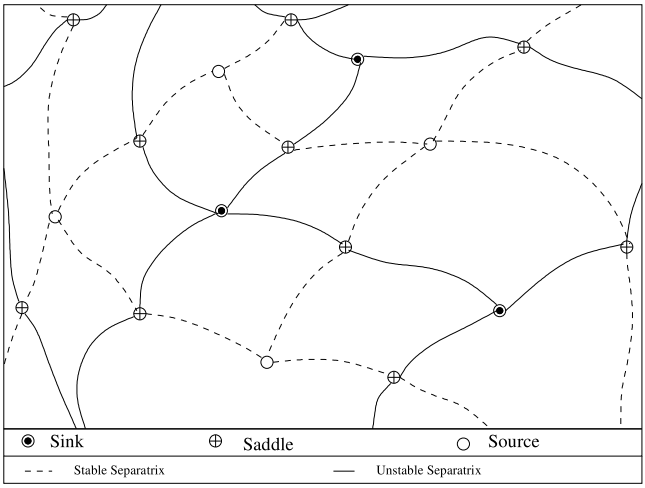

Fig. 2. Morse–Smale complex.

the α-limit of p and is denoted by $\alpha(p)$—again provided this limit exists. The stable manifold of a singular point p is the set $W^s(p) = \{q \in \mathcal{D} \mid \omega(q) = p\}$. Similarly, the unstable manifold of a singular point p is the set $W^u(p) = \{q \in \mathcal{D} \mid \alpha(q) = p\}$. Here we note that both $W^s(p)$ and $W^u(p)$ contain the singular point p itself (Hirsch and Smale, 1974).

Now, the stable and unstable manifolds of a saddle point are 1-dimensional manifolds. A stable manifold of a saddle point consists of two integral curves converging to the saddle point. Each of these integral curves (not including the saddle point) are called the stable separatrices of the saddle point. Similarly, an unstable manifold of a saddle point consists of two integral curves diverging from the saddle point and each of these integral curves (not including the saddle point) are called the unstable separatrices of the saddle point.

The Morse–Smale complex. (See Fig. 2.) A Morse function on $\mathcal{D}$ is called a Morse–Smale (MS) function if its stable and unstable separatrices are disjoint. In particular, a Morse–Smale function on a two-dimensional domain has no integral curve connecting two saddle points, since in that case a stable separatrix of one of the saddle points would coincide with an unstable separatrix of the other saddle point. The MS-complex associated with a MS-function h on $\mathcal{D}$ is the subdivision of $\mathcal{D}$ formed by the connected components of the intersections $W^s(p) \cap W^u(q)$, where p, q range over all singular points of h. If $\mathcal{D} = \mathbb{R}^2$, then, according to the Quadrangle Lemma (Edelsbrunner et al., 2003a), each region of the MS-complex is a quadrangle with vertices of index 0, 1, 2, 1, in this order around the region.

Stability of equilibrium points. We note that a gradient vector field of a MS-function $h : \mathcal{D} \to \mathbb{R}$ can have three kinds of equilibria or *singular points*, namely, sinks (corresponding to maxima of h), saddles (saddles of h) and sources (corresponding to minima of h). These singular points can be distinguished based on the local behavior of the integral curves around those points. Locally, a sink has a neighborhood, which is a stable 2-manifold. Similarly, locally a source has a neighborhood, which is an unstable 2-manifold. Locally, a saddle has a stable 1-manifold and an unstable 1-manifold crossing each other at the saddle point. A sink is called a stable equilibrium point, whereas a source or a saddle is called unstable equilibrium point. We note that, a source corresponding to a MS-function h is a sink corresponding to the function $-h$.

2.2. Computational model

Our computational model has two simple foundations: (1) BigFloat packages and (2) interval arithmetic (IA) (Moore, 1996). Together, these are able to deliver efficient and guaranteed results in the

implementation of our algorithms. A BigFloat package is a software implementation of exact ring $(+, -, \times)$ operations, division by 2, and exact comparisons, over the set $\mathbb{F} = \{m2^n : m, n \in \mathbb{Z}\}$ of *dyadic numbers*. In practice, we can use IEEE machine arithmetic as a filter for BigFloat computation to speed up many computations. *Range functions* form a basic tool of IA: given any function $F : \mathbb{R}^m \to \mathbb{R}^n$, a *range function* $\square F$ for F computes for each m-dimensional interval I (i.e., an m-box) an n-dimensional interval $\square F(I)$, such that $F(I) \subset \square F(I)$. A range function is said to be *convergent* if the diameter of the output interval converges to 0 when the diameter of the input interval shrinks to 0. Convergent range functions exist for the basic operators and functions, so all range functions are assumed to be convergent. Moreover, we assume that the sign of functions can be evaluated exactly at dyadic numbers. All our boxes are *dyadic boxes*, meaning that their corners have dyadic coordinates.

Interval implicit function theorem. To introduce a useful tool from IA, we recall some notation for *interval matrices*. An $n \times n$ interval matrix $\square M$ is defined as

$$\square M = \{M | M_{ij} \in \square M_{ij}, i, j \in \{1 \dots n\}\}.$$

Also note that we write

$$0 \notin \det \square M$$

if there exists no matrix $M \in \square M$ such that $\det M = 0$ where det represents the determinant of the corresponding matrix.

If $I = I_x \times I_y$ is a 2D-interval (box) in $\mathbb{R}^2$, the *interval Jacobian determinant* $\square \frac{\partial(f,g)}{\partial(x,y)}$ is the 2×2 interval determinant given by

$$\square \frac{\partial(f,g)}{\partial(x,y)}(I) = \begin{vmatrix} \square \frac{\partial f}{\partial x}(I) & \square \frac{\partial f}{\partial y}(I) \\ \square \frac{\partial g}{\partial x}(I) & \square \frac{\partial g}{\partial y}(I) \end{vmatrix}.$$

Proposition 2.1 *(Interval Implicit Function Theorem, Snyder, 1992a, 1992b). Let $F : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ be a C^1-map with components f and g. If $I \subset \mathbb{R} \times \mathbb{R}$ is a box for which $0 \notin \square \frac{\partial(f,g)}{\partial(x,y)}(I)$, then the system $f(x, y) = 0$, $g(x, y) = 0$ has at most one solution in I.*

Remark 2.1. In this paper, the domain $\mathbb{D}$ of h is a finite union of axis-aligned dyadic boxes. Furthermore, we (have to) assume that all stable and unstable separatrices of the saddle points are transversal to the boundary. Computationally this means that any sufficiently close approximation of these separatrices is transversal to the boundary as well.

3. Isolating boxes for singularities of gradient fields

As a first step towards the construction of the Morse–Smale complex of h we construct disjoint isolating boxes for the singular points of ∇h. To this end, we first show how to compute isolating boxes for the solutions of a generic system of two equations in two unknowns, which are confined to a bounded domain in the plane. This domain is a finite union of dyadic boxes. Applying this general method to the case in which the two equations are defined by the components of the gradient vector field ∇h we obtain isolating boxes for the singularities of this gradient field.

3.1. Certified solutions of systems of equations

We consider a system of equations

$$f(x, y) = 0, \quad g(x, y) = 0, \tag{1}$$

where f and g are C^1-functions defined on a bounded axisparallel box $\mathbb{D} \subset \mathbb{R}^2$ with *dyadic vertices*. Furthermore, we assume that the system has only non-degenerate solutions, i.e., the Jacobian determinant is non-zero at a solution. In other words,

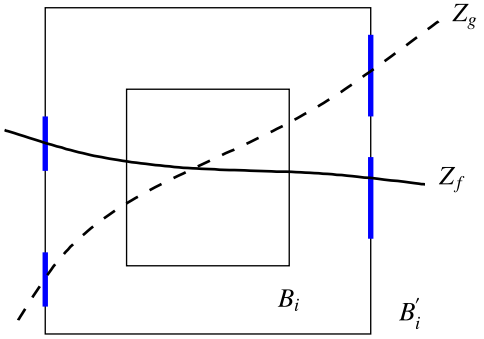

Fig. 3. A certified box-pair (B_i, B_i'), with isolating intervals in its boundary for $Z_f \cap \partial B_i'$ and $Z_g \cap \partial B_i'$. The curves Z_f and Z_g intersect inside B_i' iff these intervals are interleaved.

$$\left.\frac{\partial(f,g)}{\partial(x,y)}\right|_{(x_0,y_0)} \neq 0 \tag{2}$$

for (x_0, y_0) satisfying (1). Geometrically, this means that the curves given by $f(x, y) = 0$ and $g(x, y) = 0$ are regular near a point of intersection (x_0, y_0) and the intersection is transversal. We will denote these curves by Z_f and Z_g, respectively. Note that this condition is satisfied by Morse–Smale systems, since in that case $f = h_x$ and $g = h_y$, so the Jacobian determinant is precisely the Hessian determinant.

Since the domain $\mathbb{D}$ of f and g is compact and we assume (2), the system (1) has finitely many solutions in $\mathbb{D}$. Our goal is to construct a collection of axis-aligned boxes $B_1, \ldots, B_m$ and $B_1', \ldots, B_m'$ such that (i) box B_i is concentric with and strictly contained in B_i', (ii) the boxes B_i' are disjoint, (iii) each solution of (1) is contained in one of the boxes B_i, and (iv) each box B_i' contains exactly one solution (contained inside the enclosed box B_i). The box pair (B_i, B_i') is *certified*: B_i contains a solution, and B_i' provides positive clearance to other solutions. In fact, the sequence of boxes will satisfy the following stronger conditions; See also Fig. 3.

1. The curves Z_f and Z_g each intersect the boundary of B_i' transversally in two points.
2. There are disjoint intervals $I_i^0(f)$, $I_i^0(g)$, $I_i^1(f)$ and $I_i^1(g)$ in the boundary of B_i' (in this order), where the first and the third interval each contain one point of intersection of Z_f and $\partial B_i'$, and the second and fourth interval each contain one point of intersection of Z_g and $\partial B_i'$.
3. The (interval) Jacobian determinant of f and g does not vanish on B_i', i.e.,

$$0 \notin \square \frac{\partial(f,g)}{\partial(x,y)}(B_i').$$

3.2. Construction of certified box pairs

We first subdivide the domain $\mathbb{D}$, using a quadtree, until all leaves I satisfy certain conditions to be introduced now. This quadtree based subdivision is shown is Fig. 7. For a (square, axis-aligned) box I, let $N_\varrho(I)$ be the box obtained by multiplying box I from its center by a factor of $1 + \varrho$, where $\frac{1}{2} \leq \varrho \leq 1$. In other words, if a is the length of an edge of I, then $N_\varrho(I)$ is the box with the same center as I and the distance of each edge of $N_\varrho(I)$ from its center is $a(1 + \varrho)$. We also denote $N_1(I)$ by $N(I)$; this is the box formed by the union of I and its eight neighbor grid-boxes. We shall call $N(I)$ the surrounding box of I. The algorithm subdivides $\mathbb{D}$ until all grid-boxes (leaves of the quadtree) I satisfy $C_0(I) \vee \big(C_1(I) \wedge C_2(I)\big)$, where the clauses $C_i(I)$, $i = 0, 1, 2$, are the following predicates:

$$C_0(I): \quad 0 \notin \square f(I) \vee 0 \notin \square g(I)$$

$$C_1(I): \quad 0 \notin \square \frac{\partial(f,g)}{\partial(x,y)}(N(I))$$

$$C_2(I): \quad C_2(I,f) \wedge C_2(I,g),$$

where

$$C_2(I,f) = \langle \square \frac{\nabla f}{\| \nabla f \|}(N(I)), \square \frac{\nabla f}{\| \nabla f \|}(N(I)) \rangle \geq \cos \frac{\pi}{30}.$$

If $C_0(I)$ holds, box I does not contain a solution, so it is discarded. The second predicate guarantees that $N(I)$ contains at most one solution. This is a consequence of Interval Implicit Function Theorem (Snyder, 1992c, 1992d). Condition $C_2(I)$ is a *small angle variation condition*, guaranteeing that the variation of the unit normals of the curves $Z_f \cap N(I)$ and $Z_g \cap N(I)$ do not vary too much, so these curves are regular, and even 'nearly linear' (the unit normal of Z_f is the normalized gradient of f). Here $\langle \cdot, \cdot \rangle$ denotes the interval version of the standard inner product on $\mathbb{R}^2$. The lower bound for the angle variation of ∇f is generated by the proof of Lemma 3.2.

Remark 3.1. Condition $C_1(I)$ implies that there is a computable positive lower bound $\alpha(I)$ on the angle between $\nabla f(p)$ and $\nabla g(q)$ where p, q range over the surrounding box of I. More precisely, to compute $\alpha(I)$, we first compute a lower bound L on the quantity $\frac{\partial(f,g)}{\partial(x,y)} \frac{1}{\|\nabla f\| \cdot \|\nabla g\|}$. This L may be obtained by an interval evaluation of this quantity at $N(I)$; note that $L > 0$ iff condition $C_1(I)$ holds. We define $\alpha(I)$ as $\arcsin(L)$.

Our algorithm will construct disjoint certified boxes surrounding a box I. As observed earlier, the surrounding boxes $N(I)$ and $N(J)$ of disjoint boxes I and J may intersect. Since our algorithm will construct disjoint certified boxes surrounding a box I, its surrounding box should be smaller than the box $N(I)$. To achieve this, note that $N(I) = N_1(I)$. Now if $N(J) \cap I = \emptyset$, then $N_{1/2}(I)$ and $N_{1/2}(J)$ have disjoint interiors. This is a key observation with regard to the correctness of our algorithm.

Lemma 3.2. *Let I be a box such that conditions $\neg C_0(I)$, $C_1(I)$, and $C_2(I)$ hold. Let d be the length of its edges, and let $\frac{1}{2} \leq \varrho \leq 1$.*

1. *If Z_f intersects I, it intersects the boundary of $N_\varrho(I)$ transversally at exactly two points. At a point of intersection of Z_f and an edge e of $\partial N_\varrho(I)$ the angle between Z_f and e is at least $\frac{1}{15}\pi$.*
2. *If I contains a point of intersection of Z_f and Z_g, then the points of intersection of Z_f and $\partial N_\varrho(I)$ are at distance at least $2\varrho d \tan \frac{1}{2}\alpha(I)$ from the points of intersection of Z_g and $\partial N_\varrho(I)$. On $\partial N_\varrho(I)$, the points of intersection with Z_f and with Z_g are alternating.*

Proof. 1. Assume that Z_f intersects a vertical edge of $\partial N_\varrho(I)$ at $\overline{q}$. Let $\overline{p}$ be a point on $Z_f \cap I$. See Fig. 4. Then there is a point s on the curve segment $\overline{pq}$ at which the gradient of f is perpendicular to this line segment $\overline{pq}$. Let β be the (smallest) angle between $\nabla f(s)$ and the horizontal direction. Referring to the rightmost picture in Fig. 4, we see that this angle is not less than $\angle pqr = \frac{1}{2}\pi - \angle qpr$. Since $\| q - p \| = d\sqrt{1 + 2\varrho + 2\varrho^2}$, it follows that

$$\beta \geq \frac{\pi}{2} - \arccos \frac{\varrho}{\sqrt{1 + 2\varrho + 2\varrho^2}} > \frac{\pi}{10},$$

where the last inequality holds since $\frac{1}{2} \leq \varrho \leq 1$. Since condition $C_2(N(I), f)$ holds, the angle between the gradients of f at s and $\overline{q}$ is less than $\frac{\pi}{30}$. Therefore, the angle between Z_f and the vertical edge of $\partial N_\varrho(I)$ at $\overline{q}$ is at least $\frac{\pi}{10} - \frac{\pi}{30} = \frac{\pi}{15}$.

