Transitions and phase transitions

Nati Linial, Hebrew U. with
L. Aronshtam, T. Luczak, R. Meshulam, Y. Peled

Avi60, October 2016

In the beginning

Let us recall the classical (early 60's) work of Erdős and Rényi on random graphs. They come in three flavors, all being probability distributions on n-vertex graphs:

In the beginning

Let us recall the classical (early 60's) work of Erdős and Rényi on random graphs. They come in three flavors, all being probability distributions on n-vertex graphs:

- $G(n, M)$ - Uniform choice among all graphs with M edges.

In the beginning

Let us recall the classical (early 60's) work of Erdős and Rényi on random graphs. They come in three flavors, all being probability distributions on n-vertex graphs:

- $G(n, M)$ - Uniform choice among all graphs with M edges.
- $G(n, p)$ - For every pairs of vertices x, y pick the edge xy independently with probability p.

In the beginning

Let us recall the classical (early 60's) work of Erdős and Rényi on random graphs. They come in three flavors, all being probability distributions on n-vertex graphs:

- $G(n, M)$ - Uniform choice among all graphs with M edges.
- $G(n, p)$ - For every pairs of vertices x, y pick the edge xy independently with probability p.
- The evolution of random graphs. Start with no edges. Sequentially add edges at random.

The connectivity threshold

The same question for all three variants.

The connectivity threshold

The same question for all three variants. What is the critical parameter for graph connectivity?

The connectivity threshold

The same question for all three variants. What is the critical parameter for graph connectivity?

- What is the critical number of edges M in the $G(n, M)$ model for the transition between connected/disconnected?

The connectivity threshold

The same question for all three variants. What is the critical parameter for graph connectivity?

- What is the critical number of edges M in the $G(n, M)$ model for the transition between connected/disconnected?
- What is the critical edge density p in the $G(n, p)$ model?

The connectivity threshold

The same question for all three variants. What is the critical parameter for graph connectivity?

- What is the critical number of edges M in the $G(n, M)$ model for the transition between connected/disconnected?
- What is the critical edge density p in the $G(n, p)$ model?
- Hitting time version: At which moment does the evolving graph become connected?

Theorem (Bollobas and Thomason '85)

With probability $1 - o_n(1)$ *the random graph process becomes connected at exactly that step when the last isolated vertex disappears.*

Theorem (Bollobas and Thomason '85)

With probability $1 - o_n(1)$ the random graph process becomes connected at exactly that step when the last isolated vertex disappears.

Corollary

The critical density for graph connectivity is $p = \frac{\log n}{n}$.

Theorem (Bollobas and Thomason '85)

With probability $1 - o_n(1)$ the random graph process becomes connected at exactly that step when the last isolated vertex disappears.

Corollary

The critical density for graph connectivity is $p = \frac{\log n}{n}$.

The easy part - If $p < (1 - \epsilon)\frac{\log n}{n}$, then with almost certainty the graph is not only disconnected, it even has isolated vertices.

Theorem (Bollobas and Thomason '85)

With probability $1 - o_n(1)$ the random graph process becomes connected at exactly that step when the last isolated vertex disappears.

Corollary

The critical density for graph connectivity is $p = \frac{\log n}{n}$.

The easy part - If $p < (1 - \epsilon)\frac{\log n}{n}$, then with almost certainty the graph is not only disconnected, it even has isolated vertices.

It takes some work to show that for $p > (1 + \epsilon)\frac{\log n}{n}$, the graph is asymptotically almost surely connected.

Enter simplicial complexes

A major contact points between combinatorics and geometry (more specifically - with topology).

Enter simplicial complexes

A major contact points between combinatorics and geometry (more specifically - with topology). Combinatorially speaking, this is just a down-closed family of sets.

