*The Threshold for d-Collapsibility in Random Complexes**

L. Aronshtam, N. Linial
Department of Computer Science, Hebrew University, Jerusalem 91904, Israel; e-mail: liobsar@gmail.com; nati@cs.huji.ac.il

Received 10 July 2013; revised 19 September 2014; accepted 4 December 2014
Published online 27 March 2015 in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/rsa.20585

ABSTRACT: In this paper we determine the threshold for d-collapsibility in the probabilistic model $X_d(n,p)$ of d-dimensional simplicial complexes. A lower bound for this threshold $p = \frac{\eta_d}{n}$ was established in (Aronshtam and Linial, Random Struct. Algorithms 46 (2015) 26–35), and here we show that this is indeed the correct threshold. Namely, for every $c > \eta_d$, a complex drawn from $X_d(n, \frac{c}{n})$ is asymptotically almost surely not d-collapsible. © 2015 Wiley Periodicals, Inc. Random Struct. Alg., 48, 260–269, 2016

Keywords: topology; combinatory; simplicial complex; collapsibility

1. INTRODUCTION

This paper is part of an ongoing research effort to develop a combinatorial perspective of basic topological objects. One of the main objectives of this endeavour is to explore the power of the probabilistic method in topology. We work within the model $X_d(n,p)$ of random n-vertex d-dimensional simplicial complexes that was introduced in [6]. Recall that every complex X in $X_d(n,p)$ has a full $(d-1)$-dimensional skeleton, and each d-dimensional face is placed in X independently and with probability p. Note that $X_1(n,p)$ coincides with the Erdős-Rényi random graph model $G(n,p)$. Indeed many of the motivating questions in the study of $X_d(n,p)$ concern the similarities and differences between the one and higher-dimensional situations.

Perhaps the two most dramatic phenomena in the Erdős-Rényi's model occur at $p = \frac{1}{n}$ and at $p = \frac{\log n}{n}$. The latter is the threshold for graph connectivity (and more). Indeed [6] and subsequent work [5, 8] have resolved most (but still not all) of the problems related to

---

Correspondence to: L. Aronshtam
Supported by ISF (to N.L.); ERC (to N.L.).

the vanishing of the $(d-1)$-st homology of $X \in X_d(n,p)$. Major progress towards understanding simple connectivity of random 2-complexes was made in [3], for more references see [4].

A great deal of past research was dedicated to the so-called *phase transition* at $p = \frac{1}{n}$. Most famously, this is when the *giant component* emerges. We emphasize, however, another transition that occurs at $p = \frac{1}{n}$, namely, the almost sure (a.s.) emergence of a cycle. To wit, if $p = o(\frac{1}{n})$ then G is a.s. a forest. If $p = \frac{c}{n}$ with $1 > c > 0$, then $\Pr(G \text{ is a forest}) = p + o_n(1)$ for some $1 > p = p(c) > 0$. Finally, for $p \geq \frac{1}{n}$, there is a.s. at least one cycle in G. But at this junction the one-dimensional and the high-dimensional story part ways, since for graphs acyclicity and collapsibility are equivalent. While collapsibility implies acyclicity in all dimensions, the reverse implication no longer holds for $d \geq 2$. Here are the relevant definitions. A $(d-1)$-face τ in a d-complex X is *free* if it is contained in exactly one d-dimensional face σ of X. In the corresponding *elementary collapse* step, τ and σ are removed from X. We say that X is *d-collapsible* if it is possible to eliminate all its d-faces by a series of elementary collapses. Clearly a graph (i.e., a 1-dimensional complex) is 1-collapsible iff it is acyclic, i.e., a forest.

This naturally suggests two questions in the study of the $X_d(n,p)$ analog of the $p = \frac{1}{n}$ phase transition: (i) What is the threshold for d-collapsibility? and (ii) What is the threshold for the non-vanishing of the d-th homology? On question (ii) we only mention that the best known *upper bound* on this threshold was found in [1]. We suspect that this bound is the actual threshold for acyclicity, and this remains a subject of current study. The objective of this paper is to answer Question (i). Lower bounds on the threshold for d-collapsibility were previously found in [2,4]. Here we prove an upper bound that coincides with the lower bound from [2], thus establishing the threshold for d-collapsibility.

