# Robust Morse Decompositions of Piecewise Constant Vector Fields Supplementary Material

Andrzej Szymczak and Eugene Zhang

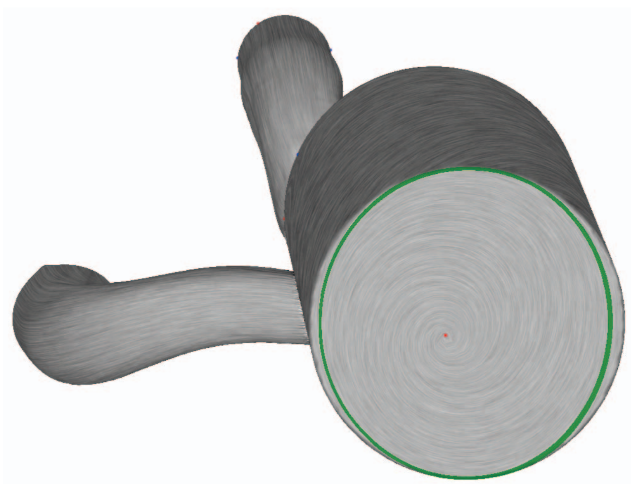

Fig. 1. The Morse set containing a large periodic orbit in the diesel engine dataset (7 refinement steps).

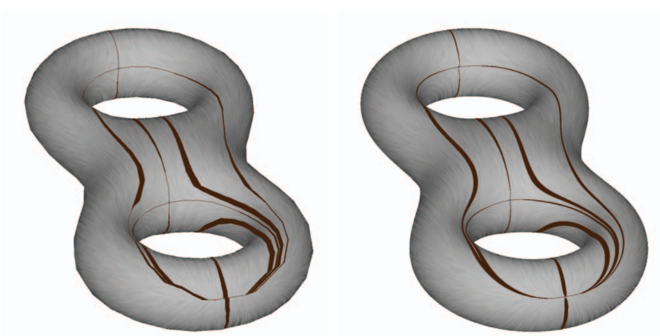

Fig. 2. Results for a height field on the original 1536-triangle mesh (left) and its subdivided version (three subdivision iterations). Note that only a source and two saddles are visible from this viewpoint.


- *Andrzej Szymczak is with Colorado School of Mines E-mail: aszymcza@mines.edu*
- *Eugene Zhang is with Oregon State University E-mail: zhange@eecs.oregonstate.edu*


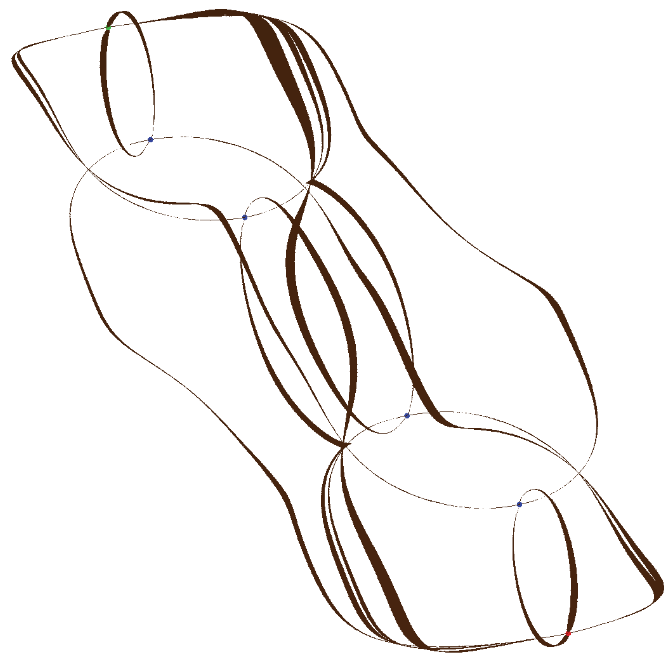

Fig. 3. Morse sets and connecting regions for the figure eight model subdivided three times.

## Appendix A
## Approximation properties

In this section, we show that trajectories of a PC vector field, defined on a fine enough mesh and constructed from a good enough vertex-based approximation of a smooth vector field $g$ are close to the trajectories of $g$.

Let $g$ be a smooth vector field on a smooth compact manifold $\mathscr{M}$ embedded in $\mathbf{R}^3$ and a triangulated manifold surface $M$ be a $G^1$-approximation of $\mathscr{M}$, i.e. the vertices of $M$ are close to $\mathscr{M}$ and the triangle normals approximate $\mathscr{M}$'s normals at nearby points. The approximate vertex based vector field on $M$ assigns the vector $\bar{g}(v) = g(\pi(v)) + \mu(v)$ to a vertex $v$, where $\pi$ is a function that maps vertices of $M$ into nearby points of $\mathscr{M}$ and $\mu(v)$ represents a perturbation, that can include the noise or measurement uncertainty (Figure 4).

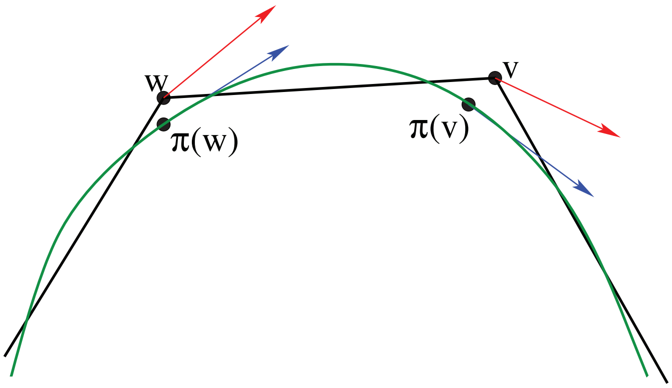


Fig. 4. 2D counterpart of the setting of the convergence proof. The black polygonal line is an approximation of the green smooth curve (1D manifold). For vertices $v$ and $w$, $\pi(v)$ and $\pi(w)$ are nearby points on the smooth curve. When constructing a vertex-based vector field on the polygonal line, one strives to assign vectors close to $f(\pi(v))$ and $f(\pi(w))$ (blue arrows) to $v$ and $w$. However, in practice, that vector would be slightly different because of measurement or numerical issues, therefore the vectors assigned to $v$ and $w$ (red lines) are slightly different. We model the error using the function $\mu$, which represents the difference between a red vector and its corresponding blue vector.

