



Hannah Alpert · Alexey Balitskiy · Larry Guth

Systolic almost-rigidity modulo 2

Received June 7, 2022; revised September 25, 2023

Abstract. No power law systolic freedom is possible for the product of mod 2 systoles of dimension 1 and codimension 1. This means that any closed n -dimensional Riemannian manifold M of bounded local geometry obeys the following systolic inequality: the product of its mod 2 systoles of dimensions 1 and $n - 1$ is bounded from above by $c(n, \varepsilon) \text{Vol}(M)^{1+\varepsilon}$, if finite (if $H_1(M; \mathbb{Z}/2)$ is non-trivial).

Keywords: systole, systolic freedom.

1. Results

Let M be a finite simplicial complex whose geometric realization is a closed n -dimensional manifold. Denote by $\text{Vol}(M)$ the number of n -simplices in M . The k -dimensional mod 2 systole $\text{sys}_k(M)$ is defined as the minimal number of k -simplices in a simplicial k -cycle representing a non-trivial homology class in $H_k(M; \mathbb{Z}/2)$. If $H_k(M; \mathbb{Z}/2)$ is trivial, then $\text{sys}_k(M) = +\infty$ by convention.

The following systolic inequality (which goes back to Loewner in the case of two-dimensional torus) seems natural to conjecture:

$$\text{sys}_k(M) \cdot \text{sys}_{n-k}(M) \leq c_n \text{sys}_n(M) = c_n \text{Vol}(M), \quad (\star)$$

where $H_k(M)$ is assumed to be non-trivial to make the left-hand side finite. The phenomenon of *systolic freedom* is the failure of $(\star)$ for $n \geq 3$. For the *integral* systoles (defined similarly to the mod 2 systoles, but with $\mathbb{Z}$ instead of $\mathbb{Z}/2$), this phenomenon

Hannah Alpert: Department of Mathematics and Statistics, Auburn University, Auburn, AL 36849, USA; hcalpert@auburn.edu

Alexey Balitskiy: Institute for Advanced Study, Princeton, NJ 08540, USA;
Department of Mathematics, University of Luxembourg, 4364 Esch-sur-Alzette, Luxembourg;
alexey.balitskiy@uni.lu

Larry Guth: Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; lguth@math.mit.edu

Mathematics Subject Classification 2020: 53C23.

was pioneered by Gromov [13]. In the integral case, there are numerous counterexamples [2, 16] to $(\star)$ in all dimensions $n \geq 3$, even if the constant c_n is allowed to depend on the topological type of M . For the mod 2 systoles, the same question is much more subtle, as the mod 2 systoles are often much smaller than the integral ones. It was Freedman [11, 12] who constructed the first counterexamples to $(\star)$ over $\mathbb{Z}/2$, with $k = 1$ and all $n \geq 3$. In his family of examples, the *systolic ratio* $\frac{\text{sys}_1(M) \cdot \text{sys}_{n-1}(M)}{\text{Vol}(M)}$ has the magnitude of $\sqrt{\log(\text{Vol}(M))}$; these are examples of *polylog systolic freedom*. We prove that they are (almost) tight in the sense that *no power law systolic freedom* can be observed for $k = 1$.

Theorem 1. *For each $\varepsilon > 0$, there is a constant $c = c(n, \varepsilon)$ such that any triangulated n -dimensional manifold M with non-trivial first homology $H_1(M; \mathbb{Z}/2)$ obeys the following systolic estimate:*

$$\text{sys}_1(M) \cdot \text{sys}_{n-1}(M) \leq c \text{Vol}(M)^{1+\varepsilon}.$$

This contrasts with the case of $\mathbb{Z}$ coefficients, in which there is *no almost-rigidity* as in Theorem 1: Gromov’s examples already exhibit *power law systolic freedom*, and there are examples of *exponential systolic freedom* [19].

As for systolic inequalities over $\mathbb{Z}/2$ with $k \geq 2$, there is an estimate

$$\text{sys}_k(M) \cdot \text{sys}_{n-k}(M) \leq c_n \text{Vol}(M)^2,$$

following from the trivial observation that the number of k -simplices in M is at most $\binom{n+1}{k+1} \text{Vol}(M)$. The outstanding result of [10] yields families of 11-dimensional triangulated manifolds in which $\text{sys}_4(M) \cdot \text{sys}_7(M)$ grows faster than $\text{Vol}(M)^{2-\varepsilon}$. The idea behind this result owes a lot to the relation between the systolic freedom over $\mathbb{Z}/2$ and quantum error correction codes. Having a cellulated manifold exhibiting systolic freedom at the level of k th homology, one can consider the three-term portion of its cellular chain complex around the k th term, and use it to build a quantum code with error correction properties depending on the systolic freedom of the manifold (see [12] for the details). The idea of [10] is to “reverse-engineer” manifolds starting from the breakthrough quantum error correction codes of [15, 17]. In terms of this quantum analogy, Theorem 1 says that for $k = 1$, manifolds do not produce any outstanding codes. With our methods, we cannot say anything for $k = 2, 3$.

Systolic inequalities in the discrete setting

We will prove a discrete systolic estimate for general simplicial complexes rather than triangulated manifolds. The systoles will be substituted by certain larger systolic invariants, making the inequality stronger.

- For a non-zero homology class $a \in H_k(M; \mathbb{Z}/2)$, define $\text{sys}_a(M)$ as the minimal number of k -simplices in a simplicial k -cycle representing a .
- For a non-zero cohomology class $\alpha \in H^k(M; \mathbb{Z}/2)$, define $\text{sys}^\alpha(M)$ as the minimal number of k -simplices in a simplicial k -cycle detected by α .