Note: By an abuse of language we occasionally say *collapsibility* where d-collapsibility is meant. We hope that this does not create too much confusion.

1.1. The Main Result and an Overview of the Proof

For fixed $\eta > 0$, let us consider the following equation in x:

$$0 = e^{(1-x)^d(-\eta)} - x. \tag{1}$$

It is not hard to verify[1] that for $\eta > 0$ small, the only root is $x = 1$. As η grows, there emerges some root $x < 1$. We denote by η_d the infimum of all $\eta > 0$ for which Eq. (1) has a root $1 > x > 0$ (notice in [2] we used the notation γ_d and not η_d). The following was proved in [2, Theorem 1.4]:

Theorem 1.1. *Fix $c < \eta_d$. A d-dimensional complex $X \in X_d(n, \frac{c}{n})$, is a.s.$d$-collapsible or it contains a copy of $\partial\Delta_{d+1}$.*

The present article shows that this lower bound is tight. Namely:

Theorem 1.2. *For every $c > \eta_d$, a d-dimensional complex $X \in X_d(n, \frac{c}{n})$, is a.s.not d-collapsible.*

[1]The reader may find it more convenient to verify this in the following equivalent form of the equation $(1-x)^d + \frac{\ln x}{\eta} = 0$.

We are inspired by work on *cores* in random graphs and in particular by Molloy's work [9]. (Similar ideas can be found in the literature at least as early as [10]). The strategy of the proof is to split the collapsing process into two *epochs* as follows. Fix a positive integer r (to be specified below) and start by carrying out r *phases* of collapses. Note that the random variables D_k, D'_k, A_k, that we later define, depend on r, although we suppress this fact in our notations. In each phase we simultaneously collapse *all* $(d-1)$-faces that are presently free. After r such phases we move on to the second epoch of the process in which we collapse free $(d-1)$-faces one by one in some order. Notice that there is no real difference between the two epochs, and the second epoch can be viewed as a slow-motion version of the first epoch. This slowed-down process allows us, however, to go into more refined analysis of the process, and this turns out to be crucial for the proof. In [1] we carried out only the first epoch. However, it is hard to analyze more than a constant number of phases of the first epoch, and the slowed-down second epoch makes it possible for us to analyze the entire collapsing process.

The order of collapsing in the second epoch is as follows: (i) Choose a set W of currently free $(d-1)$ faces that is maximal w.r.t. the condition that no d-face contains more than one $(d-1)$-face in W, and mark the faces in W. (ii) Pick a uniformly chosen random ordering of the marked faces, and collapse them one by one in this order. (iii) Repeat stages (i), (ii) until no free $(d-1)$-faces are left. A face that gets marked during the k-th run through steps (i), (ii) is said to have *mark* k. When we speak of the i-th step of the second epoch, we refer to the i-th elementary collapse, carried out in the second epoch, regardless of the run. As explained in the next paragraph, the proof is based on a detailed analysis of the resulting random process of collapses.

If ψ be a face in a finite simplicial complex Z, we define the *degree* of ψ in Z as the number of faces of dimension $\dim(\psi)+1$ in Z that contain ψ. For a $(d-1)$-face τ, we denote by $\Delta(\tau)$ the degree of τ at the end of the first epoch (i.e., the number of d-faces that contain τ at this point). For a given d-face σ we let $\Delta^\sigma(\tau) := \Delta(\tau) - 1$ (resp. $\Delta^\sigma(\tau) := \Delta(\tau)$) if $\tau \subset \sigma$ (resp. $\tau \not\subset \sigma$). The random variable X_0 counts the number of $(d-1)$-faces τ that are free, i.e., $\Delta(\tau) = 1$ at the end of the first epoch. Moreover, we define $L = L_r$ as the number of $(d-1)$-faces τ for which $\Delta(\tau) > 0$ (In the notation of [1], this number equals $f_{d-1}(R_r(X)) - \zeta_r(X)$.) We denote by X_i the number of free $(d-1)$-faces at step i of the second epoch. Our plan is to show that the expected drop in this number $\mathbb{E}(X_i - X_{i+1})$ is sufficiently large to a.s. guarantee that at some moment the complex still has some $(d-1)$-faces, but none of them are free. This clearly means that the original complex is not d-collapsible.