Enter simplicial complexes

A major contact points between combinatorics and geometry (more specifically - with topology).
Combinatorially speaking, this is just a down-closed family of sets.
Graphs are the ideal tool for modeling large systems that are governed by pairwise interactions,

Enter simplicial complexes

A major contact points between combinatorics and geometry (more specifically - with topology).
Combinatorially speaking, this is just a down-closed family of sets.
Graphs are the ideal tool for modeling large systems that are governed by pairwise interactions, and simplicial complexes can play a similar role in dealing with systems whose constituents exhibit multiway interactions.

Definition

Let V be a finite set of vertices. A collection of subsets $X \subseteq 2^V$ is called a simplicial complex if it satisfies the following condition:

$$A \in X \quad \text{and} \quad B \subseteq A \Rightarrow B \in X.$$

Definition

Let V be a finite set of vertices. A collection of subsets $X \subseteq 2^V$ is called a simplicial complex if it satisfies the following condition:

$$A \in X \text{ and } B \subseteq A \Rightarrow B \in X.$$

A member $A \in X$ is called a simplex or a face of dimension $|A| - 1$. A face of dimension k is often called a k-face.

Definition

Let V be a finite set of vertices. A collection of subsets $X \subseteq 2^V$ is called a simplicial complex if it satisfies the following condition:

$$A \in X \text{ and } B \subseteq A \Rightarrow B \in X.$$

A member $A \in X$ is called a simplex or a face of dimension $|A| - 1$. A face of dimension k is often called a k-face.
The dimension of X is the largest dimension of a face in X

Definition

Let V be a finite set of vertices. A collection of subsets $X \subseteq 2^V$ is called a simplicial complex if it satisfies the following condition:

$$A \in X \text{ and } B \subseteq A \Rightarrow B \in X.$$

A member $A \in X$ is called a simplex or a face of dimension $|A| - 1$. A face of dimension k is often called a k-face.

The dimension of X is the largest dimension of a face in X and a d-dimensional simplicial complex is often called a d-complex.

A simple but crucial observation

A graph is the synonymous with

A simple but crucial observation

A graph is the synonymous with a one-dimensional simplicial complex.

A simple but crucial observation

A graph is the synonymous with a one-dimensional simplicial complex.
So - Essentially everything that we know about graphs makes sense for simplicial complexes.

A simple but crucial observation

A graph is the synonymous with a one-dimensional simplicial complex.
So - Essentially everything that we know about graphs makes sense for simplicial complexes.
Here we tell some of the story on random simplicial complexes.

Simplicial complexes as geometric objects

We view $A \in X$ and $|A| = k+1$ as a k-dimensional simplex.

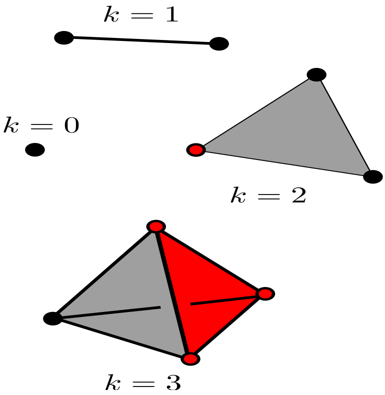

Putting simplices together properly

The intersection of every two simplices in X is a common face.

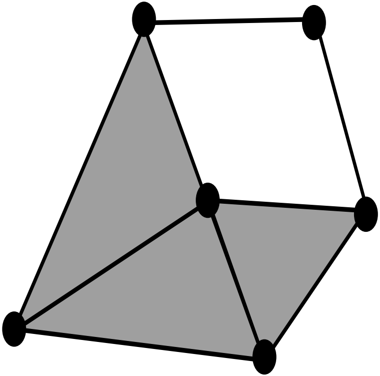

Geometric equivalence

Combinatorially different complexes may correspond to the same geometric object (e.g. via subdivision)