- For a non-zero cohomology class $\alpha \in H^k(M; \mathbb{Z}/2)$, define $\text{cut}^\alpha(M)$ as the minimal number of $(n - k)$ -simplices in an $(n - k)$ -dimensional subcomplex H that *cuts* α in the sense that α vanishes when restricted to $M \setminus H$ (in this difference M and H are understood as geometric realizations, not the combinatorial complexes); in other words, no singular k -cycle detected by α survives if we cut M along H (hence the name).

Theorem 2. *For each $\varepsilon > 0$, there is a constant $c = c(n, \varepsilon)$ such that any finite n -dimensional simplicial complex M and any non-zero class $\alpha \in H^1(M; \mathbb{Z}/2)$ obey the following estimate:*

$$\text{sys}^\alpha(M) \cdot \text{cut}^\alpha(M) \leq c \text{Vol}(M)^{1+\varepsilon}.$$

To relate cut^α to better known systolic quantities (and deduce Theorem 1), we make two remarks.

- (1) If we know that there is a non-trivial $(n - 1)$ -cycle inside the minimizing complex H from the definition of cut^α , then we get a systolic estimate. This can be achieved by assuming that there is a cohomology class $\beta \in H^{n-1}(M; \mathbb{Z}/2)$ such that $\alpha \smile \beta \neq 0$. The Lusternik–Schnirelmann lemma applies: since α vanishes on the complement of H , β restricts to H non-trivially. Therefore, there is a non-trivial $(n - 1)$ -cycle in H detected by β , and Theorem 2 implies that

$$\text{sys}^\alpha(M) \cdot \text{sys}^\beta(M) \leq c \text{Vol}(M)^{1+\varepsilon}.$$

We comment that the inequality $\text{sys}^\beta(M) \leq \text{cut}^\alpha(M)$ can be strict, even for manifolds (see Example 14).

- (2) When M is a (triangulated) manifold, it is easy to see that any cycle representing $\alpha^* \in H_{n-1}(M; \mathbb{Z}/2)$ —the Poincaré dual class to α —serves as an admissible cutting complex in the definition of cut^α , and $\text{cut}^\alpha(M) \leq \text{sys}_{\alpha^*}(M)$. Interestingly, the minimal H turns out to be a cycle representing α^* , as we will show in Lemma 13. Therefore, $\text{cut}^\alpha(M) = \text{sys}_{\alpha^*}(M)$, and Theorem 2, rephrased for manifolds, says that

$$\text{sys}^\alpha(M) \cdot \text{sys}_{\alpha^*}(M) \leq c \text{Vol}(M)^{1+\varepsilon}.$$

This estimate clearly implies Theorem 1.

Systolic inequalities in the continuous setting

Let M be a closed n -dimensional Riemannian manifold. The k -dimensional mod 2 systole is defined as the infimum of the k -volumes of all (piecewise smooth or Lipschitz) k -cycles representing non-trivial homology $H_k(M; \mathbb{Z}/2)$.

Following Freedman, we measure the *systolic freedom* of a sequence of manifolds M_i with volumes $V_i = \text{Vol}(M_i)$ and non-trivial H_k , as follows. We scale the Riemannian metric of each M_i , if needed, to make its geometry *locally bounded*: this means that the sectional curvatures are between -1 and 1 everywhere, and the injectivity radius is at

least 1. After that we measure the *systolic ratio*

$$\frac{\text{sys}_k(M_i) \cdot \text{sys}_{n-k}(M_i)}{V_i}$$

and express its growth as $i \rightarrow \infty$ as a function of V_i . The faster this function grows the more freedom the manifolds exhibit. A constant function corresponds to *systolic rigidity*. Freedman's examples had the growth rate of $\sqrt{\log(\cdot)}$. Our main result implies that for $k = 1$ the systolic ratio cannot grow faster than $(\cdot)^\varepsilon$, for arbitrarily small $\varepsilon > 0$.

Theorem 3. *For each $\varepsilon > 0$, there is a constant $c = c(n, \varepsilon)$ such that any closed Riemannian n -manifold M of bounded local geometry and with non-trivial $H_1(M; \mathbb{Z}/2)$ obeys the following systolic estimate:*

$$\text{sys}_1(M) \cdot \text{sys}_{n-1}(M) \leq c \text{Vol}(M)^{1+\varepsilon}.$$

To deduce Theorem 3 from Theorem 1, we triangulate M carefully, and replace the continuous systoles with the discrete ones. It is known [5–9, 20] that a Riemannian manifold M of bounded local geometry can be triangulated so that if we endow each simplex with the standard euclidean metric with edge-length 1, then the resulting metric on M is bi-Lipschitz to the original one (with the Lipschitz constant depending only on n). In particular, the number of n -simplices approximately equals $\text{Vol}(M)$ (maybe off by a dimensional factor). For a sequence of Riemannian manifolds of bounded local geometry, we have two ways to measure its systolic freedom: via the growth rate of the continuous or discrete systolic ratio. It turns out they have the same magnitude [10, Theorem 1.1.1]. The non-obvious part of this claim is that a systolic cycle in the continuous setting can be efficiently approximated by a discrete one in the triangulated manifold, without increasing its volume too much. This is done via a “Federer–Fleming” type of argument: inside each cell of the triangulation, the cycle is projected to the boundary radially from a random point; then this is repeated for cells of lower dimensions, until the cycle is approximated by a discrete one (see [10] for details). Therefore, Theorem 1 implies Theorem 3. We remark that manifolds of bounded local geometry, when triangulated as above, have an additional property of bounded “degree”, in the sense that every vertex of the triangulation is incident to a uniformly bounded number of top-dimensional simplices; this property is not needed in Theorem 1, making it somewhat stronger than the continuous version.