2. Let $p \in I$ be the point of intersection of Z_f and Z_g. Suppose Z_f and Z_g intersect the vertical edge e of $\partial N_\varrho(I)$ in q and r, respectively. Then there is a point s on Z_f between p and q at which the gradient of f is perpendicular to pq, and there is a point on Z_g between p and r at which the gradient of g is perpendicular to pr. Let $\alpha(I)$ be the lower bound on the angle between $\nabla f(p)$ and $\nabla g(q)$ where p, q range over the box I. See Fig. 5, where $\overline{\alpha} \geq \alpha(I)$. For fixed $\overline{\alpha}$, the distance between q and r is minimal if the projection $\overline{p}$ of p on the edge e is the midpoint of qr, in which

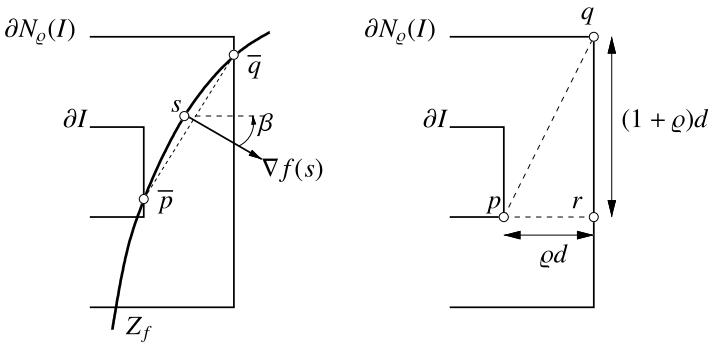

Fig. 4. Lower bound on angle of intersection of Z_f and the boundary of the surrounding box $N_\varrho(I)$ in case Z_f intersects I.

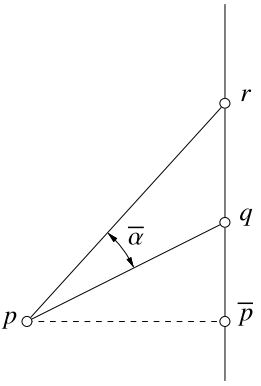

Fig. 5. A lower bound for the distance between points of Z_f and Z_g on the boundary of $N_\varrho(I)$.

case $\| r - q \| = 2 \| p - \overline{p} \| \tan \frac{1}{2}\overline{\alpha}$. Since $\| p - \overline{p} \| \geq \varrho d$ and $\overline{\alpha} \geq \alpha(I)$, the distance between q and r is at least $2\varrho d \tan \frac{1}{2}\alpha(I)$. □

The following result gives an estimate on the position of the points at which Z_f intersects the boundary of the surrounding box of I in case Z_f intersects an edge of ∂I in at least two points. For an edge e of the inner box I let l and r be the points of intersection of the line through e and the edges of the surrounding box $N_\varrho(I)$, perpendicular to e. See also Fig. 6. The dyadic intervals on the boundary of this surrounding box with length at most $\frac{2}{\sqrt{3}}(1+\varrho)d$, centered at l and r, respectively, are denoted by $L_\varrho(e)$ and $R_\varrho(e)$, where d is the length of the edges of I, and $\frac{1}{2} \leq \varrho \leq 1$. Here, we note that in the following Lemma 3.3 the length $\frac{2}{\sqrt{3}}(1+\varrho)d$ might be irrational but it should be noted that any dyadic lower bound will suffice for the widths of $L_\varrho(e)$ and $R_\varrho(e)$.

Lemma 3.3. *Let I be a box such that $\neg C_0(I)$, $C_1(I)$ and $C_2(I)$ hold, and let e be one of its edges. Let $\frac{1}{2} \leq \varrho \leq 1$.*

1. *If Z_f intersects an edge e of the boundary of I in at least two points, then it transversally intersects $\partial N_\varrho(I)$ in exactly two points, one in each of the dyadic intervals $L_\varrho(e)$ and $R_\varrho(e)$.*
2. *If Z_f intersects $\partial N_\varrho(I)$ in the dyadic intervals $L_\varrho(e)$ and $R_\varrho(e)$, then these intersections are transversal, and Z_f intersects $\partial N_\varrho(I)$ at exactly two points, one in each of these intervals.*

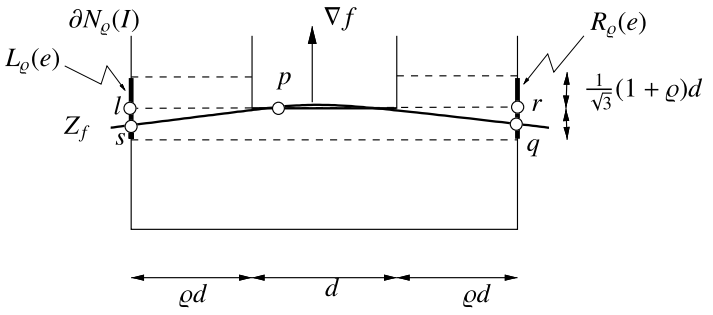

Fig. 6. Intervals containing points $Z_f \cap \partial N_\varrho(I)$.

Proof. 1. There is a point on Z_f between the points of intersection with e at which the gradient of f is perpendicular to e. See Fig. 6. The small normal variation condition $C_2(N(I), f)$ implies that Z_f does not intersect any of the two edges of $\partial N_\varrho(I)$ parallel to e, and that it intersects each of the edges of $\partial N_\varrho(I)$ perpendicular to e transversally in exactly one point. Let q be the point of intersection with the edge containing r. Then there is a point on Z_f at which the gradient of f is perpendicular to the line segment pq. Since the angle between the gradients of f at two points of $N_\varrho(I)$ does not differ by more than $\frac{1}{30}\pi$, we have $\angle qpr < \frac{1}{30}\pi$. Therefore,

$$\| q - r \| = \| p - r \| \tan \angle qpr < (1+\varrho)d \tan \tfrac{\pi}{30} < \tfrac{1}{\sqrt{3}}(1+\varrho)d.$$

In other words, Z_f intersects R_ϱ. Similarly, Z_f intersects L_ϱ. The small normal variation condition $C_2(N(I))$ implies that there are no other intersections with the edges of $\partial N_\varrho(I)$.

2. Let the points of intersection of Z_f and $R_\varrho(e)$ and $L_\varrho(e)$ be q and s, respectively. Then there is a point on Z_f at which ∇f is perpendicular to qs. Since the angle of qs and the vertical direction is at most

$$\arctan \frac{(1+\varrho)/8}{1+\varrho} = \arctan \frac{1}{8} < \frac{\pi}{25},$$

it follows from the small normal variation condition $C_2(I)$ that the gradient of f at any point of $N_\varrho(I)$ makes an angle of at most $\frac{\pi}{25} + \frac{\pi}{15} < \frac{\pi}{10}$ with the vertical direction. This rules out multiple intersections with the vertical edges of $\partial N_\varrho(I)$. It also implies that Z_f lies above the polyline qms, where m is the intersection of the line through q with slope $\tan \frac{\pi}{10}$ and the line through s with slope $-\tan \frac{\pi}{10}$. Therefore, all points of Z_f lie at distance at most $\frac{1}{4}(1+\varrho)d + \frac{1}{2}(1+\varrho)d \tan \frac{\pi}{10} < \varrho d$ from the line through e, so Z_f does not intersect the edges of $\partial N_\varrho(I)$ parallel to e. □

3.3. Towards an algorithm

After the first subdivision step, we have constructed a finite set $\mathcal{B}$ of boxes, all of the same size, such that $\neg C_0(I) \wedge (C_1(N(I)) \wedge C_2(N(I)))$ holds for each box I in $\mathcal{B}$. Each such box is a possible candidate for containing a solution and are shown by the black boxes in Fig. 7 (called, candidate critical box). Rest of the grid-boxes in $\mathbb{D}$ do not contain any solution and are discarded. For each such candidate critical box I the algorithm calls one of the following:

- DISCARD(I), if it decides that I does not contain a solution. It marks box I as processed.
- REPORTSOLUTION(I). It returns the certified pair $(N_{1/2}(I), N(I))$, and marks all boxes contained in $N(I)$ as processed.

In the latter case a solution is found inside $N_{1/2}(I)$, but, as will become clear later, it may not be contained in the smaller box I. In view of $C_1(I)$ none of the grid-boxes in $N(I)$ contain a solution different from the one reported, so they are marked as being processed.

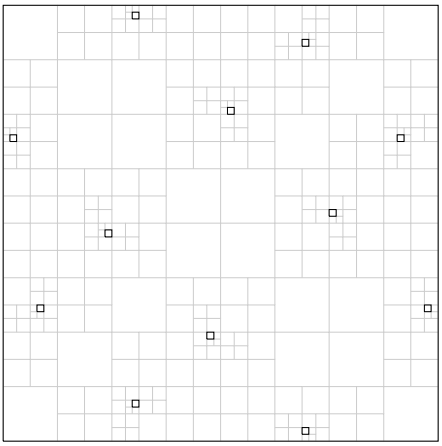

Fig. 7. A quadtree-based domain subdivision (in gray) corresponding to the function in Fig. 1(a). The candidate critical boxes (in black) satisfy $\neg C_0(I) \wedge C_1(I) \wedge C_2(I)$.

Decisions are based on evaluation of the signs of f and g at the vertices of the grid-boxes (or at certain dyadic points on edges of grid-boxes). An edge of a box is called *bichromatic* (*monochromatic*) for f if the signs of the value of f at its vertices are opposite (equal).

Algorithm, case 0: degenerate case when f and g both vanish at a vertex of I The algorithm calls ReportSolution(I).

Algorithm, case 1: I has a bichromatic edge for f and a bichromatic edge for g Then Z_f and Z_g intersect I, and, according to Lemma 3.2, part 1, both curves intersect the boundary of $N_{1/2}(I)$ transversally in exactly two points. For each of the two points in $\partial N_{1/2}(I)$ the algorithm computes an isolating interval—called an f-interval—on $\partial N_{1/2}(I)$ of length $\frac{1}{2}d\tan\frac{1}{2}\alpha(I)$. The two g-intervals are computed similarly. If the f- and g-intervals are not interleaving, there is no solution of (1) in box I—even though there may be a solution in $N_{1/2}(I)$—and Discard(I) is called. This follows from Lemma 3.2, part 2. If the intervals are interleaving, then there is a point of intersection inside $N_{1/2}(I)$, so the algorithm calls ReportSolution(I).

Algorithm, case 2: I contains no bichromatic edge for f (g), and at least one bichromatic edge for g (f, respectively) We only consider the case in which all edges of I are monochromatic for f. Then the algorithm also evaluates the sign of f at the vertices of the box $N_{1/2}(I)$. If $N_{1/2}(I)$ has no *disjoint* bichromatic edges (as in the fourth and fifth configuration of Fig. 8), the isocurve Z_f does not intersect I, so the algorithm calls Discard(I). To deal with the remaining case, in which $N_{1/2}(I)$ has two *disjoint* bichromatic edges (as in the sixth configuration in Fig. 8) we need to evaluate the sign of f at certain dyadic points of these bichromatic edges, followed from Lemma 3.3. By evaluating the signs of f at the endpoints of the interval $L_\varrho(e)$ and $R_\varrho(e)$ the algorithm decides whether they contain a point of intersection with Z_f. As noted before, for exact computation we consider suitable dyadic intervals corresponding to $L_\varrho(e)$ and $R_\varrho(e)$. If at least one of these intervals is disjoint from Z_f, then Discard(I) is called. Otherwise, the algorithm computes isolating f- and g-intervals of length $\frac{1}{2}d\tan\frac{1}{2}\alpha(I)$. As in case 1, the algorithm calls ReportSolution(I) if these intervals are interleaving, and Discard(I) otherwise.

Algorithm, case 3: all edges of I are monochromatic for both f and g Again, let e be the (unique) edge of I closest to the edge of $N_{1/2}(I)$ which is monochromatic for f, at whose vertices the sign of f is the opposite of the sign of f at the vertices of I. Edge e' of I is defined similarly for g.
Case 3.1: $e = e'$. In this case Z_f or Z_g does not intersect I. Indeed, if Z_f intersects I, it intersects e in at least two points, so there is a point $p \in Z_f$ at which $\nabla f(p)$ is perpendicular to e. Condition $C_1(I)$

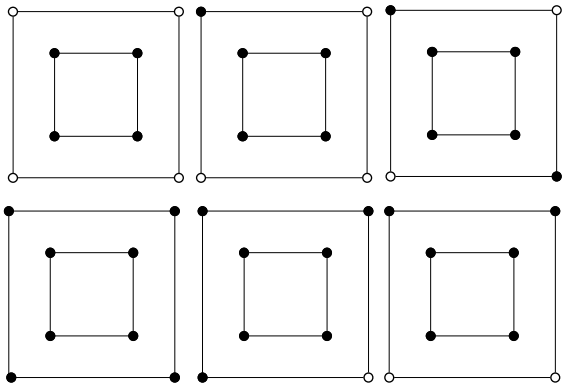

Fig. 8. Sign patterns of the box $N_{1/2}(I)$ enclosing the grid-box I with monochromatic edges for f. The three top configurations are ruled out by the small normal variation condition $C_2(I)$. The fourth, fifth and sixth configuration are all possible, but only in the sixth situation Z_f may intersect the inner box.

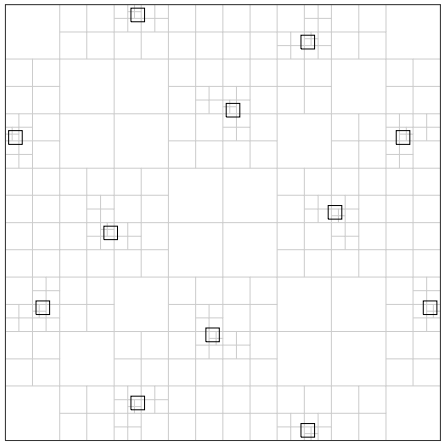

Fig. 9. The certified critical boxes (in black) and the subdivide domain (in grey) corresponding to the function in Fig. 1(a).

guarantees that ∇g is nowhere parallel to $\nabla f(p)$, so Z_g does not intersect e, and, hence, does not intersect I. Therefore, DISCARD(I) is called.

Case 3.2: $e \neq e'$. If Z_f does not intersect the dyadic interval $L_\varrho(e)$ or $R_\varrho(e)$, or if Z_g does not intersect the dyadic interval $L_\varrho(e')$ or $R_\varrho(e')$, then, as in case 2, the algorithm calls DISCARD(I). Otherwise, $L_\varrho(e)$ or $R_\varrho(e)$ are isolating f-intervals which are disjoint from the isolating g-intervals $L_\varrho(e')$ or $R_\varrho(e')$. If e and e' are perpendicular, then these f- and g-intervals are interleaving, and, hence, REPORTSOLUTION(I) is called. Otherwise, there is no solution in I, so DISCARD(I) is called.

Refinement: disjoint surrounding boxes We would like distinct isolating boxes I, J to have disjoint surrounding boxes $N(I)$, $N(J)$. There is a simple way to ensure this: we just use the predicate $C_1(N(I))$ to instead of $C_1(I)$ in the above subdivision process. Then, if the interior of $N(I) \cap N(J)$ is non-empty, we can discard any one of I or J. Fig. 9 shows an example of certified critical boxes for the function in Fig. 1(a).

4. Isolating boxes for sinks, sources and saddles

In a first step, described in Section 3, we have constructed certified disjoint isolating boxes $B'_1 \ldots, B'_m$ the singular points of ∇h in the domain $\mathbb{D}$ of h. Let $\mathbb{D}^*$ be the closure of $\mathbb{D} \setminus (B'_1 \cup \cdots \cup B'_m)$. Obviously, $\mathbb{D}^*$ is a compact subset of $\mathbb{R}^2$.

In a second step towards the construction of the MS-complex, we refine the saddle-, sink- and sourceboxes. Fig. 13 shows an example of such refined critical boxes after detecting their types for the function in Fig. 1(a). In Section 4.1 we show how to augment each saddlebox by computing four arbitrarily small disjoint intervals in its boundary, one for each intersection of a stable or unstable separatrix with the box boundary. Subsequently, in Section 4.2, we show how to construct for each source or sink of ∇h (minimum or maximum of h) a box on the boundary of which the gradient field is pointing outward or inward, respectively. These boxes are contained in the source- and sinkboxes constructed in the previous section, but are not necessarily axis-aligned.

