ORIGINAL PAPER

Collapsibility of CAT(0) spaces

Karim Adiprasito[1] · Bruno Benedetti[2]

Abstract

Collapsibility is a combinatorial strengthening of contractibility. We relate this property to metric geometry by proving the collapsibility of any complex that is CAT(0) with a metric for which all vertex stars are convex. This strengthens and generalizes a result by Crowley. Further consequences of our work are:

(1) All CAT(0) cube complexes are collapsible.
(2) Any triangulated manifold admits a CAT(0) metric if and only if it admits collapsible triangulations.
(3) All contractible d-manifolds ($d \neq 4$) admit collapsible CAT(0) triangulations. This discretizes a classical result by Ancel–Guilbault.

Keywords CAT(0) spaces · Collapsibility · Discrete Morse theory · Convexity Evasiveness · Triangulations

Mathematics Subject Classification 52B70 · 52A20 · 57Q10 · 53C23

1 Introduction

Whitehead's "simple homotopy theory", a combinatorial approach to homotopy, was partly motivated by Poincaré's conjecture that homotopy recognizes spheres among all closed manifolds. In contrast, Whitehead discovered that homotopy does not recognize $\mathbb{R}^n$ among all (open) manifolds. In fact, for each $d \geq 4$, there are compact contractible smooth d-manifolds whose boundary is not simply-connected [43,44,50].

In 1939, Whitehead introduced *collapsibility*, a combinatorial version of the notion of contractibility. All collapsible complexes are contractible, but the converse is false [16,52]. The collapsibility property lies at the very core of PL topology and is of particular interest when applied to triangulations of manifolds, cf. [12,14]. Using the idea of regular neighbor-

✉ Bruno Benedetti
  bruno@math.miami.edu

  Karim Adiprasito
  adiprasito@math.huji.ac.il

[1] Einstein Institute for Mathematics, Hebrew University of Jerusalem, 91904 Jerusalem, Israel

[2] Department of Mathematics, University of Miami, Coral Gables, Florida 33146, USA

hoods, Whitehead proved that the only manifold admitting collapsible PL triangulations is the ball [51].

Recently, collapsibility has regained interest in areas ranging from extremal combinatorics [10] to the study of subspace arrangements [1,46]; similar notions such as *nonevasiveness* and *dismantlability* have been explored in connection with theoretical computer science, cf. e.g. [23,27,40]. Unfortunately, there are only a few available criteria to predict the collapsibility of a given complex:

(a) All cones are collapsible.
(b) *Chillingworth's criterion* [26]: All subdivisions of convex 3-polytopes are collapsible.
(c) *Crowley's criterion* [28]: All 3-dimensional pseudomanifolds that are CAT(0) with the equilateral flat metric (see below for the meaning), are collapsible.

The main goal of this paper is to show that criterion (c) extends to all dimensions. In fact, in a separate companion paper, we will show that criterion (b) extends too [3]. Breaking the barrier of dimension three will allow us to list a few consequences at the interplay of combinatorics and metric geometry.

Crowley's criterion was advertised as a first "*combinatorial analog of Hadamard's theorem*" [28, p. 36]. Being CAT(0) is a property of metric spaces, popularized in the Eighties by Gromov [37]. Roughly speaking, CAT(0) spaces are metric spaces where any triangle formed by three geodesic segments looks thinner than any triangle in the euclidean plane formed by three straight lines of the same lengths. The classical Hadamard–Cartan theorem guarantees that simply connected, complete, locally CAT(0) spaces are (globally) CAT(0), and hence contractible.

Being "CAT(0) with the equilateral flat metric" is instead a property of simplicial complexes. It means that if we give the complex a piecewise-euclidean metric by assigning unit length to all its edges, then the complex becomes a CAT(0) metric space. Our first result is to prove the collapsibility of such complexes in all dimensions. In fact, we show something stronger:

Theorem I (Theorem 3.2.1) *Let C be a simplicial complex that is* CAT(0) *with a metric for which all vertex stars are convex. Then C is collapsible.*

Convexity of vertex stars means that for each vertex v, with respect to the metric introduced on C, the segment between any two points of St (v, C) lies in St (v, C). This condition cannot be removed, because not all triangulated 3-balls are collapsible (although all of them are CAT(0) with a suitable metric). It is automatically satisfied by any complex in which all simplices are acute or right-angled. The same condition plays also a central role in the authors' proof of the Hirsch conjecture for flag polytopes [4].

Corollary II *Every complex that is* CAT(0) *with the equilateral flat metric is collapsible. Similarly, every* CAT(0) *cube complex is collapsible.*

For example, the "space of phylogenetic trees" introduced by Billera, Holmes and Vogtmann [15] turns out to be collapsible (Corollary 3.2.10). A further application to enumerative combinatorics is the discrete Cheeger theorem in [2].

The converses of Theorem I and Corollary II are false. Being CAT(0) with the equilateral flat metric is a more restrictive property than being collapsible; for example, any vertex-transitive complex with this property is a simplex (Proposition 3.2.11), while Oliver provided collapsible simplicial complexes with vertex-transitive symmetry group that are not the simplex, cf. [40]. Nonetheless, a partial converse is possible if we are allowed to change the triangulation, as explained below.

Theorem I allows us to understand the topology of manifolds that admit collapsible triangulations. This was investigated by Whitehead only in the PL case, where he proved that collapsible PL triangulations are balls [51]. Revisiting previous work of Ancel and Guilbault [7,8], we are able to prove the following:

Theorem III (Theorem 3.3.2, cf. [7,8]) *For any integer d, the d-manifolds admitting a collapsible triangulation are precisely those that admit a* CAT(0) *(polyhedral) metric. In particular, all contractible d-manifolds admit a collapsible triangulation, except for d = 4.*

For simplicial complexes, the situation changes slightly: paired with a result by the first author and Funar [5, Proposition 15], Theorem I still implies that the complexes admitting a collapsible subdivision are precisely those that admit a polyhedral CAT(0) metric (Remark 3.3.3). However, these are some, but not all, of the contractible complexes: For example, the Dunce Hat has no collapsible subdivision.

2 Basic notions

2.1 Geometric and intrinsic polytopal complexes

By $\mathbb{R}^d$, $\mathbb{H}^d$ and $\mathbb{S}^d$ we denote the euclidean d-space, the hyperbolic d-space, and the unit sphere in $\mathbb{R}^{d+1}$, respectively. A *(euclidean) polytope* in $\mathbb{R}^d$ is the convex hull of finitely many points in $\mathbb{R}^d$. Similarly, a *hyperbolic polytope* in $\mathbb{H}^d$ is the convex hull of finitely many points of $\mathbb{H}^d$. A *spherical polytope* in $\mathbb{S}^d$ is the convex hull of a finite number of points that all belong to some open hemisphere of $\mathbb{S}^d$. Spherical polytopes are in natural one-to-one correspondence with euclidean polytopes, just by taking radial projections; the same is true for hyperbolic polytopes. A *geometric polytopal complex* in $\mathbb{R}^d$ (resp. in $\mathbb{S}^d$ or $\mathbb{H}^d$) is a finite collection of polytopes in $\mathbb{R}^d$ (resp. $\mathbb{S}^d$, $\mathbb{H}^d$) such that the intersection of any two polytopes is a face of both. An *intrinsic polytopal complex* is a collection of polytopes that are attached along isometries of their faces (cf. Davis–Moussong [30, Sect. 2]), so that the intersection of any two polytopes is a face of both.

Two polytopal complexes C, D are *combinatorially equivalent*, denoted by $C \cong D$, if their face posets are isomorphic. Any polytope combinatorially equivalent to the d-simplex, or to the regular unit cube $[0, 1]^d$, shall simply be called a *d-simplex* or a *d-cube*, respectively. A polytopal complex is *simplicial* (resp. *cubical*) if all its faces are simplices (resp. cubes). The set of k-dimensional faces of a polytopal complex C is denoted by $\mathrm{F}_k(C)$, and the cardinality of this set is denoted by $f_k(C)$.

The *underlying space* $|C|$ of a polytopal complex C is the topological space obtained by taking the union of its faces. If two complexes are combinatorially equivalent, their underlying spaces are homeomorphic. We will frequently abuse notation and identify a polytopal complex with its underlying space, as is common in the literature. For instance, we do not distinguish between a polytope and the complex formed by its faces. If C is simplicial, C is sometimes called a *triangulation* of $|C|$ (and of any topological space homeomorphic to $|C|$). If $|C|$ is isometric to some metric space M, then C is called a *geometric triangulation* of M.