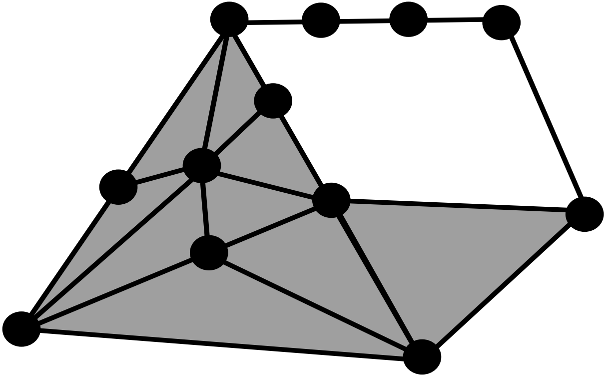

Geometric equivalence

So

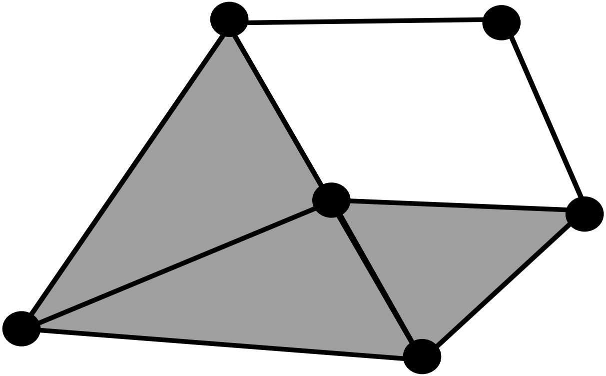

Geometric equivalence

and

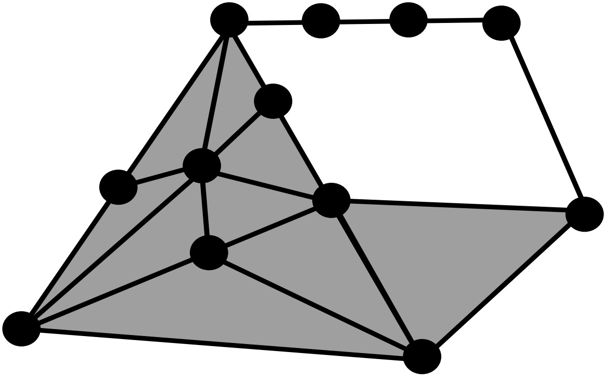

Geometric equivalence

are two different combinatorial descriptions of the same geometric object

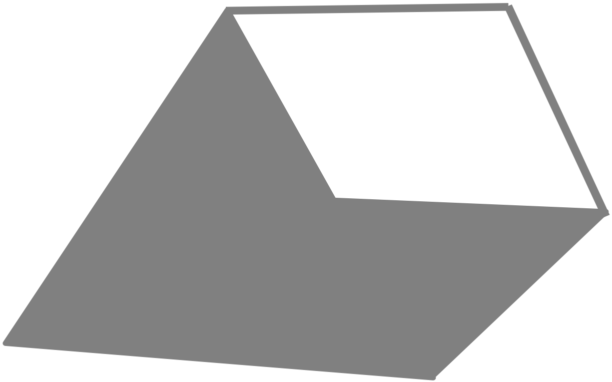

A sketch for a road-map

We want to develop a theory of random complexes, in the general spirit of random graph theory.

A sketch for a road-map

We want to develop a theory of random complexes, in the general spirit of random graph theory. In order to get started we need

- A higher-dimensional analog to $G(n, p)$.

A sketch for a road-map

We want to develop a theory of random complexes, in the general spirit of random graph theory. In order to get started we need

- A higher-dimensional analog to $G(n, p)$.
- A dictionary to translate basic graph-theoretic terms to the realm of high-dimensional simplicial complexes.

A sketch for a road-map

We want to develop a theory of random complexes, in the general spirit of random graph theory. In order to get started we need

- A higher-dimensional analog to $G(n,p)$.
- A dictionary to translate basic graph-theoretic terms to the realm of high-dimensional simplicial complexes.