4.1. Refining saddle boxes: isolating separatrix intervals

To compute *disjoint* certified separatrix intervals we consider wedge shaped regions with apex at the saddle point, enclosing the unstable and stable manifolds of the saddle point. Even though the saddle point is not known exactly, we will show how to determine certified intervals for the intersection of these wedges and the boundary of a saddle box.

First we determine the eigenvalues and eigenvectors of the Hessian of h at a point (x_0, y_0) in the interior of the saddlebox I—its center point, say—and consider these as good approximations to the eigenvalues and eigenvectors of the Hessian, i.e., the linear part of ∇h, at the saddle point. Let H be the Hessian, i.e.,

$$H = \begin{pmatrix} h_{xx} & h_{xy} \\ h_{xy} & h_{yy} \end{pmatrix}, \tag{3}$$

and let H^0 be the Hessian evaluated at (x_0, y_0). The eigenvalues λ_u and λ_s of H^0 are given by

$$\lambda_u = \tfrac{1}{2}\,(h^0_{xx} + h^0_{yy} + \sqrt{(h^0_{xx} - h^0_{yy})^2 + 4(h^0_{xy})^2})$$

$$\lambda_s = \tfrac{1}{2}\,(h^0_{xx} + h^0_{yy} - \sqrt{(h^0_{xx} - h^0_{yy})^2 + 4(h^0_{xy})^2}),$$

and the corresponding eigenvectors are

$$V^u = \begin{pmatrix} h^0_{xy} \\ \lambda_u - h^0_{xx} \end{pmatrix}, \qquad V^s = \begin{pmatrix} h^0_{xx} - \lambda_u \\ h^0_{xy} \end{pmatrix}. \tag{4}$$

The singular point is a saddle, so we have $\lambda_s < 0 < \lambda_u$. Since H^0 is a symmetric matrix, its eigenvectors are orthogonal. More precisely,

$$V^s = \begin{pmatrix} -V^u_2 \\ V^u_1 \end{pmatrix} = R_{\pi/2}(V^u).$$

Here R_α denotes counterclockwise rotation over an angle α. Therefore,

$$\| V^s \| = \| V^u \|, \text{ and } \det(V^u, V^s) = \| V^u \|^2. \tag{5}$$

The stable and unstable eigenvectors V^s and V^u are good approximations of the tangent vectors of the stable and unstable manifolds of the saddle point. These invariant manifolds are contained in wedge-shaped regions, which are defined as follows.

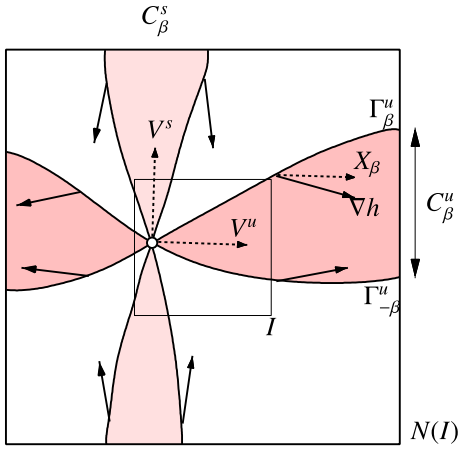

Fig. 10. The unstable wedge C^u_β enclosing the unstable separatrix, and the stable wedge C^s_β enclosing the stable separatrix. The gradient vector field ∇h, represented by solid arrows, is transversally pointing inward along the boundary of the unstable wedge, and outward along the boundary of the stable wedge. At points of the unstable wedge boundary $\Gamma^u_\beta \cup \Gamma^u_{-\beta}$ the vector field X_β is parallel to V^u or $-V^u$, so ∇h makes an angle $-\beta$ with V^u or $-V^u$ there.

Definition 4.1. *Let the (orthogonal) vectors V^u and V^s be the stable and unstable eigenvectors of the Hessian of h at the center of a saddle box I, and let $\beta \in (0, \frac{\pi}{8})$. The unstable wedge C^u_β is the set of points in the surrounding box $N(I)$ at which the (unsigned) angle between ∇h and V^u is at most β. See Fig. 10. Similarly, the stable wedge C^s_β is the set of points in the surrounding box $N(I)$ at which the (unsigned) angle between ∇h and V^s is at most β.*

The saddle point belongs to both the stable and the unstable wedge. Since V^u and V^s are orthogonal, and $0 < \beta < \frac{\pi}{8}$, this is the only common point of the stable and unstable wedge.

Conditions We now introduce additional conditions, which guarantee that each of the wedge-boundaries consist of two regular curves, cf. Lemma 4.2. In fact, these conditions guarantee that C^u_β and C^s_β are really wedge-shaped. Fix $a > 1$, and let $\delta > 0$ be an arbitrarily small constant (to be specified later). At the point (x_0, y_0) we have $HV^u = \lambda_u V^u$, $HV^s = \lambda_s V^s$, so $N(I)$ can be taken small enough to guarantee that the following condition is satisfied at all points of $N(I)$:

Condition I(a, I). At every point of the box $N(I)$ the following inequalities hold:

$$\frac{1}{a}\lambda_u \cdot \| V^u \| \le \| HV^u \| \le a\lambda_u \cdot \| V^u \|,$$
$$\frac{1}{a}|\lambda_s| \cdot \| V^s \| \le \| HV^s \| \le a|\lambda_s| \cdot \| V^s \|.$$

At the point (x_0, y_0) we also have $\langle HV^u, V^u \rangle = \lambda_u \| V^u \|^2$, $\langle HV^s, V^s \rangle = \lambda_s \| V^s \|^2$, and $\langle HV^s, V^u \rangle = 0$. Therefore, for any $\delta > 0$, the box $N(I)$ can be taken small enough such that the following condition holds:

Condition II(δ, I). At every point of the box $N(I)$ the following inequalities hold:

$$\langle HV^u, V^u \rangle \ge \tfrac{1}{2}\lambda_u \| V^u \|^2 \tag{6}$$
$$\langle HV^s, V^s \rangle \ge \tfrac{1}{2}|\lambda_s| \| V^u \|^2 \tag{7}$$
$$|\langle HV^s, V^u \rangle| \le \delta \| V^u \|^2. \tag{8}$$

Since H is symmetric, (8) also implies $|\langle HV^u, V^s \rangle| \le \delta \| V^u \|^2$.

On the boundary of C^u_β the gradient field makes a (signed) angle $\pm\beta$ or $\pi \pm \beta$ with V^u, or, in other words, $X_{\pm\beta}$ is (anti)parallel to V^u. Again, X_β is the vector field obtained by rotating ∇h over an angle β. So let $\Gamma^u_{\pm\beta}$ be the curve along which the vector field $X_{\pm\beta}$ is (anti)parallel to the unstable eigenvector V^u. Then the boundary of the unstable wedge is the union of the two curves Γ^u_β and $\Gamma^u_{-\beta}$. The curve Γ^u_β is defined by the equation

$$\psi^u_\beta(x, y) := \det(V^u, X_\beta(x, y)) = 0. \tag{9}$$

Obviously, the saddle point lies on Γ^u_β. The function $\psi^u_{-\beta}$ is defined similarly.

Similarly, the boundary of the stable wedge is the union of curves $\Gamma^s_{\pm\beta}$, along which the vector field $X_{\pm\beta}$ is (anti)parallel to V^s. The curves $\Gamma^s_{\pm\beta}$ are defined by the equation

$$\psi^s_{\pm\beta}(x, y) := \det(V^s, X_{\pm\beta}(x, y)) = 0.$$

The following technical result provides computable upper bounds for the angle variation of the normals of the boundary curves of the stable and unstable wedges.

Lemma 4.2. *Let $\omega_1 \in (0, \frac{\pi}{4})$ (to be specified later), let $a > 1$, and let I be such that Condition* I(a, I) *holds. Let $0 < \beta < \frac{\pi}{4}$ and $\delta > 0$ such that*

$$\sin\beta \le \frac{\sin\omega_1}{4a^2\sqrt{2}} \min(\left|\frac{\lambda_s}{\lambda_u}\right|, \left|\frac{\lambda_u}{\lambda_s}\right|), \tag{10}$$

$$\delta \le \frac{\sin\omega_1}{8a} \min(|\lambda_s|, |\lambda_u|), \tag{11}$$

$$\delta \le \frac{\tan\beta}{4} \min(|\lambda_s|, |\lambda_u|). \tag{12}$$

If Condition II(δ, I) *also holds, then at any point of $N(I)$*

$$\frac{\pi}{2} - \omega_1 \le \text{angle}(\nabla\psi^u_\beta, V^u) < \frac{\pi}{2} < \text{angle}(\nabla\psi^u_{-\beta}, V^u) \le \frac{\pi}{2} + \omega_1 \tag{13}$$

and

$$\frac{\pi}{2} - \omega_1 \le \text{angle}(\nabla\psi^s_\beta, V^s) < \frac{\pi}{2} < \text{angle}(\nabla\psi^s_{-\beta}, V^s) \le \frac{\pi}{2} + \omega_1. \tag{14}$$

In particular, the angle variation of any of the gradients $\nabla\psi^u_{\pm\beta}$ and $\nabla\psi^s_{\pm\beta}$ over $N(I)$ is less than $2\omega_1$.

Proof. We only show that the angle variation of $\nabla\psi^u_\beta$ over $N(I)$ is less than $2\omega_1$. Since

$$X_\beta = h_x \begin{pmatrix} \cos\beta \\ \sin\beta \end{pmatrix} + h_y \begin{pmatrix} -\sin\beta \\ \cos\beta \end{pmatrix},$$

the function ψ^u_β satisfies $\psi^u_\beta = A_\beta h_x + B_\beta h_y$, where

$$A_\beta = \det(V^u, \begin{pmatrix} \cos\beta \\ \sin\beta \end{pmatrix}) \text{ and } B_\beta = \det(V^u, \begin{pmatrix} -\sin\beta \\ \cos\beta \end{pmatrix}),$$

so

$$\begin{pmatrix} A_\beta \\ B_\beta \end{pmatrix} = \begin{pmatrix} \cos\beta & \sin\beta \\ -\sin\beta & \cos\beta \end{pmatrix} \begin{pmatrix} -V^u_2 \\ V^u_1 \end{pmatrix} = (\cos\beta)\, V^s + (\sin\beta)\, V^u.$$

Therefore,

$$\nabla\psi^u_\beta = \begin{pmatrix} h_{xx} & h_{xy} \\ h_{xy} & h_{yy} \end{pmatrix} \begin{pmatrix} A_\beta \\ B_\beta \end{pmatrix} = \cos\beta\,(HV^s) + \sin\beta\,(HV^u). \tag{15}$$

Condition I(a, I) implies that $HV^u \neq 0$ and $HV^s \neq 0$, and $\{HV^u, HV^s\}$ are independent vectors, at all points of $N(I)$, so the gradient of ψ^u_β is nonzero at every point of $N(I)$, so Γ^u_β is a regular curve.

Expression (15) for $\nabla\psi^u_\beta$ implies that

$$\| \nabla\psi^u_\beta \|^2 = \cos^2\beta \, \| HV^s \|^2 + 2\sin\beta\cos\beta \, \langle HV^s, HV^u\rangle + \sin^2\beta \, \| HV^u \|^2.$$

Using the Cauchy–Schwarz inequality $|\langle HV^s, HV^u\rangle| \leq \| HV^s \| \cdot \| HV^u \|$ and the fact that $\beta > 0$ we get

$$\begin{aligned} \| \nabla\psi^u_\beta \|^2 &\geq \cos^2\beta \, \| HV^s \|^2 - 2\sin\beta\cos\beta \, \| HV^s \| \cdot \| HV^u \| \\ &\quad + \sin^2\beta \, \| HV^u \|^2 \\ &= \cos^2\beta \left(\| HV^s \| - \| HV^u \| \tan\beta \right)^2. \end{aligned} \tag{16}$$

Since $\sin\beta \leq \dfrac{\sin\omega_1}{4a^2\sqrt{2}} \left|\dfrac{\lambda_s}{\lambda_u}\right|$ and $0 < \beta < \frac{\pi}{4}$, it follows from Condition I(a, I) that

$$\| HV^u \| \tan\beta \leq a\lambda_u \| V^u \| \frac{\sin\beta}{\sqrt{2}} < \frac{|\lambda_s|}{2a} \| V^u \|.$$

Using Condition I again we get

$$\| HV^s \| - \| HV^u \| \tan\beta \geq \frac{|\lambda_s|}{a} \cdot \| V^u \| - \frac{|\lambda_s|}{2a} \| V^u \| = \frac{|\lambda_s|}{2a} \| V^u \|.$$

In view of (16) we get, using $\cos\beta \geq \frac{1}{\sqrt{2}}$:

$$\| \nabla\psi^u_\beta \| \geq \frac{|\lambda_s|}{2a\sqrt{2}} \| V^u \|. \tag{17}$$

Expression (15) for $\nabla\psi^u_\beta$ also implies that

$$\begin{aligned} \langle \nabla\psi^u_\beta, V^u\rangle &= \cos\beta \, \langle HV^s, V^u\rangle + \sin\beta \, \langle HV^u, V^u\rangle \\ &= \cos\beta \langle HV^u, V^u\rangle \, \Big(\frac{\langle HV^s, V^u\rangle}{\langle HV^u, V^u\rangle} + \tan\beta \Big). \end{aligned}$$

Condition II and (12) imply

$$\Big| \frac{\langle HV^s, V^u\rangle}{\langle HV^u, V^u\rangle} \Big| \leq \frac{2\delta}{\lambda_u} \leq \tfrac{1}{2}\tan\beta.$$

Since $\beta > 0$, this implies $\langle \nabla\psi^u_\beta, V^u\rangle > 0$ on $N(I)$.

According to Condition II we have $|\langle HV^s, V^u\rangle| \leq \delta \| V^u \|^2$, whereas the Cauchy–Schwarz inequality implies $|\langle HV^u, V^u\rangle| \leq \| HV^u \| \cdot \| V^u \| \leq a\lambda_u \| V^u \|^2$. Therefore,

$$0 \leq \langle \nabla\psi^u_\beta, V^u\rangle \leq (\delta\cos\beta + a\lambda_u \sin\beta) \| V^u \|^2.$$

Together with (17) this implies

$$\Big\langle \frac{\nabla\psi^u_\beta}{\| \nabla\psi^u_\beta \|}, \frac{V^u}{\| V^u \|} \Big\rangle \leq \frac{2a\sqrt{2}}{|\lambda_s|} (\delta\cos\beta + a\lambda_u \sin\beta).$$

Given the upper bounds (11) for δ and (10) for $\sin\beta$, we get

$$\frac{2a\sqrt{2}}{|\lambda_s|} \delta\cos\beta \leq \frac{4a}{|\lambda_s|} \delta \leq \tfrac{1}{2} \sin\omega_1 \quad \text{and} \quad \frac{2a\sqrt{2}}{|\lambda_s|} a\lambda_u \sin\beta \leq \tfrac{1}{2} \sin\omega_1.$$

Therefore,

$$0 < \langle \frac{\nabla \psi^u_\beta}{\| \nabla \psi^u_\beta \|}, \frac{V^u}{\| V^u \|} \rangle \leq \sin \omega_1 = \cos(\frac{\pi}{2} - \omega_1).$$

At all points of $N(I)$ we then have

$$\frac{\pi}{2} - \omega_1 \leq \text{angle}(\nabla \psi^u_\beta, V^u) < \frac{\pi}{2}.$$

Since V^u is constant, the angle variation of $\nabla \psi^u_\beta$ over $N(I)$ does not exceed $2\omega_1$. □

The main result of this subsection states that, under suitable conditions, the intersection of the boundary of a saddle box and the stable and unstable wedges can be computed. Moreover, at all points of these intersections the gradient vector field is transversal to the boundary of the saddle box, and, even stronger, at these points there is a computable positive lower bound for the angle of the gradient vector field and the boundary of the saddle box.