A *subdivision* of a polytopal complex C is a polytopal complex C' with the same underlying space of C, such that for every face F' of C' there is some face F of C for which $F' \subset F$. Two polytopal complexes C and D are called *PL equivalent* if some subdivision C' of C is combinatorially equivalent to some subdivision D' of D. In case $|C|$ is a topological manifold (with or without boundary), we say that C is *PL* (short for Piecewise-Linear) if the star of every face of C is PL equivalent to the simplex of the same dimension.

A *derived subdivision* sd C of a polytopal complex C is any subdivision of C obtained by stellarly subdividing at all faces in order of decreasing dimension of the faces of C, cf. [38]. An example of a derived subdivision is the *barycentric subdivision*, which uses as vertices the barycenters of all faces of C.

If C is a polytopal complex in $\mathbb{R}^d$ (or $\mathbb{S}^d$, or $\mathbb{H}^d$) and A is a subset of $\mathbb{R}^d$ (resp. $\mathbb{S}^d$, resp. $\mathbb{H}^d$), we define the *restriction* R (C, A) *of* C to A as the inclusion-maximal subcomplex D of C such that D lies in A. The *star* of σ in C, denoted by St (σ, C), is the minimal subcomplex of C that contains all faces of C containing σ. The *deletion* $C - D$ of a subcomplex D from C is the subcomplex of C given by R $(C, C\backslash\text{relint}\, D)$.

Next, we define the notion of *link* with a metric approach. (Compare also Charney [22] and Davis–Moussong [30, Sect. 2.2].) Let p be any point of a metric space X. By $\text{T}_p X$ we denote the tangent space of X at p. Let $\text{T}^1_p X$ be the restriction of $\text{T}_p X$ to unit vectors. If Y is any subspace of X, then $\text{N}_{(p,Y)} X$ denotes the subspace of the tangent space $\text{T}_p X$ spanned by the vectors orthogonal to $\text{T}_p Y$. If p is in the interior of Y, we define $\text{N}^1_{(p,Y)} X := \text{N}_{(p,Y)} X \cap \text{T}^1_p Y$.

If τ is any face of a polytopal complex C containing a nonempty face σ of C, then the set $\text{N}^1_{(p,\sigma)} \tau$ of unit tangent vectors in $\text{N}^1_{(p,\sigma)} |C|$ pointing towards τ forms a spherical polytope $P_p(\tau)$, isometrically embedded in $\text{N}^1_{(p,\sigma)} |C|$. The family of all polytopes $P_p(\tau)$ in $\text{N}^1_{(p,\sigma)} |C|$ obtained for all $\tau \supset \sigma$ forms a polytopal complex, called the *link* of C at σ; we will denote it by Lk $_p(\sigma, C)$. If C is a geometric polytopal complex in $X^d = \mathbb{R}^d$ (or $X^d = \mathbb{S}^d$), then Lk $_p(\sigma, C)$ is naturally realized in $\text{N}^1_{(p,\sigma)} X^d$. Obviously, $\text{N}^1_{(p,\sigma)} X^d$ is isometric to a sphere of dimension $d - \dim \sigma - 1$, and will be considered as such. Up to ambient isometry Lk $_p(\sigma, C)$ and $\text{N}^1_{(p,\sigma)} \tau$ in $\text{N}^1_{(p,\sigma)} |C|$ or $\text{N}^1_{(p,\sigma)} X^d$ do not depend on p; for this reason, p will be omitted in notation whenever possible.

If C is simplicial, and v is a vertex of C, then Lk (v, C) is combinatorially equivalent to

$$(C - v) \cap \text{St}\,(v, C) = \text{St}\,(v, C) - v.$$

By convention, Lk $(\emptyset, C) := C$. If C is a simplicial complex, and σ, τ are faces of C, then $\sigma * \tau$ is the minimal face of C containing both σ and τ (assuming it exists). If σ is a face of C, and τ is a face of Lk (σ, C), then $\sigma * \tau$ is the face of C with Lk $(\sigma, \sigma * \tau) = \tau$. In both cases, the operation $*$ is called the *join*.

2.2 CAT(k) spaces and convex subsets

All of the metric spaces we consider here are *compact, connected length spaces*. For a more detailed introduction, we refer the reader to the textbook by Burago–Burago–Ivanov [20]. Given two points a, b in a length space X, we denote by $|ab|$ the *distance* between a and b, which is also the infimum (and by compactness, the minimum) of the lengths of all curves from a to b. Rectifiable curves that connect a with b and realize the distance $|ab|$ are *(geodesic) segments* of X. A *geodesic* from a to b is a curve γ from a to b which is locally a segment. A *geodesic triangle* in X is given by three vertices a, b, c connected by segments $[a, b]$, $[b, c]$ and $[a, c]$.

Let k be a real number. Depending on the sign of k, by *the k-plane* we mean either the euclidean plane (if $k = 0$), or the sphere of radius $k^{-\frac{1}{2}}$ with its length metric (if $k > 0$), or the hyperbolic plane of curvature k (if $k < 0$). A *k-comparison triangle* for $[a, b, c]$ is a triangle $[\bar{a}, \bar{b}, \bar{c}]$ in the k-plane such that $|\bar{a}\bar{b}| = |ab|$, $|\bar{a}\bar{c}| = |ac|$ and $|\bar{b}\bar{c}| = |bc|$. A length space X "*has curvature* $\leq k$" if locally (i.e. in some neighborhood of each point of X) the following holds:

TRIANGLE CONDITION: For each geodesic triangle $[a, b, c]$ inside X and for any point d in the relative interior of $[a, b]$, one has $|cd| \leq |\bar{c}\bar{d}|$, where $[\bar{a}, \bar{b}, \bar{c}]$ is any k-comparison triangle for $[a, b, c]$ and $\bar{d}$ is the unique point on $[\bar{a}, \bar{b}]$ with $|ad| = |\bar{a}\bar{d}|$.

A CAT(k) *space* is a length space of curvature $\leq k$ in which the triangle condition holds globally, i.e. for any geodesic triangle (with edgelengths $< k^{-1/2}\pi$ if k is positive). Obviously, any CAT(0) space has curvature ≤ 0. The converse is false: The circle S^1 is non-positively curved (because it is locally isometric to $\mathbb{R}$) but not CAT(0), as shown by any geodesic triangle that is not homotopy equivalent to a point. Also, all CAT(0) spaces are contractible, while S^1 is not even simply connected. This is not a coincidence, as explained by the Hadamard–Cartan theorem, which provides a crucial local-to-global correspondence:

Theorem 2.2.1 (Hadamard–Cartan theorem, cf. Alexander–Bishop [6]) *Let X be any complete space with curvature ≤ 0. The following are equivalent:*

(1) *X is simply connected;*
(2) *X is contractible;*
(3) *X is* CAT(0).

In a CAT(0) space, any two points are connected by a unique geodesic. The same holds for CAT(k) spaces ($k > 0$), as long as the two points are at distance $< \pi k^{-\frac{1}{2}}$.

Let K be a subset of a metric space X. A set K is called *convex* if any two points of K are connected by some segment in X that lies entirely in K.

Let c be a point of X and let K be a closed subset of X, not necessarily convex. We denote by $\pi_c(K)$ the subset of the points of K at minimum distance from c, the *closest-point projection* of c to K. In case $\pi_c(K)$ contains a single point, with abuse of notation we write $\pi_c(K) = x$ instead of $\pi_c(K) = \{x\}$. This is always the case when K is convex, as the following well-known lemma shows.

Lemma 2.2.2 ([18, Prp. 2.4]) *Let X be a connected* CAT(k)-*space, $k \leq 0$. Let c be a point of X. Then the function "distance from c" has a unique local minimum on each closed convex subset K of X. Similarly, if $k > 0$ and K is at distance less than $\frac{1}{2}\pi k^{-1/2}$ from c, then there exists a unique local minimum of distance at most $\frac{1}{2}\pi k^{-1/2}$ from c.*

2.3 Discrete Morse theory, collapsibility and non-evasiveness

The *face poset* $(C, \subseteq)$ of a polytopal complex C is the set of nonempty faces of C, ordered with respect to inclusion. By $(\mathbb{R}, \leq)$ we denote the poset of real numbers with the usual ordering. A *discrete Morse function* is an order-preserving map f from $(C, \subseteq)$ to $(\mathbb{R}, \leq)$, such that the preimage $f^{-1}(r)$ of any number r consists of either one element, or of two elements one contained in the other. A *critical cell* of C is a face at which f is strictly increasing.