Our starting point is that if the integer r is large enough, then the results from [1] give us a good idea about several of the relevant random variables. The main new ingredient in our proof is a detailed accounting of the free $(d-1)$-faces. In every step this number gets decreased by one due to the loss of the free $(d-1)$-face that we are presently collapsing. What complicates matters is that as we carry out the collapse step, we may be creating some new free $(d-1)$-faces. We define the event S_i^j that our i-th collapsing step creates j new free $(d-1)$-faces. Note that $0 \le j \le d$, since the collapsed d-face contains $d+1$ distinct $(d-1)$-faces, and at least one of them was free prior to the collapsing step. In particular $X_i \le X_{i-1} + j - 1$. We further introduce the random variable Y_i that counts the number of new free $(d-1)$-faces added in the i-th collapsing step.

Clearly

$$X_i \le X_{i-1} - 1 + Y_i \le X_0 - i + \sum_{l=1}^{i} Y_l \tag{2}$$

To recap: when $X_i = 0$, we run out of free faces, and if this happens before all $(d-1)$-faces get eliminated, the original complex is not d-collapsible.

We need the following definition: Two $(d-1)$-faces in a d-complex are considered *neighbors* if they are both contained in some d-face of the complex. Construct a graph whose vertices are all the currently free $(d-1)$-faces in the complex, with this adjacency relation.

We now recall a basic technique from Section 2 of [1]. A *rooted d-tree* is defined recursively as follows : (i) A single $(d-1)$-face is a rooted d-tree, its root being that single $(d-1)$-face. (ii) A rooted d-tree with N distinct d-faces is constructed by taking T, a rooted d-tree with $N-1$ distinct d-faces, choosing a $(d-1)$-face τ in T, and a new vertex u, and adding the d-face $\tau \cup u$ and all its subfaces to T. The root of the new tree remains unchanged. With the same concept of adjacency between $(d-1)$-faces, we can use the notion of graph metric and talk about the distance between pairs of $(d-1)$-faces in a rooted d-tree. We consider the probability space, $\mathcal{T}(c,t)$, of rooted d-trees of radius at most t (i.e., every $(d-1)$-face in the tree is at distance at most t from the root). The probability space $\mathcal{T}(c,0)$ clearly contains only the d-tree that is just the root. A tree is sampled from $\mathcal{T}(c,t)$ by (i) Sample a tree T from $\mathcal{T}(c,t-1)$. (ii) For every $(d-1)$-face, τ in T, at distance $t-1$ from the root, create l new vertices $v_1,\ldots,v_l$ and add the d-faces $\{\tau \cup v_1,\ldots,\tau \cup v_l\}$ and their subfaces to T. The integer l is sampled from the Poisson distribution with parameter c. We are interested in τ-*collapsing* of rooted d-trees. These are processes where we collapse faces of the d-tree but not the root τ.

Let $t \geq 0$ be a fixed integer. As shown in Section 3 of [1], for every $(d-1)$-face τ in $X \in X_d(n, \frac{c}{n})$ the t-th neighborhood of τ is a.s. a d-tree. Moreover, the distributions of such d-trees is very close to the distribution of $\mathcal{T}(c,t)$ of d-trees rooted at τ. This simplifies the analysis of the τ-collapse process on a local scale. We denote by $\gamma_t(d,c)$ (resp. $\beta_t(d,c)$) the probability that τ becomes isolated[2] in fewer than (resp. more than) t collapse phases in $\mathcal{T}(c,t+1)$. It is interesting to notice that this event can happen even if the tree, rooted in τ, has radius much larger than t. Obviously $\beta_t(d,c) = 1-\gamma_{t+1}(d,c)$. The following relations are proved in [1, Eq.(4)]:

$$\gamma_0(d,c) = 0, \quad \gamma_{t+1}(d,c) = \exp(-c(1-\gamma_t(d,c))^d) \ \text{ for } \ t = 0,1,\ldots.$$

To simplify notations, and as long as everything is clear, we suppress the dependence on c and d and write γ_t, β_t rather than $\gamma_t(d,c), \beta_t(d,c)$.