Corollary 4.3. *Let a and ω_1 be constants such that $a > 1$, and $\omega_1 = \frac{1}{3} \arctan \frac{1}{2}$. Let $\beta \in (0, \omega_1)$ and $\delta > 0$ such that (10), (11) and (12) hold.*

If I is a saddle box with concentric surrounding box $N(I)$ satisfying Condition $I(a, I)$ and Condition $\text{II}(\delta, I)$, then

1. *The saddle point is the only common point of the stable wedge C^u_β and the unstable wedge C^s_β.*
2. *The gradient vector field ∇h is transversal at points on the boundary of these wedges, different from the saddle point: on the boundary of the unstable wedge it points inward, except at the saddle point, and on the boundary of the stable wedge it points outward, except at the saddle point.*
3. *The unstable wedge C^u_β contains the unstable separatrices of the saddle point, and the stable wedge C^s_β contains the stable separatrices.*
4. *The unstable wedge intersects the boundary of $N(I)$ in two intervals, called the unstable intervals. Similarly, the stable wedge intersects the boundary of $N(I)$ in two intervals, called the stable intervals. These four intervals are disjoint, and the unstable and stable intervals occur alternatingly on the boundary of $N(I)$. At each point of a stable or unstable interval the (unsigned) angle between ∇h and the boundary edge containing this point is at least ω_1. Moreover, there are computable isolating intervals for each stable and unstable interval.*

Proof. 1. If $p \in C^u_\beta \cap C^s_\beta$, then $\nabla h(p)$ makes an angle $\beta \in (0, \frac{\pi}{4})$ with both V^u and V^s. Therefore, $\nabla h(p) = 0$, since these vectors are orthogonal. Hence, p is a singular point of ∇h inside $N(I)$, which is the saddle point.

2. Recall from Lemma 4.2 that $\langle \nabla \psi^u_\beta, V^u \rangle$ is positive on $N(I)$. Let p be the saddle point of ∇h in I. Since X_β is parallel to V^u on one component of $\Gamma^u_\beta \setminus \{p\}$, and parallel to $-V^u$ on the other component, it follows that X_β is pointing into the unstable wedge along both components. See again Fig. 10. Since ∇h is obtained by rotating X_β over $-\beta$, and $X_{-\beta}$ over β, also ∇h is pointing into the unstable wedge. Similarly, $\langle \nabla \psi^u_{-\beta}, V^u \rangle$ is negative along both components of $\Gamma^u_{-\beta} \setminus \{p\}$, so ∇h is also pointing into the unstable wedge along each of these components.

A similar argument shows that ∇h is pointing outward along each of the boundary components of the stable wedge C^s_β, except at the saddle point.

3. The second part of the lemma implies that the unstable wedge is forward invariant under the flow of the gradient vector field ∇h. In particular, it contains the unstable separatrices of the saddle point. Similarly, the stable wedge C^s_β is backward invariant, so it contains the stable separatrices of the saddle point.

4. Suppose Γ^u_β intersects an edge e of the surrounding box at a point $\overline{q}$, see Fig. 11. We first show that the angle of $\nabla h(\overline{q})$ and e is bounded away from zero. To see this, observe that there is a point $s \in \Gamma^u_\beta$ at which $\nabla \psi^u_\beta(s)$ is perpendicular to $\overline{p}\,\overline{q}$. Therefore, the angle between $\nabla \psi^u_\beta(s)$ and the normal of e is at least ω_0, where $\omega_0 = \arctan \frac{1}{2} = 3\omega_1$.

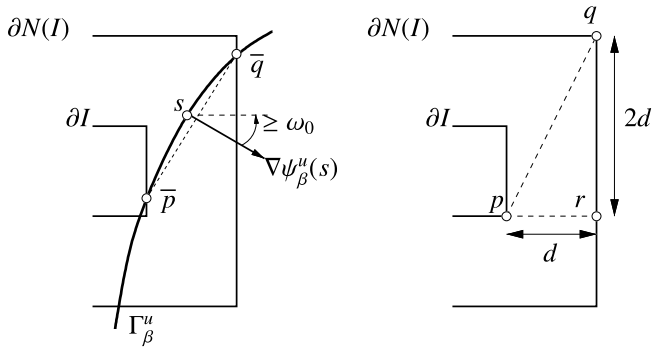

Fig. 11. Lower bound on angle of intersection of Γ^u_β and the boundary of the surrounding box $N(I)$.

The angle between V^u and $\nabla\psi^u_\beta$ lies in the interval $[\frac{\pi}{2} - \omega_1, \frac{\pi}{2})$, cf. Lemma 4.2, so the angle between V^u and e is at least $\omega_0 - \omega_1 = 2\omega_1$.

At any point of C^u_β the angle between ∇h and V^u is at most β—by the definition of C^u_β—so at any point of $\in C^u_\beta \cap e$ the angle between ∇h and e is at least $2\omega_1 - \beta \geq \omega_1$.

To find isolating intervals for the intersection of the stable and unstable wedges with the boundary of the surrounding box $N(I)$, we compute isolating intervals for the intersection of each of the four curves $\psi^u_{\pm\beta} = 0$, $\psi^s_{\pm\beta} = 0$ with this boundary. The normal to each of the curves $\psi^u_{\pm\beta} = 0$ makes an angle of at least $\frac{\pi}{2} - 4\omega_1$ with each of the curves $\psi^s_{\pm\beta} = 0$. This follows from (13) and (14), and the fact that V^u and V^s are perpendicular. Since $\omega_1 = \frac{1}{3}\arctan\frac{1}{2} < \frac{\pi}{20}$, so the angle between each of the curves $\psi^u_{\pm\beta} = 0$ and each of the curves $\psi^s_{\pm\beta} = 0$ is at least $\frac{\pi}{2} - 4\frac{\pi}{20} = \frac{3\pi}{10}$, which is bounded away from zero. Therefore, the method of Section 3 provides such isolating intervals. □

As will become clear in the certified construction of the MS-complex, we need to be able to provide certified separatrix intervals of arbitrarily small width, without refining the saddle box:

Lemma 4.4. *Let I be a separatrix box satisfying the conditions of Corollary 4.3. Then the isolating separatrix intervals in the boundary of $N(I)$ can be made arbitrarily small.*

The proof of this result is rather technical. For a sketch we refer to Appendix B.

4.2. Refining boxes for maxima and minima

To construct the MS-complex, the algorithm needs to determine when an unstable (stable) separatrix will have a given maximum (minimum) of h as its ω-limit (α-limit). For each maximum (minimum) the algorithm determines a certified box such that the gradient vector field points inward (outward) on the boundary of the box. Unfortunately, we cannot always choose an axis-aligned box, as will become clear from the following example.

Let $h(x, y) = -5x^2 - 4xy - y^2$, then the origin is a sink of the gradient vector field

$$\nabla h(x, y) = \begin{pmatrix} -10x - 4y \\ -4x - 2y \end{pmatrix}.$$

This vector field is horizontal along the line $y = -2x$, which intersects the horizontal edges of every axis aligned box centered at the sink $(0, 0)$. In other words, the vector field is not transversal on the boundary of any such box.

However, there is a box aligned with (approximations) of a pair of eigenvectors of the linear part of the gradient vector field at (or, near) its singular point, for which the vector field is transversal

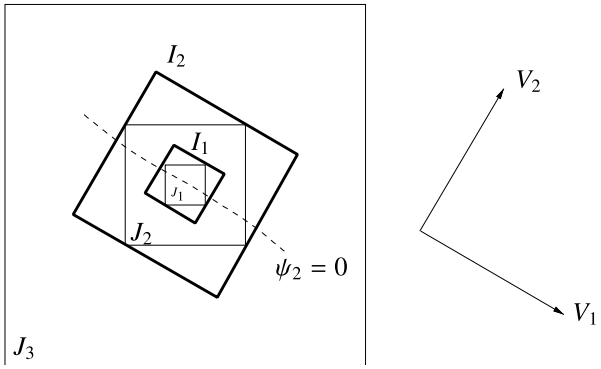

Fig. 12. Construction of a sinkbox.

to the boundary. To see this, we refine the sink-box to obtain three concentric axis aligned boxes $J_1 \subset J_2 \subset J_3$, such that

(i) the edge length of J_i, $i = 2, 3$, is three times the edge length of J_{i-1}, and
(ii) the sink is contained in the inner box J_1. See also Fig. 12. Moreover, let V_1 and V_2 be the (orthogonal) eigenvectors of the Hessian matrix H^0 at the center of the boxes. These eigenvectors, corresponding to the eigenvectors λ_1 and λ_2, are computed as in Section 4.1, cf. (4).

We require that

(iii) the gradients of the two functions ψ_1 and ψ_2, defined by

$$\psi_i(x, y) = \langle \nabla h(x, y), V_i \rangle, \quad i = 1, 2,$$

have small angle variation over the outer box J_3. This condition is made precise in Lemma 4.5 below. Note that

$$\nabla \psi_i(x, y) = H(x, y) V_i, \tag{18}$$

where $H(x, y)$ is again the Hessian matrix of h at (x, y). Since this matrix is non-singular, we can find a triple of boxes $J_1 \subset J_2 \subset J_3$, satisfying conditions (i), (ii) and (iii), such that $H(x, y)$ is nearly constant over the outer box (again, this is made precise in Lemma 4.5). In particular, $\nabla \psi_i$ is nearly parallel to V_i, since $H^0 V_i = \lambda_i V_i$. Now construct boxes I_1 and I_2, which are the smallest boxes enclosing J_1 and J_2, respectively, with edges parallel to V^1 or V^2.

Lemma 4.5. *Suppose on the outer box J_3 the following conditions hold:*

1. $\| HV_i \| \geq \frac{1}{2} |\lambda_i| \cdot \| V_i \|$, *for* $i = 1, 2$;
2. $|\langle HV_1, V_2 \rangle| = |\langle HV_2, V_1 \rangle| \leq \frac{1}{4} \| V_1 \|^2 \arctan \frac{1}{2}$.

Then the gradient vector field is transversal to the boundary of I_2.

Proof. The second condition limits the variation of the angle of $\nabla \psi_i = HV_i$ and the basis vectors V_1 and V_2 over J_3. Using this bound, we use the same arguments as in Section 3, applied to the pair of boxes I_1, I_2, to show that the curve $\psi_i(x, y) = 0$ does not intersect the edges of I_2 perpendicular to V_i. Therefore, $\psi_i = \langle \nabla h, V_i \rangle$ is nowhere zero on these edges, so ∇h is nowhere tangent to these edges. Therefore, I_2 is the desired sink box, on the boundary of which ∇h is pointing inward. In other words, if an unstable separatrix intersects the boundary of this box, the part of the separatrix beyond this point of intersection lies inside the sink-box. Certified source-boxes are constructed similarly. □

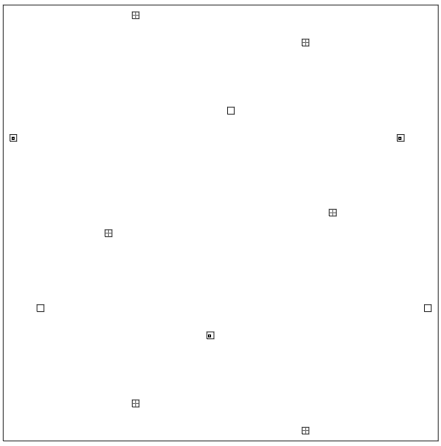

Fig. 13. Refined critical boxes (after detecting types) corresponding to the function in Fig. 1(a)—6 saddles (each box with a cross), 3 sinks (each box with a small box inside) and 3 sources (empty boxes).

5. Isolating funnels around separatrices

To compute isolating strips (funnels) for unstable separatrices of ∇h we approximate, for each isolating unstable separatrix interval, the forward orbits of its endpoints. Isolating funnels for stable separatrices are computed similarly. If the forward orbits of the endpoints of an unstable separatrix interval have the same sink of ∇h as ω-limit, these forward orbits bound a region around the unstable separatrix leaving the saddle box via this unstable segment. This region is called a *funnel* for the separatrix (this terminology is borrowed from Hubbard and West, 1991).

In this section we provide the details of the construction of a certified funnel for each separatrix. First we show, in Section 5.1 how to construct two polylines per separatrix interval those are candidates for the funnel around the corresponding separatrix. Then, in Section 5.2, we introduce the notion of *width of a funnel*, and derive upper bounds for its growth. These upper bounds are the ingredients for a certified algorithm computing these funnels. The algorithm computes the Morse–Smale complex by providing disjoint certified funnels for each stable and unstable separatrix. The proof of correctness and termination is presented in Section 5.3.

5.1. Construction of fences around a separatrix

Let X_ϑ be the vector field obtained by rotating the vector field $X = \nabla h$ over an angle ϑ, i.e.,

$$X_\vartheta = \begin{pmatrix} \cos\vartheta & -\sin\vartheta \\ \sin\vartheta & \cos\vartheta \end{pmatrix} \begin{pmatrix} h_x \\ h_y \end{pmatrix}.$$

We compute an *isolating funnel* for the forward orbit of ∇h through a point p by enclosing it between (approximations of) the forward orbits of X_ϑ and $X_{-\vartheta}$ through p. See Fig. 14.

Small angle variation We first determine some bounds on the angle variation of the gradient vector field ∇h over $\mathbb{D}^*$. We subdivide the region $\mathbb{D}^*$ into square boxes over which the angle variation of ∇h is at most ϑ, where ϑ is to be determined later. Let w be the edge length of the boxes. If $X = (f, g)$ is a vector field on $\mathbb{R}^2$, then the angle variation over a regular curve Γ is given by (Arnol'd, 2006, Section 36.7):

$$\int_\Gamma \frac{g\,df - f\,dg}{f^2 + g^2}.$$

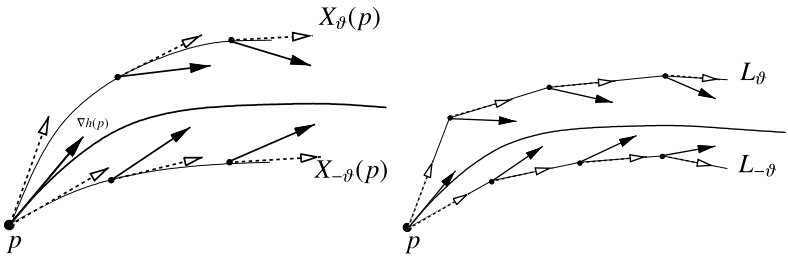

Fig. 14. Orbits of the rotated vector fields X_ϑ and $X_{-\vartheta}$ through a point p enclose the forward orbit of ∇h through p. On the right polygonal lines approximating these orbits.

If $X = \nabla h$, this angle variation is equal to

$$\int_\Gamma \frac{(h_x h_{xy} - h_y h_{xx})\,dx + (h_x h_{yy} - h_y h_{xy})\,dy}{h_x^2 + h_y^2}.$$

Let C_0 and C_1 be constants such that

$$\max_{\mathbb{D}^*} \left| \frac{h_x h_{xy} - h_y h_{xx}}{h_x^2 + h_y^2} \right| \le C_0 \quad \text{and} \quad \max_{\mathbb{D}^*} \left| \frac{h_x h_{yy} - h_y h_{xy}}{h_x^2 + h_y^2} \right| \le C_1. \tag{19}$$

Then the angle variation over a curve Γ is less than

$$\int_\Gamma (C_0\,dx + C_1\,dy) \le (C_0 + C_1)\ \text{length}(\Gamma).$$

This inequality provides an upper bound for the maximal angle variation over a square box:

Lemma 5.1. *Let C_0 and C_1 satisfy (19). Then the total angle variation over a square box in $\mathbb{D}^*$ with edge length w does not exceed $(C_0 + C_1)w\sqrt{2}$.*

The grid boxes have edge length w such that the angle variation of ∇h over any box in $\mathbb{D}^*$ is less than ϑ.