Then we can take whatever we know about $G(n,p)$ graphs and seek the high-dimensional counterparts.

Our lingua franca - Linear algebra spoken here

- To say that $G = (V, E)$ is connected we use $A = A_G$, the incidence matrix of G.

Our lingua franca - Linear algebra spoken here

- To say that $G = (V, E)$ is connected we use $A = A_G$, the incidence matrix of G.
- It is a $V \times E$ matrix, indexed by vertices resp. edges. If $e = [i, j] \in E$, then

Our lingua franca - Linear algebra spoken here

- To say that $G = (V, E)$ is connected we use $A = A_G$, the incidence matrix of G.
- It is a $V \times E$ matrix, indexed by vertices resp. edges. If $e = [i, j] \in E$, then

$$a_{i,e} = 1, \quad a_{j,e} = -1 \quad \forall k \neq i, j \;\; a_{i,k} = 0.$$

- Clearly, $\mathbf{1}A = 0$, since every column of A contains one $+1$ and one -1 and all other entries are zero.

- Clearly, $\mathbf{1}A = 0$, since every column of A contains one $+1$ and one -1 and all other entries are zero.
- Likewise, if S is the vertex set of a connected component of G, then $\mathbf{1}_S A = 0$.

- Clearly, $\mathbf{1}A = 0$, since every column of A contains one $+1$ and one -1 and all other entries are zero.
- Likewise, if S is the vertex set of a connected component of G, then $\mathbf{1}_S A = 0$.
- It is not hard to see that G is connected if and only if the left kernel of A is one-dimensional.

- Clearly, $\mathbf{1}A = 0$, since every column of A contains one $+1$ and one -1 and all other entries are zero.
- Likewise, if S is the vertex set of a connected component of G, then $\mathbf{1}_S A = 0$.
- It is not hard to see that G is connected if and only if the left kernel of A is one-dimensional.

This brings us back to topology

Now with subtitles in topologese

The linear transformation corresponding to the matrix A is the boundary operator, usually denoted by ∂ and the condition that it has just the trivial left kernel means that the zeroth homology of G vanishes.

The many roads not taken...

Very often scenarios and concepts that are simple, even obvious, for $d = 1$ (graph theory) are much richer in higher dimension.

The many roads not taken...

Very often scenarios and concepts that are simple, even obvious, for $d = 1$ (graph theory) are much richer in higher dimension. The d-dimensional analog of graph connectivity that we discuss here is:

The many roads not taken...

Very often scenarios and concepts that are simple, even obvious, for $d = 1$ (graph theory) are much richer in higher dimension. The d-dimensional analog of graph connectivity that we discuss here is: The $(d-1)$-st homology vanishes.

The many roads not taken...

Very often scenarios and concepts that are simple, even obvious, for $d = 1$ (graph theory) are much richer in higher dimension. The d-dimensional analog of graph connectivity that we discuss here is: The $(d-1)$-st homology vanishes.
and only mention in passing

The many roads not taken...

Very often scenarios and concepts that are simple, even obvious, for $d = 1$ (graph theory) are much richer in higher dimension. The d-dimensional analog of graph connectivity that we discuss here is:
The $(d-1)$-st homology vanishes.
and only mention in passing beautiful work of Babson Hoffman and Kahle on the threshold for the vanishing of the fundamental group in random 2-dimensional complexes.

Setting the ground

Our paper with Roy Meshulam ('06) introduces $X_d(n, p)$, a d-dimensional analog of $G(n, p)$.

Setting the ground

Our paper with Roy Meshulam ('06) introduces $X_d(n,p)$, a d-dimensional analog of $G(n,p)$. This is a d-dimensional complex on n vertices. It has a full $(d-1)$-dimensional skeleton. Namely, every face of dimension $\leq d-1$ is present.