Our goal is to show that the trajectories of the PC vector field obtained from $\bar{g}$ using the procedure described in Section 3.3 converge to the trajectories of $g$ as $M$ converges to $\mathscr{M}$ in the $G^1$ sense, and the maximum diameter of a triangle of $M$ and maxima of $|\pi(v)-v|$ and $|\mu(v)|$ over all vertices of $M$ go to zero.

By the assumptions, $g(x)$ is tangent to $\mathscr{M}$ at $x$. Also, $g$ is bounded and Lipschitz continuous. Let $L$ be the Lipschitz constant. For a triangle $\Delta$ of $M$ and $x \in \mathscr{M}$, $P_\Delta$ and $Q_x$ are the orthogonal projections to the 2D linear space parallel to $\Delta$ and to the tangent plane to $\mathscr{M}$ at $x$ (respectively). By $D$ we denote the common upper bound on $|g(x)|$ for $x \in \mathscr{M}$ and $|\bar{g}(v)|$ over all vertices of $M$. Let $\delta$ be the maximum diameter of the set $\{a,b,c,\pi(a),\pi(b),\pi(c)\}$ over all triangles $\Delta abc$ of $M$ and $\kappa$ be the maximum value of $|\mu(v)|$ over all vertices of $M$. By our assumptions, $\delta$ and $\kappa$ converge to 0 as $M$ gets closer to $\mathscr{M}$. Pick any $\varepsilon > 0$. Take a triangle $\Delta$ in $M$ with vertices $a$, $b$ and $c$. If $M$ is a good enough approximation of $\mathscr{M}$, then $||P_\Delta - Q_{\pi(p)}|| < \varepsilon$ for $p \in \{a,b,c\}$. Since $f(\Delta) = P_\Delta(\frac{1}{3}(\bar{g}(a)+\bar{g}(b)+\bar{g}(c)))$, $|f(\Delta) - g(\pi(a))| \le \frac{1}{3}(|P_\Delta(\bar{g}(a)) - g(\pi(a))| + |P_\Delta(\bar{g}(b)) - g(\pi(a))| + |P_\Delta(\bar{g}(c)) - g(\pi(a))|)$ by the triangle inequality. Notice that $|P_\Delta(\bar{g}(a)) - g(\pi(a))| \le \kappa + |P_\Delta(g(\pi(a))) - Q_{\pi(a)}(g(\pi(a)))| \le \kappa + D\varepsilon$. Since $g(\pi(b))$ belongs to the tangent plane to $\mathscr{M}$ at $\pi(b)$, $Q_{\pi(b)}(g(\pi(b))) = g(\pi(b))$. Therefore, by the triangle inequality, $|P_\Delta(\bar{g}(b)) - g(\pi(a))| \le \kappa + |P_\Delta(g(\pi(b)) - Q_{\pi(b)}(g(\pi(b)))| + |g(\pi(b)) - g(\pi(a))| \le \kappa + D\varepsilon + L\delta$. Similar inequality holds if $b$ is replaced by $c$. We conclude that $|f(\Delta) - g(\pi(a))| \le \eta := 3\kappa + 3D\varepsilon + 2L\delta$. Note that $\eta \to 0$ as $M$ converges to $\mathscr{M}$.

If $\mathbf{e}$ is an exploding or imploding mesh edge with incident triangles $\Delta_0$ and $\Delta_1$, then $f(\mathbf{e}) = P_\mathbf{e}(w_0 f(\Delta_0) + w_1 f(\Delta_1))$ where $P_\mathbf{e}$ is the orthogonal projection to the 1D vector space $W$ parallel to $\mathbf{e}$. Since $\Delta_0$ and $\Delta_1$ share a vertex, $|f(\Delta_0) - f(\Delta_1)| < 2\eta$. If $M$ is a good enough approximation of $\mathscr{M}$, then the angle between the outward normal vectors of any two adjacent triangles is below $\pi/2$. Since $\mathbf{e}$ is exploding or imploding, both $f(\Delta_0)$ and $f(\Delta_1)$ are no more than $2\eta$ away from $W$. Therefore, the distance between a weighted average of these two vectors and its projection to $W$ is bounded by $2\eta$. We conclude that, for any vertex $p$ of $\mathbf{e}$, $|f(\mathbf{e}) - g(\pi(p))| = |P_\mathbf{e}(w_0 f(\Delta_0) + w_1 f(\Delta_1)) - g(\pi(p))| \le w_0|f(\Delta_0) - g(\pi(p))| + w_1|f(\Delta_1) - g(\pi(p))| + 2\eta \le 3\eta$.

Now, let $v$ be a mesh vertex that our algorithm makes stationary. If $M$ is a good enough approximation of $\mathscr{M}$, the outward normals of all triangles incident upon $v$ are contained in a cone of opening angle $\pi/2$. Let $P$ be the orthogonal projection to the linear 2D space $U$ perpendicular to the cone's axis. $P$ is a one-to-one mapping of the star of $v$ (union of all triangles incident to $v$) into $U$. Let $\mathscr{V}$ be the set of projections of vectors assigned to triangles and edges incident to $v$ by $f$. The convex hull of $\mathscr{V}$ contains the zero vector because $v$ is stationary. Otherwise, one could rotate the domain so that all vectors in $V$ point into the upper half-plane. Thus, there would be only one stable sector (in the lower half-plane) and only one unstable sector (in the upper half-plane) and two hyperbolic sectors (since the flow moves up near $v$): therefore, $v$ would not be stationary. Since the diameter of $\mathscr{V}$ is bounded by $6\eta$ by the previous estimates, $\mathscr{V}$ consists of vectors shorter than $6\eta$. Hence the magnitude of any vector assigned to a triangle or edge incident to $v$ is below $6\sqrt{2}\eta$ (because the angle between any vector in $\mathscr{V}$ and $U$ is bounded by $\pi/4$). In particular, since $|f(\Delta) - g(\pi(v))| < \eta$, where $\Delta$ is a triangle incident to $v$, $|g(\pi(v))| < \xi := (6\sqrt{2}+1)\eta$

To sum up, if $z \in M$ is contained in a triangle $\Delta$, then for any vector $v \in F_*(z)$, $|v - g(\pi(a))| < \xi$ for some vertex $a$ of $\Delta$. Hence, for any $x \in \mathscr{M}$, $|v - g(x)| \le |v - g(\pi(a))| + |g(\pi(a)) - g(x)| \le \xi + L|\pi(a) - x| \le \xi + L(|\pi(a) - a| + |a - z| + |z - x|) \le \xi + 2L\delta + L|z-x| = \tau + L|z-x|$, where $\tau = \xi + 2L\delta$.