Lemma 5.2. *Let $0 < \vartheta < \frac{\pi}{2}$, and let the grid boxes in $\mathbb{D}^*$ have width w satisfying*

$$w \le \frac{\vartheta}{(C_0 + C_1)\sqrt{2}}. \tag{20}$$

Then the following properties hold.

1. *The angle variation of ∇h over any gridbox is less than ϑ.*
2. *Let p be a point on an edge of a gridbox, and let q_ϑ be the point on the boundary of the gridbox into which $X_\vartheta(p)$ is pointing, such that the line segment pq_ϑ has direction $X_\vartheta(p)$. The point $q_{-\vartheta}$ is defined similarly. See Fig. 15. Then ∇h is pointing rightward along pq_ϑ and leftward along $pq_{-\vartheta}$.*
3. *The function h is strictly increasing on the line segments from p to q_ϑ and from p to $q_{-\vartheta}$.*

Proof. The first claim follows from Lemma 5.1, using the fact that the diameter of a grid box is $w\sqrt{2}$.

With regard to the second claim, the small angle variation condition implies that the orientation of $\{\nabla h(q), X_\vartheta(p)\}$ does not change as q ranges over I. Since this orientation is positive for $q = p$, it is positive for all $q \in I$. Similarly, the orientation of $\{\nabla h(q), X_{-\vartheta}(p)\}$ is negative for all $q \in I$. Therefore, the second claim also holds.

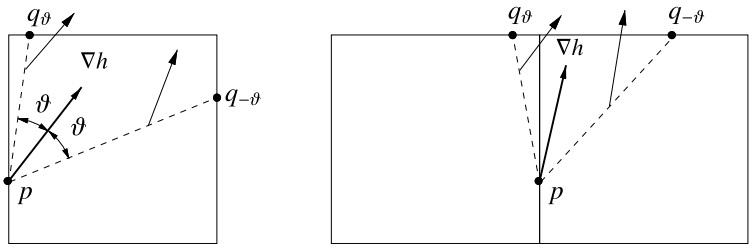

Fig. 15. The orientation of ∇h with respect to $X_\vartheta(p)$ does not change over a grid box.

At a point r of the line segment $pq_{\pm\vartheta}$ the directional derivative of h in the direction of this line segment is $\langle \nabla h(r), X_{\pm\vartheta}(p)\rangle$, which is positive since the angle between $\nabla h(r)$ and $X_{\pm\vartheta}(p)$, is less than ϑ, and $\vartheta < \frac{\pi}{2}$. This proves the third part. □

Fencing in the separatrices For each isolating unstable separatrix interval J on the boundary of a saddle box we construct two polylines $L_{-\vartheta}(J)$ and $L_\vartheta(J)$ as follows. The initial points of these polylines are the endpoints of J, q_- and q_+, where q_- comes before q_+ in the counterclockwise orientation of the boundary of the saddle box. The polyline $L_\vartheta(J)$ is uniquely defined by requiring that its vertices $q_+ = p_0, p_1, \ldots, p_n$ lie on grid edges, with the property that

1. The line segment $p_{i-1}p_i$, $0 < i \leq n$, lies in a (closed) grid box of $\mathbb{D}^*$, and has direction $X_\vartheta(p_{i-1})$.
2. p_n, the last vertex, lies on the boundary of $\mathbb{D}^*$.

The polyline $L_{-\vartheta}(J)$ is defined similarly, with the obvious changes: its initial vertex is q_-, and each edge has direction equal to the value of the vector field $X_{-\vartheta}$ at the initial point of this edge. The polylines $L_{\pm\vartheta}(J)$ are called *fences* of the (unique) unstable separatrix of ∇h intersecting J.

It is not hard to see that a grid box contains at most two *consecutive* edges of each of these polylines, but it is not obvious a priori that each box cannot contain more than two edges of each polyline in total. It follows from the next result that the intersection of a grid box with any of these polylines is connected, and, hence, that these polylines are finite.

The following results states that, when walking along the polylines L_ϑ and $L_{-\vartheta}$ in the direction of increasing h-values, each grid box is passed at most once.

Lemma 5.3. *Let I be a box such that the angle variation of ∇h over the surrounding box $N(I)$ is at most ϑ. Then the intersection of $L_\vartheta(J)$ and I ($L_{-\vartheta}(J)$ and I) is either empty or a connected polyline (consisting of one or two segments).*

Proof. Let q be a point at which $L_\vartheta(J)$ leaves I, i.e., the segment of $L_\vartheta(J)$ ending at q lies inside I and the segment qr beginning at q lies outside I. Let p be a point on the boundary of I at which h attains its maximum value M.

Case 1: p is a vertex of I, incident to the edge of I containing q. See Fig. 16, top row. Let l be the line through the edge of I containing q, let α be the angle between l and $\nabla h(p)$, and let β be the angle between l and the segment of $L_\vartheta(J)$ with initial point q. The angles α and β are both positive, since p is a vertex of I. The angle between $\nabla h(p)$ and $\nabla h(q)$ is at most ϑ, since the angle variation of ∇h over I is less than ϑ. Therefore, $|\alpha - \beta| \leq 2\vartheta$.

Let pp_0 and pp_1, with p_0 and p_1 on the boundary of $N(I)$, be the line segments that make an angle of $\frac{\pi}{2} - \vartheta$ with $\nabla h(p)$. Since the angle variation of ∇h over $N(I)$ is at most ϑ, the value of h at any point of these line segments is at least M. We shall prove that the connected component of $L_\vartheta(J) \cap N(I)$ containing q intersects one of the line segments pp_0 and pp_1.

First assume $\alpha \geq \vartheta$. Then the line segment pp_0 lies in the grid box J containing segment qr of $L_\vartheta(J)$. If r lies on an edge of J incident to p, then qr intersects pp_0. So assume r lies on the edge of

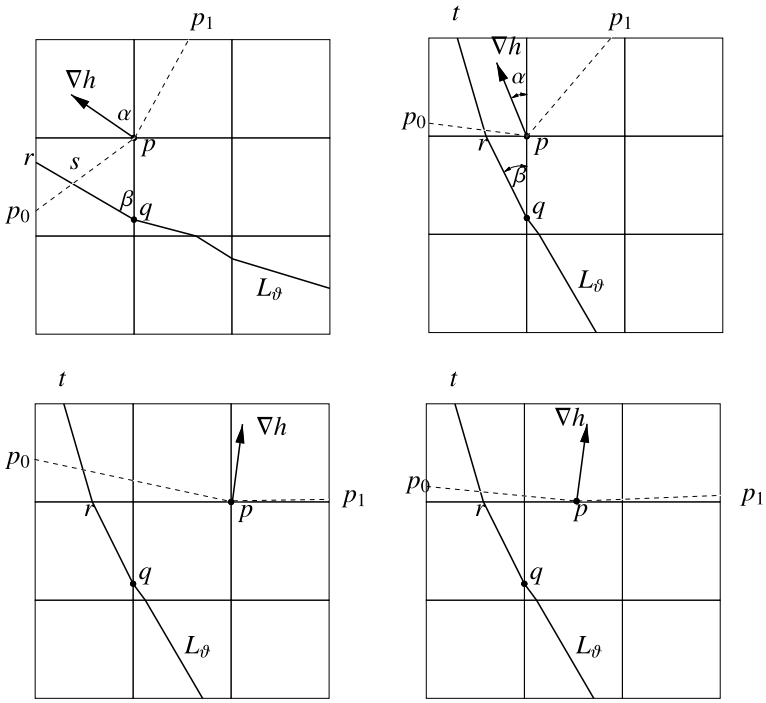

Fig. 16. The value of h at the point where polyline $L_\vartheta(J)$ leaves the surrounding box $N(I)$ is greater than the maximum value of h on I.

J contained in the boundary of $N(I)$ (Fig. 16, leftmost picture). Let s be the point of intersection of the line through pp_0 and the line through qr. This point lies on the same side of l as p_0 and r, since $\angle p_0pq = \frac{\pi}{2} - \alpha + \vartheta < \frac{\pi}{2}$ and $\angle pqr = \beta \leq \frac{\pi}{2}$. Furthermore, $\angle psq = \pi - (\frac{\pi}{2} - \alpha + \vartheta) - \beta \geq \frac{\pi}{2} - \vartheta > \frac{\pi}{4}$. Therefore, x lies inside $N(I)$, in other words, qr intersects pp_0 also in this case.

Now consider the case $\alpha < \vartheta$. Then p_0 lies on the side of $N(I)$ parallel to the line through p and q. Furthermore, $\beta \leq \alpha + 2\vartheta < 3\vartheta < \frac{3\pi}{40}$, so $L_\vartheta(J)$ 'leaves' $N(I)$ at a point t on the side of $N(I)$ perpendicular to the line through p and q. See Fig. 16, rightmost picture. It follows that the part of $L_\vartheta(J)$ between q and t intersects pp_0. In particular, $h(t) > M$.

Case 2: p is not a vertex of I, incident to the edge of I containing q. Then either p is a vertex of I, not incident to the edge of I, containing q, as in Fig. 16, bottom-left picture, or p lies on the relative interior of an edge of I, as in Fig. 16, bottom-right picture.

In this case $\nabla h(p)$ is nearly vertical, as are the edges of $L_\vartheta(J)$. Similarly, the line segments pp_0 and pp_1 are nearly horizontal, so $L_\vartheta(J)$ intersects pp_0. The details are similar to those of Case 1 of this proof. □

If the endpoints of the fences $L_\vartheta(J)$ and $L_{-\vartheta}(J)$ lie on the same connected component of the boundary of $\mathbb{D}^*$, then these fences split $\mathbb{D}^*$ into two connected regions. See Fig. 17. In this case, the region containing the separatrix interval J in its boundary is called the *funnel* of J (with angle ϑ) denoted by $F_\vartheta(J)$. Its boundary consists of J, the two fences $L_\vartheta(J)$ and $L_{-\vartheta}(J)$, and a curve J^* on the boundary of $\partial\mathbb{D}^*$ connecting the endpoints of these fences. If the funnel is simply connected, it contains the part of the unstable separatrix through J lying inside $\mathbb{D}^*$, which enters the funnel through J and leaves it through J^*. Note that J^* is a curve either on the outer boundary of $\mathbb{D}^*$ or on a sink box.

Similarly, each *stable* separatrix interval has two fences (for an angle ϑ). If the endpoints of these fences lie on the same connected component of $\partial\mathbb{D}^*$, the enclosed region is again called a funnel for

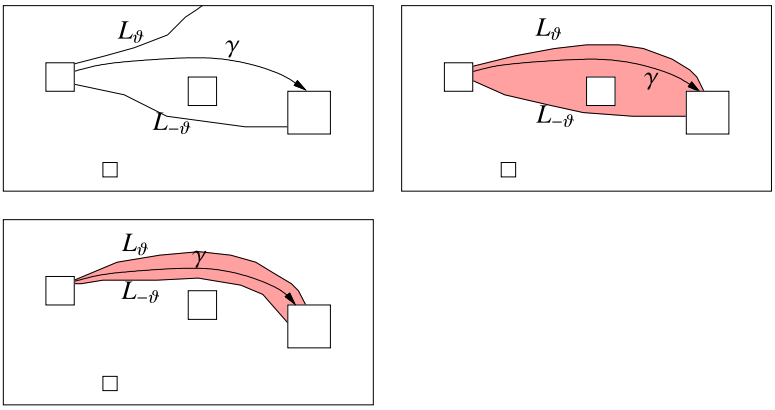

Fig. 17. Fences around a separatrix γ. If the fences end in the same connected component of the boundary of $\mathbb{D}^*$, then they enclose a funnel (top right picture). If the funnel is simply connected, it isolates the separatrix from the source-, sink- and saddle-boxes (bottom picture).

the stable separatrix interval. Our goal is to construct disjoint, simply connected funnels for the stable and unstable separatrix intervals. If these funnels are disjoint, then they form, together with the sink boxes, source boxes and saddle boxes, a (fattened) Morse–Smale complex for ∇h.

It is intuitively clear that a funnel $F_\vartheta(J)$ is simply connected if ϑ, the length of the separatrix interval J, and the edge length w of the grid boxes are sufficiently small. The next subsection presents computable upper bounds on these quantities, guaranteeing that the endpoints of two fences of a separatrix interval lie on the same boundary component. It is then easy to check whether the enclosed funnel is simply connected.

5.2. Controlling the width of the funnel

If the width of a funnel is sufficiently small, in a sense to be made more precise, it encloses a simply connected region in $\mathbb{D}^*$. The width of a funnel is, roughly speaking, the number of grid boxes between its bounding fences in the vertical direction, in regions where the fences are nearly horizontal, and in the horizontal direction, in regions where the fences are nearly vertical. To define the width of a funnel more precisely, we distinguish quasihorizontal and quasivertical parts of a funnel, and show that the width of a funnel does not increase substantially at transitions between these quasihorizontal and quasivertical parts.

Quasihorizontal and quasivertical parts of a funnel A nonzero vector $v = (v_1, v_2)$ is called *quasihorizontal* if $|v_2| \leq 2|v_1|$, and *quasivertical* if $|v_1| \leq 2|v_2|$. Note that each nonzero vector is quasihorizontal, quasivertical, or both. Consider a subdivision of $\mathbb{D}^*$ into boxes of equal width, where non-disjoint boxes share either an edge or a vertex. A *horizontal ε-strip* is the union of a sequence of boxes where successive boxes share a vertical edge, such that the horizontal edge of the rectangle thus obtained has length at most ε. A *vertical ε-strip* is defined similarly. An *ε-box* is a square box with edge length at most ε which is the union of a number of boxes. Two polygonal curves L_+ and L_- form an ε-funnel if there is a set $\mathcal{H}_\varepsilon$ of horizontal ε-strips, a set $\mathcal{V}_\varepsilon$ of vertical ε-strips, and a set $\mathcal{B}_\varepsilon$ of ε-boxes such that the following holds:

1. The vertices of L_- and L_+ lie on the edges of the grid-boxes; L_- intersects a grid box in at most one vertex or in at most one edge; the same holds for L_+;
2. Both L_- and L_+ lie in the union of the rectangles in $\mathcal{H}_\varepsilon$, $\mathcal{V}_\varepsilon$ and $\mathcal{B}_\varepsilon$;
3. An edge of $L_\pm$ contained in a horizontal ε-strip is quasi-vertical, and an edge contained in a vertical ε-strip is quasihorizontal. Moreover, neither L_+ nor L_- intersect the vertical sides of a

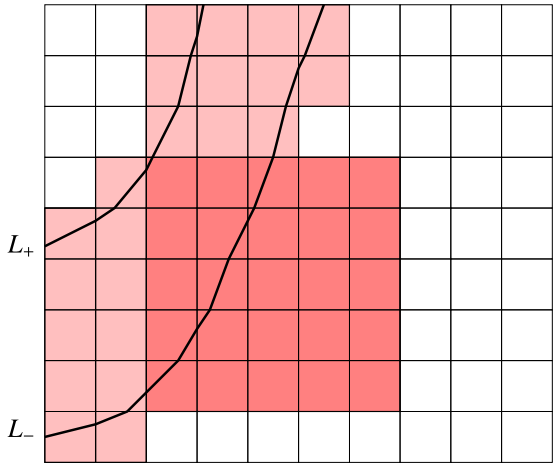

Fig. 18. A funnel formed by two polylines covered by two vertical ε-strips, one ε-box and three horizontal ε-strips. Here ε is six times the width of a grid box. L_+ intersects the ε-box in a single edge, which is quasivertical but not quasihorizontal. All edges of L_- inside the ε-box are quasivertical as well (and some of them are also quasihorizontal).

horizontal ε-strip, or the horizontal sides of a vertical ε-strip. Each ε-strip and each ε-box is intersected by both polylines.