Setting the ground

Our paper with Roy Meshulam ('06) introduces $X_d(n, p)$, a d-dimensional analog of $G(n, p)$. This is a d-dimensional complex on n vertices. It has a full $(d-1)$-dimensional skeleton. Namely, every face of dimension $\leq d-1$ is present. Every d-face is joined in independently with probability p.

Setting the ground

Our paper with Roy Meshulam ('06) introduces $X_d(n, p)$, a d-dimensional analog of $G(n, p)$. This is a d-dimensional complex on n vertices. It has a full $(d-1)$-dimensional skeleton. Namely, every face of dimension $\leq d-1$ is present. Every d-face is joined in independently with probability p. Note that $X_1(n, p)$ is identical with $G(n, p)$.

More groundwork

The full d-dimensional boundary operator ∂_d is an $\binom{n}{d} \times \binom{n}{d+1}$ matrix indexed by subsets of $[n]$ of cardinalities d and $d+1$ resp.

More groundwork

The full d-dimensional boundary operator ∂_d is an $\binom{n}{d} \times \binom{n}{d+1}$ matrix indexed by subsets of $[n]$ of cardinalities d and $d+1$ resp.
If $R \subset S$, with $S \setminus R = \{x\}$ the (R, S) entry is $(-1)^i$ where x is the i-th largest element of S.

In linear algebra terms, to sample a graph from $G(n, p)$ we start from the above-mentioned $n \times \binom{n}{2}$ matrix and pick each column independently and with probability p.

In linear algebra terms, to sample a graph from $G(n,p)$ we start from the above-mentioned $n \times \binom{n}{2}$ matrix and pick each column independently and with probability p.

In general, to sample a d-complex from $X_d(n,p)$ we start from the above $\binom{n}{d} \times \binom{n}{d+1}$ and pick each column independently with probability p.

In linear algebra terms, to sample a graph from $G(n, p)$ we start from the above-mentioned $n \times \binom{n}{2}$ matrix and pick each column independently and with probability p.

In general, to sample a d-complex from $X_d(n, p)$ we start from the above $\binom{n}{d} \times \binom{n}{d+1}$ and pick each column independently with probability p.

We can now spell out in elementary terms what it means that the $(d-1)$-st homology vanishes.

It is easy to verify that

$$\partial_{d-1}\partial_d = 0$$

It is easy to verify that

$$\partial_{d-1}\partial_d = 0$$

and the left kernel of ∂_d is the row space of ∂_{d-1}.

It is easy to verify that

$$\partial_{d-1}\partial_d = 0$$

and the left kernel of ∂_d is the row space of ∂_{d-1}. Clearly if $p \to 1$, then the same should also hold for Y's boundary operator. On the other hand, if $p \to 0$, the left kernel gets larger.

Specifically, if p is small enough so there is a all-zero row, then we clearly get new vectors in the left kernel. Note that for $d = 1$ an all-zero row corresponds to an isolated vertex. So we are considering the critical density for the existence of isolated vertices.

Specifically, if p is small enough so there is a all-zero row, then we clearly get new vectors in the left kernel. Note that for $d = 1$ an all-zero row corresponds to an isolated vertex. So we are considering the critical density for the existence of isolated vertices.

In this respect the high-dimensional answer is consistent with the one-dimensional case. Namely,

and indeed

Theorem (Linial-Meshulam '06, Meshulam-Wallach '09)

The threshold for the vanishing of the $(d-1)$-st homology in $X_d(n,p)$ over any finite ring of coefficients is

$$p = (1+o(1))\frac{d\ln n}{n}.$$

Again it's easy to see that for $p < (1-\epsilon)\frac{d \ln n}{n}$ the $(d-1)$-st homology is nonzero, since we get an all-zero row in Y's boundary operator.

Again it's easy to see that for $p < (1-\epsilon)\frac{d \ln n}{n}$ the $(d-1)$-st homology is nonzero, since we get an all-zero row in Y's boundary operator.