Let $\sigma_E$ be the trajectory of $g$ starting at a point $x_0 \in \mathscr{M}$ and $\sigma_{PC}$ be a trajectory the PC variant starting at a nearby point $x_1 \in M$. Since $\dot{\sigma}_{PC}(t) \in F_*(\sigma_{PC}(t))$ almost everywhere, the inequality $|\dot{\sigma}_E(t) - \dot{\sigma}_{PC}(t)| \le \tau + L|\sigma_E(t) - \sigma_{PC}(t)|$ also holds almost everywhere. By the Gronwall's lemma [2], $|\sigma_E(t) - \sigma_{PC}(t)| \le |x_0 - x_1|\exp(Lt) + \tau/L(\exp(Lt) - 1)$. Since $\tau \to 0$ as $M$ converges to $\mathscr{M}$ in the $G^1$ sense, this means that trajectories of the flow defined by the PC vector fields on $M$ converge to the trajectories of $g$.

## Appendix B
## Admissibility

In this section, we show that the flow constructed in Section 3.3 is admissible. First, we prove upper semicontinuity, then – acyclicity of the trajectory set.

*Theorem 1:* The flow induced by a PC vector field built as described in Section 3.3 is upper semicontinuous.

*Proof:* Take a convergent sequence of trajectories $\sigma_n : [0,T] \to M$. We need to prove that $\sigma_* := \lim_{n\to\infty} \sigma_n$ is also a valid trajectory.

Take $t \in [0,T]$. If $\sigma_*(t)$ is in the interior of triangle $\Delta$ then $\sigma_n(t+h) = \sigma_n(t) + hf(\Delta)$ for small enough $h$ and sufficiently large $n$ (since $\sigma_n(t) \to \sigma_*(t)$). We conclude that also $\sigma_*(t+h) = \sigma_*(t) + hf(\Delta)$ and therefore $\sigma_*$ has velocity $f(\Delta)$ in the interior of a triangle $\Delta$.

The rest of the proof is based on the following simple observation related to general properties of piecewise linear functions. Let $g_n : I \to \mathbf{R}^n$ be a sequence of piecewise linear functions, each with no more than $N$ linear segments with derivatives in a finite set of vectors $V \subset \mathbf{R}^n$. If $g_n$ converges to $g_*$, then $g_*$ is also a piecewise linear function with no more than $N$ linear segments with derivatives in $V$.

Assume $\sigma_*(t)$ is in the interior of an edge $\mathbf{e}$. For some $h > 0$ and sufficiently large $n$, $\sigma_n(t+s)$ belongs to $\mathbf{e}$ or the interior of one of its incident triangles, $\Delta_0$ and $\Delta_1$, for any $s \in [-h,h]$. This means that $\sigma_{n|[t-h,t+h]}$ is a piecewise linear function with at most 2 segments with derivatives $f(\Delta_0)$, $f(\Delta_1)$ or $f(\mathbf{e})$ if $\mathbf{e}$ is exploding or imploding. The same is true about the limit function $\sigma_{*|[t-h,t+h]}$. Since the velocity of $\sigma_*$ is $f(\Delta_i)$ in the interior of $\Delta_i$, $\sigma_{*|[t-h,t+h]}$ is a valid trajectory.

Finally, assume that $\sigma_*(t)$ is a non-stationary mesh vertex $v$. We show that $\sigma_*$ does not stay at $t$ for a positive time. For small enough $h$ and large enough $n$, $\sigma_n([t,t+h])$ is contained in the interior of the union of triangles incident to $v$. Since $v$ is a non-stationary vertex, $\sigma_{n|[t,t+h]}$ cannot leave a mesh edge and return to the same edge later or reenter the interior of the same triangle. Thus, it a piecewise linear function of no more than twice the degree of $v$ segments. The derivatives of each segment is either $f(\Delta)$, where $\Delta$ is a triangle incident to $v$ or $f(\mathbf{e})$, where $\mathbf{e}$ is an exploding or imploding edge incident to $v$. In particular, all of them are nonzero. Hence $\sigma_*$ leaves $v$ immediately at time $t$. Note that it can leave $v$ only along an unstable direction because we have already showed that $\sigma_*$ is a valid trajectory at all points other than mesh vertices.

To sum up, we have proven that $\sigma_*$ is a (locally) valid trajectory at all points except possibly stationary vertices (which have not been considered above). In fact, this means that $\sigma_*$ is a valid trajectory. Note that $\sigma_*$ is Lipschitz continuous and therefore also absolutely continuous since it is a limit of Lipschitz-continuous functions with the same Lipschitz constant. In particular, it has a well-defined derivative almost everywhere. Because of what we have proven earlier, the derivative satisfies $\dot{\sigma}_*(t) \in F_*(\sigma_*(t))$ at all points where it is defined. In particular, for a time $t_0$ such that $\sigma_*(t_0)$ is a stationary vertex $v$ and $\sigma_*(t) \neq v$ for $t$ in a neighborhood of $t_0$, the trajectory has to leave that vertex along its unstable direction (which is then its derivative if it exists). Otherwise, $\dot{\sigma}_*(t) = 0$ if it exists. □

*Theorem 2:* There exists a positive number $h$ such that for every point $x_0 \in M$, the set $S(x_0,h)$ of trajectories starting at $x_0$ and defined over time interval $[0,h]$ is acyclic for the flow induced by a PC vector field built as described in Section 3.3. The topology on $S(x_0,h)$ is induced by the standard maximum norm on the space of continuous functions defined on $[0,h]$ and with values in the three-dimensional Euclidean space.