4. Either L_- or L_+ intersects an ε-box in exactly one of its edges, which is contained in a grid box at the corner of the ε-box. This single edge is either quasivertical or quasihorizontal (but not both). If this edge is quasihorizontal (quasivertical), all edges of the other polyline inside the ε-box are quasihorizontal (quasivertical) as well—and possibly also quasivertical (quasihorizontal). The other polygonal curve intersects the same edges of the ε-box, each in exactly one point, and is disjoint from the other edges of the ε-box.

See also Fig. 18.

We determine $\vartheta > 0$ later, but for now we assume that

$$\vartheta \leq \frac{\pi}{40}. \tag{21}$$

We start with a simple observation.

Lemma 5.4. *Let L be a polyline with quasihorizontal edges and with vertices on the edges of a grid with edge length w satisfying (20). If L lies in a vertical strip of width w, where each of the vertical lines bounding the strip contains one of its endpoints, then L intersects at most three grid boxes contained in this vertical strip.*

A similar property holds for a polyline with quasivertical edges intersecting a horizontal strip.

Proof. We only prove the first part, in which L lies in a vertical strip and has quasihorizontal edges. The slope of the line segment connecting the endpoints of L does not exceed the maximum slope of any of the edges of L, so this slope is at most $\arctan 2$. Hence the projection of this line segment on any of the vertical lines bounding the strip has length at most $2w$, so it intersects at most three boxes. □

The next result shows that the width of a funnel does not grow substantially at a transition between a quasihorizontal and a quasivertical part. We take $\varepsilon > 0$ such that the angle variation of ∇h over a box with edge length ε is at most $\frac{\pi}{20}$. Again, by Lemma 5.1, this is guaranteed by taking

$$\varepsilon \leq \frac{\pi}{20(C_0 + C_1)\sqrt{2}}. \tag{22}$$

Lemma 5.5. *Let J be an ε-box intersected by both $L_{-\vartheta}$ and $L_{+\vartheta}$, with an edge e which contains the initial vertex of both $L_{-\vartheta} \cap J$ and $L_{+\vartheta} \cap J$. Assume that at least one of the polylines has an edge which is either quasihorizontal or quasivertical, but not both. Then both polylines intersect the boundary of J in exactly two points, and there is an edge e' of J, adjacent to e, containing the terminal vertices of both $L_{-\vartheta} \cap J$ and $L_{+\vartheta} \cap J$. See Fig. 18.*

Proof. Assume that $L_{+\vartheta}$ has an edge e_+ which is quasivertical but not quasihorizontal. We first show that all edges of $L_{+\vartheta}$ are quasivertical (and possibly quasihorizontal). The angle between e_+ and the horizontal direction is at least $\arctan 2$, which is greater than $\frac{\pi}{4} + \frac{\pi}{10}$. Since the slope of e_+ is the slope of the vector field X_ϑ at the initial vertex of e_+, and the angle variation of X_ϑ over J is at most $\frac{\pi}{20}$, the slope of X_ϑ at any point of J is at least $\frac{\pi}{4} + \frac{\pi}{20}$. Since the slope of an edge of $L_{+\vartheta}$ is the slope of X_ϑ at the initial vertex of this edge, we conclude that all edges of $L_{+\vartheta}$ are quasivertical.

All edges of $L_{-\vartheta}$ are also quasivertical (and possibly quasihorizontal). To see this, observe that the slope of an edge of $L_{-\vartheta}$ is the slope of $X_{-\vartheta}$ at the initial vertex of this edge, and, hence, the slope of X_ϑ at this initial vertex, minus 2ϑ. In other words, the slope of any edge of $L_{-\vartheta}$ is at least $\frac{\pi}{4} + \frac{\pi}{20} - 2\vartheta$. Since $\vartheta \le \frac{\pi}{40}$, this slope is at least $\frac{\pi}{4}$. Therefore, all edges of $L_{-\vartheta}$ are quasivertical.

The polylines $L_{+\vartheta}$ and $L_{-\vartheta}$ do not intersect the edge of J opposite e, since then at least one of the edges of these polylines would have a slope less than $\frac{\pi}{4}$. Let e' be the edge containing the endpoint of $L_\vartheta(J) \cap J$. Then e' is adjacent to e. Given the bounds on the slope variation of the edges of the polylines, it is easy to see that

(i) the endpoint of $L_\vartheta(J)$ is the only point of this polyline on e';
(ii) the endpoint $L_{-\vartheta} \cap J$ also lies on e', and this is the only point of this polyline on e';
(iii) none of the polylines intersects the edge opposite e'.

This concludes the proof of Lemma 5.5. □

Growth of the width of quasihorizontal and quasivertical funnel parts The width of the funnel may grow exponentially in the number of grid boxes it is traversing. The next result gives an upper bound for the growth of this width. Even though the bounds are conservative, they provide the tools for the construction of certified funnels for all separatrices.

A gridbox is called *quasihorizontal* (*quasivertical*) if it contains a point at which ∇h is quasihorizontal (quasivertical). Again, a gridbox may be both quasihorizontal and quasivertical.

An integral curve of ∇h in a quasihorizontal gridbox $[x_0, x_1] \times [y_0, y_1]$ is the graph of a function $x \mapsto y(x)$, where $y(x)$ is a solution of the differential equation

$$\begin{aligned} y'(x) &= F(x, y(x)), \\ y(x_0) &= y_0, \end{aligned} \tag{23}$$

where $F(x, y) = \dfrac{h_y(x, y)}{h_x(x, y)}$. Here x ranges over the full interval $[x_0, x_1]$ if $y_0 \le y(x) \le y_1$. Otherwise, the range of x is restricted to a suitable maximal subinterval $[\xi_0, \xi_1]$, such that $(\xi_0, y(\xi_0))$ and $(\xi_1, y(\xi_1))$ are points on the boundary of the gridbox. Similarly, a trajectory of X_ϑ in a quasihorizontal gridbox $[x_0, x_1] \times [y_0, y_1]$ is the graph of a function $x \mapsto y(x)$, where $y(x)$ is a solution of the differential equation

$$\frac{dy}{dx} = F_\vartheta(x, y), \tag{24}$$

with

$$F_\vartheta(x, y) = \frac{h_x(x, y)\,\sin\vartheta + h_y(x, y)\,\cos\vartheta}{h_x(x, y)\,\cos\vartheta - h_y(x, y)\,\sin\vartheta}.$$

Similarly, a trajectory of $X_{-\vartheta}$ is the graph of a function $y \mapsto x(y)$, where $x(y)$ is a solution of the differential equation

$$\frac{dx}{dy} = G_\vartheta(x, y),$$

with

$$G_\vartheta(x, y) = \frac{1}{F_\vartheta(x, y)} = \frac{h_x(x, y)\cos\vartheta - h_y(x, y)\sin\vartheta}{h_x(x, y)\sin\vartheta + h_y(x, y)\cos\vartheta}.$$

Here y ranges over the full interval $[y_0, y_1]$ if $x_0 \le x(y) \le x_1$, or a suitable maximal subinterval otherwise.

The union of all quasihorizontal gridboxes in $\mathbb{D}^*$ is denoted by ${\mathbb{D}^*}_{\mathrm{qh}}$, and the union of all quasivertical gridboxes by ${\mathbb{D}^*}_{\mathrm{qv}}$.

Even though the width of a funnel may grow exponentially in the number of grid boxes it traverses, this growth is controlled. To this end, we introduce several *computable* constants that only depend on the function h and (the size of) its domain $\mathbb{D}^*$. Let A_{qh}, A_{qv}, B_{qh}, B_{qv}, C_{qh} and C_{qv} be positive constants such that

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qh}}} |F(x, y)| \le A_{\mathrm{qh}}, \quad \max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qv}}} |G(x, y)| \le A_{\mathrm{qv}},$$

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qh}}} |\frac{\partial F}{\partial x}(x, y)| \le B_{\mathrm{qh}}, \quad \max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qv}}} |\frac{\partial G}{\partial y}(x, y)| \le B_{\mathrm{qv}},$$

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qh}}} |\frac{\partial F}{\partial y}(x, y)| \le C_{\mathrm{qh}}, \quad \max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qv}}} |\frac{\partial G}{\partial x}(x, y)| \le C_{\mathrm{qv}}.$$

Note that

$$F_\vartheta(x, y) - F(x, y) = \frac{(h_x(x, y)^2 + h_y(x, y)^2)\sin\vartheta}{h_x(x, y)^2\cos\vartheta - h_x(x, y)\,h_y(x, y)\sin\vartheta}.$$

Let $M^{(1)}_{\mathrm{qh}}$ be a dyadic number such that

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qh}}} \left|\frac{h_y(x, y)}{h_x(x, y)}\right| \le M^{(1)}_{\mathrm{qh}}. \tag{25}$$

Take $\vartheta_{\mathrm{qh}} \in (0, \frac{1}{2}\pi)$ such that $\tan\vartheta_{\mathrm{qh}} \le \dfrac{1}{2M^{(1)}_{\mathrm{qh}}}$. Finally, let $M^{(2)}_{\mathrm{qh}}$ be a constant such that

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qh}}} \left|\frac{h_x(x, y)^2 + h_y(x, y)^2}{h_x(x, y)^2\cos\vartheta_{\mathrm{qh}}}\right| \le M^{(2)}_{\mathrm{qh}}. \tag{26}$$

Taking $M_{\mathrm{qh}} = \dfrac{M^{(2)}_{\mathrm{qh}}}{2M^{(1)}_{\mathrm{qh}}}$, we have, for $|\vartheta| \le \vartheta_{\mathrm{qh}}$:

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qh}}} |F_\vartheta(x, y) - F(x, y)| \le M_{\mathrm{qh}}\sin\vartheta. \tag{27}$$

Similarly, there are (computable) constants M_{qv} and ϑ_{qv} such that

$$\max_{(x,y)\in{\mathbb{D}^*}_{\mathrm{qv}}} |G_\vartheta(x, y) - G(x, y)| \le M_{\mathrm{qv}}\sin\vartheta, \tag{28}$$

for $|\vartheta| \le \vartheta_{\mathrm{qv}}$. Finally, let the constants c_0, c_1 and ϑ_0 be defined by

$$c_0 = 2\max(C_{\mathrm{qh}} + A_{\mathrm{qh}}B_{\mathrm{qh}}, C_{\mathrm{qv}} + A_{\mathrm{qv}}B_{\mathrm{qv}}) \tag{29}$$

$$c_1 = \max(\frac{1}{2M_{\mathrm{qh}}K_{\mathrm{qh}}}, \frac{1}{2M_{\mathrm{qv}}K_{\mathrm{qv}}}) \tag{30}$$

$$\vartheta_0 = \max(\vartheta_{\mathrm{qh}}, \vartheta_{\mathrm{qv}}). \tag{31}$$

The next result provides an upper bound for the growth of the funnel width along a quasihorizontal part of its bounding polylines. We assume that the funnel runs from left to right, so its initial points are on the line with smallest x-coordinate. If the funnel runs from right to left, a similar result is obtained.

Lemma 5.6. *Let $y_{\vartheta,w}, y_{-\vartheta,w} : [a,b] \to [c,d]$ be the piecewise linear functions the graphs of which are quasihorizontal parts of the polylines L_ϑ and $L_{-\vartheta}$ for a grid with edge length w, respectively. Let Δ be an upper bound for the distance of the initial points of these polylines, i.e.,*

$$|y_{\vartheta,w}(a) - y_{-\vartheta,w}(a)| \le \Delta.$$

Then the width of the fence, bounded by L_ϑ and $L_{-\vartheta}$, is bounded:

$$|y_{\vartheta,w}(x) - y_{-\vartheta,w}(x)| \le \Delta\, e^{C_{\mathrm{qh}}(x-a)} + (c_0 w + c_1 \sin\vartheta)\, \frac{e^{C_{\mathrm{qh}}(x-a)} - 1}{C_{\mathrm{qh}}}.$$

Proof. Let $y_{\pm\vartheta}(x)$ be the exact solution of the rotated system with initial condition $y_{\pm\vartheta}(a)$. In particular, $|y_\vartheta(a) - y_{-\vartheta}(a)| \le \Delta$. Then (27) implies

$$\Big|\frac{dy_{\pm\vartheta}}{dx}(x) - F(x, y_{\pm\vartheta}(x))\Big| = \big|F_{\pm\vartheta}(x, y_{\pm\vartheta}(x) - F(x, y_{\pm\vartheta}(x))\big| \le M_{\mathrm{qh}} \sin\vartheta.$$

Therefore, according to the Fundamental Inequality (Hubbard and West, 1991, Theorem 4.4.1)—See also Appendix A—we have

$$|y_\vartheta(x) - y_{-\vartheta}(x)| \le \Delta\, e^{C_{\mathrm{qh}}(x-a)} + \frac{2M_{\mathrm{qh}} \sin\vartheta}{C_{\mathrm{qh}}}\,(e^{C_{\mathrm{qh}}(x-a)} - 1). \tag{32}$$

The interval $[a,b]$ is subdivided into a finite number of subintervals of length at most w, where the endpoints correspond to the x-coordinates of the breakpoints of the fences L_ϑ and $L_{-\vartheta}$. Let $y_{\vartheta,w}$ be the Euler approximation to the ordinary differential equation (24). Its graph is (a quasihorizontal) part of the fence L_ϑ. Theorem 4.5.2 in Hubbard and West (1991)—See also Appendix A—gives the following explicit bound for the error in Euler's method:

$$|y_{\vartheta,w}(x) - y_\vartheta(x)| \le w\, \frac{B_{\mathrm{qh}} + A_{\mathrm{qh}} C_{\mathrm{qh}}}{C_{\mathrm{qh}}}\,(e^{C_{\mathrm{qh}}(x-a)} - 1). \tag{33}$$

We get a similar inequality for the Euler approximation $y_{-\vartheta,w}$ of $y_{-\vartheta}$. Combining (32) and (33), and using (29) and (30), yields

$$\begin{aligned} |y_{\vartheta,w}(x) - y_{-\vartheta,w}(x)| &\le \Delta\, e^{C_{\mathrm{qh}}(x-a)} + \\ &\quad 2(w(B_{\mathrm{qh}} + A_{\mathrm{qh}} C_{\mathrm{qh}}) + M_{\mathrm{qh}} \sin\vartheta)\, \frac{e^{C_{\mathrm{qh}}(x-a)} - 1}{C_{\mathrm{qh}}} \\ &= \Delta\, e^{C_{\mathrm{qh}}(x-a)} + (c_0 w + c_1 \sin\vartheta)\, \frac{e^{C_{\mathrm{qh}}(x-a)} - 1}{C_{\mathrm{qh}}}. \end{aligned}$$

□

A similar result holds for quasivertical trajectories. Next we need to control the increase of the funnel width upon transition from a quasihorizontal to a quasivertical part of its bounding polylines (or from a quasivertical to a quasihorizontal part).

Transitions: bounded increase of funnel width Transition from a quasihorizontal to a quasivertical, or from a quasivertical to a quasihorizontal part of the funnel takes place at an ε-box. If the width of the funnel at the 'entry' of the box is less than the width w of a grid box, then the width may increase, but it will not be greater than $2w$ at the exit. This is made more precise by the following result.

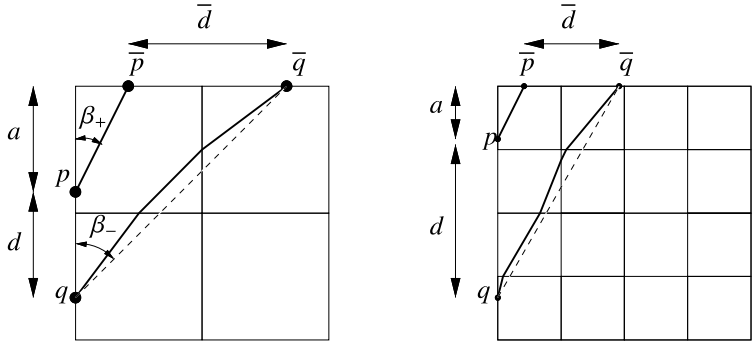

Fig. 19. The distance between the two polylines upon entry and exit of a box. Left: If the distance d between the initial points p and q of the polylines is less than the edge-length w of a grid box, then the distance $\overline{d}$ between the terminal points $\overline{p}$ and $\overline{q}$ is less than $2w$. Right: Otherwise, the distance $\overline{d}$ between the terminal points is less than d.