We now use this model and consider a normal run of the first epoch with the change that we forbid collapsing via τ (i.e. a τ-collapsing). Denote by $\mathcal{G}_k^\tau$ the event (in this modified process) that τ has degree k at the end of the first epoch. For, $k \geq 0$

$$\begin{aligned}
\Pr(\mathcal{G}_k^\tau) &= \sum_{j=k}^{\infty} \frac{c^j}{j!} e^{-c} \binom{j}{k} (1-(1-\gamma_r)^d)^{j-k} (1-\gamma_r)^{kd} \\
&= \frac{(1-\gamma_r)^{kd} c^k}{k!} \sum_{j=0}^{\infty} \frac{c^j}{j!} e^{-c} (1-(1-\gamma_r)^d)^j \\
&= \frac{((\beta_{r-1})^d c)^k}{k!} \exp(-c(1-\gamma_r)^d) \\
&= \frac{((\beta_{r-1})^d c)^k}{k!} \gamma_{r+1}. \qquad (3)
\end{aligned}$$

[2] A $(d-1)$-dimensional face of a d-complex is *isolated* if it is not contained in any d-face.

Consequently, the degree of τ is Poisson-distributed with parameter $(\beta_{r-1})^d c$.

2. PROOF OF THEOREM 1.2

As mentioned, the first epoch proceeds for r phases. As shown in [1, Eq.(7)]:

$$\mathbb{E}[L_r] = (1 - o(1))\binom{n}{d}(1 - (\gamma_{r+1} + c\gamma_r\beta_{r-1}^d)). \tag{4}$$

A concentration of measure argument similar to the one used in [1, Section (4.2)] yields that a.s.

$$|L_r - \mathbb{E}[L_r]| \le o(n^d). \tag{5}$$

A more detailed proof of this appears in the Appendix.

For a $(d-1)$-face to be free at the end of the first epoch, it must have degree ≥ 2 after collapsing phase $r-1$ and degree 1 after collapsing phase r. Therefore the probability for a $(d-1)$-face to be free at the beginning of the second epoch is:

$$\sum_{j=1}^{\infty}\frac{c^j}{j!}e^{-c}j(1-(1-\gamma_r)^d)^{j-1}(1-\gamma_r)^d - \sum_{j=1}^{\infty}\frac{c^j}{j!}e^{-c}j(1-(1-\gamma_{r-1})^d)^{j-1}(1-\gamma_r)^d$$
$$= (\beta_{r-1})^d c\gamma_{r+1} - (\beta_{r-1})^d c\gamma_r = (\beta_{r-1})^d c(\gamma_{r+1} - \gamma_r).$$

Notice this expression is different than the expression in (3) with $k = 1$, since here we permit collapsing of the $(d-1)$-face (i.e. this is not a τ-collapsing). Consequently, the expectation of X_0 is:

$$\mathbb{E}[X_0] = \binom{n}{d}(\beta_{r-1})^d c(\gamma_{r+1} - \gamma_r). \tag{6}$$

Let $\sigma \supset \tau$ be faces that we collapse at some step of the second epoch. The $(d-1)$-faces *affected* by that step are all the $(d-1)$-subfaces of σ other than τ. As mentioned, we control the number of free $(d-1)$-faces throughout the process. Clearly the change in this number at a given step is determined by the current degrees of the affected faces. We consider the count of steps in the second epoch as "time". In particular, time zero means the end of the first epoch and the beginning of the second one.

Let D_k be the random variable that counts the $(d-1)$-faces τ, that have degree k, at the end of the first epoch, of the τ-collapsing. In order to calculate the expectation of D_k we use the events $\mathcal{G}_k$. Notice these events are defined in the probability space $\mathcal{T}(c,t)$, but the proximity of the distributions yields:

$$\mathbb{E}[D_k] = (1 - o(1))\sum_{\tau}\Pr(\mathcal{G}_k^\tau) = (1 - o(1))\binom{n}{d}\frac{((\beta_{r-1})^d c)^k}{k!}\gamma_{r+1}. \tag{7}$$

Let B_j be the random variable that counts $(d-1)$-faces which are not isolated after the j-th collapsing phase of the first epoch. Recall that β_j is the probability that the root of the d-tree becomes isolated, in more than t collapse phases, in $\mathcal{T}(c, j+1)$. Hence $\mathbb{E}[B_j] = (1 - o(1))\binom{n}{d}\beta_j$ (the event B_j is defined in $X_d(n, \frac{c}{n})$, but the proximity of the distributions yields this expectation). Again a concentration of measure argument in the spirit of [1, Section (4.2)] yields that a.s. $|\mathbb{E}[D_k] - D_k| \leq o(n^d)$ and $|\mathbb{E}[B_j] - B_j| \leq o(n^d)$.