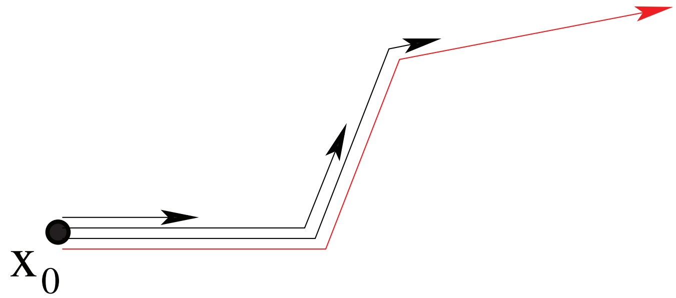


Fig. 5. Sucking a red trajectory into a stationary point $x_0$. The black trajectories stay in $x_0$ for increasing an amount of time and then follow the red trajectory. The black trajectories represent samples of $H(\sigma,s)$ where $\sigma$ is the red trajectory and $H$ is the deformation used in the case of $x_0$ being a stationary point.

*Proof:* First of all, notice that since the vector field is bounded, there exists $a > 0$ such that for every $x_0 \in M$, any trajectory in $S(x_0,a)$ is contained in the star of some vertex $v$, i.e. in the interior (relative to $M$) of the union of all triangles incident upon $v$. We shall prove that for any $x_0 \in M$ and $h \in [0,a]$, $S(x_0,h)$ is acyclic.

First, assume that $x_0$ is a stationary mesh vertex (note that then $v = x_0$). Let $S = S(x_0,h)$. It is easy to see that the function $H : S \times [0,1] \to S$ defined by $H(\sigma,s)(t) = \sigma(\max(0,t-sh))$ is well-defined and continuous. Also, $H(.,0)$ is the identity on $S$ and $H(.,1)$ maps every point into the constant trajectory that stays at $x_0$ for all times. We conclude that $S$ is contractible and therefore also acyclic. Intuitively, $H$ sucks trajectories in $S$ into $x_0$ as shown in Figure 5.

Now, let $x_0$ be a non-stationary vertex, $x_0 = v$. This means that $v$ has exactly one unstable parabolic sector, one stable parabolic sector and two hyperbolic sectors. Trajectories in $S$ leave $v$ in one of the unstable directions, pointing into the unstable parabolic sector. Let $d_i$, $i = 1,2,\ldots,n$ be all unstable directions, in counterclockwise order in that sector. This means that by moving from $d_1$ to $d_n$ counterclockwise around $v$ one sweeps the entire sector. Let $S_i$ be all trajectories in $S$ contained in the part of the sector between $d_i$ and $d_n$. By definition, $S_n$ consists of exactly one trajectory and therefore is contractible. We will show that $S_{i+1}$ is a strong deformation retract of $S_i$. The trajectories in $S_i \setminus S_{i+1}$ either move clockwise or counterclockwise relative to $v$. Consider these two cases.

If they move counterclockwise, a trajectory $\sigma \in S_i \setminus S_{i+1}$ leaves $v$ in the direction $d_i$ and then turns toward $d_{i+1}$ (as seen from $v$) at some time $t \in (0,h]$, unless it does not turn at all. In the former case, we denote $\sigma$ by $\sigma_t$ and in the latter case – by $\sigma_h$ (Figure 6). Note that $\sigma_t$ is uniquely defined since it cannot reach a mesh vertex or an exploding edge. Let $\sigma_0$ be the trajectory that moves along $d_{i+1}$ all the time. It is not hard to see that $\sigma_t$ depends continuously on $t$. The strong deformation retraction $H : S_i \times [0,1] \to S_{i+1}$ is defined by $H(\sigma_t,s) = \sigma_{(1-s)t}$ and $H(\sigma,s) = \sigma$ for any $\sigma \notin \{\sigma_t | t \in [0,h]\}$.

If they move clockwise, a trajectory $\sigma \in S_i \setminus S_{i+1}$ leaves $v$ in the direction $d_{i+1}$ and then turns toward $d_i$ at some time $t \in (0,h]$ (then we call it $\sigma_t$), unless it does not turn at all ($\sigma_h$).

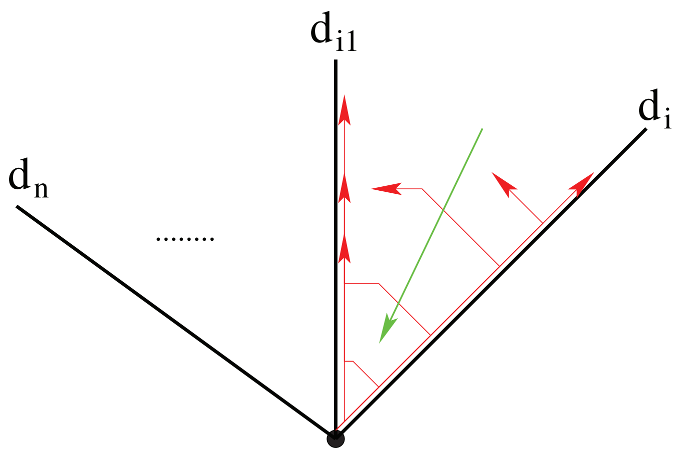


Fig. 6. The case of trajectories moving counterclockwise between unstable directions $d_i$ and $d_{i+1}$. $\sigma_h$ is the red trajectory that does not leave $d_i$. For $t < h$, $\sigma_t$ is the red trajectory that leaves $d_i$ at time $t$. The deformation $H$ gradually scales down the parameter (time a trajectory exits $d_i$) to zero. In other words, trajectories are mapped into trajectories that stay on $d_i$ for a shorter and shorter time (as indicated by the green arrow). Ultimately, all trajectories between $d_i$ and $d_{i+1}$ are collapsed to the one that follows $d_{i+1}$.

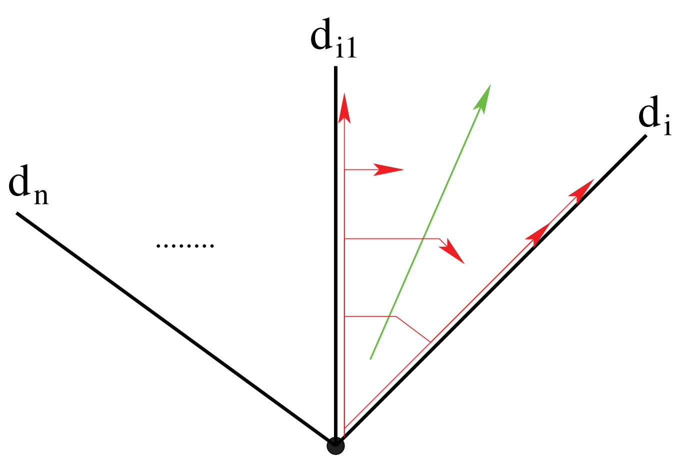


Fig. 7. The case of trajectories moving clockwise between unstable directions $d_i$ and $d_{i+1}$. $\sigma_h$ is the red trajectory that does not leave $d_{i+1}$. For $t < h$, $\sigma_t$ is the red trajectory that leaves $d_{i+1}$ at time $t$. The deformation $H$ gradually increases the parameter (time a trajectory exits $d_{i+1}$) to one. In other words, trajectories are mapped into trajectories that stay on $d_{i+1}$ for a gradually longer time (as indicated by the green arrow). Ultimately, all trajectories between $d_i$ and $d_{i+1}$ are collapsed to the one that follows $d_{i+1}$.