The proof shows that in the complementary range $p > (1+\epsilon)\frac{d \ln n}{n}$ there is nothing additional in the left kernel of Y's boundary operator.

Another good story we must skip (intended for the mavens)

It is conjectured that $p = \frac{d \ln n}{n}$ is also the threshold for the vanishing of the $(d-1)$-st homology over $\mathbb{Z}$.

Another good story we must skip (intended for the mavens)

It is conjectured that $p = \frac{d \ln n}{n}$ is also the threshold for the vanishing of the $(d-1)$-st homology over $\mathbb{Z}$.

This is presently known for $d = 2$ (Luczak and Peled '16).

Another good story we must skip (intended for the mavens)

It is conjectured that $p = \frac{d \ln n}{n}$ is also the threshold for the vanishing of the $(d-1)$-st homology over $\mathbb{Z}$.

This is presently known for $d = 2$ (Luczak and Peled '16).

It is also known up to a constant factor for all d (Hoffman, Kahle, Paquette '14+).

Phase transition in $G(n, p)$

Perhaps the most famous discovery of Erdős and Rényi refers to $p = \frac{1}{n}$ ($\approx$time $\frac{n}{2}$). At this point several dramatic changes occur.

Phase transition in $G(n,p)$

Perhaps the most famous discovery of Erdős and Rényi refers to $p = \frac{1}{n}$ ($\approx$time $\frac{n}{2}$). At this point several dramatic changes occur.

This is the first time when the graph is almost surely no longer a forest.

Phase transition in $G(n, p)$

Perhaps the most famous discovery of Erdős and Rényi refers to $p = \frac{1}{n}$ ($\approx$time $\frac{n}{2}$). At this point several dramatic changes occur.

This is the first time when the graph is almost surely no longer a forest.

Namely, for every $\epsilon > 0$ there is $1 > q(\epsilon) > 0$, such that the probability for a $G(n, \frac{1-\epsilon}{n})$ graph to be a forest is $(1 + o_n(1))q(\epsilon)$, but a graph in $G(n, \frac{1+\epsilon}{n})$ is almost surely not a forest.

At around this time a giant component emerges.

At around this time a giant component emerges. Namely, almost every graph graph in $G(n, \frac{1+\epsilon}{n})$ has a connected component on $c(\epsilon) \cdot n$ vertices, where $c(\cdot)$ is some well-specified function.

Can we say these things in high-dimenese?

At least some of them we do recognize. Right?

Can we say these things in high-dimenese?

At least some of them we do recognize. Right?

- A forest is just an acyclic graph. We know what cycles are in all dimensions. Agree?

Can we say these things in high-dimenese?

At least some of them we do recognize. Right?

- A forest is just an acyclic graph. We know what cycles are in all dimensions. Agree? Well, not so fast. We'll see.
- But one thing is certainly a problem - There is no obvious notion of connected components in dimensions $d \geq 2$.

Can we say these things in high-dimenese?

At least some of them we do recognize. Right?

- A forest is just an acyclic graph. We know what cycles are in all dimensions. Agree? Well, not so fast. We'll see.
- But one thing is certainly a problem - There is no obvious notion of connected components in dimensions $d \geq 2$. So how should we proceed?

Cycles in all dimensions

Definition

The cycle space of a simplicial complex is the right kernel of the corresponding boundary operator.

A theorem from Discrete Math 101

Theorem

TFAE for an n-vertex graph $G = (V, E)$ with $n - 1$ edges.

A theorem from Discrete Math 101

Theorem

TFAE for an n-vertex graph $G = (V, E)$ with $n - 1$ edges.

- *It is connected.*

A theorem from Discrete Math 101

Theorem

TFAE for an n-vertex graph $G = (V, E)$ with $n - 1$ edges.