Let $\sigma \supset \tau$ be the faces that we collapse at time i. We are interested in the event that some affected $(d-1)$-face τ' becomes free due to this step. This can happen if $\Delta^\sigma(\tau') = k > 1$ and $k-1$ of the k distinct τ'-containing d-faces other than σ got collapsed at some time $< i$ (one of them must not be collapsed before step i, since the degree of τ after step i is one). The other way in which this can happen is the event that $\Delta^\sigma(\tau') = 1$ and the other d-face containing τ' was not collapsed in time $< i$. This event is obviously contained in the event $C^{\sigma,\tau'}$ that $\Delta^\sigma(\tau') = 1$.

We next define the random variable D'_k exactly as D_k, except that "degree" is interpreted a little differently. Namely a $(d-1)$-face of degree s that is contained in σ is considered as having degree $s-1$. Clearly $D_k - (d+1) \leq D'_k \leq D_k + d + 1$ and the same concentration of measure arguments work for both random variables. Since σ was not collapsed in the first epoch, we know that τ' was not collapsed, hence the degree of τ' in $X \setminus \sigma$ after the first epoch is the degree in $X \setminus \sigma$ if we forbid collapsing on τ'.

Let us consider the d-tree, rooted at τ', that consists of the r-neighborhood of τ' excluding the d-face σ. Before the last phase of the first epoch, τ' does not become isolated in this tree. Otherwise τ' would have become a free face of σ, and consequently σ would have been collapsed in the first epoch. The number of $(d-1)$ faces τ that do not become isolated before the last phase of the first epoch is B_{r-1}. Therefore

$$\Pr(C^{\sigma,\tau'}) = (1 \pm o(1))\frac{\mathbb{E}[D'_1]}{\mathbb{E}[B_{r-1}]} = (1 \pm o(1))(\beta_{r-1})^{d-1} c\gamma_{r+1}. \tag{8}$$

We denote this probability by $x := (1 \pm o(1))(\beta_{r-1})^{d-1} c\gamma_{r+1}$.

Let $Q_i^{\tau'}$ be the event that i is the first time at which τ' is an affected face. The number of $(d-1)$-faces that get affected before time i cannot exceed $(i-1)d$. Therefore,

$$\Pr(Q_i^{\tau'}) \geq 1 - \frac{(i-1)d}{\beta_{r-1}\binom{n}{d}} \tag{9}$$

We define Q_i as the event that each face affected in step i is affected for the first time. We prove

Lemma 1. *For every i there holds*

$$\Pr(Q_i) \geq \prod_{\tau_j \in T} \Pr(Q_i^{\tau_j}).$$

Here $T = \{\tau_1 \ldots, \tau_d\}$ is the set of faces that are affected in collapsing step i.

Proof. Let τ be the $(d-1)$-face collapsed in the i-th collapsing step, and let k be its mark. If some face $\tau_j \in T$ has a neighbor with mark $< k$ then $\Pr(Q_i^{\tau_j}) = 0$, $\Pr(Q_i) = 0$ and the inequality clearly holds. Otherwise, every face in T is affected in a collapsing step where the collapsed $(d-1)$-face has mark k. Hence the probability of the event $Q_i^{\tau_j}$ depends only on the random order selected for marking step k. Specifically let α_j be the number of neighbors of τ_j with mark k. The probability that i is the first time in which τ' is an affected face, is the probability that the permutation sampled in marking step k, is one where τ precedes all the other neighbors of τ_j with mark k. Hence $\Pr(Q_i^{\tau_j}) = \frac{1}{\alpha_j}$. Let J be the set of indices j for which $\alpha_j > 1$, if $J = \emptyset$ then the inequality of the lemma easily follows. Otherwise $\Pr(\mathcal{Q}_i)$ is the probability that the permutation sampled in marking step k, is one where τ precedes all the neighbors of faces τ_j, for which $j \in J$. Hence,

$$\Pr(\mathcal{Q}_i) \geq \frac{1}{1 + \sum_J (\alpha_j - 1)} \geq \frac{1}{\sum_J \alpha_j}.$$