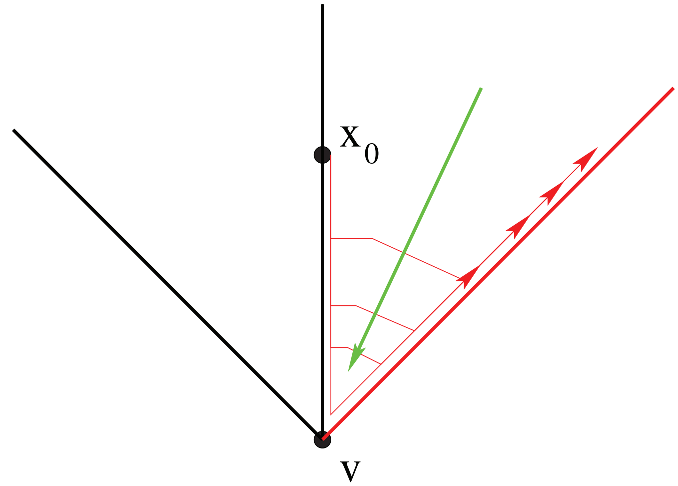


Fig. 8. Trajectories leaving an exploding edge containing $x_0$ to the same side. In this case, the limit trajectory $\sigma_{m_*}$ moves from $x_0$ to $v$ and then leaves $v$ along an unstable direction (red line). All trajectories that leave the edge in the clockwise direction are collapsed by the deformation retraction $H$ that gradually increases the time for which a trajectory follows the edge, pushing them toward the limit trajectory (as indicated by the green arrow).

Let $\sigma_0$ be the trajectory that follows $d_i$ all the time. A suitable deformation retraction, illustrated in Figure 7, can be defined by $H(\sigma_t, s) = \sigma_{h-(1-s)(h-t)}$ and $H(\sigma, s) = \sigma$ for any $\sigma \in S_{i+1}$.

Contractibility of $S(v,h)$ for a non-stationary vertex follows by induction on $n$. At this point, we have proven that $S(v,h)$ is contractible (and hence acyclic) for each mesh vertex $v$.

Now, assume that $x_0$ belongs to an exploding edge $\mathbf{e}$ with incident triangles $\Delta_0$ and $\Delta_1$. Consider a trajectory $\sigma \in S$. It starts at $x_0$ and either stays in the interior $\mathrm{int}(\mathbf{e})$ of $\mathbf{e}$ for all times or leaves it at some time $t < h$. It can leave either by arriving at $v$ or by entering a triangle $\Delta_0$ or $\Delta_1$ at time $t$. In what follows, we denote trajectories that stay in $\mathrm{int}(\mathbf{e})$ all the time or leave by arriving at $v$ by $S_*$. Let $m_*$ be the time needed for a trajectory in $S_*$ to reach $v$, or $h$ if $S_* = \emptyset$. $S_{i,t}$ is the set of trajectories that leave into $\Delta_i$ at time $t < m_*$ for $i \in \{0,1\}$. Let $S'_i = \bigcup_{t\in[0,m_*)} S_{i,t}$ and $S_i = S'_i \cup S_*$.

Let $A$ consists of all $i \in \{0,1\}$ such that $S_{i,t}$ consists of exactly one trajectory $\sigma_t$ for each $t \in [0,m_*)$. Notice that $S_*$ is a strong deformation retract of $S_i$ if $i \in A$. The strong deformation retraction can be defined by $H(\sigma_t, s) = \sigma_{sm_*+(1-s)t}$, where $\sigma_{m_*} = \lim_{t\to m_*^-} \sigma_t$ and $H(\sigma, s) = \sigma$ if $\sigma \notin S'_i$. This is illustrated in Figure 8. Note that the same idea can be used on both sides of $e$ if $A = \{0,1\}$ to construct a deformation retraction from $S$ to $S_*$. Therefore, it remains to focus on showing that the set $S_c := S_* \cup \bigcup_{i\in A^c} S'_i$ is acyclic, where $A^c = \{0,1\} \setminus A$.

Define a function $m$ on $S_c$ as follows. For a trajectory $\sigma$, let $m(\sigma)$ be the smallest $t$ such that $\sigma(t) = v$ or $h$ if no such $t$ exists. $m$ is continuous, since trajectories in $S_c$ leaving $e$ move toward a stable direction of $v$ and therefore $m$ is a continuous function of the time at which they exit $e$ (Figure 9).

Pick any trajectory $\tau$ out of $v$. For a trajectory $\sigma \in S_c$, let $\hat{\sigma}$ be defined by $\hat{\sigma}(s) = \sigma(s)$ for $s \le m(\sigma)$ and $\hat{\sigma}(s) =$

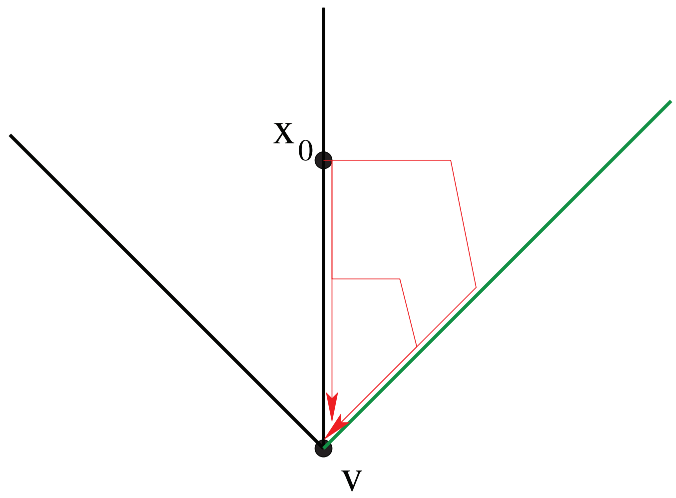


Fig. 9. Trajectories that leave $x_0$ and move toward a stable direction (green) of $v$. Time at which the trajectory arrives at $v$ depends continuously on the trajectory.

