Perfect discrete Morse functions

Neža Mramor Kosta
joint work with
Hanife Vari, Mehmetcik Pamuk

University of Ljubljana, Slovenia
Middle Eastern Technical University, Ankara, Turkey

GETCO 2015, Aalborg

Discrete Morse functions

M regular (locally) finite cell complex (simplicial, cubical, ...)

Discrete Morse functions

M regular (locally) finite cell complex (simplicial, cubical, ...)

A *discrete Morse function* F on M is a labelling of the cells such that

- $F(\tau) \leq F\sigma$ for all $\tau < \sigma$
- $F^{-1}(t)$ contains at most two cells $\tau < \sigma$ for all $t \in \mathbb{R}$.

Discrete Morse functions

M regular (locally) finite cell complex (simplicial, cubical, ...)

A *discrete Morse function* F on M is a labelling of the cells such that

- $F(\tau) \leq F\sigma$ for all $\tau < \sigma$
- $F^{-1}(t)$ contains at most two cells $\tau < \sigma$ for all $t \in \mathbb{R}$.

Introduced by Robin Forman in the 1990's

Discrete Morse functions

M regular (locally) finite cell complex (simplicial, cubical, ...)

A *discrete Morse function* F on M is a labelling of the cells such that

- $F(\tau) \leq F\sigma$ for all $\tau < \sigma$
- $F^{-1}(t)$ contains at most two cells $\tau < \sigma$ for all $t \in \mathbb{R}$.

Introduced by Robin Forman in the 1990's

If $F^{-1}(F(\sigma)) = \{\sigma\}$ then σ is *critical*,

Discrete Morse functions

M regular (locally) finite cell complex (simplicial, cubical, ...)

A *discrete Morse function* F on M is a labelling of the cells such that

- $F(\tau) \leq F\sigma$ for all $\tau < \sigma$
- $F^{-1}(t)$ contains at most two cells $\tau < \sigma$ for all $t \in \mathbb{R}$.

Introduced by Robin Forman in the 1990's

If $F^{-1}(F(\sigma)) = \{\sigma\}$ then σ is *critical*,

if $F(\tau) = F(\sigma)$, and $\tau < \sigma$ then τ is of codimension 1 in σ, and (τ, σ) is a *regular pair*, denoted by an arrow $\tau \to \sigma$.

A discrete Morse function defines an *acyclic matching* on the cells of M.

A discrete Morse function defines an *acyclic matching* on the cells of M.

That is, the cells of M are partitioned into three classes:

A discrete Morse function defines an *acyclic matching* on the cells of M.

That is, the cells of M are partitioned into three classes:

critical cells C,

arrow tails A,

arrow heads B,

A discrete Morse function defines an *acyclic matching* on the cells of M.

That is, the cells of M are partitioned into three classes:

critical cells C,

arrow tails A,

arrow heads B,

with a bijection $r : A \to B$,

A discrete Morse function defines an *acyclic matching* on the cells of M.

That is, the cells of M are partitioned into three classes:

critical cells C,

arrow tails A,

arrow heads B,

with a bijection $r : A \to B$,

such that adjoining sequences of arrows (*gradient paths*)

$$\tau_0 \to \sigma_0 > \tau_1 \to \sigma_1 \cdots \tau_n \to \sigma_n > \tau_{n+1}$$

do not form cycles.

A discrete Morse function defines an *acyclic matching* on the cells of M.

That is, the cells of M are partitioned into three classes:

critical cells C,

arrow tails A,

arrow heads B,

with a bijection $r : A \to B$,

such that adjoining sequences of arrows (*gradient paths*)

$$\tau_0 \to \sigma_0 > \tau_1 \to \sigma_1 \cdots \tau_n \to \sigma_n > \tau_{n+1}$$

do not form cycles.

Any acyclic matching comes from a discrete Morse function.

The homotopy type of the cell complex is determined by the critical data:

Theorem (Forman)

A cell complex with a discrete Morse function has the homotopy type of a CW complex with one cell of dimension p for every critical cell of the function of dimension p.

The homotopy type of the cell complex is determined by the critical data:

Theorem (Forman)

A cell complex with a discrete Morse function has the homotopy type of a CW complex with one cell of dimension p for every critical cell of the function of dimension p.