Lemma 5.7. *Let J be a ε-box as in Lemma 5.5, where, moreover, the initial points p and q of $L_\vartheta \cap J$ and $L_{-\vartheta} \cap J$, respectively, are on the boundary of the gridboxes containing the vertices of edge e of J. If the distance between p and q is at least w, then the distance between the terminal points $\overline{p}$ and $\overline{q}$ of $L_\vartheta \cap J$ and $L_{-\vartheta} \cap J$, respectively, is less than the distance of p and q. If the distance between p and q is less than w, then the distance of $\overline{p}$ and $\overline{q}$ is at most $2w$. (See Fig. 19.)*

Proof. Assume that the first edge of L_ϑ is quasivertical, but not quasihorizontal. Edge e of J is then vertical. Assume that this polyline consists of a single edge, namely the line segment $p\overline{p}$. Let β_+ be the angle between $p\overline{p}$ and edge e, then $\arctan\frac{1}{2} \le \beta_+ \le \vartheta + \arctan\frac{1}{2}$. Let β_- be the angle between the line segment $q\overline{q}$ and edge e, then β_- is inbetween the smallest and largest slope of any edge of $L_{-\vartheta}$. Since the angle variation of X over J is less than $\frac{\pi}{20}$, the angle β_- is greater than $\beta_+ - \frac{\pi}{20}$. Let a be the distance of p to the nearest vertex of e, then $a \le w$. If $d \ge w$, the distance $\overline{d}$ between $\overline{p}$ and $\overline{q}$ satisfies

$$\begin{aligned}\overline{d} &= (d+a)\tan\beta_- - a\tan\beta_+\\ &\le d\tan(\beta_+ + \tfrac{\pi}{20}) + a(\tan(\beta_+ + \tfrac{\pi}{20}) - \tan\beta_+)\\ &< \tfrac{3}{4}d + \tfrac{1}{4}a\\ &\le d,\end{aligned}$$

since $a \le w \le d$. Here we used $\tan\beta_+ \le \frac{1}{2}$ to get

$$\tan(\beta_+ + \tfrac{\pi}{20}) = \frac{\tan\beta_+ + \tan\frac{\pi}{20}}{1 - \tan\beta_+\tan\frac{\pi}{20}} \le \frac{\frac{1}{2} + \tan\frac{\pi}{20}}{1 - \frac{1}{2}\tan\frac{\pi}{20}} \le \tfrac{3}{4}.$$

Since $\arctan\frac{1}{2} - \frac{\pi}{40} \le \arctan\frac{1}{2} - \vartheta \le \beta_+ \le \arctan\frac{1}{2}$, a short computation shows that $\tan(\beta_+ + \frac{1}{20}\pi) - \tan\beta_+ < \frac{1}{4}$.

If $d < w$, then q lies in the same gridbox as p, or in a gridbox adjacent to it. Then it is easy to see that $\overline{p}$ lies in the same grid box as p, and $\overline{q}$ also lies in this box, or in a box adjacent to it. Therefore, $\overline{d} \le 2w$ in this case.

If $L_{-\vartheta}$ consists of a single edge, then the argument is similar. □

Lemmas 5.6 and 5.7 provide the following result on the upper bound on the funnel width of a separatrix with M transitions between quasihorizontal and quasivertical parts.

Corollary 5.8. *Let T be the (computable) edge length of a bounding square of the domain $\mathbb{D}$ of the function h, and let M be the total number of quasihorizontal and quasivertical parts of the polylines bounding a separatrix funnel. Let $C = \max(C_{\mathrm{qh}}, C_{\mathrm{qv}})$ and let $D = \min(C_{\mathrm{qh}}, C_{\mathrm{qv}})$. Then the width of the funnel does not exceed*

$$(c_1\vartheta + c_2 w)\,\frac{e^{CMT}}{D},$$

provided $\vartheta \le \vartheta_0$ where $c_2 = 2 + \dfrac{c_0}{D}$, with c_0 and c_1 given by (29) and (30), respectively.

In particular, this width is at most ε if

$$c_1\vartheta + c_2 w \le \frac{D}{e^{CMT}}\,\varepsilon. \tag{34}$$

Proof. Let $\mathbb{D}^* \subset [a,b]\times[c,d]$, then $T \le \max(b-a, d-c)$. There are $M-1$ transitions from quasihorizontal to quasivertical parts of the funnel, each occurring at an ε-box. Let Δ_0 be the width of the initial separatrix interval, and let $\Delta_1, \ldots, \Delta_{M-1}$ be the width of the funnel at the entry of the corresponding boxes, in other words, Δ_k is the width at the end of the k-th part of the funnel. Using induction, we will prove that, for $k = 1, \ldots, M$:

$$\Delta_k \le 2w\,e^{kCT} + \frac{c_0 w + c_1 \sin\vartheta}{D}\,(e^{kCT} - 1). \tag{35}$$

So assume (35) holds for $k = n-1$. If $\Delta_{n-1} > w$, the initial width of the n-th part of the funnel does not exceed Δ_{n-1}, cf. Lemma 5.7. Assume that the n-th part of the funnel is quasihorizontal, then Lemma 5.6 implies that the width of this part at a point with horizontal coordinate x is at most

$$\Delta_{n-1}\,e^{C_{\mathrm{qh}}(x-a)} + (c_0 w + c_1 \sin\vartheta)\,\frac{e^{C_{\mathrm{qh}}(x-a)} - 1}{C_{\mathrm{qh}}},$$

so in particular, since $D \le C_{\mathrm{qh}} \le C$ and $0 \le x - a \le T$:

$$\Delta_n \le \Delta_{n-1}\,e^{CT} + \frac{c_0 w + c_1 \sin\vartheta}{D}\,(e^{CT} - 1).$$

Therefore, (35) holds for $k = n$. If $\Delta_{n-1} \le w$, then the initial width of the n-th part of the funnel is at most $2w$, cf. Lemma 5.7. Therefore, Lemma 5.6 implies that the width of this part at a point with horizontal coordinate x is at most

$$2w\,e^{C_{\mathrm{qh}}(x-a)} + (c_0 w + c_1 \sin\vartheta)\,\frac{e^{C_{\mathrm{qh}}(x-a)} - 1}{C_{\mathrm{qh}}},$$

so in particular

$$\begin{aligned}\Delta_n &\le 2w\,e^{CT} + \frac{c_0 w + c_1 \sin\vartheta}{D}\,(e^{CT} - 1)\\ &\le 2w\,e^{nCT} + \frac{c_0 w + c_1 \sin\vartheta}{D}\,(e^{nCT} - 1).\end{aligned}$$

Therefore, for $n = M$, we have

$$\begin{aligned}\Delta_M &\le (c_1 \sin\vartheta + c_2 w)\,\frac{e^{CMT} - 1}{D}\\ &\le (c_1\vartheta + c_2 w)\,\frac{e^{CMT}}{D},\end{aligned}$$

which proves the corollary. □

Remark 5.9. The computable constants ϑ_0, c_1, c_2, C and D depend only on $\mathbb{D}^*$ and h.

In the next section, we assemble the bits and pieces into a certified algorithm for the construction of the MS-complex, and show how the upper bounds on the funnel width are used to prove that this algorithm terminates.

5.3. Construction of the MS-complex

The algorithm The construction of the MS-complex of h is a rather straightforward application of the preceding results. It uses a parameter M, the (a priori unknown) number of transitions (at ε-boxes) between quasihorizontal and quasivertical parts of a funnel. Let T be the edge length of a bounding square of the domain $\mathbb{D}$ of h. Then the algorithm performs the following steps.

Step 1. Construct certified isolating boxes $B'_1, \ldots, B'_m$ for the singularities of ∇h (cf. Section 3).

Step 2. Let $\mathbb{D}^*$ be the closure of $\mathbb{D} \setminus (B'_1 \cup \cdots \cup B'_m)$. Compute the constants ϑ_0, c_1, c_2, C and D, which depend only on h and $\mathbb{D}^*$. Set ε to the minimum of the width of the source-, sink- and saddleboxes.

Step 3. Let ϑ and w be such that $w \leq \frac{\vartheta}{2C_0\sqrt{2}}$, $\vartheta \leq \min(\frac{\pi}{40}, \vartheta_0)$, and $c_1\vartheta + c_2 w \leq \varepsilon\, D\, e^{-CMT}$ (cf. Corollary 5.8). Subdivide $\mathbb{D}^*$ until all gridboxes have maximum width w. For each saddle box, compute four separatrix intervals on its boundary, of width at most w.

Step 4. For each stable and unstable separatrix interval do the following. Start the computation of a funnel for a separatrix by setting M to a small number M_0 (say 4). Compute the fences $L_{-\vartheta}$ and L_{ϑ}, keeping track of the width of the enclosed funnel under construction and of the number m of transitions between quasihorizontal and quasivertical parts of this funnel.

If *the width of the funnel exceeds ε or the number of transitions m exceeds M*, then abort the computation of the current funnel, discard all funnels constructed so far, set M to twice its current value and goto Step 3.

If *the funnel intersects an already constructed funnel, or a source- or sinkbox on which it does not terminate* (i.e., if only one of its fences intersects this box), then set ε to half its current value, discard all funnels constructed so far, and goto Step 3.

If *the funnel intersects a saddlebox B'_i*, then decrease the size of B'_i by a factor of two via *subdivision*, discard all funnels constructed so far, set ε to half its current value, and goto Step 2. (Note that $\mathbb{D}^*$ gets larger, so the constants in Step 2 have to be recomputed.)

Otherwise, the fences end on the same component of the boundary of $\partial\mathbb{D}^*$. The enclosed funnel is simply connected, and does not intersect any of the funnels constructed so far. Add the funnel to the output, and reset M to M_0 (and repeat until all separatrices have been processed).

5.4. Termination

Since the gradient field ∇h is a $2D$ Morse–Smale system, its separatrices are disjoint. Their intersections with $\mathbb{D}^*$ are compact, and have positive distance (although this distance is not known a priori). In the main loop of the algorithm, the maximal funnel width ε is bisected if funnels intersect, and saddleboxes intersected by the funnel are subdivided, so after a finite number of iterations of the main loop its value is less than half the minimum distance between any pair of distinct separatrices, and funnels stay clear from saddleboxes (apart from the one containing the α- or ω-limit of the enclosed separatrix).

Separatrices that intersect $\partial\mathbb{D}$ do so transversally, cf. Remark 2.1. Therefore, after a finite number of subdivision steps, both fences around such separatrices will intersect $\partial\mathbb{D}$ transversally. Hence, eventually all funnels become disjoint, at which point the algorithm terminates after returning a topologically correct MS-complex for ∇h.

6. Implementation and experimental results

The algorithm has been implemented using the `Boost` library (http://www.boost.org) for IA. All experiments have been performed on a 3 GHz Intel Pentium 4 machine under Linux with 1 GB RAM

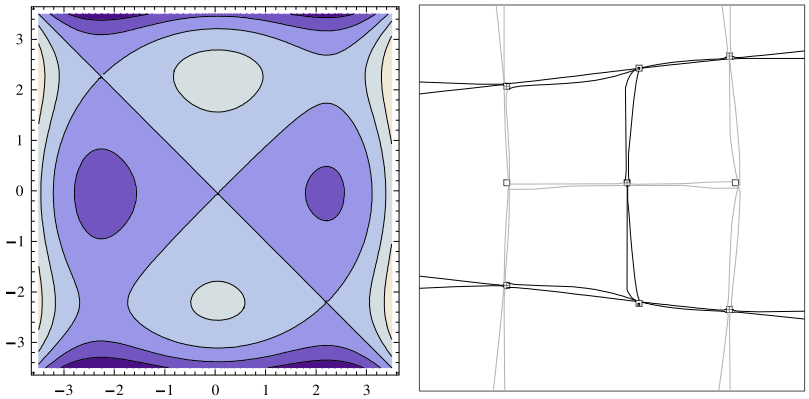

(a) Contour plot (left) and MS-Complex (right) of $h(x, y) = -10\,x^2 + x^4 + 10\,y^2 - y^4 + x + xy^2$, on the box $[-4, 3.5] \times [-4, 3.5]$. CPU-time: 11 seconds.

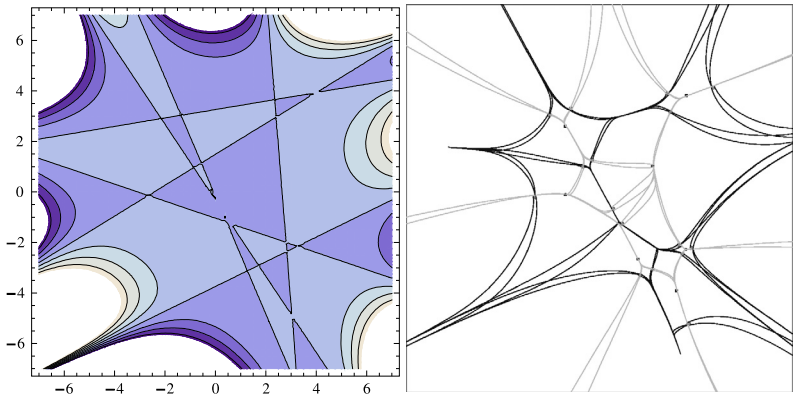

(b) Contour plot (left) and MS-complex (right) of a product of seven linear functions, on the box $[-7, 7] \times [-7, 7]$. CPU-time: 11.5 minutes.

Fig. 20. Contour plots of MS-functions and their Morse–Smale complexes.

using the g++ compiler, version 3.3.5. Figs. 1(a)–(b) and 20(a)–(b) depict the output of our algorithm, for several Morse–Smale functions. In our implementation the parameter ϑ, used in the construction of separatrix-funnels, is $\frac{\pi}{30}$, which is larger than the theoretical bound given by Corollary 5.8. The algorithm halves this angle several times, depending on the input, until the funnels are simply connected, mutually disjoint, and connect a saddle-box to a source-box (for stable separatrices) or sink-box (for unstable separatrices), in which case a topologically correct MS-complex has been computed. Each of the funnels with deep black boundaries contains an unstable separatrix, whereas a funnel with light black boundaries contains a stable separatrix. The CPU-time for computing a MS-system increases with the number of critical points and the complexity of the vector field, as indicated in the captions of the figures.

7. Conclusion

The outcome of our research is two-fold. Firstly, we compute the topologically correct MS-complex of a Morse–Smale system. The correct saddle–sink or saddle–source connectivity can also be represented as a graph, which is of special interest from different application point of view. On the other hand, depending on a user-specified width of funnel one can compute a geometrically close approximation of the MS-complex. We give the proof of convergence of our algorithms. However, the

complexity of the given algorithm depends on the input function and the complexity of the interval arithmetic library used in the algorithm. As we discussed some of the separatrices inside a bounding box B may have discontinuous components. The algorithm we propose here is able to compute only the part of the separatrices which are connected to the corresponding saddle. Therefore one open question is how to compute all the components of separatrices inside a bounding box.

Acknowledgements

The work reported in this paper was funded by the Dutch National Science Foundation (NWO), project code 435038 and partially supported by the University of Groningen, project code 135170.

Appendix A. Mathematical results used in the text

Error in Euler's method Error bounds for approximate solutions of ordinary differential equation play a crucial role in the construction of certified funnels for separatrices. We quote the relevant parts of the book (Hubbard and West, 1991).