$\tau(s-m(\sigma))$ for $s \in [m(\sigma),h]$. Thus, $\hat{\sigma}$ follows $\sigma$ until it reaches $v$ and then follows $\tau$ for later times. Clearly, $\hat{\sigma} \in S_c$. Let $\hat{S}_c = \{\sigma \in S_c | \hat{\sigma} = \sigma\}$. The mapping $F : S_c \to \hat{S}_c$ defined by $F(\sigma) = \hat{\sigma}$ is continuous. For $\theta \in \hat{S}_c$, $F^{-1}(\theta)$ is equal to $S_*$ or $S_{i,t}$ for some $i \in \{0,1\}$ and $t \in [0,m_*)$. Therefore, it either consists of a single point (if no trajectory in $S_{i,t}$ reaches $v$) or is homeomorphic to $S(v,h-m(\sigma))$ (a homeomorphism is simply the restriction of a trajectory to $[m(\sigma),h]$, followed by a time shift of $-m(\sigma)$). In all cases, it is acyclic and therefore, by the Vietoris-Begle theorem [3], $F$ induces an isomorphism on homology. Hence $S_c$ is acyclic if $\hat{S}_c$ is. It remains to show that $\hat{S}_c$ is acyclic. Let $\hat{\sigma}_{i,t}$ be the unique trajectory in $S_{i,t} \cap \hat{S}_c$ for $t < m_*$ and $i \in A^c$ or the unique trajectory in $S_* \cap \hat{S}_c$ if $t = m_*$. A deformation retraction of $\hat{S}_c$ to a single point is given by $H(\sigma_{i,t},s) := \sigma_{i,m_*+(1-s)(t-m_*)}$. We conclude that $\hat{S}_c$ is acyclic.

The last case left to consider is $x_0$ that does not belong to an exploding edge and is not a mesh vertex. In this case, the trajectory of $x_0$ can be continued in the unique manner until it reaches a vertex. Thus, if the trajectory does not reach $v$ for a time $t < h$, $S(x_0,h)$ consists of one point. Otherwise, if $v$ is reached at time $t$, $S$ is homeomorphic to $S(x_0,h-t)$, that has already been proven to be acyclic. □

## Appendix C
## Morse sets

In this section, we provide a rigorous definition of a Morse set represented by a strongly connected component of a transition graph and prove its correctness. The most challenging task is to define Morse sets so that they are disjoint. The definition given in this section is similar to the definition of pseudo-Morse sets, but it uses trajectory segments represented by long enough paths in a strongly connected components (as opposed to single arcs as in the definition of a pseudo-Morse set). This has the effect of 'shrinking' the sets so that they are disjoint.

We start with a number of technical definitions. In this section, we allow trajectories to be defined on a zero-length interval. In other words, any function $\sigma : [t,t] \to M$ is also a trajectory.

*Definition 1:* A path $\pi = a_1 \to a_2 \to \ldots \to a_k$ in the transition graph $\mathscr{G}$ *captures a trajectory* $\sigma$ if $\sigma$ can be obtained by concatenating simple trajectory segments (possibly of zero length) that start in $a_i$ and end in $a_{i+1}$ for $i = 1,2,\ldots,k-1$. In what follows, we call the segment that starts in $a_i$ and ends in $a_{i+1}$ the *i-th segment* of $\pi$. If $k$ is even, the *middle segment* of $\pi$ is its $\frac{k}{2}$-th segment.

*Definition 2:* By a *trivial trajectory* we mean a trajectory that stays at the same point over the entire time interval it is defined on.

A trivial trajectory can either be defined on a zero-length time interval or stay at a stationary point over the entire time interval it is defined on. Note that trivial trajectories may be captured by nontrivial (non-constant) paths in $\mathscr{G}$. For example, a path formed by edge pieces incident to a spiral sink captures the trivial trajectory that stays at the sink. Here are the key technical lemmas that will lead to the proof of correctness of the Morse decomposition later on.

*Lemma 1:* Let $\sigma$ be a nontrivial trajectory captured by a path $\pi$. $\sigma$ visits a point $p$ such that a minimum n-set containing $p$ is on $\pi$.

*Proof:* Since $\sigma$ is nontrivial and captured by a path in $\mathscr{G}$, it has to visit a point $p$ on $M_1$ (the one-skeleton of $M$) that is not a mesh vertex. The point $p$ has to be on a simple trajectory segment that starts in an n-set $a$ and ends in an n-set $b$, with the arc $a \to b$ belonging to the path $\pi$. Either $a$ or $b$ contains $p$. Since $p$ is not a mesh vertex, that n-set has to be an edge piece and therefore is the minimal n-set containing $p$. □

*Lemma 2:* Let $D$ be the maximum degree of a vertex of $M$, $\pi$ – a path of length $D+1$ in $\mathscr{G}$ and $\sigma$ – a trajectory captured by $\pi$. If $\sigma$ is trivial, then there are two possibilities:

(i) $\sigma$ stays at a center, spiral sink or source or
(ii) $\sigma$ stays at a mesh vertex $v$ that is not a spiral sink or source; in this case, $v$ has to appear as one of the n-sets along $\pi$.

*Proof:* Let $p$ be the point on $\sigma$ ($p$ is unique since $\sigma$ is trivial). Notice that $p$ cannot be in the interior of a triangle (otherwise, $\sigma$ would not be trivial). Thus, $p$ has to be in $M_1$.

Now, we show that $p$ is a mesh vertex. Assume the contrary. Then, $p$ is on an edge $\mathbf{e}$ but is not at its endpoint. Edge pieces in $\mathbf{e}$ are either not linked by arcs at all if the edge is crossing or are linked in order along $\mathbf{e}$ otherwise. $\pi$ can only visit edge pieces containing $p$ and therefore its length is at most 2 nodes – this is a contradiction.

