Optimal discrete Morse vectors are not unique

Bruno Benedetti
University of Miami

Memphis, Oct. 17, 2015

Conventions

Free face = a face contained in only one other face. Like a leaf in a graph.

Conventions

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.

Conventions

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex = a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex = a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Clearly this number is at least one (without free faces, the collapsing sequence cannot even start).

Conventions

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex = a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Clearly this number is at least one (without free faces, the collapsing sequence cannot even start).

Conventions

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex = a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Clearly this number is at least one (without free faces, the collapsing sequence cannot even start).

For $d = 1$ this number is 2 (collapsible graphs are just trees).

Conventions

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex = a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Clearly this number is at least one (without free faces, the collapsing sequence cannot even start).

For $d = 1$ this number is 2 (collapsible graphs are just trees). Maybe it grows linearly with d?

Free face $=$ a face contained in only one other face. Like a leaf in a graph.
Elementary collapse $=$ the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex $=$ a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Clearly this number is at least one (without free faces, the collapsing sequence cannot even start).

For $d = 1$ this number is 2 (collapsible graphs are just trees). Maybe it grows linearly with d? Polynomially?

Free face = a face contained in only one other face. Like a leaf in a graph.
Elementary collapse = the deletion of a free face (and of the unique face containing it). Like, stripping a graph of a leaf.
Collapsible complex = a complex that can be reduced to a point via some sequence of elementary collapses. That is, by repeatedly deleting one free face.

How many free faces does a collapsible d-dimensional complex have, *at least*?

Clearly this number is at least one (without free faces, the collapsing sequence cannot even start).

For $d = 1$ this number is 2 (collapsible graphs are just trees). Maybe it grows linearly with d? Polynomially? Exponentially?

Actually, it doesn't grow at all

Actually, it doesn't grow at all

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 2$, there is a collapsible d-complex G_d with exactly one free face.

Actually, it doesn't grow at all

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 2$, there is a collapsible d-complex G_d with exactly one free face. The complex we construct is also **shellable**, and has $2^d + d + 1$ vertices.

Actually, it doesn't grow at all

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 2$, there is a collapsible d-complex G_d with exactly one free face. The complex we construct is also **shellable**, and has $2^d + d + 1$ vertices.

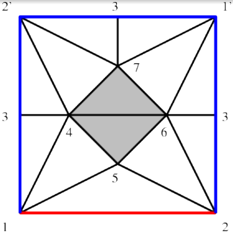

Actually, it doesn't grow at all

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 2$, there is a collapsible d-complex G_d with exactly one free face. The complex we construct is also **shellable**, and has $2^d + d + 1$ vertices.

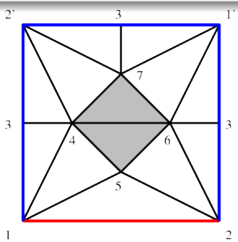

Idea: put a minicube (2^d vertices) inside the cross-polytope,

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 2$, there is a collapsible d-complex G_d with exactly one free face. The complex we construct is also **shellable**, and has $2^d + d + 1$ vertices.

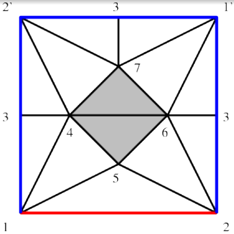

Idea: put a minicube (2^d vertices) inside the cross-polytope, and triangulate the space in between using the **antiprism subdivision**.

Actually, it doesn't grow at all

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 2$, there is a collapsible d-complex G_d with exactly one free face. The complex we construct is also **shellable**, and has $2^d + d + 1$ vertices.

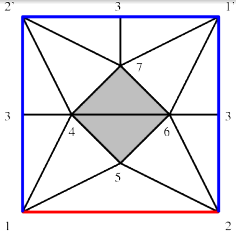

Idea: put a minicube (2^d vertices) inside the cross-polytope, and triangulate the space in between using the **antiprism subdivision**. Fold the boundary in a clever order, given by a line shelling, until all boundary facets are identified to a single one ($d + 1$ vertices).

How to use the gadget above

How to use the gadget above

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

How to use the gadget above

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

Exercise if you don't care about discrete Morse theory (which comes next):

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

Exercise if you don't care about discrete Morse theory (which comes next): let e be an edge in the boundary of a 2-disk D.