- *It is connected.*
- *It is acyclic.*

A theorem from Discrete Math 101

Theorem

TFAE for an n-vertex graph $G = (V, E)$ with $n - 1$ edges.

- *It is connected.*
- *It is acyclic.*
- *It is collapsible.*

A theorem from Discrete Math 101

Theorem

TFAE for an n-vertex graph $G = (V, E)$ with $n - 1$ edges.

- *It is connected.*
- *It is acyclic.*
- *It is collapsible.*

Actually collapsibility is not usually discussed in undergraduate class, although it's completely elementary.

Collapsibility

Definition

- $d = 1$, graphs. An elementary collapse is a step where we eliminate a vertex of degree 1 and the single edge that contains it.

Collapsibility

Definition

- $d = 1$, graphs. An elementary collapse is a step where we eliminate a vertex of degree 1 and the single edge that contains it.
- General d. Let σ be a $(d-1)$-dimensional face that is contained in a single d-dimensional face τ. In the corresponding elementary collapse step we delete both σ and τ from the complex.

d-Collapsibility

Definition

A d-dimensional simplicial complex is said to be d-collapsible if it is possible to eliminate all its d-dimensional faces by a series of elementary collapses.

d-Collapsibility

Definition

A d-dimensional simplicial complex is said to be d-collapsible if it is possible to eliminate all its d-dimensional faces by a series of elementary collapses.

It is easy to see, using either very elementary topological or combinatorial arguments that:

Proposition

A d-collapsible d-dimensional complex is acyclic.

A little surprise

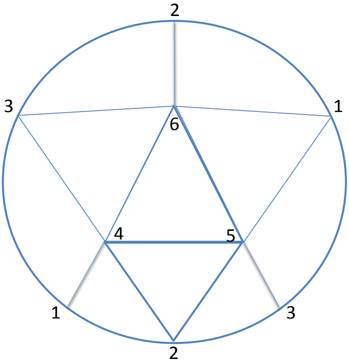

A little surprise

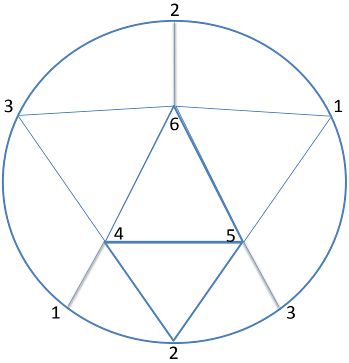

The triangulation of the projective plane is non-collapsible but is $\mathbb{R}$-acyclic

A little surprise

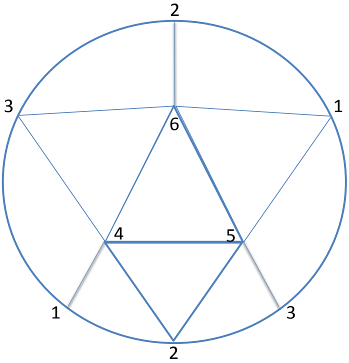

The triangulation of the projective plane is non-collapsible but is $\mathbb{R}$-acyclic (and $\mathbb{F}_2$-cyclic...).

Questions that suggest themselves

- What is the threshold for d-collapsibility in $X_d(n, p)$?
- For acyclicity ($=$ the vanishing of the d-th homology)?
- Clearly $p_{\text{collapsibility}} \leq p_{\text{acyclicity}}$, but is the inequality strict?

The short answer

Theorem

There are explicit constants γ_d and c_d for all $d \geq 2$, such that in $X_d(n,p)$

- *The threshold for d-collapsibility is $p = \frac{\gamma_d}{n}$.*
- *The threshold for acyclicity is $p = \frac{c_d}{n}$.*

The asymptotics in d are

$$\gamma_d = (1 + o(1)) \log d$$

and

$$c_d = d + 1 - o(1)$$

A taste of the ideas that go into the proofs - d-collapsibility

The k-core of a graph G is the largest induced subgraph in which every vertex has degree $\geq k$.