At this point, we know that $p$ is a mesh vertex. Assume it is not a center or spiral sink or source and the n-set corresponding to $p$ does not appear in $\pi$. Then, $\sigma$ consists only of edge pieces incident to $p$. However, there are at most $D$ such edge pieces. The arcs connecting them cannot form loops (since $p$ is not a spiral sink or source). Hence the path $\pi$ satisfying all of the above requirements cannot exist. □

*Definition 3:* For a path $\pi = a_1 \to a_2 \to \ldots \to a_{2D+2}$ in $\mathscr{G}$, let $\mathscr{U}(\pi)$ be the union of middle segments of all trajectories captured by $\pi$. Let $A_i$ ($i = 1,2,\ldots,m$) be a strongly connected component of $\mathscr{G}$. Let $C_i$ be the union of sets $\mathscr{U}(\pi)$ over all paths $\pi$ of length $2D+2$ in $A_i$.

The rest of this section is devoted to showing that the sets $C_i$ form a valid Morse decomposition.

Since any path $\pi$ in Definition 3 is contained in $A_i$, its middle segment is an arc connecting nodes in $A_i$. Therefore, $C_i$ is contained in the corresponding pseudo-Morse set $\mathscr{R}(A_i)$. Note that in some cases, $C_i$ is empty but $\mathscr{R}(A_i)$ is not. However, this can only happen for a trivial Morse set by the general results of [1].

$C_i$ contains all stationary vertices $v$ in $A_i$ (they are on the middle segment of the trivial trajectory captured by the path $\pi$ that stays at $v$). Also, any vertex $v$ that is a center or a spiral sink or source and has an incident edge piece in $A_i$ belongs to $C_i$. A path $\pi$ whose middle segment contains $v$ can be obtained by following the loop in the transition graph formed by edge pieces incident upon $v$. This proves that the classification procedure (Section 5.2) yields correct results.

*Lemma 3:* $C_i$ is closed.

*Proof:* It suffices to show that $\mathscr{U}(\pi)$ is closed for any path $\pi$. Take a sequence $(x_i)$ of points in $\mathscr{U}(\pi)$ converging to $x_*$. For each $i$, there is a trajectory $\sigma_i$ captured by $\pi$ with the middle segment passing through $x_i$. Select a subsequence of trajectories $(\sigma_{i_j})$ such that the sequences of endpoints of their $k$-th segments converge for each $k \in \{1,2,\ldots,2D+1\}$. By connecting the limit points with line segments, we obtain a trajectory captured by $\pi$ whose middle segment passes through $x_*$. Hence $x_* \in \mathscr{U}(\pi)$ and $\mathscr{U}(\pi)$ is closed. $\square$

*Lemma 4:* $C_i$ and $C_j$ do not intersect for $i \neq j$.

*Proof:* First, notice that if $s$ is a center, a spiral sink or a spiral source and $s$ belongs to a simple trajectory segment connecting n-sets $a$ and $b$ connected by an arc $a \to b$ in $\mathscr{G}$, then either $a$ or $b$ is an edge piece incident upon $s$ (recall that there is no node corresponding to $s$ in the graph). The edge pieces incident upon $s$ form a loop in $\mathscr{G}$ and therefore all of them are either in $A_i$ or in $A_j$. But this means that $s$ can belong to only one of the Morse sets $C_i$, $C_j$.

Now, assume $C_i$ and $C_j$ intersect and $p$ be a point in the intersection. There are paths $\pi_i$ in $A_i$ and $\pi_j$ in $A_j$ such that the sets $\mathscr{U}(\pi_i)$ and $\mathscr{U}(\pi_j)$ contain $p$. By the argument in the first paragraph, $p$ is not a spiral sink or source. Let $\sigma_i$ and $\sigma_j$ be trajectories through $p$ captured by $\pi_i$ and $\pi_j$ (respectively). Consider the trajectory $\sigma_{i\to j}$ obtained by following the initial section (call it section 1) of $\sigma_i$ until $p$ and then following $\sigma_j$, starting at $p$ (section 2) until the end. By Lemmas 1 and 2, section 1 passes through a point $q_1$ such that the minimum n-set containing it is on $\pi_i$ (in particular, in $A_i$). Similarly, section 2 passes through a point $q_2$ such that the minimum n-set containing it is on $\pi_j$ (i.e. in $A_j$). Now, follow $\sigma_{i\to j}$ from $q_1$ to $q_2$, recording the minimum n-sets of the encountered points. The result is a path in $\mathscr{G}$ from a node in $A_i$ to a node in $A_j$. By the same argument (with $A_i$ and $A_j$ exchanged), we obtain a path in the graph connecting $A_j$ in $A_i$. But this means that $A_i$ and $A_j$ are contained in the same strongly connected component. $\square$

*Theorem 3:* The sets $C_i$ form a valid Morse decomposition.

*Proof:* It remains to show that any trajectory that is not contained in the union of all Morse sets links two different Morse sets and that linkage graph is acyclic.

Consider a trajectory $\sigma : (-\infty,\infty) \to M$ that is not contained in the union of the Morse sets. By recording the minimum n-sets of points along $\sigma$ we obtain a path $\pi = (n_k)_{k=-\infty}^{\infty}$ in $\mathscr{G}$. $\pi$ is not contained in any strongly connected component. There are strongly connected components $A_i$ and $A_j$ such that $n_k \in A_i$ for sufficiently large negative $k$ and $n_k \in A_j$ for sufficiently large positive $k$. But then, by Definition 3, $\sigma(t) \in C_i$ for sufficiently large negative times $t$ and in $\sigma(t) \in C_j$ for sufficiently large positive times $t$. Hence $\sigma$ links $C_i$ and $C_j$.

Acyclicity of the linkage graph follows from the standard property of strongly connected components: collapsing each of them to a single supernode produces an acyclic graph. $\square$

## Appendix D
## Intersections of pseudo-Morse sets

In this section, we show that pseudo-Morse sets can only intersect along the boundary. This means that there is no significant overlap between Morse sets so that ambiguities in visualization are generally easy to avoid.

*Theorem 4:* Let $A_i$ and $A_j$ be strongly connected components of $\mathscr{G}$. Then $\mathscr{R}(A_j)$ can intersect the interior of $\mathscr{R}(A_i)$ only at mesh vertices if $i \neq j$. In particular, $\mathscr{R}(A_i)$ and $\mathscr{R}(A_j)$ have disjoint interiors.

**Proof.** $\mathscr{R}(A)$ is the union of all sets represented by arcs that start and end in $A$. These sets can be of one of three types that we call *sticks*, *blocks* and *dots*. Blocks are sets with non-empty interior represented by type T arcs. Sticks are line segments (represented by type T or E arcs). Dots consist of a single point (represented by type S arcs).