But for all j,

$$\prod_{1 \leq j \leq d} \Pr(\mathcal{Q}_i^{\tau_j}) = \prod_J \Pr(\mathcal{Q}_i^{\tau_j}) = \prod_J \frac{1}{\alpha_j}$$

and the conclusion follows, since $\frac{1}{\sum_J \alpha_j} \geq \prod_J \frac{1}{\alpha_j}$. ∎

In view of (9) and Lemma 1, we see that if $i \ll \beta_{r-1}\binom{n}{d}/d$, then $\Pr(\mathcal{Q}_i^{\tau'}) = 1 - o(1)$ and hence $\Pr(\mathcal{Q}_i) = 1 - o(1)$.

We can finally calculate $\mathbb{E}[Y_i]$, the expected number of new free $(d-1)$-faces added in a collapsing step. We do this in terms of the events S_i^j. It is convenient for us to condition on the almost sure event $\mathcal{Q}_i$. For every $1 \leq j \leq d$ we write

$$\begin{aligned}
\Pr(S_i^j) &= \Pr(S_i^j | \mathcal{Q}_i) \Pr(\mathcal{Q}_i) + \Pr(S_i^j | \bar{\mathcal{Q}}_i) \Pr(\bar{\mathcal{Q}}_i) \\
&\leq \Pr(S_i^j | \mathcal{Q}_i) + 1 - \Pr(\mathcal{Q}_i) \\
&\leq (1 + o(1)) \binom{d}{j} (1-x)^{d-j} x^j
\end{aligned}$$

To see this, notice that conditioned on $\mathcal{Q}_i$ the event S_i^j happens if $\Delta^\sigma(\tau') = 1$ for exactly j of the affected faces τ'.

Notice $\Pr(Y_i = j) = \Pr(S_i^j)$, hence

$$\begin{aligned}
\mathbb{E}[Y_i] &= \sum_{i=1}^{j} j \cdot \Pr(S_i^j) \leq (1 + o_n(1)) \sum_{i=1}^{j} j \binom{d}{j} (1-x)^{d-j} x^j \\
&= (1 + o_n(1)) dx = (1 + o_n(1)) d (\beta_{r-1})^{d-1} c \gamma_{r+1}. \qquad (10)
\end{aligned}$$

Claim 2.1. *For $c > \eta_d$ and $i \ll \beta_{r-1}\binom{n}{d}/d$ there holds $\mathbb{E}[Y_i] < 1$.*

Proof. Recall that $d \cdot \beta^{d-1} \eta_d \gamma = 1$ where $\beta = \lim_{r\to\infty} \beta_r(\eta_d)$ and $\gamma = 1 - \beta$ [2, Section (3)]. We turn to show that $d \cdot \beta_{r-1}^{d-1} c \gamma_{r+1} < 1$ for $c > \eta_d$. To this end we recall and slightly extend some analysis from [1,2]. The role of the parameter η_d is revealed by analyzing the function $f_c(t) = 1 - e^{-ct^d} - t$.

For $c > \eta_d$ the function f_c has exactly two zeroes in the open interval $(0, 1)$, which we denote by $0 < b(c) < B(c) < 1$. Note that $B(c) = \lim_{k\to\infty} \beta_k(c)$. The fact that $b(c), B(c)$ are roots of f_c yields $c = g(B(c)) = g(b(c))$ where $g(t) = \frac{-\ln(1-t)}{t^d}$. Therefore $g'(b(c))b'(c) = g'(B(c))B'(c) = 1$. (Prime denotes derivative w.r.t. c.) We claim that B (resp. b) is an increasing (resp. decreasing) function of c, which would follow if $g'(b(c)) < 0 < g'(B(c))$. This is so, since, as is easily verified, for every integer $d \geq 2$ the function g has exactly one critical point in $(0, 1)$, which is a minimum, and $0 < b < B < 1$.