In particular, the homology of M can be computed from the *Morse chain complex* with:

The homotopy type of the cell complex is determined by the critical data:

Theorem (Forman)

A cell complex with a discrete Morse function has the homotopy type of a CW complex with one cell of dimension p for every critical cell of the function of dimension p.

In particular, the homology of M can be computed from the *Morse chain complex* with:

chain groups generated by critical cells and

The homotopy type of the cell complex is determined by the critical data:

Theorem (Forman)

A cell complex with a discrete Morse function has the homotopy type of a CW complex with one cell of dimension p for every critical cell of the function of dimension p.

In particular, the homology of M can be computed from the *Morse chain complex* with:

chain groups generated by critical cells and

boundary homomorphisms given by gradient paths starting in the boundary of a critical p-cell and ending in a critical $p-1$ cell ($p > 0$).

This gives rise to efficient algorithms for computing homology, homology generators, as well as homology homomorphisms generated by smooth maps

This gives rise to efficient algorithms for computing homology, homology generators, as well as homology homomorphisms generated by smooth maps

Harker, Mischaikow, Mrozek, Nanda, Dłotko, Wagner, Kaczynski,...

This gives rise to efficient algorithms for computing homology, homology generators, as well as homology homomorphisms generated by smooth maps

Harker, Mischaikow, Mrozek, Nanda, Dłotko, Wagner, Kaczynski,. . .

Good implementations of these algorithms exist as well:

This gives rise to efficient algorithms for computing homology, homology generators, as well as homology homomorphisms generated by smooth maps

Harker, Mischaikow, Mrozek, Nanda, Dłotko, Wagner, Kaczynski,…

Good implementations of these algorithms exist as well:

Perseus, CHomP, RedHom, …

This gives rise to efficient algorithms for computing homology, homology generators, as well as homology homomorphisms generated by smooth maps

Harker, Mischaikow, Mrozek, Nanda, Dłotko, Wagner, Kaczynski,...

Good implementations of these algorithms exist as well:

Perseus, CHomP, RedHom, ...

The complexity of the computation of homology after reducing the complex depends on the number of critical cells.

Discrete Morse theory provides a strategy: cancelling critical cells.

Discrete Morse theory provides a strategy: cancelling critical cells.

Given two critical cells σ^p and τ^{p-1} of consecutive dimensions with only one gradient path connecting them

$$\sigma = \sigma_0 > \tau_1 \to \sigma_1 \cdots \tau_n \to \sigma_n > \tau_{n+1} = \tau,$$

Discrete Morse theory provides a strategy: cancelling critical cells.

Given two critical cells σ^p and τ^{p-1} of consecutive dimensions with only one gradient path connecting them

$$\sigma = \sigma_0 > \tau_1 \to \sigma_1 \cdots \tau_n \to \sigma_n > \tau_{n+1} = \tau,$$

switch the arrows along the gradient path:

Discrete Morse theory provides a strategy: cancelling critical cells.

Given two critical cells σ^p and τ^{p-1} of consecutive dimensions with only one gradient path connecting them

$$\sigma = \sigma_0 > \tau_1 \to \sigma_1 \cdots \tau_n \to \sigma_n > \tau_{n+1} = \tau,$$

switch the arrows along the gradient path:

$$\sigma = \sigma_0 \to \tau_1 < \sigma_1 \cdots \tau_n < \sigma_n \to \tau_{n+1} = \tau$$

Discrete Morse theory provides a strategy: cancelling critical cells.

Given two critical cells σ^p and τ^{p-1} of consecutive dimensions with only one gradient path connecting them

$$\sigma = \sigma_0 > \tau_1 \to \sigma_1 \cdots \tau_n \to \sigma_n > \tau_{n+1} = \tau,$$

switch the arrows along the gradient path:

$$\sigma = \sigma_0 \to \tau_1 < \sigma_1 \cdots \tau_n < \sigma_n \to \tau_{n+1} = \tau$$

The beginning and ending cells are not critical any more, no cycles are produced so we still have an acylic matching.

If the values of F on vertices correspond to some data measurements then,

If the values of F on vertices correspond to some data measurements then,

these can be extended to a discrete Morse function on the complex (which can be some reconstruction from the data points),

If the values of F on vertices correspond to some data measurements then,

these can be extended to a discrete Morse function on the complex (which can be some reconstruction from the data points),

cancelling up to a certain threshold ε provides a smoothing algorithms for the data.