The function f induces a perfect matching on the non-critical cells: two cells are matched whenever they have identical image under f. This is called *Morse matching*, and it is usually represented by a system of arrows: Whenever $\sigma \subsetneq \tau$ and $f(\sigma) = f(\tau)$, one draws an arrow from the barycenter of σ to the barycenter of τ. We consider two discrete Morse functions *equivalent* if they induce the same Morse matching. Since any Morse matching pairs together faces of different dimensions, we can always represent a Morse matching by its associated partial function Θ from C to itself, defined on a subset of C as follows:

$$\Theta(\sigma) := \begin{cases} \sigma & \text{if } \sigma \text{ is unmatched,} \\ \tau & \text{if } \sigma \text{ is matched with } \tau \text{ and } \dim \sigma < \dim \tau. \end{cases}$$

A *discrete vector field* V on a polytopal complex C is a collection of pairs (σ, Σ) of faces such that σ is a codimension-one face of Σ, and no face of C belongs to two different pairs of V. A *gradient path* in V is a concatenation of pairs of V

$$(\sigma_0, \Sigma_0), (\sigma_1, \Sigma_1), \ldots, (\sigma_k, \Sigma_k), \ k \geq 1,$$

so that for each i the face σ_{i+1} is a codimension-one face of Σ_i different from σ_i. A gradient path is *closed* if $\sigma_0 = \sigma_k$ for some k (that is, if the gradient path forms a closed loop). A discrete vector field V is a Morse matching if and only if V contains no closed gradient paths [32]. The main result of discrete Morse Theory is the following:

Theorem 2.3.1 (Forman [33, Theorem 2.2]) *Let C be a polytopal complex. For any Morse matching on C, the complex C is homotopy equivalent to a CW complex with one i-cell for each critical i-simplex.*

Inside a polytopal complex C, a *free face* σ is a face strictly contained in only one other face of C. An *elementary collapse* is the deletion of a free face σ from a polytopal complex C. We say that *C (elementarily) collapses* onto $C - \sigma$, and write $C \searrow_e C - \sigma$. We also say that the complex *C collapses* to a subcomplex C', and write $C \searrow C'$, if C can be reduced to C' by a sequence of elementary collapses. A *collapsible* complex is a complex that collapses onto a single vertex. Collapsibility is, clearly, a combinatorial property (i.e. it only depends on the combinatorial type), and does not depend on the geometric realization of a polytopal complex. The connection to discrete Morse theory is highlighted by the following simple result:

Theorem 2.3.2 (Forman [32]) *Let C be a polytopal complex. The complex C is collapsible if and only if C admits a discrete Morse function with only one critical face.*

Collapsible complexes are contractible; collapsible PL manifolds are necessarily balls [51]. The following facts are easy to verify.

Lemma 2.3.3 *Let C be a simplicial complex, and let C' be a subcomplex of C. Then the cone over base C collapses to the cone over C'.*

Lemma 2.3.4 *Let v be any vertex of any simplicial complex C. If $\mathrm{Lk}\,(v, C)$ collapses to some subcomplex S, then C collapses to $(C - v) \cup (v * S)$. In particular, if $\mathrm{Lk}\,(v, C)$ is collapsible, then $C \searrow C - v$.*

Lemma 2.3.5 *Let C denote a simplicial complex that collapses to a subcomplex C'. Let D be a simplicial complex such that $D \cup C$ is a simplicial complex. If $D \cap C = C'$, then $D \cup C \searrow D$.*

Proof It is enough to consider the case $C \searrow_e C' = C - \sigma$, where σ is a free face of C. For this, notice that the natural embedding $C \mapsto D \cup C$ takes the free face $\sigma \in C$ to a free face of $D \cup C$.   □

Non-evasiveness is a further strengthening of collapsibility that emerged in theoretical computer science [40]. A 0-dimensional complex is *non-evasive* if and only if it is a point. Recursively, a d-dimensional simplicial complex $(d > 0)$ is *non-evasive* if and only if there is some vertex v of the complex whose link and deletion are both non-evasive. Again, non-evasiveness is a combinatorial property only.

The notion of non-evasiveness is rather similar to vertex-decomposability, a notion defined only for *pure* simplicial complexes [45]; to avoid confusions, we recall the definition and

explain the difference in the lines below. A 0-dimensional complex is *vertex-decomposable* if and only if it is a finite set of points. In particular, not all vertex-decomposable complexes are contractible. Recursively, a *d*-dimensional simplicial complex ($d > 0$) is *vertex-decomposable* if and only if it is pure and there is some vertex v of the complex whose link and deletion are both vertex-decomposable. Note that in particular the link and the deletion in a vertex decomposition have to be pure. All vertex-decomposable contractible complexes are easily seen to be non-evasive, similar to the basic fact that contractible shellable complexes are collapsible [17]. An important difference arises when considering cones. It is easy to see that the cone over a simplicial complex C is vertex-decomposable if and only if C is. In contrast,

Lemma 2.3.6 (cf. Welker [49]) *The cone over any simplicial complex is non-evasive.*

By Lemma 2.3.4 every non-evasive complex is collapsible. As a partial converse, we also have the following lemma

Lemma 2.3.7 (cf. Welker [49]) *The derived subdivision of every collapsible complex is non-evasive. In particular, the derived subdivision of any non-evasive complex is non-evasive.*

A *non-evasiveness step* is the deletion from a simplicial complex C of a single vertex whose link is non-evasive. Given two simplicial complexes C and C', we write $C \searrow_{\mathrm{NE}} C'$ if there is a sequence of non-evasiveness steps which lead from C to C'. We will need the following lemmas, which are well known and easy to prove:

Lemma 2.3.8 *If $C \searrow_{\mathrm{NE}} C'$, then $\mathrm{sd}^m C \searrow_{\mathrm{NE}} \mathrm{sd}^m C'$ for all non-negative m.*

Lemma 2.3.9 *Let v be any vertex of any simplicial complex C. Let $m \geq 0$ be an integer. Then $(\mathrm{sd}^m C) - v \searrow_{\mathrm{NE}} \mathrm{sd}^m (C - v)$. In particular, if $\mathrm{sd}^m \mathrm{Lk}\,(v, C)$ is non-evasive, then $\mathrm{sd}^m C \searrow_{\mathrm{NE}} \mathrm{sd}^m (C - v)$.*

Proof The case $m = 0$ is trivial. We treat the case $m = 1$ as follows: The vertices of $\mathrm{sd}\,C$ correspond to faces of C, the vertices that have to be removed in order to deform $(\mathrm{sd}\,C) - v$ to $\mathrm{sd}\,(C - v)$ correspond to the faces of C strictly containing v. The order in which we remove the vertices of $(\mathrm{sd}\,C) - v$ is by increasing dimension of the associated face. Let τ be a face of C strictly containing v, and let w denote the vertex of $\mathrm{sd}\,C$ corresponding to τ. Assume all vertices corresponding to faces of τ have been removed from $(\mathrm{sd}\,C) - v$ already, and call the remaining complex D. Denote by $\mathrm{L}(\tau, C)$ the set of faces of C strictly containing τ, and let $\mathrm{F}(\tau - v)$ denote the set of nonempty faces of $\tau - v$. Then $\mathrm{Lk}\,(w, D)$ is combinatorially equivalent to the order complex of $\mathrm{L}(\tau, C) \cup \mathrm{F}(\tau - v)$, whose elements are ordered by inclusion. Every maximal chain contains the face $\tau - v$, so $\mathrm{Lk}\,(w, D)$ is a cone, which is non-evasive by Lemma 2.3.6. Thus, we have $D \searrow_{\mathrm{NE}} D - w$. The iteration of this procedure shows $(\mathrm{sd}\,C) - v \searrow_{\mathrm{NE}} \mathrm{sd}\,(C - v)$, as desired. The general case follows by induction on m: Assume that $m \geq 2$. Then

$$(\mathrm{sd}^m C) - v = (\mathrm{sd}\,(\mathrm{sd}^{m-1} C)) - v \searrow_{\mathrm{NE}} \mathrm{sd}\,((\mathrm{sd}^{m-1} C) - v)$$
$$\searrow_{\mathrm{NE}} \mathrm{sd}\,(\mathrm{sd}^{m-1}(C - v)) = \mathrm{sd}^m(C - v),$$