Let $h(c) = d \cdot B(c)^{d-1} \cdot c \cdot (1 - B(c))$, and recall that $h(\eta_d) = 1$. As we show next, $h(c) < 1$ for c that is a little larger than η_d. We denote derivatives w.r.t. B by an upper dot, and we prove this claim by showing that $\dot{h} < 0$. By choosing r large enough we obtain $d(\beta_{r-1})^{d-1}c\gamma_{r+1} < 1$. Together with (10) this implies the claim.

We take the derivative w.r.t.B of the relation $c = g(B)$:

$$\dot{c} = \frac{1}{B(c)^d(1-B(c))} + \frac{d \cdot \ln(1-B(c))}{B(c)^{d+1}} = \frac{1}{B(c)^d(1-B(c))} - \frac{d \cdot c}{B(c)}.$$

But

$$\begin{aligned} \dot{h} &= d\left(((d-1)B(c)^{d-2} - dB(c)^{d-1})c + (B(c)^{d-1} - B(c)^d)\dot{c}\right) \\ &= d\left(-c \cdot B(c)^{d-2} + \frac{1}{B(c)}\right) \\ &= d\left(\frac{\ln(1-B(c))}{B(c)^2} + \frac{1}{B(c)}\right) \end{aligned}$$

which is negative, since $-\ln(1-u) > u$ for all u. ■

Final steps in proof of Theorem 1.2. Fix $c > \eta_d$. As shown in Claim 2.1, $\mathbb{E}[Y_i] = \mathbb{E}[Y_i'] \leq 1-\epsilon$ for some $\epsilon > 0$. We select $\delta \ll \frac{\beta_{r-1}}{2d}\epsilon$ and we want to guarantee that $\mathbb{E}[X_0] \leq \delta\binom{n}{d}$. By (6), this can be guaranteed by choosing r, the number of phases in the first epoch, a large enough integer. Let ρ be a number in the range $\frac{2\delta}{\epsilon} \leq \rho \ll \frac{\beta_{r-1}}{d}$. To complete our proof we need to show that for some $i \leq \rho\binom{n}{d}$ small enough, a.s. $X_i = 0$. As was done in [9, Proof of Lemma 4], we define a series of random variables $Z_0, Z_1, \ldots$ (that are very similar to $X_0, X_1, \ldots$) as follows. We let $Z_0 = X_0$. For every $i > 0$, $Z_i := Z_{i-1} - 1 + Y_i'$, where Y_i' has the same distribution as Y_i. In other words, for every $0 \leq j \leq d$ the random variable Z_i takes the value $Z_{i-1} - 1 + j$ with probability $\Pr(S_i^j)$. As long as X_i is larger than 0 it behaves like Z_i, however once it vanishes it stays 0, while Z_i can become negative. Hence the simple but crucial fact is that if $Z_i < 0$ then $X_i = 0$. Thus for $i > \frac{2\delta}{(1+o_n(1))\epsilon}\binom{n}{d}$ by (10) the displayed equation follows:

$$\mathbb{E}[Z_i] = \mathbb{E}[X_0] - i + (1+o_n(1))dxi < -\delta\binom{n}{d}.$$

The desired conclusion follows now by applying a version of Azuma's inequality from [7, Lemma (1.2)]:

Proposition 2.2. *Let $Y_1, \ldots, Y_m$ be random variables taking values from $\{0,1\}$. Suppose $\Phi : \{0,1\}^m \to \mathbb{R}$ satisfies*

$$|\Phi(x) - \Phi(x')| \leq \epsilon_k$$

for x and x' that differ only at their k-th coordinate. Then for every $t > 0$:

$$\Pr[|\Phi(Y_1, \ldots, Y_m) - \mathbb{E}[\Phi(Y_1, \ldots, Y_m)]| \geq t] \leq 2e^{-\frac{2t^2}{\sum_k \epsilon_k^2}} \tag{11}$$

Fix an i bigger than $\frac{2\delta}{\epsilon}\binom{n}{d}$ and smaller than $\beta_{r-1}\binom{n}{d}/d$ and L_r (by (4) and (5) such an i exists). The random variables $Y_1', \ldots, Y_i'$ take values in $\{0, \ldots, d\}$. Notice $Z_i = f(Y_1', \ldots, Y_i') = X_0 - i + \sum_{j=0}^{i} Y_j'$ and for two vectors t, t' that differ in only the k-th coordinate $|f(t) - f(t')| \leq d$, hence for $l = \delta\binom{n}{d}$

$$\Pr[|f(Y_1', \ldots, Y_i') - \mathbb{E}[f(Y_1', \ldots, Y_i')]| \geq l] \leq 2e^{-\frac{2 \cdot l^2}{i \cdot d^2}} = o(1).$$