Theorem 3 can also be strengthened in a different direction: we stay in the continuous realm, but replace the condition of bounded local geometry by its “macroscopic” cousin as follows. Given $0 < v_1 < v_2$, let us say that a Riemannian manifold M has (v_1, v_2) -bounded macroscopic geometry if

- $\text{sys}_1(M) \geq 2$,
- every ball of radius $1/4$ has volume at least v_1 ,
- every ball of radius 1 has volume at most v_2 .

These three conditions should be viewed as the analogues of a lower bound on injectivity radius, upper bound on sectional curvature, and lower bound on sectional curvature, respectively.

Theorem 4. *For each $\varepsilon, v_1, v_2 > 0$, there is a constant $c = c(\varepsilon, v_1, v_2)$ such that any compact Riemannian n -manifold M of (v_1, v_2) -bounded macroscopic geometry exhibits systolic almost-rigidity:*

$$\text{sys}_1(M) \cdot \text{sys}_{n-1}(M) \leq c \text{Vol}(M)^{1+\varepsilon}.$$

This theorem is more general than Theorem 3. Indeed, on a manifold of bounded local geometry, the systolic bound $\text{sys}_1 \geq 2$ follows trivially from the bound on the injectivity radius, and volumetric estimates are mere corollaries of the Rauch comparison theorem (see, e.g., [4, Section 7.1.1]).

Remark 5. The referee asked us about the minimal set of conditions implying almost-rigidity. Namely, what assumptions can one impose on a class of manifolds to prohibit the following scenario: there exists $\varepsilon > 0$ and a sequence of manifolds M_i satisfying these assumptions, with $\text{Vol } M_i \rightarrow \infty$, and with $\text{sys}_1(M_i) \cdot \text{sys}_{n-1}(M_i) > c \text{Vol}(M_i)^{1+\varepsilon}$? We do not know if any of the three conditions of bounded macroscopic geometry in Theorem 4 can be omitted. We remark that the condition “every ball of radius $1/4$ has volume at least v_1 ” is not sufficient by itself. To see that, recall that Freedman constructed a sequence of three-dimensional manifolds M_g (with an integer parameter g , which can be arbitrarily large), which have bounded local geometry, $\text{diam}(M_g) \sim \log g$, $\text{sys}_1(M_g) \sim \sqrt{\log g}$, $\text{sys}_2(M_g) \sim g$, $\text{Vol}(M_i) \sim g$. They exhibit polylog systolic freedom. Scaling down the metric on M_g by a factor of $g^{1/3}/\log g$, we obtain a sequence of manifolds M'_g with volumes going to infinity, and diameters going to zero. So a lower bound on the volumes of balls of radius $1/4$ is present, but the sequence now has power law systolic freedom: $\text{sys}_1(M'_g) \cdot \text{sys}_2(M'_g) \sim \text{Vol}(M'_g)^{7/6}$.

Question 6. *Does it suffice to assume that $\text{sys}_1 \geq 2$ to guarantee systolic almost-rigidity? In other words, is there a sequence of manifolds M_i with $\text{sys}_1(M_i) \geq 2$ and $\text{Vol } M_i \rightarrow \infty$ such that $\text{sys}_1(M_i) \cdot \text{sys}_{n-1}(M_i) > c \text{Vol}(M_i)^{1+\varepsilon}$ (for some $\varepsilon > 0$)? If the bound on the 1-systole alone is not enough, does it suffice to assume both $\text{sys}_1 \geq 2$ and a lower bound on the volumes of balls of radius $1/4$?*

Structure of the paper

The proofs of discrete systolic estimates (mainly, of Theorem 2) are explained in Section 2. The core idea is the “Schoen–Yau minimal surface” approach, after [14, 18].

Some further generalizations are considered in Section 3. There we deal with the product of more than two systoles, and with a “macroscopic” systolic estimate, implying Theorem 4.

2. Proofs in the discrete case

Let M be an n -dimensional simplicial complex. The distance between vertices is measured as the edge-length of the shortest path between them in the 1-skeleton of M .

Definition 7. Let x be a vertex of M , and let r be a non-negative integer.

- (1) The *ball* $B(x, r)$ is the subcomplex of M defined as follows. If $r = 0$, $B(x, 0) = \{x\}$, so assume $r \geq 1$. An inclusion-maximal k -simplex, $0 \leq k \leq n$, is in $B(x, r)$ if all its vertices are distance at most r away from x and at least one of the vertices is at distance less than r from x . An arbitrary simplex is in $B(x, r)$ if it is a face of an inclusion-maximal simplex as in the previous sentence.
- (2) The *sphere* $S(x, r)$ is the subcomplex M consisting of all simplices of $B(x, r)$ with vertices at distance exactly r from x .

This definition might seem confusing, but it essentially defines spheres and balls as the level and sublevel sets of the distance function from x measured in the geometric realization of M with an auxiliary piecewise-Finsler metric F , described as follows. For each $k \leq n$, consider a k -simplex Δ^k in a k -dimensional linear space and endow it with the norm whose unit ball is the Minkowski sum $\Delta^k + (-\Delta^k)$. In particular, the distance between any disjoint faces of Δ^k is 1. Now we define F by saying that its restriction to any k -simplex of M is a flat Finsler metric isometric to Δ^k .