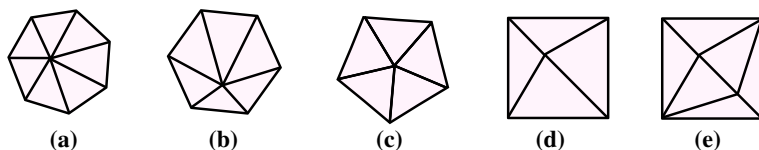by applying the inductive assumption twice, and Lemma 2.3.8 for the second deformation.
□

Fig. 1 Complex **a** is CAT(-1) with an equilateral hyperbolic metric. Complex **b** is $CAT(0)$ with the equilateral flat metric. Complex **c** is not CAT(0) with the equilateral flat metric; it is CAT(0) with the acute piecewise euclidean metric that assigns length 1 resp. 1.4 to all interior resp. boundary edges. Complex **d** is flat with the non-obtuse piecewise euclidean metric that assigns length 1 resp. $\sqrt{2}$ to all interior resp. boundary edges. Finally, only an obtuse metric can make complex **e** CAT(0); but the complex is CAT(1) with a non-obtuse metric

2.4 Acute and non-obtuse triangulations

A simplex is called *acute* (resp. *non-obtuse*) if the dihedral angle between any two facets is smaller than $\frac{\pi}{2}$ (resp. smaller or equal than $\frac{\pi}{2}$). In any acute simplex, all faces are themselves acute simplices. In particular, all triangles in an acute simplex are acute in the classical sense. The same holds for non-obtuse simplices. A simplex is called *equilateral* or *regular* if all edges have the same length. Obviously, all equilateral simplices are acute and all acute simplices are non-obtuse. The next, straightforward lemma characterizes these notions in terms of orthogonal projections.

Lemma 2.4.1 *A d-simplex Δ is acute (resp. non-obtuse) if and only if, for each facet F of $\partial \Delta$, the closest-point projection $\pi_{\Delta - F}$ span F of the vertex $\Delta - F$ to the affine span of F intersects the relative interior of F (resp. intersects F).*

In 2004, Eppstein–Sullivan–Üngör showed that $\mathbb{R}^3$ can be tiled into acute tetrahedra [31]. This was strengthened by Van der Zee et al. [48] and Kopczyński–Pak–Przytycki [41], who proved that the unit cube in $\mathbb{R}^3$ can be tiled into acute tetrahedra. In contrast, there is no geometric triangulation of the 4-cube into acute 4-simplices [41]. For $d \geq 5$, neither $\mathbb{R}^d$ nor the $(d+1)$-cube have acute triangulations, cf. [41]. In contrast, by subdividing a cubical grid, one can obtain non-obtuse triangulations of $\mathbb{R}^d$ and of the d-cube for any d. So, acute is a much more restrictive condition than non-obtuse.

Let C be an intrinsic simplicial complex in which every face of C is isometric to a regular euclidean simplex. If such C is CAT(0), we say that it is CAT(0) *with the equilateral flat metric*. More generally, a CAT(k) intrinsic simplicial complex C is CAT(k) *with an acute metric* (resp. CAT(k) *with a non-obtuse metric*) if every face of C is acute (resp. non-obtuse), see Fig. 1.

3 Collapsibility and curvature bounded above

3.1 Gradient matchings and star-minimal functions

Here we obtain non-trivial Morse matchings on a simplicial (or polytopal) complex, by studying real-valued continuous functions on the complex.

Definition 3.1.1 (Star-minimal) Let C be an intrinsic simplicial complex. A (nonlinear) function $f : |C| \to \mathbb{R}$ is called *star-minimal* if it satisfies the following three conditions:

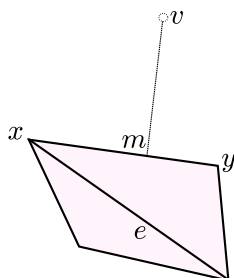**Fig. 2** How to match an edge e with respect to the function $f =$ distance from some point v. This (nonlinear) function has a unique minimum m on $|\mathrm{St}\,(e, C)|$. The inclusion-minimal face containing m in its interior is spanned by the vertices x and y. Between these two vertices, we choose the one with minimal distance from v. If it is x, since x is contained in e, we do nothing. If it is y, which does not belong to e, we match e with $e * y$

  (i) f is continuous,

 (ii) On the star of each face of C, the function f has a unique absolute minimum, and

(iii) No two vertices have the same value under f.

Condition (iii) is just a technical detail, since it can be forced by 'wiggling' the complex a bit, or by perturbing f. Alternatively, we can perform a 'virtual wiggling' by choosing a strict total order $\vartriangleleft$ on the vertices of C such that so that if $f(x) < f(y)$, then $x \vartriangleleft y$.

Conditions (i) and (ii) are also not very restrictive: Every generic continuous function on any simplicial complex is star-minimal.

Our next goal is to show that any star-minimal function on a complex C naturally induces a certain type of Morse matching, called *gradient matching*. The key is to define a "pointer function" $y_f : C \to C$, which intuitively maps each face into the "best" vertex of its star. (How good a vertex is, is decided simply by looking at its image under f.) Unlike f, which is defined on $|C|$, the map y_f is purely combinatorial. Later, we will obtain a matching from y_f basically by pairing every still unmatched face σ together with the face $\sigma * y_f(\sigma)$ (Fig. 2).

In detail: Let σ be a face of C. On the star of σ, the function f has a unique minimum m. We denote by $\mu(\sigma)$ the inclusion-minimal face among all the faces of $\mathrm{St}\,(\sigma, C)$ that contain the point m. We set

$$y_f(\sigma) \; = \; y,$$

where y is the f-minimal vertex of $\mu(\sigma)$. Since f is injective on the vertices of C, the pointer function y_f is well-defined.

Next, we define a matching $\Theta_f : C \longrightarrow C$ associated with the function f. The definition is recursive on the dimension: In other words, we start defining Θ_f for the faces of lower dimension, working all the way up to the facets. Let σ be a face of C. Set $\Theta(\emptyset) := \emptyset$. If for all faces $\tau \subsetneq \sigma$ one has $\Theta_f(\tau) \neq \sigma$, we define

$$\Theta_f(\sigma) \; := \; y_f(\sigma) * \sigma.$$

Note that every face of C is either in the image of Θ_f, or in its domain, or both. In the latter case, we have $\Theta_f(\sigma) = \sigma$.

Definition 3.1.2 (Gradient matching) Let $\Theta : C \longrightarrow C$ be (the partial function associated to) a matching on the faces of a simplicial complex C. We say that Θ is a *gradient matching* if $\Theta = \Theta_f$ for some star-minimal function $f : |C| \to \mathbb{R}$.

Next we show that all gradient matchings are indeed Morse matchings.

Theorem 3.1.3 *Let C be a simplicial complex. Let $f : |C| \to \mathbb{R}$ be any star-minimal continuous function on the underlying space of C. Then the induced gradient matching Θ_f is a Morse matching. Moreover, the map*

$$\sigma \mapsto (y_f(\sigma), \sigma)$$

yields a bijection between the set $\mathfrak{C}$ of the (nonempty) critical faces and the set

$$\mathfrak{P} := \left\{ (v, \tau) \mid v \subset \tau, \ y_f(\tau) = v \text{ and } y_f(\tau - v) \neq v \right\}.$$

In particular, the complex C admits a discrete Morse function with c_i critical i-simplices, where

$$c_i = \# \left\{ (v, \tau) \mid v \subset \tau, \ \dim \tau = i, \ y_f(\tau) = v \text{ and } y_f(\tau - v) \neq v \right\}.$$

Proof Our proof has four parts:

  (I) The pairs $(\tau, \Theta_f(\tau))$ form a discrete vector field;
 (II) This discrete vector field contains no closed gradient path;
(III) The image of $\sigma \mapsto (y_f(\sigma), \sigma)$ is in $\mathfrak{P}$.
(IV) The map $\sigma \mapsto (y_f(\sigma), \sigma)$ is a bijection.

Parts (I) and (II) guarantee that the pairs $(\tau, \Theta_f(\tau))$ form a Morse matching; parts (III) and (IV) determine how many of the faces are critical. The latter count will turn out to be crucial for establishing Theorem 3.2.1 below.