Thus a.s. $Z_i < 0$, as claimed. ■

APPENDIX

In this section we prove Eq. (5) namely:

$$|L_r - \mathbb{E}[L_r]| \leq o(n^d).$$

We prove this concentration of measure claim using Proposition 2.2. We fix a numbering of all d-faces on vertex set $[n]$ and define indicator random variables $\mathcal{Y}_1, \ldots, \mathcal{Y}_{\binom{n}{d+1}}$ where $\mathcal{Y}_i = 1$ iff the i-th d-face belongs to X. Let X and X' be two d-complexes that are identical except for the d-face σ, i.e. $\sigma \in X$ and $\sigma \notin X'$. A $(d-1)$-face τ is said to *linger longer* in X than in X' if and only if the phase it becomes isolated in is different in X and X'.

We show that $|L_r(X) - L_r(X')| \leq (d+1)d^r$. First notice a few crucial observations:

- $L_r(X) \geq L_r(X')$, namely for some $(d-1)$-face, τ, the fact that σ does not belong to X' can only make $\Delta(\tau)$ in X' smaller than in X.
- Suppose that a $(d-1)$-face τ, is isolated in X' at the end of the first epoch, but not isolated in X at the end of the first epoch. This can happen only if τ is within distance r of σ, in X, and lingers longer in X than in X'.
- A simple inductive argument shows that there are at most $(d+1) \cdot d^{k-1}$ distinct $(d-1)$-faces at distance k from σ that linger longer in X than in X'. Therefore the number of $(d-1)$ faces at distance $\leq r$ of σ, that linger longer in X than in X', is at most $\sum_{i=1}^{r}(d+1)d^{i-1} \leq (d+1)d^r$.

Consequently $|L_r(X) - L_r(X')| \leq (d+1)d^r$, and Proposition 2.2 yields:

$$\Pr[|L_r(X) - \mathbb{E}[L_r(X)]|] \geq t \cdot n^d] \leq 2e^{-\frac{2\cdot(t\cdot n^d)^2}{\binom{n}{d+1}(d+1)d^r}} = o(1)$$

as needed.

ACKNOWLEDGMENTS

We are grateful to the referees for thorough reading and insightful comments which have considerably improved the quality of presentation of this paper.

REFERENCES

[1] L. Aronshtam and N. Linial, When does the top homology of a random simplicial complex vanish?, Random Struct Algorithms 46 (2015), 26–35.

[2] L. Aronshtam, N. Linial, T. Łuczak, and R. Meshulam, Collapsibility and vanishing of top homology in random simplicial complexes, Discrete Comput Geom 49 (2013), 317–334.

[3] E. Babson, C. Hoffman, and M. Kahle, The fundamental group of random 2-complexes, J Am Math Soc 24 (2011), 1–28.

[4] D. Cohen, A. Costa, M. Farber, and T. Kappeler, Topology of random 2-complexes, Discrete Comput Geom 47 (2012), 117–149.

[5] C. Hoffman, M. Kahle, and E. Paquette, The threshold for integer homology in random d-complexes, http://arxiv.org/abs/1308.6232.

[6] N. Linial and R. Meshulam, Homological connectivity of random 2-complexes, Combinatorica 26 (2006), 475–487.

[7] C. McDiarmid, On the Method of Bounded Differences, Surveys in combinatorics 1989, London Mathematical Society Lecture Note Series, Vol. 141, Cambridge University Press, Cambridge, pp. 148–188, 1989.

[8] R. Meshulam and N. Wallach, Homological connectivity of random k-dimensional complexes, Random Struct Algorithms 34 (2009), 408–417.

[9] M. Molloy, Cores in random hypergraphs and Boolean formulas, Random Struct Algorithms 27 (2005), 124–135.

[10] B. Pittel, J. Spencer, and N. Wormald, Sudden emergence of a giant *k*-core in a random graph, J Combin Theory (Ser B) 67 (1996), 111–151.