From this interpretation it is seen that every sphere actually separates the inside from the outside: every curve starting in $B(x, r)$ and ending outside $B(x, r)$ must intersect $S(x, r)$.

We write Vol for the number of n simplices in an n -complex, and Area for the number of $(n - 1)$ -simplices in an $(n - 1)$ -complex.

Everywhere below, α denotes a non-trivial class in $H^1(M; \mathbb{Z}/2)$, and H denotes an $(n - 1)$ -dimensional subcomplex of M .

Lemma 8 (Co-area inequality). $\text{Vol } B(x, r) \geq \sum_{i=1}^r \text{Area } S(x, i)$.

Proof. To each $(n - 1)$ -simplex in $S(x, i)$ we can assign an incident n -simplex in $B(x, i)$, and all those n -simplices are distinct. ■

Lemma 9 (Vitali cover). *Suppose we have a ball $B(x_i, 2\rho_i)$ at each vertex of an $(n - 1)$ -complex H inside M . Then one can select a subcollection of those balls that covers H but such that the corresponding balls $B(x_i, \rho_i)$ do not overlap (no n -simplex belongs to two balls from the subcollection).*

Proof. Greedily add balls $B(x_i, \rho_i)$ to the non-overlapping subcollection, in the non-decreasing order of radii. ■

Lemma 10 (Curve factoring). *If a 1-cycle ℓ lies in a ball of radius $r < \frac{\text{sys}^\alpha - 1}{2}$, then $\alpha([\ell]) = 0$.*

Proof. The cycle ℓ can be split as a sum of loops of length at most $2r + 1$, and none of those is detected by α . ■

Lemma 11 (Cut-and-paste trick). *Suppose H cuts α , $x \in H$, and $r < \frac{\text{sys}^\alpha - 1}{2}$. Then the complex $H' = [H \setminus (H \cap B(x, r))] \cup S(x, r)$ cuts α as well.*

Proof. Every connected component of a 1-cycle in the complement of H' either avoids H (and then it is not detected by α since H cuts α) or falls inside $B(x, r)$ (and then it is not detected by α by Lemma 10). ■

Proof of Theorem 2. Let $V = \text{Vol } M$, $R = \text{sys}^\alpha(M)$, and let H be a cutting subcomplex on which $\text{cut}^\alpha(M)$ is attained.

Consider all singular (that is, not necessarily simplicial) loops detected by α . There are two cases.

- (1) Every such loop meets the $(n-2)$ -skeleton of M . Then H can be taken to be this skeleton, $\text{Area } H = 0$, and there is nothing to prove.
- (2) There is a loop avoiding the $(n-2)$ -skeleton of M . Deform it to the 1-skeleton of M . Each of the edges of the deformed loop is incident to an n -dimensional simplex of M . If these n -simplices are not all distinct, we can replace the loop with a shorter one (still detected by α). Once we make it as short as possible, all incident n -simplices are distinct, and we conclude that $R \leq V$. We use this inequality in what follows.

The second inequality we observe is $\text{Area } H \leq (n+1)V$. Indeed, consider an $(n-1)$ -simplex in H . If it is not incident to n -simplices in M , we can safely remove it from H (but leave its boundary) preserving the cutting property of H . Therefore, to every $(n-1)$ -simplex of H we can assign an incident n -simplex of M , and it follows that $\text{Area } H \leq (n+1)V$.

If $\text{Area } H \geq R^{1/\varepsilon}$, then

$$R \cdot \text{Area } H \leq (\text{Area } H)^{1+\varepsilon} \leq (n+1)^{1+\varepsilon} V^{1+\varepsilon},$$

so in the rest of the argument we assume that $\text{Area } H < R^{1/\varepsilon}$.

For brevity, we write $v(x, r) = \text{Vol } B(x, r)$ and $h(x, r) = \text{Area}(H \cap B(x, r))$.

A ball centered at $x \in H$ of positive radius divisible by 2 will be called *good* if

$$h(x, r) \leq \frac{4}{R^{1-\varepsilon}} v(x, r/2).$$

A sufficient condition for a ball of radius $0 < r \leq R$ divisible by 4 to be good is

$$h(x, r) \leq \frac{r}{R^{1-\varepsilon}} h(x, r/4).$$

Indeed, by the co-area inequality (Lemma 8) one has

$$\begin{aligned} h(x, r) &\leq \frac{r}{R^{1-\varepsilon}} \frac{1}{r/2 - r/4} \sum_{t=r/4}^{r/2-1} h(x, t) \\ &\leq \frac{r}{R^{1-\varepsilon}} \frac{4}{r} \sum_{t=r/4}^{r/2-1} \text{Area } S(x, t) \\ &\leq \frac{4}{R^{1-\varepsilon}} v(x, r/2). \end{aligned}$$

Here we used the minimality of H , which implies $h(x, t) \leq \text{Area } S(x, t)$, since otherwise we can apply the cut-and-paste trick (Lemma 11), decreasing the area of H and preserving its cutting property.