Assume $p$ is in the interior of $\mathscr{R}(A_i)$ and in $A_j$ but is not a mesh vertex. Then, all blocks containing $p$ are represented by arcs with both endpoints in $A_i$. First, assume $p$ is in the interior of a triangle $\Delta$. Let $q$ and $r$ be points on the boundary of $\Delta$ such that the line segment connecting them is a simple trajectory segment containing $p$. One of these points (say, $q$) is not a mesh vertex. All (one or two) edge pieces containing $q$ are in $A_i$ (since each of them defines a block containing $p$). But this means that $p$ cannot belong to $A_j$. This is a contradiction.

At this point, we know that $p$ is on a mesh edge. All edge pieces containing $p$ are in $A_i$. On the other hand, since $p \in A_j$, there is a stick represented by an arc $\alpha$ in $A_j$ containing $p$. At least one endpoint of $\alpha$ is an edge piece in $A_j$ containing $p$. This is a contradiction. **Q.E.D.**

## Appendix E
## Effect of refinement on Morse sets

The goal of this section is to prove that finer transition graphs lead to finer Morse decompositions (i.e. smaller Morse sets). This means that our refinement scheme in fact produces a hierarchy of Morse decompositions.

*Lemma 5:* Assume that the transition graph $\mathscr{G}'$ is obtained from the transition graph $\mathscr{G}$ by means of a refinement operation that replaces node $\mathbf{f}$ into two nodes $\mathbf{f}_1$ and $\mathbf{f}_2$. Then:

(i) An arc $\mathbf{f} \to \mathbf{g}$ ($\mathbf{g} \to \mathbf{f}$) is in $\mathscr{G}$ if and only if at least one of the arcs $\mathbf{f}_i \to \mathbf{g}$ (respectively, $\mathbf{g} \to \mathbf{f}_i$) is in $\mathscr{G}'$ for $i \in \{1,2\}$.

(ii) Let $\beta$ be a path in $\mathscr{G}'$. Define $\hat{\beta}$ as the sequence of nodes of $\mathscr{G}$ obtained by replacing each occurrence of $\mathbf{f}_1$ or $\mathbf{f}_2$ in

$\beta$ by $\mathbf{f}$ and then replacing any subsequence of consecutive $\mathbf{f}$'s in the resulting sequence with a single $\mathbf{f}$. $\hat{\beta}$ is a path in $\mathscr{G}$.

(iii) Any strongly connected component of $\mathscr{G}$ that does not contain $\mathbf{f}$ is also a strongly connected component in $\mathscr{G}'$,

(iv) Assume there is a strongly connected component $A'$ of $\mathscr{G}'$ containing $\mathbf{f}_1$ or $\mathbf{f}_2$. Then, $\mathbf{f}$ belongs to a strongly connected component $A$ of $\mathscr{G}$ that contains $A' \setminus \{\mathbf{f}_1, \mathbf{f}_2\}$.

(v) Under the assumptions of **(iv)**, the refined pseudo-Morse (Morse) set defined by $A'$ is a subset of the pseudo-Morse (repsectively, Morse) set defined by $A$.

**Proof.** **(i)** follows immediately from the description of refinement operation in Section 4.2.2. **(ii)** is a simple consequence of **(i)**.

**Proof of part (iii).** Take a strongly connected component $A$ in $\mathscr{G}$ that does not contain $\mathbf{f}$. Pick any node $\mathbf{g} \in A$. $A$ is the union of all loops passing through $\mathbf{g}$. These loops are also loops in $\mathscr{G}'$ and therefore $A$ is a subset of its strongly connected component $A'$. It remains to show that $A$ and $A'$ are the same. If they are not, there is a loop $\beta$ in $\mathscr{G}'$ through $\mathbf{g}$ that is not a loop in $\mathscr{G}$. It has to pass through $\mathbf{f}_1$ or $\mathbf{f}_2$. But then $\hat{\beta}$ is a loop in $\mathscr{G}$ passing through both $\mathbf{g}$ and $\mathbf{f}$. This is a contradiction.

**Proof of part (iv).** For any loop $\beta$ in $\mathscr{G}'$ passing through $\mathbf{f}_1$ or $\mathbf{f}_2$, $\hat{\beta}$ is a loop in $\mathscr{G}$ passing through $\mathbf{f}$. Therefore, if there is a strongly connected component of $\mathscr{G}'$ containing $\mathbf{f}_1$ or $\mathbf{f}_2$, then there is a strongly connected component $A$ of $\mathscr{G}$ containing $\mathbf{f}$.

Let $A'$ be a strongly connected component of $\mathscr{G}'$ containing $\mathbf{f}_i$. Any node $\mathbf{g} \in A' \setminus \{\mathbf{f}_1, \mathbf{f}_2\}$ is on a loop $\beta$ passing through $\mathbf{f}_i$ and $\mathbf{g}$. But then, $\mathbf{g}$ is also in $A$ since it is on the loop $\hat{\beta}$ (that passes through $\mathbf{f}$ and $\mathbf{g}$).

**Proof of part (v).** This follows immediately from the definitions, part **(iv)** and the containment $\mathbf{f}_i \subset \mathbf{f}$ for $i \in \{1,2\}$. **Q.E.D.**

*Corollary 1:* Assume that the transition graph $\mathscr{G}'$ is obtained from the transition graph $\mathscr{G}$ by means of refinement operations. Any Morse (pseudo-Morse) set obtained from $\mathscr{G}'$ is a subset of a Morse (pseudo-Morse) set obtained from $\mathscr{G}$. Moreover, Morse (pseudo-Morse) sets defined by strongly connected components of $\mathscr{G}$ whose nodes are not refined are identical to Morse (pseudo-Morse) sets obtained from $\mathscr{G}'$.

**Proof.** Use Lemma 5 and induction with respect to number of refinement operations. **Q.E.D.**

## References


[1] C. Conley. *Isolated Invariant Sets and Morse Index*, volume 38 of *CBMS Lecture Notes*. AMS Providence, RI, 1978.

[2] M. Schatzman. *Numerical Analysis: A Mathematical Introduction*. Oxford University Press, USA, 2002.

[3] E. H. Spanier. *Algebraic Topology*. Springer, 1966.