If the values of F on vertices correspond to some data measurements then,

these can be extended to a discrete Morse function on the complex (which can be some reconstruction from the data points),

cancelling up to a certain threshold ε provides a smoothing algorithms for the data.

Cancelling with $\varepsilon = \infty$ produces a smaller, in some cases even the minimal, number of critical cells.

Perfect discrete Morse functions

Morse inequalities:

Perfect discrete Morse functions

Morse inequalities:

M a manifold of dimension n with Betti numbers (with respect to some coefficient ring K) b_k, $k = 0, 1, \ldots, n$,
F is a discrete Morse function on M with c_k critical cells of dimension k, $k = 0, 1, \ldots, n$. Then

Perfect discrete Morse functions

Morse inequalities:

M a manifold of dimension n with Betti numbers (with respect to some coefficient ring K) b_k, $k = 0, 1, \ldots, n$,
F is a discrete Morse function on M with c_k critical cells of dimension k, $k = 0, 1, \ldots, n$. Then

1. $c_k \geq b_k$ for all k,
2. $c_k - c_{k-1} + \cdots \pm c_0 \geq b_k - b_{k-1} + \cdots \pm b_0$, for all k,
3. $c_0 - c_1 + \cdots + (-1)^n c_n = b_0 - b_1 + \cdots + (-1)^n b_n = \chi(M)$.

Perfect discrete Morse functions

Morse inequalities:

M a manifold of dimension n with Betti numbers (with respect to some coefficient ring K) b_k, $k = 0, 1, \ldots, n$,
F is a discrete Morse function on M with c_k critical cells of dimension k, $k = 0, 1, \ldots, n$. Then

1. $c_k \geq b_k$ for all k,
2. $c_k - c_{k-1} + \cdots \pm c_0 \geq b_k - b_{k-1} + \cdots \pm b_0$, for all k,
3. $c_0 - c_1 + \cdots + (-1)^n c_n = b_0 - b_1 + \cdots + (-1)^n b_n = \chi(M)$.

A discrete Morse function on M is *perfect* (with respect to the coefficient ring K) if the Morse inequalities are equalities.

Perfect discrete Morse functions

Morse inequalities:

M a manifold of dimension n with Betti numbers (with respect to some coefficient ring K) b_k, $k = 0, 1, \ldots, n$,
F is a discrete Morse function on M with c_k critical cells of dimension k, $k = 0, 1, \ldots, n$. Then

1. $c_k \geq b_k$ for all k,
2. $c_k - c_{k-1} + \cdots \pm c_0 \geq b_k - b_{k-1} + \cdots \pm b_0$, for all k,
3. $c_0 - c_1 + \cdots + (-1)^n c_n = b_0 - b_1 + \cdots + (-1)^n b_n = \chi(M)$.

A discrete Morse function on M is *perfect* (with respect to the coefficient ring K) if the Morse inequalities are equalities.

A perfect discrete Morse function has in each dimension the number of critical cells equal to the Betti number of the complex.

An example of a $\mathbb{Z}_2$-perfect discrete Morse function on the projective plane $\mathbb{R}P^2$ ...

An example of a $\mathbb{Z}_2$-perfect discrete Morse function on the projective plane $\mathbb{R}P^2$ ...

Some ovious nonexistence results:

An example of a $\mathbb{Z}_2$-perfect discrete Morse function on the projective plane $\mathbb{R}P^2$ ...

Some ovious nonexistence results:

No nonorientable manifold can admit a $\mathbb{Z}$-perfect discrete Morse function. More generally, any torsion in homology is an obstruction to existence of $\mathbb{Z}$-perfect discrete Morse functions.

An example of a $\mathbb{Z}_2$-perfect discrete Morse function on the projective plane $\mathbb{R}P^2$ …

Some ovious nonexistence results:

No nonorientable manifold can admit a $\mathbb{Z}$-perfect discrete Morse function. More generally, any torsion in homology is an obstruction to existence of $\mathbb{Z}$-perfect discrete Morse functions.

Complexes that are contractible but not collapsible do not admit perfect discrete Morse functions with respect to any coefficients, since:

An example of a $\mathbb{Z}_2$-perfect discrete Morse function on the projective plane $\mathbb{R}P^2$ ...