How to use the gadget above

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

Exercise if you don't care about discrete Morse theory (which comes next): let e be an edge in the boundary of a 2-disk D. Let us glue D to G_2 by identifying e with the free edge of G_2.

How to use the gadget above

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

Exercise if you don't care about discrete Morse theory (which comes next): let e be an edge in the boundary of a 2-disk D. Let us glue D to G_2 by identifying e with the free edge of G_2.

Is $D \cup G_2$ collapsible? How?

How to use the gadget above

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

Exercise if you don't care about discrete Morse theory (which comes next): let e be an edge in the boundary of a 2-disk D. Let us glue D to G_2 by identifying e with the free edge of G_2.

Is $D \cup G_2$ collapsible? How? What does a shelling order look like?

Note: given a contractible 3-complex, attaching copies of G_3 will disallow freeness of *triangles*. There is also another (simpler) trick to disallow freeness of *edges*, while maintaining contractibility:

How to use the gadget above

This gadget G_d works as a relay: by attaching it to another d-complex, we can disallow some collapsing/shelling sequences.

Exercise if you don't care about discrete Morse theory (which comes next): let e be an edge in the boundary of a 2-disk D. Let us glue D to G_2 by identifying e with the free edge of G_2.

Is $D \cup G_2$ collapsible? How? What does a shelling order look like?

Note: given a contractible 3-complex, attaching copies of G_3 will disallow freeness of *triangles*. There is also another (simpler) trick to disallow freeness of *edges*, while maintaining contractibility: namely, we can just sew together two adjacent two boundary edges. (The resulting edge will belong to four triangles instead of two, and hence will "be freed later"). Let's call this **sewing trick**.

In discrete Morse theory (DMT), we can dismantle a complex not only via elementary collapses, but also via deletions of top-dimensional faces (called "critical").

Discrete Morse theory in two minutes

In discrete Morse theory (DMT), we can dismantle a complex not only via elementary collapses, but also via deletions of top-dimensional faces (called “critical”). The **discrete Morse vector** counts how many of these exceptional steps were necessary (dimension by dimension).

Discrete Morse theory in two minutes

In discrete Morse theory (DMT), we can dismantle a complex not only via elementary collapses, but also via deletions of top-dimensional faces (called "critical"). The **discrete Morse vector** counts how many of these exceptional steps were necessary (dimension by dimension).

For example: that a triangulated 3-sphere "has discrete Morse vector $(1, 0, 0, 1)$" means that, after removing a tetrahedron, is collapsible. By the way, being collapsible is the same as "having discrete Morse vector $(1, 0, \ldots, 0)$".

The name comes from Morse theory in differential geometry, which is source of inspirations for most results.

Discrete Morse theory in two minutes

In discrete Morse theory (DMT), we can dismantle a complex not only via elementary collapses, but also via deletions of top-dimensional faces (called "critical"). The **discrete Morse vector** counts how many of these exceptional steps were necessary (dimension by dimension).

For example: that a triangulated 3-sphere "has discrete Morse vector $(1, 0, 0, 1)$" means that, after removing a tetrahedron, is collapsible. By the way, being collapsible is the same as "having discrete Morse vector $(1, 0, \dots, 0)$".

The name comes from Morse theory in differential geometry, which is source of inspirations for most results. One result that hasn't been discretized yet, is that any smooth manifold has a unique optimal Morse vector (i.e., the poset of smooth Morse vectors, partially ordered by their norm, has a unique minimum).

Discrete Morse theory in two minutes

In discrete Morse theory (DMT), we can dismantle a complex not only via elementary collapses, but also via deletions of top-dimensional faces (called "critical"). The **discrete Morse vector** counts how many of these exceptional steps were necessary (dimension by dimension).

For example: that a triangulated 3-sphere "has discrete Morse vector $(1,0,0,1)$" means that, after removing a tetrahedron, is collapsible. By the way, being collapsible is the same as "having discrete Morse vector $(1,0,\ldots,0)$".

The name comes from Morse theory in differential geometry, which is source of inspirations for most results. One result that hasn't been discretized yet, is that any smooth manifold has a unique optimal Morse vector (i.e., the poset of smooth Morse vectors, partially ordered by their norm, has a unique minimum).