Part (I). *The pairs $(\tau, \Theta_f(\tau))$ form a discrete vector field.* We have to show that Θ_f is injective. Suppose Θ_f maps two distinct k-faces σ_1, σ_2 to the same $(k+1)$-face Σ. By definition of Θ_f,

$$y_f(\sigma_1) * \sigma_1 = \Sigma = y_f(\sigma_2) * \sigma_2.$$

All vertices of Σ are contained either in σ_1 or in σ_2. So $\text{St}(\sigma_1, C) \cap \text{St}(\sigma_2, C) = \text{St}(\Sigma, C)$. Moreover, each $y_f(\sigma_i)$ belongs to Σ but not to σ_i; so $y_f(\sigma_i)$ must belong to σ_{3-i}. This means that the function f attains its minimum on $\text{St}(\sigma_i, C)$ at a point m_i which lies in σ_{3-i}. Since on $\text{St}(\sigma_1, C) \cap \text{St}(\sigma_2, C) = \text{St}(\Sigma, C)$, the function f has a unique minimum, we obtain $m_1 = m_2$. But then $y_f(\sigma_1) = y_f(\sigma_2) = y_f(\Sigma)$. Since $y_f(\sigma_1) * \sigma_1 = y_f(\sigma_2) * \sigma_2$, it follows that $\sigma_1 = \sigma_2$, a contradiction.

Part (II). *The discrete vector field contains no closed gradient path.* Consider the function $\min_{\text{St}(\sigma, C)} f$ from C to $\mathbb{R}$. Choose an arbitrary gradient path $(\sigma_1, \Sigma_1), \ldots, (\sigma_k, \Sigma_k)$. In this path, consider a pair (σ_i, Σ_i). By definition of gradient path, the faces σ_i and Σ_i are matched with one another, and σ_{i+1} is matched with Σ_{i+1} and so on. In particular $\Theta_f(\sigma_i) = \Sigma_i$ and $\Theta_f(\sigma_{i+1}) = \Sigma_{i+1}$, and hence $\min_{\text{St}(\sigma_i, C)} f = \min_{\text{St}(\Sigma_i, C)} f$. We claim that in this case

$$\min_{\text{St}(\Sigma_i, C)} f > \min_{\text{St}(\sigma_{i+1}, C)} f.$$

In fact, suppose by contradiction that $\min_{\text{St}(\Sigma_i, C)} f = \min_{\text{St}(\sigma_{i+1}, C)} f$. Note that $\sigma_{i+1} \subset \Sigma_i$. By the star-minimality of f, the local minimum of f must be unique. So

$$y_f(\sigma_i) = y_f(\Sigma_i) = y_f(\sigma_{i+1}).$$

So $y_f(\sigma_{i+1})$ is the unique vertex in Σ_i but not in σ_i. In particular, we have that $y_f(\sigma_{i+1})$ belongs to σ_{i+1}. Yet this contradicts the fact that σ_{i+1} is matched with coface Σ_{i+1}. In

conclusion, for any gradient path $\sigma_1, \ldots, \sigma_m$ one has $\min_{\mathrm{St}(\sigma_i, C)} f = \min_{\mathrm{St}(\Sigma_i, C)} f > \min_{\mathrm{St}(\sigma_{i+1}, C)} f$. This excludes the possibility that the gradient path is closed.

Part (III). *The image of* $\sigma \mapsto (y_f(\sigma), \sigma)$ *is in* $\mathfrak{P}$. Recall that

$$\mathfrak{P} := \big\{ (v, \tau) : v \subset \tau, \ y_f(\tau) = v \text{ and } y_f(\tau - v) \neq v \big\}.$$

Consider an arbitrary vertex w of C. Either $y_f(w) \neq w$, or $y_f(w) = w$. If $y_f(w)$ is a vertex x different from w, then w is in the domain of Θ_f; moreover, we see that $\Theta_f(w)$ is the edge $[x, w]$ and w is not critical. If $y_f(w) = w$, again w is in the domain of Θ_f; one has $\Theta_f(w) = w$, so w is critical. In this case it is easy to verify that (w, w) belongs to $\mathfrak{P}$.

   Now, consider a critical face σ of dimension $k \geq 1$. Since σ is critical, there is no $(k-1)$-face τ such that $\Theta_f(\tau) = \sigma$. By the definition of Θ_f, the face σ is in the domain of Θ_f; so since σ is critical we must have $\Theta_f(\sigma) = \sigma$. So the vertex $y_f(\sigma)$ belongs to σ. In order to conclude that $(y_f(\sigma), \sigma) \in \mathfrak{P}$, we only have to prove that $y_f(\sigma - y_f(\sigma)) \neq y_f(\sigma)$. Let us adopt the abbreviations $v := y_f(\sigma)$ and $\delta := \sigma - v$. Suppose by contradiction that $y_f(\delta) = v$. If δ were not in the image of any of its facets under the map Θ_f, then we would have $\Theta_f(\delta) = \sigma$, which would contradict the assumption that σ is a critical face. So, there has to be a codimension-one face ρ of δ such that $\Theta_f(\rho) = \delta$. In other words, we see that $y_f(\rho)$ is the unique vertex that belongs to δ but not to ρ. Clearly $\operatorname{argmin}_{\mathrm{St}(\rho, C)} f \in \mathrm{St}(y_f(\rho), C)$, whence

$$\operatorname{argmin}_{\mathrm{St}(\rho, C)} f = \operatorname{argmin}_{\mathrm{St}(\rho * y_f(\rho), C)} f = \operatorname{argmin}_{\mathrm{St}(\delta, C)} f.$$

So $y_f(\rho) = y_f(\delta)$. Recall that we are assuming $y_f(\delta) = v$, where $\delta := \sigma - v$. Hence

$$y_f(\rho) = y_f(\delta) = y_f(\sigma - v) = v \notin \sigma - v.$$

This contradicts the fact that $\Theta_f(\rho) = \sigma - v$. Thus, the assumption $y_f(\delta) = v$ must be wrong.

Part (IV). *The map* $\sigma \mapsto (y_f(\sigma), \sigma)$ *from* $\mathfrak{C}$ *to* $\mathfrak{P}$ *is a bijection.*

   The map $s : \sigma \mapsto (y_f(\sigma), \sigma)$ is clearly injective. Let us verify surjectivity. Consider a pair (v, τ) in $\mathfrak{P}$. By definition of s, $y_f(\tau) = v$ and $y_f(\tau - v) \neq v$. Assume that $s(\tau) \notin \mathfrak{P}$, or, equivalently, that τ is not critical.

   Since $y_f(\tau) = v \subset \tau$, there must be a facet η of τ such that $\Theta_f(\eta) = \tau$. By definition of Θ_f, we have $\operatorname{argmin}_{\mathrm{St}(\eta, C)} \in \mathrm{St}(\tau, C)$ and hence $\operatorname{argmin}_{\mathrm{St}(\eta, C)} = \operatorname{argmin}_{\mathrm{St}(\tau, C)}$. Since y_f only depends on the point where f attains the minimum, $y_f(\tau) = y_f(\eta) = y_f(\tau - v) = v$. This contradicts the assumption that (v, τ) is in $\mathfrak{P}$. Thus τ is critical, as desired.   □

Corollary 3.1.4 *Every gradient matching is a Morse matching.*

   In fact, one can characterize gradient matchings within Morse matchings as follows: A Morse matching is a gradient matching if and only if for every matching pair $(\sigma, \Theta(\sigma))$ there is a facet $\Sigma \supset \Theta(\sigma)$ of C such that, for any face τ of Σ that contains σ but does not contain $\Theta(\sigma)$, we have $\Theta(\tau) = \tau * \Theta(\sigma)$.

3.2 Discrete Hadamard–Cartan Theorem

In this section, we prove our discrete version of the Hadamard–Cartan theorem, namely, that CAT(0) complexes with convex vertex stars are collapsible (Theorem 3.2.1). Our theorem strengthens a result by Crowley, whose proof technique we shall briefly discuss.

Crowley's approach uses the classical notion of *minimal disk*, which is the disk of minimum area spanned by a 1-sphere in a simply connected complex. Gromov [37] and Gersten [35,36] have studied minimal disks in connection with group presentations and the word problem; later Chepoi and others [11,25] have used them to relate CAT(0) cube complexes to median graphs, and systolic complexes to bridged graphs. To collapse a complex onto a fixed vertex v, Crowley's idea is to reduce the problem to the two-dimensional case, by studying the minimal disk(s) spanned by two geodesics that converge to v. Her argument is based on two observations:

(1) If the complex is three-dimensional, these minimal disks are CAT(0) with the equilateral flat metric, as long as the starting complex is. This was proven in [28, Theorem 2]. A similar result was obtained for "systolic" complexes in [24] and [39], and for "weakly systolic" complexes in [25, Claim 1], [28, Theorem 2] implies that all three-dimensional simplicial complex with equilateral simplices are systolic.
(2) Any contractible 3-dimensional pseudomanifold is collapsible if and only if it admits a discrete Morse function without critical edges.