Fundamental Inequality. (See Hubbard and West, 1991, Theorem 4.4.1.) Consider the differential equation

$$\frac{dy}{dx} = F(x, y)$$

on a box $B = [a, b] \times [c, d]$, and let C be a constant such that

$$\max_{(x,y)\in B} |\frac{\partial F}{\partial y}(x, y)| \leq C.$$

If $y_1(x)$ and $y_2(x)$ are two approximate piecewise differentiable solutions satisfying

$$|y_1'(x) - F(x, y_1(x))| \leq \varepsilon_1,$$
$$|y_2'(x) - F(x, y_2(x))| \leq \varepsilon_2$$

for all $x \in [a, b]$ at which $y_1(x)$ and $y_2(x)$ are differentiable, and if, for some $x_0 \in [a, b]$

$$|y_1(x_0) - y_2(x_0)| \leq \delta,$$

then, for all $x \in [a, b]$

$$|y_1(x) - y_2(x)| \leq \delta e^{C|x-x_0|} + \varepsilon \frac{e^{C|x-x_0|} - 1}{C},$$

where $\varepsilon = \varepsilon_1 + \varepsilon_2$.

The well-known Euler method for constructing approximate solutions to ordinary differential equations is also useful for the construction of certified strips. It proceeds as follows. For a given initial position (x_0, y_0), define the sequence of points (x_n, y_n) by

$$x_n = x_{n-1} + \eta = x_0 + n\eta$$
$$y_n = y_{n-1} + \eta\, F(x_{n-1}, y_{n-1}),$$

as long as $(x_n, y_n) \in B$. Then the *Euler approximate solution* $y_\eta(x)$ through (x_0, y_0) with step η is the piecewise linear function the graph of which joins the points (x_n, y_n), so

$$y_h(x) = y_n + (x - x_n)\, F(x_n, y_n) \quad \text{for } x \in [x_n, x_{n+1}].$$

The following result states that the Euler approximate solution converges to the actual solution as the step tends to zero, and gives a bound for the error.

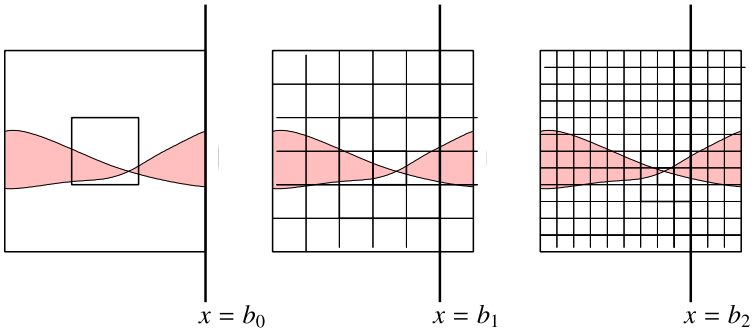

Fig. 21. Zooming in on the saddle point by subdivision.

Error in Euler's method. (See Hubbard and West, 1991, Theorem 4.5.2.) Consider the differential equation

$$\frac{dy}{dx} = F(x, y)$$

on a box $B = [a, b] \times [c, d]$, where F is a C^2-function on B. Let the constants A, B and C satisfy

$$\max_{(x,y)\in B} |F(x, y)| \le A, \quad \max_{(x,y)\in B} |\frac{\partial F}{\partial x}(x, y)| \le B, \quad \max_{(x,y)\in B} |\frac{\partial F}{\partial y}(x, y)| \le C.$$

The deviation of the Euler approximate solution y_η with step η from a solution y of the differential equation with $|y(a) - y_\eta(a)| \le \Delta$ satisfies

$$|y_\eta(x) - y(x)| \le \Delta\, e^{C|x-a|} + \eta\,(B + AC)\,\frac{e^{C|x-a|-1}}{C},$$

for all $x \in [a, b]$.

The preceding result also holds if, as in the current chapter, η is not the exact step, but an upper bound for a possibly varying step.

Appendix B. Narrowing separatrix intervals

We first sketch the algorithm for narrowing the separatrix intervals. To this end we subdivide the box I, and hence the box $N(I)$, yielding a nested sequence of boxes $I = I_0 \supset I_1 \supset \ldots$, with surrounding boxes $N(I) = N(I_0) \supset N(I_1) \supset \ldots$, such that

1. $\text{width}(I_{n+1}) = \frac{1}{2}\,\text{width}(I_n)$
2. the saddle point p is contained in box I_n, for all n.

See Fig. 21. Let s be the x-coordinate of the saddle point p, and let b_n be the x-coordinate of the rightmost vertical boundary edge of $N(I_n)$. Let w_n be the width of I_n, and let c_n be the x-coordinate of its center. Then $b_n = c_n + \frac{3}{2}w_n$, and $w_{n+1} = \frac{1}{2}w_n$. Since then $b = b_0 > b_1 > \ldots$, since

$$b_{n+1} = c_{n+1} + \tfrac{3}{2}w_{n+1} \le c_n + \tfrac{1}{4}w_n + \tfrac{3}{4}w_n = b_n - \tfrac{1}{2}w_n.$$

Since $|c_n - s| \le \frac{1}{2}w_n$, we get

$$w_n \le b_n - s \le 2w_n. \tag{B.1}$$

Consider the forward integral curves of the vector field ∇h through the points of intersection $q_n^{\pm}$ of the line $x = b_n$ and the boundary curves $\Gamma^u_{\pm\beta}$. See Fig. 22. These curves intersect the rightmost edge

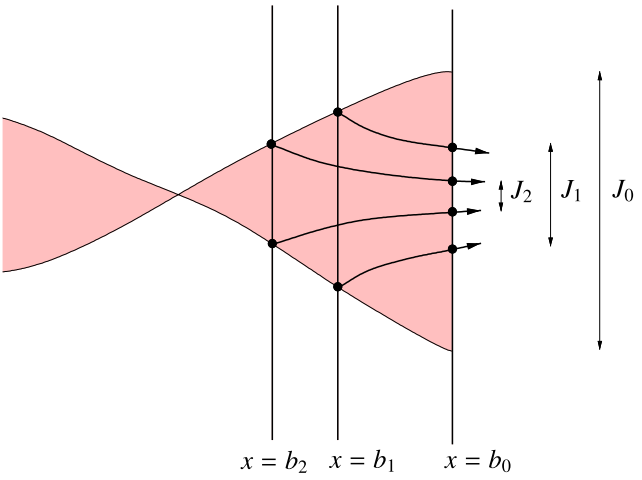

Fig. 22. Narrowing separatrix intervals.

of $N(I)$ in two points bounding an interval J_n on this edge. Arbitrarily good approximations of these integral curves are obtained as follows. Let ϑ_n be (an upper bound on) the maximum angle variation of ∇h over any of the boxes of the n-th subdivision of $N(I)$. Since h is C^2, the angle variation is a Lipschitz function, so $\lim_{n\to\infty} \vartheta_n = 0$. In particular, the rotated vector fields $X_{\pm\vartheta_i}$ converge to ∇h. We construct an upper fence with angle $\frac{1}{2}\vartheta_n$ for the upper integral curve, and a lower fence with angle $-\frac{1}{2}\vartheta_n$ for the lower integral curve. See also Section 5.1 for the construction of a fence. These fences are disjoint, since the angle variation of ∇h over a grid box is less than ϑ_n.

Since $\lim_{n\to\infty} b_n = s$, cf (B.1), the points $q_n^{\pm}$ converge to the saddle point. Therefore, the intervals $J_0 \supset J_1 \supset \ldots$, contained in the intersection of the unstable wedge C_β^u and the rightmost vertical edge of $N(I)$, converge to the intersection of the unstable separatrix and the rightmost vertical edge of $N(I)$.

A proof of Lemma 4.4 can be given along the lines of Coddington and Levinson (1955, page 330ff) or Chow and Hale (1982, Chapter 3.6). Rather than giving a complete proof we give an example illustrating the main ideas. Consider the function $h(x, y) = \frac{1}{2}\lambda_u x^2 + \frac{1}{2}\lambda_s y^2$, with $\lambda_s < 0 < \lambda_u$. The gradient vector field is given by $\nabla h(x, y) = (\lambda_u x, \lambda_s y)^T$. Obviously, the origin is a saddle point, with positive eigenvalue λ_u and negative eigenvalue λ_s, and eigenvectors $(1, 0)^T$ and $(0, 1)^T$, respectively. The unstable cone of this saddle point is bounded by the curves $\Gamma^u_{\pm\beta}$, defined implicitly by $\psi^u_{\pm\beta}(x, y) = 0$, where

$$\begin{aligned}\psi^u_{\pm\beta}(x, y) &= \det(V^u, X_\beta(x, y)) \\ &= \begin{vmatrix} 1 & (\lambda_u \cos\beta)\, x \mp (\lambda_s \sin\beta)\, y \\ 0 & (\lambda_u \sin\beta)\, x \pm (\lambda_s \cos\beta)\, y \end{vmatrix} \\ &= (\lambda_u \sin\beta)\, x \pm (\lambda_s \cos\beta)\, y.\end{aligned}$$

Therefore, the equation of $\Gamma^u_{\pm\beta}$ is $y = \pm a x$, with $a = -(\tan\beta)\frac{\lambda_u}{\lambda_s} > 0$. Let the right vertical edge of I be on the line $x = b$, $b > 0$, and consider a point $q^+ = (\xi, a\,\xi)$, with $0 < \xi < b$, on the boundary curve Γ^u_β of the unstable cone. The integral curve of ∇h through q^+ satisfies the differential equation

$$\frac{dy}{dx} = \Lambda\, \frac{y}{x},$$

with initial condition $y(\xi) = a\xi$, where $\Lambda = \dfrac{\lambda_s}{\lambda_u} < 0$. Therefore,

$$y(x) = a\,\xi\left(\frac{x}{\xi}\right)^{\Lambda}.$$

The integral curve through q^+ intersects the rightmost edge of I in the point $(b, \delta(\xi))$, where

$$\delta(\xi) = \frac{ab^{\Lambda}}{\xi^{\Lambda-1}}.$$

Similarly, the integral curve through $q^- = (\xi, -a\xi) \in \Gamma^u_{-\beta}$ intersects the rightmost edge of I in the point $(b, -\delta(\xi))$. Now let ξ range over the sequence $b_0, b_1, \ldots$. Then the interval J_n has endpoints $(b_n, \pm\delta(b_n))$, so its width is $2\delta(b_n)$. In view of (B.1), with $s = 0$, we have

$$2\frac{ab^{\Lambda}}{w_n^{\Lambda-1}} \le \text{width}(J_n) \le 2\frac{ab^{\Lambda}}{(2w_n)^{\Lambda-1}}.$$

In other words, with $K = 1 - \Lambda > 1$ and $c = ab^{\Lambda}\, w_0^K$,

$$c\left(\tfrac{1}{2}\right)^{Kn} \le \text{width}(J_n) \le 2^K\, c\left(\tfrac{1}{2}\right)^{nK}.$$

Hence,

$$\text{width}(J_{n+2}) \le \frac{1}{2^K}\,\text{width}(J_n).$$

Since $K > 1$, after two subdivision steps the size of the separatrix interval reduces by more than a factor two. Hence interval arithmetic provides an arbitrarily good approximation of the intersection of the unstable separatrix and the boundary of the saddle box. A similar observation holds for the intersection of the other separatrices and the boundary of their saddle box.

References

Arnol'd, V., 2006. Ordinary Differential Equations. Universitext. Springer-Verlag, New York, Heidelberg, Berlin.

Bajaj, C., Schikore, D., 1998. Topology preserving data simplification with error bounds. Comput. Graph. 22 (1), 3–12.

Banchoff, T.F., 1970. Critical points and curvature for embedded polyhedral surfaces. Am. Math. Mon. 77, 475–485.

Biasotti, S., Floriani, L.D., Falcidieno, B., Frosini, P., Giorgi, D., Landi, C., Papaleo, L., Spagnuolo, M., 2008. Describing shapes by geometrical-topological properties of real functions. ACM Comput. Surv. 40 (4), 12.1–12.87.

`Boost` interval arithmetic library, http://www.boost.org.

Cazals, F., Chazal, F., Lewiner, T., 2003. Molecular Shape Analysis based upon the Morse–Smale Complex and the Connolly Function. In: Proceedings of the 19th ACM Symposium on Computational Geometry. ACM Press, New York, NY, pp. 351–360.

Chattopadhyay, A., Vegter, G., Yap, C., 2012. Certified Computation of Planar Morse–Smale Complexes. In: Proceedings of the 27th ACM Symposium on Computational Geometry. Chapel Hill, pp. 259–268.

Chow, S.-N., Hale, J., 1982. Methods of Bifurcation Theory. Grundlehren der mathematischen Wissenschaften, vol. 251. Springer-Verlag, New York, Heidelberg, Berlin.

Coddington, E., Levinson, N., 1955. Theory of Ordinary Differential Equations. McGraw–Hill Book Company.

Edelsbrunner, H., Harer, J., Natarajan, V., Pascucci, V., 2003b. Morse–Smale complexes for piecewise linear 3-manifolds. In: Proc. 19th Ann. Sympos. Comput. Geom., pp. 361–370.

Edelsbrunner, H., Harer, J., Zomorodian, A., 2003a. Hierarchical Morse–Smale complexes for piecewise linear 2-manifolds. Discrete Comput. Geom. 30, 87–107.

Forman, R., 1998. Morse theory for cell complexes. Adv. Math. 134, 90–145.

Gyulassy, A., Bremer, P., Hamann, B., Pascucci, V., 2008. A practical approach to Morse–Smale complex computation. IEEE Trans. Vis. Comput. 14, 1619–1626.

Helman, J.L., Hesselink, L., 1991. Visualizing vector field topology in fluid flows. IEEE Comput. Graph. Appl. 11 (3), 36–46.

Hirsch, M.W., Smale, S., 1974. Differential Equations, Dynamical Systems, and Linear Algebra. Academic Press.

Hubbard, J., West, B., 1991. Differential Equations. A Dynamical Systems Approach. Part I. Texts Appl. Math., vol. 5. Springer Verlag, New York, Heidelberg, Berlin.

Li, C., Pion, S., Yap, C., 2004. Recent progress in Exact Geometric Computation. J. Log. Algebraic Program. 64 (1), 85–111. Special issue on "Practical Development of Exact Real Number Computation".

Lin, L., Yap, C., 2011. Adaptive isotopic approximation of nonsingular curves: the parameterizability and nonlocal isotopy approach. Discrete Comput. Geom. 45 (4), 760–795.

Meyer, F., 1994. Topographic distance and watershed lines. Signal Process. 38, 113–125.

Milnor, J., 1968. Morse Theory. Princeton University Press.

Moore, R., 1996. Interval Analysis. Prentice-Hall.

Palis, J., de Melo, W., 1982. Geometric Theory of Dynamical Systems: An Introduction. Springer-Verlag.

Plantinga, S., Vegter, G., 2007. Isotopic meshing of implicit surfaces. Vis. Comput. 23, 45–58.
Snyder, J., 1992a. Generative Modeling for Computer Graphics and CAD. Symbolic Shape Design Using Interval Analysis. Academic Press Professional, Inc., San Diego, CA, USA.
Snyder, J., 1992b. Interval analysis for computer graphics. SIGGRAPH Comput. Graph. 26 (2), 121–130.
Snyder, J.M., 1992c. Interval analysis for computer graphics. SIGGRAPH Comput. Graph. 26 (2), 121–130.
Snyder, J.M., 1992d. Generative Modeling for Computer Graphics and CAD: Symbolic Shape Design Using Interval Analysis. Academic Press Professional, Inc., San Diego, CA, USA.
Takahashi, S., Ikeda, T., Shinagawa, Y., Fujishiro, I., 1995. Algorithms for extracting correct critical points and constructing topological graphs from discrete geographic elevation data. Comput. Graph. Forum 14 (3), 181–192.
Yap, C., 2009. In praise of numerical computation. In: Albers, S., Alt, H., Näher, S. (Eds.), Efficient Algorithms. In: Lecture Notes in Computer Science, vol. 5760. Springer-Verlag, pp. 308–407. Essays dedicated to Kurt Mehlhorn on the Occasion of His 60th Birthday.
Zhang, E., Mischaikow, K., Turk, G., 2006. Vector field design on surfaces. ACM Trans. Graph. 25 (4), 1294–1326.