A complex with two main collapsing strategies

Recall that we have constructed

- a collapsible G_2, with an edge e that belongs to only one triangle T, and
- a collapsible G_3, with a triangle T' that belongs to only one tetrahedron Δ.

Recall that we have constructed

- a collapsible G_2, with an edge e that belongs to only one triangle T, and
- a collapsible G_3, with a triangle T' that belongs to only one tetrahedron Δ.

What happens if we glue them by identifying T and T'?

A complex with two main collapsing strategies

Recall that we have constructed

- a collapsible G_2, with an edge e that belongs to only one triangle T, and
- a collapsible G_3, with a triangle T' that belongs to only one tetrahedron Δ.

What happens if we glue them by identifying T and T'?
The resulting complex has **two**, very distinct ways to be collapsed, depending on whether you wish to pair T with e or with Δ. The first gives

A complex with two main collapsing strategies

Recall that we have constructed

- a collapsible G_2, with an edge e that belongs to only one triangle T, and
- a collapsible G_3, with a triangle T' that belongs to only one tetrahedron Δ.

What happens if we glue them by identifying T and T'?
The resulting complex has **two**, very distinct ways to be collapsed, depending on whether you wish to pair T with e or with Δ. The first gives $(1, 1, 0, 1)$, the second gives...

A complex with two main collapsing strategies

Recall that we have constructed

- a collapsible G_2, with an edge e that belongs to only one triangle T, and
- a collapsible G_3, with a triangle T' that belongs to only one tetrahedron Δ.

What happens if we glue them by identifying T and T'?
The resulting complex has **two**, very distinct ways to be collapsed, depending on whether you wish to pair T with e or with Δ. The first gives $(1, 1, 0, 1)$, the second gives... $(1, 0, 0, 0)$.

Many technical discussions later...

Many technical discussions later...

Indeed:

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 3$, there is a non-collapsible d-complex A_d, that admits two discrete Morse vectors $(1, 0, \ldots, 0, 1, 1, 0)$ and $(1, 0, \ldots, 0, 0, 1, 1)$.

Many technical discussions later...

Indeed:

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 3$, there is a non-collapsible d-complex A_d, that admits two discrete Morse vectors $(1, 0, \ldots, 0, 1, 1, 0)$ and $(1, 0, \ldots, 0, 0, 1, 1)$. The complex we construct in dimension three is obtained from G_2 and ten copies of G_3 (one of them slightly modified via a sewing trick), totaling up to 106 vertices.

Many technical discussions later...

Indeed:

Theorem (Adiprasito–B.–Lutz, 2015)

In each dimension $d \geq 3$, there is a non-collapsible d-complex A_d, that admits two discrete Morse vectors $(1, 0, \ldots, 0, 1, 1, 0)$ and $(1, 0, \ldots, 0, 0, 1, 1)$. The complex we construct in dimension three is obtained from G_2 and ten copies of G_3 (one of them slightly modified via a sewing trick), totaling up to 106 vertices.

It remains open (and rather interesting) if there is a **triangulated 3-manifold** with two different optimal discrete Morse vectors.

Bonus

A bonus result if you like non-evasiveness, a strengthening of collapsibility.

Bonus

A bonus result if you like non-evasiveness, a strengthening of collapsibility. It is well known that

(Folklore)

Every (finite) tree has at least two leaves, and the bound is sharp: Paths have exactly 2 leaves.

How can one generalize this to higher dimensions?

Bonus

A bonus result if you like non-evasiveness, a strengthening of collapsibility. It is well known that

(Folklore)

Every (finite) tree has at least two leaves, and the bound is sharp: Paths have exactly 2 leaves.

How can one generalize this to higher dimensions? For example, it is not true that every collapsible d-complex has at least two free faces!, since our G_d has one...

Bonus

A bonus result if you like non-evasiveness, a strengthening of collapsibility. It is well known that

(Folklore)

Every (finite) tree has at least two leaves, and the bound is sharp: Paths have exactly 2 leaves.

How can one generalize this to higher dimensions? For example, it is not true that every collapsible d-complex has at least two free faces!, since our G_d has one...

Theorem (Adiprasito–B.–Lutz, 2015)

Let d be any positive integer. Every nonevasive d-complex has at least two free faces, and the bound is sharp: We constructed nonevasive d-complexes with exactly 2 leaves, for all d.