Neither of these two facts extends to higher-dimensional manifolds:

(1) If we take the join of a triangle with a cycle of length 5, the resulting complex is obviously collapsible, but not with Crowley's argument (even though it does support a CAT(0) equilateral flat metric): The minimal disk bounded by the 5-cycle contains a degree-five vertex, and is consequently not CAT(0) with the equilateral flat metric.
(2) For any integers $m \geq 0$ and $d \geq 6$, some contractible d-manifold different from a ball admits a discrete Morse function with 1 critical vertex, m critical $(d-3)$-faces, m critical $(d-2)$-faces, and no further critical face. (In particular, no critical edges.) This follows from discretizing, as in [34] and [13], the smooth result by Sharko [47, pp. 27–28].

Our new approach consists in applying the ideas of Sect. 3.1 to the case where C is a CAT(0) complex, and the function $f : C \to \mathbb{R}$ is the distance from a given vertex v of C. Because of the CAT(0) property, this function is star-minimal—so it will induce a convenient gradient matching on the complex. The main advantage of this new approach is that it works in all dimensions.

Theorem 3.2.1 *Let C be a* CAT(k) *intrinsic simplicial complex such that*

 (i) *for each face σ in C, the underlying space of* St (σ, C) *in C is convex, and*
(ii) *if $k > 0$, every facet is contained in some ball of radius $< \frac{1}{2}\pi k^{-1/2}$.*

Then C is collapsible.

Proof Fix x in C such that, if $k > 0$, then every facet is contained in some ball of radius $< \frac{1}{2}\pi k^{-1/2}$ around x. Let d $: C \longmapsto \mathbb{R}$ be the distance from x in C, and let w denote the vertex of C that minimizes d. Let us perform on the face poset of C the Morse matching constructed in Theorem 3.1.3. The vertex w will be mapped onto itself and for each vertex $u \neq w$ we have $y_{\mathrm{d}}(u) \neq u$. So every vertex is matched with an edge, apart from w, which is the only critical vertex.

By contradiction, suppose there is a critical face τ of dimension ≥ 1. Set $v := y_{\mathrm{d}}(\tau)$. By Lemma 2.2.2, on St (τ, C) the function d attains its minimal value in the relative interior of a face $\sigma_v \in$ St (τ, C) that contains v. Let δ be any face of St $(\tau - v, C)$ containing σ_v. Clearly δ contains v. Thus St (τ, C) and St $(\tau - v, C)$ coincide in a neighborhood of v. By Lemma 2.2.2, we have argmin $_{\mathrm{St}\,(\tau-v,C)}$ d $=$ argmin $_{\mathrm{St}\,(\tau,C)}$ d. Therefore $y_{\mathrm{d}}(\tau - v) = y_{\mathrm{d}}(\tau) = v$. This means that

$$(v, \tau) \notin \{(v, \tau) : v \subset \tau,\ y_{\mathrm{d}}(\tau) = v \text{ and } y_{\mathrm{d}}(\tau - v) \neq v\};$$

hence by Theorem 3.1.3 τ is not critical, a contradiction. □

Remark 3.2.2 In a CAT(0) space, any closed connected locally convex subset is also convex [21]. So in the assumptions of Theorem 3.2.1 it suffices to check condition (i) locally: We can replace condition (i) with the request that St (σ, C) is convex for every ridge σ (ridges are facets of facets of a simplicial complex). In fact, if all ridges have convex stars, the closest-point projection to St (v, C) is a well-defined map and it is locally non-expansive for every vertex v. But then the closest-point projection to St (v, C) is also globally non-expansive. Now let x and y be any two points in St (v, C), and let us assume the geodesic γ from x to y leaves St (v, C). Let us project γ to St (v, C). The result is a curve γ' connecting x and y that is obviously lying in St (v, C), and not longer than γ. This contradicts the uniqueness of segments in CAT(0) spaces.

Corollary 3.2.3 *Let C be an intrinsic simplicial complex. Suppose that C is* CAT(0) *with a non-obtuse metric. Then C is collapsible.*

Proof By the assumption, there is a metric structure on C of non-positive curvature, such that every face of C is non-obtuse. In non-obtuse triangulations, the star of every ridge is convex. In fact, let Σ, Σ' be two facets containing a common ridge R. Since Σ and Σ' are convex, and their union is locally convex in a neighborhood of R, we have that $\Sigma \cup \Sigma'$ is locally convex. Since the embedding space is CAT(0), convexity follows as in Remark 3.2.2. □

Corollary 3.2.4 *Every intrinsic simplicial complex that is* CAT(0) *with the equilateral flat metric is collapsible.*

Corollary 3.2.5 (Crowley [28]) *Every 3-pseudomanifold that is* CAT(0) *with the equilateral flat metric is collapsible.*

Remark 3.2.6 Technical details aside, our proof is based on the idea of "distance from a basepoint". This idea has widely been used in mathematics, and in particular in graph theory, where it lies at the core of several algorithms, such as Bread-First-Search. Using a refinement of Bread-First-Search, Chepoi and Osajda proved that weakly systolic complexes are dismantlable, and in particular collapsible. Their result can thus be viewed as a "more graph-theoretic" version of Theorem 3.2.1.

Extension to polytopal complexes

Theorem 3.2.1 can be extended even to polytopal complexes that are not simplicial. The key for this is Bruggesser–Mani's rocket shelling of polytope boundaries, cf. e.g. Ziegler [53, Sect. 8.2]. The following, well-known lemma follows from the fact that there is a (rocket) shelling of ∂P in which St (σ, ∂P) is shelled first [53, Cor. 8.13]:
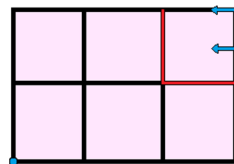
Lemma 3.2.7 (Bruggesser–Mani) *For any polytope P and for any face μ, P collapses onto* St (μ, ∂P).

Theorem 3.2.8 *Let C be any* CAT(k) *(intrinsic) polytopal complex such that*

 (i) *for each face σ in C, the underlying space of* St (σ, C) *in C is convex, and*
(ii) *if $k > 0$, every facet is contained in some ball of radius $< \frac{1}{2}\pi k^{-1/2}$.*

Then C is collapsible.

Fig. 3 How to collapse this CAT(0) cubical complex to the bottom left vertex v. We proceed by induction: Here σ is the top right square, and F is the complex highlighted in red. (Color figure online)

Proof We make the stronger claim that every polytopal complex that admits a function $f : |C| \to \mathbb{R}$ that takes a unique local minimum on each vertex star is collapsible. This situation clearly applies: Fix x in C such that, if $k > 0$, then every facet is contained in some ball of radius $< \frac{1}{2}\pi k^{-1/2}$ around x. Let $\mathrm{d} : C \longmapsto \mathbb{R}$ be the distance from x in C. This d is a function that has a unique local minimum on the star of each face by Lemma 2.2.2.

We prove the claim by induction. Let σ be a facet of C maximizing $\min_\sigma f$, and let μ denote the strict face of σ that minimizes f. Let $F \subset \mathrm{St}\,(\mu, C)$ be the subcomplex induced by the facets of C that attain their minimum at μ. By Lemma 3.2.7, we can collapse each facet P of F to $\mathrm{St}\,(\mu, \partial P)$. Hence, we can collapse F to

$$\bigcup_{P \in F} \mathrm{St}\,(\mu, \partial P) = \mathrm{St}\,(\mu, C) \cap \bigcup_{P \in F} \partial P,$$

where P ranges over the facets of F. In particular, we can collapse C to $C' := C - F$ (Fig. 3).

It remains to show the existence of a function on C' that attains a unique local minimum on each star. Let us show that the restriction $f_{|C'}$ of f to C' is the function we are looking for. Assume that for some vertex w of C the function $f_{|C'}$ attains two local minima on $\mathrm{St}\,(w, C')$. Clearly w is in F. Let x be the absolute minimum of f restricted to $\mathrm{St}\,(w, C')$; let y be the other (local) minimum. Let P be a facet of $C - F$ containing x. When restricted to C, the function f attains a unique local minimum on the star of every face. Therefore, the point y must lie in F. In particular, the facet P must contain the minimum of f on F. But y is not that local minimum since P is not in F, so f takes two local minima on P, in contradiction with the assumption on f. □

In particular, Theorem 3.2.8 holds for CAT(0) *cube complexes*, which are complexes of regular unit cubes glued together to yield a CAT(0) metric space. These complexes have been extensively studied in the literature: See for example [9,15,29,37].