Some ovious nonexistence results:

No nonorientable manifold can admit a $\mathbb{Z}$-perfect discrete Morse function. More generally, any torsion in homology is an obstruction to existence of $\mathbb{Z}$-perfect discrete Morse functions.

Complexes that are contractible but not collapsible do not admit perfect discrete Morse functions with respect to any coefficients, since:

a regular cell complex with a discrete Morse function with only one critical cell is collapsible,

An example of a $\mathbb{Z}_2$-perfect discrete Morse function on the projective plane $\mathbb{R}P^2$ ...

Some ovious nonexistence results:

No nonorientable manifold can admit a $\mathbb{Z}$-perfect discrete Morse function. More generally, any torsion in homology is an obstruction to existence of $\mathbb{Z}$-perfect discrete Morse functions.

Complexes that are contractible but not collapsible do not admit perfect discrete Morse functions with respect to any coefficients, since:

a regular cell complex with a discrete Morse function with only one critical cell is collapsible,

the collapsing strategy is given by the gradient paths.

Some (more or less) obvious existence results:

Some (more or less) obvious existence results:

Every graph admits a perfect discrete Morse function.

Some (more or less) obvious existence results:

Every graph admits a perfect discrete Morse function.

Suppose M has dimension 2 and is a subcomplex of a two dimensional manifold. Then there exists a $\mathbb{Z}_2$-perfect discrete Morse function. It is obtained by a recursive process of cancelling critical cells in lower stars of vertices.

Some (more or less) obvious existence results:

Every graph admits a perfect discrete Morse function.

Suppose M has dimension 2 and is a subcomplex of a two dimensional manifold. Then there exists a $\mathbb{Z}_2$-perfect discrete Morse function. It is obtained by a recursive process of cancelling critical cells in lower stars of vertices.

Spheres are characterized by the existence of a perfect Morse function:

Some (more or less) obvious existence results:

Every graph admits a perfect discrete Morse function.

Suppose M has dimension 2 and is a subcomplex of a two dimensional manifold. Then there exists a $\mathbb{Z}_2$-perfect discrete Morse function. It is obtained by a recursive process of cancelling critical cells in lower stars of vertices.

Spheres are characterized by the existence of a perfect Morse function:

Every complex with exactly two critical cells is a triangulated sphere, and for every sphere there exists a triangulation which admits a perfect discrete Morse function. (Froman, building on deep results from smooth Morse theory by Milnor, Smale, Sharko).

Some (more or less) obvious existence results:

Every graph admits a perfect discrete Morse function.

Suppose M has dimension 2 and is a subcomplex of a two dimensional manifold. Then there exists a $\mathbb{Z}_2$-perfect discrete Morse function. It is obtained by a recursive process of cancelling critical cells in lower stars of vertices.

Spheres are characterized by the existence of a perfect Morse function:

Every complex with exactly two critical cells is a triangulated sphere, and for every sphere there exists a triangulation which admits a perfect discrete Morse function. (Froman, building on deep results from smooth Morse theory by Milnor, Smale, Sharko).

Every sphere of dimension $d > 4$ has a triangulation which does not admit a discrete Morse function (Benedetti).

It is known that finding optimal discrete Morse functions (not necessarily perfect) is NP-complete (Joswig and Pfetsch, Lewiner)

It is known that finding optimal discrete Morse functions (not necessarily perfect) is NP-complete (Joswig and Pfetsch, Lewiner)

Several approaches exist:

It is known that finding optimal discrete Morse functions (not necessarily perfect) is NP-complete (Joswig and Pfetsch, Lewiner)

Several approaches exist:

Joswig-Pfetch (an algorithm for finding maximal acyclic matchings on the Hasse diagram),

Engström (using a form of Fourier transforms),

Lewiner, Lopes and Tavares (finding maximal hyperforests of hypergraphs, on big complexes of dimension 2 and 3)

Benedetti and Lutz (an efficient algorithm for generating a large number of random discrete Morse functions, look for the minimal Morse vector)

...

Perfect discrete Morse functions on connected sums

This is joint work with Hanife Isal and Mehmetcik Pamuk from METU, Turkey, and Jose Antonio Vilches and Rafael Ayala, University of Sevilla

Perfect discrete Morse functions on connected sums

This is joint work with Hanife Isal and Mehmetcik Pamuk from METU, Turkey, and Jose Antonio Vilches and Rafael Ayala, University of Sevilla

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented, triangulated n manifolds with perfect discrete Morse functions F_1 and F_2 on M1 and M2, respectively.