We will prove below that for every $x \in H$ there is a good ball centered at x . Having a cover of H by good balls, we pick a Vitali subcover $\bigcup B(x_i, 2r_i) \supset H$ such that the balls $B(x_i, r_i)$ do not overlap (by Lemma 9). This will conclude the proof:

$$R \cdot \text{Area } H \leq R \sum h(x_i, 2r_i) \leq 4R^\varepsilon \sum v(x, r_i) \leq 4V^{1+\varepsilon}.$$

Now we show that around every $x \in H$ one can place a good ball. In the range $[R^{1-\varepsilon}, R]$, pick a longest sequence of integers of the form $r_0 = 4^m, r_1 = 4^{m+1}, \dots, r_N = 4^{m+N}$. We assume it has at least two elements ($R^\varepsilon \geq 4^3$), since for small R the statement of the theorem holds vacuously. Observe that $4^N = r_N/r_0 > R^\varepsilon/16$. If for some $1 \leq i \leq N$ the ball $B(x, r_i)$ obeys the estimate

$$h(x, r_i) \leq \frac{r_i}{R^{1-\varepsilon}} h(x, r_i/4),$$

then we have found a good ball, so assume these estimates all fail. If $h(x, r_0) = 0$, then $B(x, r_0)$ is good. If $h(x, r_0) \geq 1$, we have the following:

$$\begin{aligned} R^{1/\varepsilon} &> \text{Area } H \geq h(x, r_N) \\ &> 4^N h(x, r_{N-1}) \\ &> 4^N 4^{N-1} h(x, r_{N-2}) \\ &> \dots \\ &> 4^N 4^{N-1} \dots 4^1 h(x, r_0) \\ &\geq (4^N)^{(N+1)/2} \\ &> \left(\frac{R^\varepsilon}{16}\right)^{(\varepsilon \log_4 R - 1)/2}. \end{aligned}$$

This inequality fails for all large $R \geq R_0(\varepsilon)$, but for all $R \leq R_0(\varepsilon)$ the statement of the theorem holds trivially. $\blacksquare$

Remark 12. A careful analysis of this proof shows that the growth rate $V^{1+\varepsilon}$ can be replaced by $V \exp(2(\log V)^{1/\sqrt{2}})$. This function grows slower than any power $V^{1+\varepsilon}$. Unfortunately, this proof cannot be refined to give the growth rate $V \text{ polylog } V$, as is tempting to conjecture.

In fact, Theorem 2 holds (with the same proof) for any coefficient ring A , not just $\mathbb{Z}/2$, if we define the systole $\text{sys}^\alpha(M)$ as the shortest edge-length of a loop detected by $\alpha \in H^1(M; A)$, and define $\text{cut}^\alpha(M)$ in the same way as with $\mathbb{Z}/2$ coefficients. For example, if $A = \mathbb{Z}/p$, then sys^α might capture some loops not detected by the $\mathbb{Z}/2$ -cohomology, but in return, the cutting complexes from the definition of cut^α will have to cut those loops as well. However, we need to assume $A = \mathbb{Z}/2$ in order to relate cut^α to the systolic invariants (and finish the proof of Theorem 1). We do it for general k , as we will need it in the next section.

Lemma 13 (Minimality forces regularity). *Let $\alpha \in H^k(M; \mathbb{Z}/2)$ be a non-trivial cohomology class, and let H be an $(n - k)$ -dimensional subcomplex in M on which $\text{cut}^\alpha(M)$ is attained. If M is a manifold, then the top-dimensional simplices of H form an $(n - k)$ -cycle representing α^* . Hence, $\text{cut}^\alpha(M) = \text{sys}_{\alpha^*}(M)$.*

Proof. It suffices to show that H contains an $(n - k)$ -cycle representing α^* ; if so, then this cycle is among the subcomplexes defining cut^α , and so by minimality it is equal to H . Let $i: H \hookrightarrow M$ denote the inclusion, and suppose for the sake of contradiction that $i_*(H_{n-k}(H))$ does not contain α^* . Then there is some hyperplane of $H_{n-k}(M)$ that contains $i_*(H_{n-k}(H))$ but not α^* . By the universal coefficient theorem, this hyperplane can be expressed as pairing with some $\beta \in H^{n-k}(M)$ with $\beta(\alpha^*) \neq 0$. In this case $\beta|_H = 0$ while $\alpha \smile \beta \neq 0$, contradicting the Lusternik–Schnirelmann lemma because we have assumed that $\alpha|_{M \setminus H} = 0$. ■

Note that the same proof shows that for any field coefficients, there is a cycle representing α^* with support H . However, unless the coefficients are $\mathbb{Z}/2$ we do not control the coefficient of each simplex in H . Only with $\mathbb{Z}/2$ coefficients does it make sense to say that H itself forms a cycle achieving $\text{sys}_{\alpha^*}(M)$. This completes the proof of Theorem 1.

Example 14. We construct a triangulated manifold M of dimension n with a cohomology class $\alpha \in H^1(M; \mathbb{Z}/2)$ such that $\text{sys}^\beta(M) < \text{cut}^\alpha(M)$ for any class $\beta \in H^{n-1}(M; \mathbb{Z}/2)$ with $\alpha \smile \beta \neq 0$. Let M_1 and M_2 be large identical triangulated copies of $S^1 \times S^{n-1}$. We form M as the connected sum of M_1 and M_2 by removing an n -simplex from each and identifying their boundaries. We assume $n \geq 3$, although the argument is easily adapted for $n = 2$. For $n \geq 3$, the Mayer–Vietoris sequence implies

$$H_{n-1}(M; \mathbb{Z}/2) \cong H_{n-1}(M_1; \mathbb{Z}/2) \oplus H_{n-1}(M_2; \mathbb{Z}/2) \cong \mathbb{Z}_2 \times \mathbb{Z}_2.$$

We select $\alpha \in H^1(M; \mathbb{Z}/2)$ such that α^* is the class $(1, 1)$ in $H_{n-1}(M; \mathbb{Z}_2)$. Given any cycle representing α^* , we may split it into an M_1 part and an M_2 part and add faces from the common boundary of the n -simplex to each part, to obtain a non-trivial cycle on each of M_1 and M_2 . Thus we have

$$\text{cut}^\alpha(M) = \text{sys}_{\alpha^*}(M) \geq 2 \text{sys}_{n-1}(M_1) - 2(n + 1).$$

On the other hand, for any $\beta \in H^{n-1}$ such that β pairs non-trivially with $\alpha^* = (1, 1)$, either β pairs non-trivially with $(1, 0)$ or β pairs non-trivially with $(0, 1)$. Each of $(1, 0)$ and $(0, 1)$ in $H_{n-1}(M; \mathbb{Z}/2)$ can be represented by a cycle of size $\text{sys}_{n-1}(M_1)$, so we have

$$\text{sys}^\beta(M) \leq \text{sys}_{n-1}(M_1) < \text{cut}^\alpha(M).$$

3. Generalizations

The following result follows from Theorem 2 by induction.