Corollary 3.2.9 *Every* CAT(0) *cube complex is collapsible.*

One instance of a CAT(0) cube complex is the space of phylogenetic trees, introduced by Billera, Holmes and Vogtmann [15]:

Corollary 3.2.10 *The space of phylogenetic trees is collapsible.*

Vertex-transitive triangulations

In this short section we include a connection of Theorem 3.2.1 between metric geometry and the evasiveness conjecture, following a suggestion by Anders Björner. A *vertex-transitive* complex is a simplicial complex with n vertices on which the symmetric group acts transitively. An important open problem in theoretical computer science is whether there is any vertex-transitive non-evasive simplicial complex, apart from the simplex. This is known as *evasiveness conjecture* [40,42].

It is known that collapsibility is not enough to force a vertex-transitive complex to be a simplex [42]. However, non-evasiveness is strictly stronger than collapsibility. We have shown in Sect. 3.2 that the property of "being CAT(0) with the equilateral flat metric" is also strictly stronger than collapsibility. Thus it makes sense to compare it with vertex-transitivity, in parallel with the statement of the evasiveness conjecture. Here is a simple observation:

Proposition 3.2.11 *Every vertex-transitive simplicial complex that is* CAT(0) *with the equilateral flat metric is a simplex.*

This follows directly from the Bruhat-Tits fixed point theorem for CAT(0) spaces, cf. [19]. A direct proof is simple enough, though, so we present it here:

Proof Let C be a vertex-transitive intrinsic simplicial complex and let $v_1, \ldots, v_n$ be the vertices of C. Let $a : |C| \longmapsto \mathbb{R}$ be the function $a(x) := \sum_{i=1}^{n} d(x, v_i)$. When C is CAT(0), the function a is strongly convex and has therefore a *unique* minimum m. Since the simplicial complex is vertex transitive, and the function a is also invariant under the symmetries of C, we claim that

$$d(m, v_1) = d(m, v_2) = \ldots = d(m, v_n),$$

so that m minimizes simultaneously all functions $x \mapsto d(x, v_i)$. In fact, were $d(m, v_k) < d(m, v_j)$ for some $k \neq j$, we could find by symmetry a point m' such that

$$m \neq m' \text{ and } \sum_i d(m', v_i) = \sum_i d(m, v_i).$$

But then we would have $a(m) = a(m')$: A contradiction since a has a unique minimum.

So, there is one point in the complex equidistant from all of the vertices. This implies that C has only one facet, because for each facet F of C, the unique point equidistant from all vertices of F is the barycenter of F itself. □

3.3 Collapsible, contractible and CAT(0) manifolds

We already mentioned the famous result obtained in 1939 by Whitehead:

Theorem 3.3.1 *(Whitehead) Let M be a (compact) manifold with boundary. If some PL triangulation of M is collapsible, then M is a ball.*

Let us focus on non-PL triangulations. Let us introduce a convenient notation:

- By CAT_d we denote the class of d-manifolds homeomorphic to CAT(0) cube complexes;
- By COLL_d we denote the d-manifolds that admit a collapsible triangulation;
- By CONT_d we denote all contractible d-manifolds.

Moreover, by *PL singular set of M* we mean the subcomplex given by the faces of M whose link is not homeomorphic to a PL sphere or a PL ball.

Theorem 3.3.2 *For each $d \geq 5$ one has*

$$\mathrm{CAT}_d = \mathrm{COLL}_d = \mathrm{CONTR}_d,$$

whereas for $d = 4$ one has

$$\mathrm{CAT}_4 = \mathrm{COLL}_4 \subsetneq \mathrm{CONTR}_4.$$

In particular, when $d \geq 5$ not all collapsible manifolds are balls.

Proof. Clearly, every d-ball can be given a CAT(0) cubical structure (consisting of a single d-cube). By Theorem 3.2.8, every CAT(0) cube complex is collapsible, and so is its first derived subdivision [49]; hence CAT(0) cube complexes admit collapsible triangulations. Finally, collapsible complexes are contractible. This proves that

$$\{d\text{-balls}\} \subset \text{CAT}_d \subset \text{COLL}_d \subset \text{CONTR}_d \text{ for all } d.$$

When $d \leq 4$, every triangulation of a d-manifold is PL (This is non-trivial; for $d = 4$, this statement relies on the Poincaré-Perelman theorem). By Whitehead's theorem, this implies $\text{COLL}_d \subset \{d\text{-balls}\}$, so the first two containments above are actually equalities. All contractible 3-manifolds are balls, so

$$\{d\text{-balls}\} = \text{CAT}_d = \text{COLL}_d = \text{CONTR}_d \text{ for } d \leq 3.$$

In contrast, some contractible 4-manifolds are not balls, like the Mazur manifold [43]; hence

$$\{4\text{-balls}\} = \text{CAT}_4 = \text{COLL}_4 \subsetneq \text{CONTR}_4.$$

When $d \geq 5$, it is a classical result that $\{d\text{-balls}\} \subsetneq \text{CONTR}_d$. So all we need to show is that $\text{CONTR}_d \subset \text{CAT}_d$, namely, that every contractible d-manifold, $d \geq 5$, admits a CAT(0) cube structure. By a result of Ancel and Guilbault [7], every contractible manifold M admits a triangulation C such that:

(1) The PL singular set S of C lies in the interior of C,
(2) The PL singular set S of C is a path (i.e. a graph homeomorphic to a curve),
(3) S is a deformation retract of C, and
(4) $|C|\backslash|S|$ is homeomorphic to $M \times (0, 1]$.

Let us hyperbolize C (as in Davis–Januskiewicz [29]), and pass to the universal cover. The resulting complex C' is a CAT(0) cube complex. Since S is a path, its image under hyperbolization and lift to universal cover is a disjoint union S' of convex compact sets in C'. Let N denote the subcomplex of C' consisting of faces that intersect a fixed connected component of S'. As in [29, Thm. (5b.1)], we have that N is a cubical complex PL homeomorphic to C and consequently homeomorphic to M. Furthermore N is a convex subcomplex of a CAT(0) cube complex; in particular, it is itself a CAT(0) cube complex. In conclusion,

$$\{d\text{-balls}\} \subsetneq \text{CAT}_d = \text{COLL}_d = \text{CONTR}_d \text{ for } d \geq 5. \quad \square$$

Remark 3.3.3 The first author and Funar have recently extended this result to arbitrary complexes, proving that any collapsible simplicial complex is PL equivalent to a CAT(0) (and even CAT(-1)) polyhedral complex [5, Proposition 15]. With our Theorem 3.2.8, this implies that a simplicial complex admits a collapsible subdivision if and only if it admits a CAT(0) cube structure. However, not all contractible simplicial complexes admit a collapsible subdivision: A well known counterexample is given by the Dunce Hat.

Here is a curious consequences concerning discrete Morse theory:

Corollary 3.3.4 *Forman's discrete Morse theory can give sharper upper bounds than smooth Morse Theory (or of PL handle theory) for the Betti numbers of a manifold.*

Proof Let M be any collapsible (non-PL) triangulation of a contractible manifold different than a ball, as given by Theorem 3.3.2. By a classical result of Gleason, every open contractible manifold admits a smooth structure, and therefore a smooth handle decomposition; from that one can always obtain a PL triangulation and a PL handle decomposition. (See e.g.[13] for the

definitions.) However, since *M* is not a ball, any of its (smooth or PL) handle decompositions must contain more than one handle. In particular the Betti vector $(1, 0, \ldots, 0)$ is not a possible smooth or PL handle vector for *M*. Yet by Theorem 3.3.2, the same vector is a possible discrete Morse vector for (a suitable triangulation of) *M*, as long as dim $M \geq 5$. □

Remark 3.3.5 The situation changes if we restrict ourselves to PL triangulations. Indeed, every PL handle decomposition naturally yields a discrete Morse function, and vice versa, as proven in [13].