Perfect discrete Morse functions on connected sums

This is joint work with Hanife Isal and Mehmetcik Pamuk from METU, Turkey, and Jose Antonio Vilches and Rafael Ayala, University of Sevilla

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented, triangulated n manifolds with perfect discrete Morse functions F_1 and F_2 on M1 and M2, respectively.

Theorem

There exists a polyhedral subdivision $\tilde{M}$ of M and a perfect discrete Morse function F on M that agrees, up to a constant on each summand with F_1 and F_2, except in a neighbourhood of the two removed cells.

The proof is constructive:

The proof is constructive:

Form the connected sum by removing an n-cell with the critical vertex in its bounday on the upper summand, and the unique top-dimensional critical cell on the lower summand.

The proof is constructive:

Form the connected sum by removing an n-cell with the critical vertex in its bounday on the upper summand, and the unique top-dimensional critical cell on the lower summand.

Attach a tube connecting the two boundaries.

The proof is constructive:

Form the connected sum by removing an n-cell with the critical vertex in its bounday on the upper summand, and the unique top-dimensional critical cell on the lower summand.

Attach a tube connecting the two boundaries.

Extend the discrete vector field so that all arows point down from the upper boundary and the vector field coincides with the original one on the lower boundary.

The opposite direction, $n = 2$

The opposite direction, $n = 2$

Theorem

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented surfaces of genus g_1 and g_2 respectively and F be a perfect discrete Morse function on M. We can find a separating circle S^1 on M such that the cells on S^1 are paired with either cells on S^1 or cells in $M - M_2$.

The opposite direction, $n = 2$

Theorem

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented surfaces of genus g_1 and g_2 respectively and F be a perfect discrete Morse function on M. We can find a separating circle S^1 on M such that the cells on S^1 are paired with either cells on S^1 or cells in $M - M_2$.

Outline of proof:

The opposite direction, $n = 2$

Theorem

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented surfaces of genus g_1 and g_2 respectively and F be a perfect discrete Morse function on M. We can find a separating circle S^1 on M such that the cells on S^1 are paired with either cells on S^1 or cells in $M - M_2$.

Outline of proof:

Collapse along gradient paths starting in the boundary of the critical 2-cell and ending in the g_1 critical 1-cells that belong to the upper summand.

The opposite direction, $n = 2$

Theorem

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented surfaces of genus g_1 and g_2 respectively and F be a perfect discrete Morse function on M. We can find a separating circle S^1 on M such that the cells on S^1 are paired with either cells on S^1 or cells in $M - M_2$.

Outline of proof:

Collapse along gradient paths starting in the boundary of the critical 2-cell and ending in the g_1 critical 1-cells that belong to the upper summand.

The separating circle S^1 is obtained after removing a neighborhood of these critical 1-cells.

The opposite direction, $n = 2$

Theorem

Let $M = M_1 \# M_2$ be the connected sum of two closed, oriented surfaces of genus g_1 and g_2 respectively and F be a perfect discrete Morse function on M. We can find a separating circle S^1 on M such that the cells on S^1 are paired with either cells on S^1 or cells in $M - M_2$.

Outline of proof:

Collapse along gradient paths starting in the boundary of the critical 2-cell and ending in the g_1 critical 1-cells that belong to the upper summand.

The separating circle S^1 is obtained after removing a neighborhood of these critical 1-cells.

If two paths from the critical 2-cell to critical 1-cells meet in a vertex or are separated only by an edge or 1-path, subdivisions are necessary.

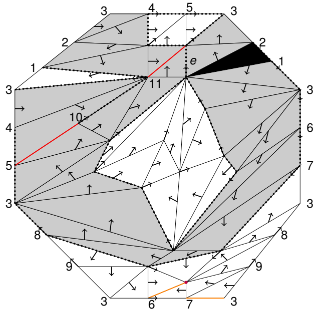

A perfect discrete Morse function on a genus 2 surface.

The shaded region represents the gradient paths from the critical triangle to the top 2 critical edges.