Theorem 15. *For each $\varepsilon > 0$, there is a constant $C = C(n, \varepsilon)$ such that the following holds. Let M be a simplicial n -complex with cohomology classes $\alpha_1, \dots, \alpha_k \in H^1(M; \mathbb{Z}/2)$ whose cup-product is non-zero. Then*

$$\text{sys}^{\alpha_1}(M) \cdot \dots \cdot \text{sys}^{\alpha_k}(M) \cdot \text{cut}^{\alpha_1 \smile \dots \smile \alpha_k}(M) \leq C \text{Vol}(M)^{1+\varepsilon}.$$

We remark that Lemma 13 applies and gives, in the case when M is a manifold,

$$\text{cut}^{\alpha_1 \smile \dots \smile \alpha_k}(M) = \text{sys}_{(\alpha_1 \smile \dots \smile \alpha_k)^*}(M).$$

Results very close to Theorem 15 for $k = n - 1$ have been known even with $\varepsilon = 0$ (cf. [14, Theorem 3], [3, Theorem 1], [1, Theorem 1.18]). For $k \leq n - 2$ one cannot plug $\varepsilon = 0$.

Proof of Theorem 15. Assume $\varepsilon < 1$. Take H to be a minimal $(n - 1)$ -subcomplex of M such that α_1 vanishes when restricted to $M \setminus H$. The proof in Section 2 implies that

$$\text{sys}^{\alpha_1}(M) \cdot \text{Area}(H) \leq c(n, \varepsilon/3) \text{Vol}(M)^{1+\varepsilon/3}.$$

The Lusternik–Schnirelmann lemma tells us that the product of the remaining α_i is non-trivial on H ; hence, the induction assumption applies for H in place of M . The systoles in H dominate the corresponding systoles in M , since the intrinsic distances in H dominate the extrinsic ones: $\text{sys}^{\alpha_i}(M) \leq \text{sys}^{\alpha_i|_H}(H)$, $2 \leq i \leq k$. Any $(n - k)$ -dimensional subcomplex of H that cuts $\alpha_2 \smile \dots \smile \alpha_k|_H$ also cuts $\alpha_1 \smile \dots \smile \alpha_k$ in M , as one can deduce from the Lusternik–Schnirelmann lemma, so $\text{cut}^{\alpha_1 \smile \dots \smile \alpha_k}(M) \leq \text{cut}^{\alpha_2 \smile \dots \smile \alpha_k|_H}(H)$. Assembling this all together, we get

$$\begin{aligned} & \text{sys}^{\alpha_1}(M) \cdot \text{sys}^{\alpha_2}(M) \cdot \dots \cdot \text{sys}^{\alpha_k}(M) \cdot \text{cut}^{\alpha_1 \smile \dots \smile \alpha_k}(M) \\ & \leq \text{sys}^{\alpha_1}(M) \cdot \text{sys}^{\alpha_2|_H}(H) \cdot \dots \cdot \text{sys}^{\alpha_k|_H}(H) \cdot \text{cut}^{\alpha_2 \smile \dots \smile \alpha_k|_H}(H) \\ & \leq \text{sys}^{\alpha_1}(M) \cdot C(n - 1, \varepsilon/3) \text{Area}(H)^{1+\varepsilon/3} \\ & \leq C(n - 1, \varepsilon/3)^{1+\varepsilon/3} (\text{sys}^{\alpha_1}(M) \cdot \text{Area}(H))^{1+\varepsilon/3} \\ & \leq C(n - 1, \varepsilon/3)^{1+\varepsilon/3} (c(n, \varepsilon/3) \text{Vol}(M)^{1+\varepsilon/3})^{1+\varepsilon/3} \\ & \leq C(n, \varepsilon) \text{Vol}(M)^{1+\varepsilon}. \end{aligned}$$

■

Another version of the same theorem holds in the continuous setting, with the assumption of locally bounded geometry replaced by its macroscopic cousin. We work in the setting of *Riemannian polyhedra*, that is, finite pure¹ simplicial complexes whose top-dimensional faces are endowed with Riemannian metrics matching on the common faces. We say that a subspace $H \subset M$ is a k -dimensional *subpolyhedron* if it admits the structure of a finite pure k -dimensional simplicial complex whose cells are embedded smoothly in the cells of M . The piecewise Riemannian metric is inherited from the ambient

¹In a *pure* simplicial n -complex, a simplex of any dimension is contained in an n -dimensional simplex.

polyhedron, allowing one to measure intrinsic distances and volumes. Any subpolyhedron H admits a thin tubular neighborhood that deformation retracts onto H , and the Lusternik–Schnirelmann lemma for cohomology is applicable in the following form: if $\alpha|_{M \setminus H} = 0$, $\beta|_H = 0$, then $\alpha \smile \beta$ vanishes as a cohomology class of M . The systole $\text{sys}^\alpha(M)$ is the infimum of the k -volumes of all (piecewise smooth or Lipschitz) k -cycles detected by $\alpha \in H_k(M; \mathbb{Z}/2)$. The cutting area $\text{cut}^\alpha(M)$ is the infimum of the k -volumes of all those k -dimensional subpolyhedra $H \subset M$ that cut α in the sense that α vanishes on $M \setminus H$.