References

1. Adiprasito, K.: Combinatorial stratifications and minimality of 2-arrangements. J. Topol. **7**, 1200–1220 (2014)
2. Adiprasito, K., Benedetti, B.: A Cheeger-type exponential bound for the number of triangulated manifolds. To appear in Annales de l'Institut Henri Poincaré D. Preprint (2017) arXiv:1710.00130
3. Adiprasito, K., Benedetti, B.: Barycentric subdivisions of convex complexes are collapsible. To appear in Discrete & Computational Geometry. Preprint (2017) arXiv:1709.07930
4. Adiprasito, K., Benedetti, B.: The Hirsch Conjecture holds for normal flag complexes. Math. Oper. Res. **39**(4), 1340–1348 (2014)
5. Adiprasito, K., Funar, L.: Hyperbolicity of contractible manifolds. Preprint (2015) arXiv:1512.06403
6. Alexander, S.B., Bishop, R.L.: The Hadamard–Cartan theorem in locally convex metric spaces. Enseign. Math. (2) **36**, 309–320 (1990)
7. Ancel, F.D., Guilbault, C.R.: Compact contractible *n*-manifolds have arc spines ($n \geq 5$). Pac. J. Math. **168**, 1–10 (1995)
8. Ancel, F.D., Guilbault, C.R.: Interiors of compact contractible *n*-manifolds are hyperbolic ($n \geq 5$). J. Differ. Geom. **45**, 1–32 (1997)
9. Ardila, F., Owen, M., Sullivant, S.: Geodesics in CAT(0) cubical complexes. Adv. Appl. Math. **48**, 142–163 (2012)
10. Aronshtam, L., Linial, N., Łuczak, T., Meshulam, R.: Collapsibility and vanishing of top homology in random simplicial complexes. Discret. Comput. Geom. **49**, 317–334 (2013)
11. Bandelt, H.-J., Chepoi, V.: Metric graph theory and geometry: a survey. In: Goodman, J.E., Pach, J., Pollack, R. (eds.) Surveys on Discrete and Computational Geometry. Contemporary Mathematics, vol. 453, pp. 49–86. American Mathematical Society, Providence (2008)
12. Benedetti, B.: Discrete Morse theory for manifolds with boundary. Trans. Am. Math. Soc. **364**, 6631–6670 (2012)
13. Benedetti, B.: Smoothing discrete Morse theory. Ann. Sci. Norm. Sup. Pisa Cl. Sci. Ser. V **16**(2), 335–368 (2016)
14. Benedetti, B., Ziegler, G.M.: On locally constructible spheres and balls. Acta Math. **206**, 205–243 (2011)
15. Billera, L.J., Holmes, S.P., Vogtmann, K.: Geometry of the space of phylogenetic trees. Adv. Appl. Math. **27**, 733–767 (2001)
16. Bing, R.H.: Some Aspects of the Topology of 3-Manifolds Related to the Poincaré Conjecture. Lectures on Modern Mathematics, vol. II, pp. 93–128. Wiley, New York (1964)
17. Björner, A.: Topological Methods. Handbook of Combinatorics, vol. 1, 2, pp. 1819–1872. Elsevier, Amsterdam (1995)
18. Bridson, M.R., Haefliger, A.: Metric Spaces of Non-Positive Curvature. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319. Springer, Berlin (1999)
19. Bruhat, F., Tits, J.: Groupes réductifs sur un corps local. Inst. Hautes Études Sci. Publ. Math. **41**, 5–251 (1972)

20. Burago, D., Burago, Yu., Ivanov, S.: A Course in Metric Geometry. Graduate Studies in Mathematics, vol. 33. American Mathematical Society, Providence (2001)
21. Bux, K.-U., Witzel, S.: Local convexity in CAT(ˇ)-spaces. Preprint arXiv:1211.1871
22. Charney, R.: Metric geometry: connections with combinatorics. In: New Brunswick, N.J. (ed.) Formal Power Series and Algebraic Combinatorics. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 24, pp. 55–69. American Mathematical Society, Providence (1996)
23. Chepoi, V.: Bridged graphs are cop-win graphs: an algorithmic proof. J. Comb. Theory Ser. B **69**, 97–100 (1997)
24. Chepoi, V.: Graphs of some CAT(0) complexes. Adv. Appl. Math. **24**, 125–179 (2000)
25. Chepoi, V., Osajda, D.: Dismantlability of weakly systolic complexes and applications. Trans. Am. Math. Soc. **367**, 1247–1272 (2015)
26. Chillingworth, D.R.J.: Collapsing three-dimensional convex polyhedra. Proc. Camb. Philos. Soc. **63**, 353–357 (1967)
27. Civan, Yusuf, Yalçın, Ergün: Linear colorings of simplicial complexes and collapsing. J. Comb. Theory Ser. A **114**, 1315–1331 (2007)
28. Crowley, K.: Simplicial collapsibility, discrete Morse theory, and the geometry of nonpositively curved simplicial complexes. Geom. Dedic. **133**, 35–50 (2008)
29. Davis, M.W., Januszkiewicz, T.: Hyperbolization of polyhedra. J. Differ. Geom. **34**, 347–388 (1991)
30. Davis, M.W., Moussong, G.: Notes on nonpositively curved polyhedra. In: Eger, (ed.) Low Dimensional Topology. Bolyai Society Mathematical Studies, vol. 8, pp. 11–94. János Bolyai Mathematical Society, Budapest (1998)
31. Eppstein, D., Sullivan, J.M., Üngör, A.: Tiling space and slabs with acute tetrahedra. Comput. Geom. **27**, 237–255 (2004)
32. Forman, R.: Morse theory for cell complexes. Adv. Math. **134**, 90–145 (1998)
33. Forman, R.: A user's guide to discrete Morse theory. Sém. Lothar. Combin. **48**, 35 (2002)
34. Gallais, É.: Combinatorial realization of the Thom–Smale complex via discrete Morse theory. Ann. Sc. Norm. Super. Pisa Cl. Sci. (5) **9**, 229–252 (2010)
35. Gersten, S.M.: Dehn functions and L_1-norms of finite presentations. Algorithms and Classification in Combinatorial Group Theory. Mathematical Sciences Research Institute Publications, vol. 23, pp. 195–224. Springer, New York (1992)
36. Gersten, S.M.: Isoperimetric and isodiametric functions of finite presentations. Geometric Group Theory. London Mathematical Society Lecture Note Series, vol. 181, pp. 79–96. Cambridge University Press, Cambridge (1993)
37. Gromov, M.: Hyperbolic groups. Essays in Group Theory. Mathematical Sciences Research Institute Publications, vol. 8, pp. 75–263. Springer, New York (1987)
38. Hudson, J.F.P.: Piecewise Linear Topology. University of Chicago Lecture Notes. W. A. Benjamin Inc, New York (1969)
39. Januszkiewicz, T., Świątkowski, J.: Simplicial nonpositive curvature. Publ. Math. Inst. Hautes Études Sci. **104**, 1–85 (2006)
40. Kahn, J., Saks, M., Sturtevant, D.: A topological approach to evasiveness. Combinatorica **4**, 297–306 (1984)
41. Kopczyński, E., Pak, I., Przytycki, P.: Acute triangulations of polyhedra and $\mathbb{R}^N$. Combinatorica **32**, 85–110 (2012)
42. Lutz, F.H.: Examples of $\mathbb{Z}$-acyclic and contractible vertex-homogeneous simplicial complexes. Discret. Comput. Geom. **27**, 137–154 (2002)
43. Mazur, B.: A note on some contractible 4-manifolds. Ann. Math. (2) **73**, 221–228 (1961)
44. Newman, M.H.A.: Boundaries of ULC sets in Euclidean n-space. Proc. Nat. Acad. Sci. USA **34**, 193–196 (1948)
45. Provan, J.S., Billera, L.J.: Decompositions of simplicial complexes related to diameters of convex polyhedra. Math. Oper. Res. **5**, 576–594 (1980)
46. Salvetti, M., Settepanella, S.: Combinatorial Morse theory and minimality of hyperplane arrangements. Geom. Topol. **11**, 1733–1766 (2007)
47. Sharko, V.V.: Functions on manifolds. Algebraic and Topological Aspects, Translated from the Russian by V. V. Minachin [V. V. Minakhin]. Translations of Mathematical Monographs, vol. 131. American Mathematical Society, Providence (1993)
48. VanderZee, E., Hirani, A.N., Zharnitsky, V., Guoy, D.: A dihedral acute triangulation of the cube. Comput. Geom. **43**, 445–452 (2010)
49. Welker, V.: Constructions preserving evasiveness and collapsibility. Discret. Math. **207**, 243–255 (1999)
50. Whitehead, J.H.C.: A certain open manifold whose group is unity. Q. J. Math. (Oxford Series) **6**, 268–279 (1935)

51. Whitehead, J.H.C.: Simplicial spaces, nuclei and m-groups. Proc. Lond. Math. Soc. **S2–45**, 243 (1939)
52. Zeeman, E.C.: On the dunce hat. Topology **2**, 341–358 (1964)
53. Ziegler, G.M.: Lectures on Polytopes. Graduate Texts in Mathematics, vol. 152. Springer, New York (1995)