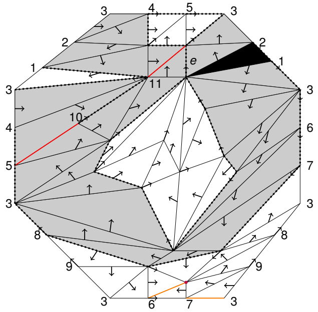

A perfect discrete Morse function on a genus 2 surface.
The shaded region represents the gradient paths from the critical triangle to the top 2 critical edges.

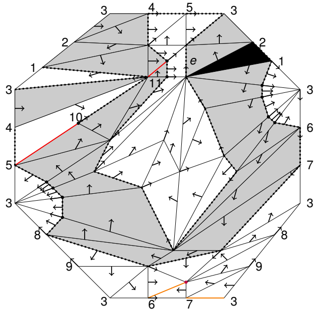

The shaded region after subdivion around the vertex 5 and the edge *e* where the gradient paths meet.
The boundary of the shaded region is a separating circle.

Theorem

If there are no arrows on the vertices and edges of the separating circle S^1 pointing upwards into $M - M_1$, perfect discrete Morse functions F_1 and F_2 exist on M_1 and M_2 which coincide on each summands with F, except on the two added discs.

Theorem

If there are no arrows on the vertices and edges of the separating circle S^1 pointing upwards into $M - M_1$, perfect discrete Morse functions F_1 and F_2 exist on M_1 and M_2 which coincide on each summands with F, except on the two added discs.

Idea of proof: in this configuration of arrows on the separating circle S^1

there exist triangulations of the discs glued to S^1 to obtain M_1 and M_2

and extensions of the discrete vector field to these two discs such that there is only one critical vertex on the top summand and only one critical 2-cell on the bottom summand,

they are relatively easy to construct.

Opposite direction, $n > 2$

Opposite direction, $n > 2$

Opposite direction, $n > 2$

Theorem (Ayala, Fernández-Ternero, Vilches)

Let M be a closed orientable 3-manifold and let coefficients be either in $\mathbb{Z}$ or in a field. Then M admits a perfect discrete Morse function if and only if there exists a spine K of M, which admits a perfect discrete Morse function.

Opposite direction, $n > 2$

Theorem (Ayala, Fernández-Ternero, Vilches)

Let M be a closed orientable 3-manifold and let coefficients be either in $\mathbb{Z}$ or in a field. Then M admits a perfect discrete Morse function if and only if there exists a spine K of M, which admits a perfect discrete Morse function.

The *spine* of a connected closed orientable n-manifold is an $n-1$-dimensional complex onto which the manifold collapses after removing an n-cell.

Opposite direction, $n > 2$

Theorem (Ayala, Fernández-Ternero, Vilches)

Let M be a closed orientable 3-manifold and let coefficients be either in $\mathbb{Z}$ or in a field. Then M admits a perfect discrete Morse function if and only if there exists a spine K of M, which admits a perfect discrete Morse function.

The *spine* of a connected closed orientable n-manifold is an $n-1$-dimensional complex onto which the manifold collapses after removing an n-cell.

The proof is based on the fact that a connected sum has a spine $K = K_1 \vee K_2$, where K_1 and K_2 are the spines of the two components.

Opposite direction, $n > 2$

Theorem (Ayala, Fernández-Ternero, Vilches)

Let M be a closed orientable 3-manifold and let coefficients be either in $\mathbb{Z}$ or in a field. Then M admits a perfect discrete Morse function if and only if there exists a spine K of M, which admits a perfect discrete Morse function.

The *spine* of a connected closed orientable n-manifold is an $n-1$-dimensional complex onto which the manifold collapses after removing an n-cell.

The proof is based on the fact that a connected sum has a spine $K = K_1 \vee K_2$, where K_1 and K_2 are the spines of the two components.

The proof works for arbitrary n.

Opposite direction, $n > 2$

Theorem (Ayala, Fernández-Ternero, Vilches)

Let M be a closed orientable 3-manifold and let coefficients be either in $\mathbb{Z}$ or in a field. Then M admits a perfect discrete Morse function if and only if there exists a spine K of M, which admits a perfect discrete Morse function.

The *spine* of a connected closed orientable n-manifold is an $n-1$-dimensional complex onto which the manifold collapses after removing an n-cell.

The proof is based on the fact that a connected sum has a spine $K = K_1 \vee K_2$, where K_1 and K_2 are the spines of the two components.

The proof works for arbitrary n.

Open question: if the collapses come from a discrete Morse function, will the resulting spine be a wedge?