Theorem 16. *For each $\varepsilon, v_1, v_2 > 0$, there is a constant $C = C(\varepsilon, v_1, v_2)$ such that the following holds. Let M be a compact Riemannian n -polyhedron such that $\text{sys}_1(M) \geq 2$, every ball of radius $1/4$ has volume at least v_1 , and every ball of radius 1 has volume at most v_2 . Let $\alpha_1, \dots, \alpha_k \in H^1(M; \mathbb{Z}/2)$ be cohomology classes with non-zero cup-product. Then*

$$\text{sys}^{\alpha_1}(M) \cdot \dots \cdot \text{sys}^{\alpha_k}(M) \cdot \text{cut}^{\alpha_1 \smile \dots \smile \alpha_k}(M) \leq C \text{Vol}(M)^{1+\varepsilon}.$$

Note that this theorem implies Theorem 4.

Proof of Theorem 16. It suffices to prove this for $k = 1$ and then induct. In the case $k = 1$, let $\alpha = \alpha_1$, $R = \text{sys}^\alpha(M)$, $V = \text{Vol } M$, and proceed along the same scheme as in the proof of Theorem 2. Trivially, $2v_1 \lfloor R \rfloor \leq V$ (think of a necklace made of beads of radius $1/4$ along the systole), so $R \lesssim V$, where $\lesssim$ indicates an inequality that holds up to a factor depending on ε, v_1, v_2 . Let A be the infimum of the areas of $(n-1)$ -subpolyhedra such that α vanishes on their complement. We need to show that $R \cdot A \lesssim V^{1+\varepsilon}$. If $A \geq R^{1/\varepsilon}$, then, as before,

$$R \cdot A \leq A^{1+\varepsilon} \lesssim V^{1+\varepsilon},$$

but the last inequality requires an explanation. Pick a set of disjoint balls of radius $1/4$ such that the concentric balls of radius $3/4$ cover M , and for each of these find a concentric sphere of radius between $3/4$ and 1 that has area at most $4v_2$ (by the co-area inequality). The union of these spheres has area at most $4v_2 V/v_1$, and α vanishes on the complement of this union by the curve factoring lemma (here we use the assumption $\text{sys}_1(M) \geq 2$). Therefore, $A \leq 4v_2 V/v_1$.

From now on assume that $A < R^{1/\varepsilon}$. Let H be an almost-minimizing subpolyhedron for which α vanishes on its complement, with $A \leq \text{Area } H < A + \delta$, where $\delta = v_1/R$.

A ball centered at $x \in H$ of radius $r \in [1, R]$ will be called *good* if

$$h(x, r) \leq \frac{13}{R^{1-\varepsilon}} v(x, r/3),$$

where we use again the notation $v(x, r) = \text{Vol } B(x, r)$ and $h(x, r) = \text{Area}(H \cap B(x, r))$.

It suffices to show that every $x \in H$ is the center of a good ball. Once that is done, we pick a Vitali cover of H by good balls $B(x_i, r_i)$ such that the balls $B(x_i, r_i/3)$ do not intersect, and conclude just as in Section 2:

$$R \cdot A \leq R \sum h(x_i, r_i) \leq 13R^\varepsilon \sum v(x_i, r_i/3) \leq 13R^\varepsilon V \lesssim V^{1+\varepsilon}.$$

We start by looking at balls $B(x, R)$, $B(x, R/4)$, $\dots$, $B(x, R/4^N)$, where $N = \lfloor \varepsilon \log_4 R \rfloor$ is the integer defined by $R/4^{N+1} < R^{1-\varepsilon} \leq R/4^N$, or equivalently $4^N \leq R^\varepsilon < 4^{N+1}$ (note that $N \geq 0$ since $R > 1$). If at least one of these balls obeys the inequality

$$h(x, r) \leq \frac{r}{R^{1-\varepsilon}} h(x, r/4),$$

then it is good, since

$$\begin{aligned} \frac{r}{R^{1-\varepsilon}} h(x, r/4) &\leq \frac{12}{R^{1-\varepsilon}} \int_{r/4}^{r/3} (\text{Area } S(x, t) + \delta) dt \\ &\leq \frac{12}{R^{1-\varepsilon}} v(x, r/3) + R^\varepsilon \delta \leq \frac{13}{R^{1-\varepsilon}} v(x, r/3). \end{aligned}$$

If none of them are good, then for each $0 \leq m \leq N-1$ we have

$$h(x, R/4^m) > \frac{R/4^m}{R^{1-\varepsilon}} h(x, R/4^{m+1}) \geq 4^{N-m} h(x, R/4^{m+1}),$$

and assembling these,

$$\begin{aligned} R^{1/\varepsilon} &> A \geq h(x, R) \\ &> 4^N h(x, R/4) \\ &> \dots \\ &> 4^N \dots 4^1 h(x, R/4^N) \\ &\geq (4^{N+1})^{N/2} h(x, R^{1-\varepsilon}) \\ &> R^{\varepsilon(\varepsilon \log_4 R - 1)/2} h(x, 1). \end{aligned}$$

The function $R^{\varepsilon(\varepsilon \log_4 R - 1)/2}$ grows (in R) faster than $R^{1-\varepsilon+1/\varepsilon}$, so for $R \geq R_0(\varepsilon, v_1)$ large enough we have

$$h(x, 1) < R^{1/\varepsilon - \varepsilon(\varepsilon \log_4 R - 1)/2} < \frac{13v_1}{R^{1-\varepsilon}} \leq \frac{13v(x, 1/3)}{R^{1-\varepsilon}},$$

proving that the ball $B(x, 1)$ is good. For $R < R_0(\varepsilon, v_1)$ the statement of the theorem follows from the above-mentioned observation $A \lesssim V$. $\blacksquare$

Acknowledgments. We thank the anonymous referee whose suggestions motivated us to highlight the version of the main theorem under “macroscopic” assumptions.

Funding. This material is based in part upon work supported by the National Science Foundation under Grant No. DMS-1926686, and by the HORIZON EUROPE Research and Innovation programme under Marie Skłodowska-Curie grant agreement No. 101107896. Hannah Alpert is supported in part by the Simons Foundation (Gift Number 965348). Larry Guth is supported by Simons Investigator Award.

References

- [1] Avvakumov, S., Balitskiy, A., Hubard, A., Karasev, R.: [Systolic inequalities for the number of vertices](#). J. Topol. Anal. (online, 2023)

- [2] Babenko, I., Katz, M.: *Systolic freedom of orientable manifolds*. Ann. Sci. École Norm. Sup. (4) **31**, 787–809 (1998) Zbl 0944.53019 MR 1664222
- [3] Balacheff, F., Karam, S.: *Length product of homologically independent loops for tori*. J. Topol. Anal. **8**, 497–500 (2016) Zbl 1342.53060 MR 3509570
- [4] Berger, M.: *A panoramic view of Riemannian geometry*. Springer, Berlin (2003) Zbl 1038.53002 MR 2002701
- [5] Boissonnat, J.-D., Dyer, R., Ghosh, A.: *Delaunay triangulation of manifolds*. Found. Comput. Math. **18**, 399–431 (2018) Zbl 1395.57032 MR 3777784
- [6] Boissonnat, J.-D., Dyer, R., Ghosh, A., Wintraecken, M.: *Local criteria for triangulating general manifolds*. Discrete Comput. Geom. **69**, 156–191 (2023) Zbl 1521.57023 MR 4527538
- [7] Bowditch, B. H.: *Bilipschitz triangulations of riemannian manifolds*. Preprint, <http://homepages.warwick.ac.uk/~masgak/papers/triangulations.pdf> (2020), visited on June 12, 2024
- [8] Cairns, S. S.: *On the triangulation of regular loci*. Ann. of Math. (2) **35**, 579–587 (1934) Zbl 0012.03605 MR 1503181
- [9] Dyer, R., Vegter, G., Wintraecken, M.: *Riemannian simplices and triangulations*. Geom. Dedicata **179**, 91–138 (2015) Zbl 1333.57036 MR 3424659
- [10] Freedman, M., Hastings, M.: *Building manifolds from quantum codes*. Geom. Funct. Anal. **31**, 855–894 (2021) Zbl 1485.53115 MR 4317505
- [11] Freedman, M. H.: *Z_2 -systolic-freedom*. In: Proceedings of the Kirbyfest (Berkeley, CA, 1998), Geom. Topol. Monogr. 2, Geometry & Topology Publications, Coventry, 113–123 (1999) Zbl 1024.53029 MR 1734404
- [12] Freedman, M. H., Meyer, D. A., Luo, F.: *Z_2 -systolic freedom and quantum codes*. In: Mathematics of quantum computation, Comput. Math. Ser., Chapman & Hall/CRC, Boca Raton, FL, 287–320 (2002) Zbl 1075.81508 MR 2007952
- [13] Gromov, M.: *Systoles and intersystolic inequalities*. In: Actes de la Table Ronde de Géométrie Différentielle (Luminy, 1992), Sémin. Congr. 1, Société Mathématique de France, Paris, 291–362 (1996) MR 1427763
- [14] Guth, L.: *Systolic inequalities and minimal hypersurfaces*. Geom. Funct. Anal. **19**, 1688–1692 (2010) Zbl 1185.53045 MR 2594618
- [15] Hastings, M. B., Haah, J., O’Donnell, R.: *Fiber bundle codes: breaking the $N^{1/2}\text{polylog}(N)$ barrier for quantum LDPC codes*. In: STOC ’21—Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 1276–1288 (2021) Zbl 07765248 MR 4398919
- [16] Katz, M. G., Suciu, A. I.: *Volume of Riemannian manifolds, geometric inequalities, and homotopy theory*. In: Tel Aviv Topology Conference: Rothenberg Festschrift (1998), Contemp. Math. 231, American Mathematical Society, Providence, RI, 113–136 (1999) Zbl 0967.53024 MR 1705579
- [17] Panteleev, P., Kalachev, G.: *Quantum LDPC codes with almost linear minimum distance*. IEEE Trans. Inform. Theory **68**, 213–229 (2022) Zbl 1489.94162 MR 4395409
- [18] Papasoglu, P.: *Uryson width and volume*. Geom. Funct. Anal. **30**, 574–587 (2020) Zbl 1453.53050 MR 4108616
- [19] Pittet, C.: *Systoles on $S^1 \times S^n$* . Differential Geom. Appl. **7**, 139–142 (1997) Zbl 0878.53038 MR 1454720
- [20] Whitehead, J. H. C.: *On C^1 -complexes*. Ann. of Math. (2) **41**, 809–824 (1940) Zbl 0025.09203 MR 0002545