A taste of the ideas that go into the proofs - d-collapsibility

The k-core of a graph G is the largest induced subgraph in which every vertex has degree $\geq k$. You can find it by repeatedly removing from G every vertex of degree $< k$.

A taste of the ideas that go into the proofs - d-collapsibility

The k-core of a graph G is the largest induced subgraph in which every vertex has degree $\geq k$. You can find it by repeatedly removing from G every vertex of degree $< k$. If this process reminds you of the concept of collapsing, you are right.

A taste of the ideas that go into the proofs - d-collapsibility

The k-core of a graph G is the largest induced subgraph in which every vertex has degree $\geq k$. You can find it by repeatedly removing from G every vertex of degree $< k$. If this process reminds you of the concept of collapsing, you are right. It was a major achievement of random graph theory to determine when the k-core emerges in the evolution of random graphs.

A taste of the ideas that go into the proofs - d-collapsibility

The k-core of a graph G is the largest induced subgraph in which every vertex has degree $\geq k$. You can find it by repeatedly removing from G every vertex of degree $< k$. If this process reminds you of the concept of collapsing, you are right. It was a major achievement of random graph theory to determine when the k-core emerges in the evolution of random graphs. Indeed ideas from that domain have played a crucial role for us.

Some ideas around the d-collapsibility threshold

A key idea in the study of k-cores and other problems in random graph theory is to consider the local structure of the graph that's being generated.

Some ideas around the d-collapsibility threshold

A key idea in the study of k-cores and other problems in random graph theory is to consider the local structure of the graph that's being generated. This is often done using the machinery of Galton-Watson trees. Here we do something analogous with complexes.

Some ideas around the d-collapsibility threshold

A key idea in the study of k-cores and other problems in random graph theory is to consider the local structure of the graph that's being generated. This is often done using the machinery of Galton-Watson trees. Here we do something analogous with complexes. More generally, the idea of weak local limits is crucial for the whole line of research.

Some ideas around the acyclicity threshold

It is a simple linear-algebra observation that a d-complex with more d-faces than $(d-1)$-faces is cyclic.

Some ideas around the acyclicity threshold

It is a simple linear-algebra observation that a d-complex with more d-faces than $(d-1)$-faces is cyclic. This implies immediately that $c_d \leq d+1$.

Some ideas around the acyclicity threshold

It is a simple linear-algebra observation that a d-complex with more d-faces than $(d-1)$-faces is cyclic. This implies immediately that $c_d \leq d+1$. But more is true: Being acyclic is a topological property, so it's not affected by collapse steps.

Some ideas around the acyclicity threshold

It is a simple linear-algebra observation that a d-complex with more d-faces than $(d-1)$-faces is cyclic. This implies immediately that $c_d \leq d+1$. But more is true: Being acyclic is a topological property, so it's not affected by collapse steps. Moreover, the above observation remains valid if we eliminate every $(d-1)$-face that is not contained in any d-face.

Some ideas around the acyclicity threshold

It is a simple linear-algebra observation that a d-complex with more d-faces than $(d-1)$-faces is cyclic. This implies immediately that $c_d \leq d+1$.
But more is true: Being acyclic is a topological property, so it's not affected by collapse steps. Moreover, the above observation remains valid if we eliminate every $(d-1)$-face that is not contained in any d-face.
This argument, when formalized, gives the right value of c_d.

And what about the giant component?

A key idea in our work, and already in even more recent papers is that of a shadow.

And what about the giant component?

A key idea in our work, and already in even more recent papers is that of a shadow.
As usual, we start with graphs. Let $G = (V, E)$ be a graph, and let e be an edge not in E.

And what about the giant component?

A key idea in our work, and already in even more recent papers is that of a shadow.
As usual, we start with graphs. Let $G = (V, E)$ be a graph, and let e be an edge not in E. We say that e is in G's shadow if its addition to G creates new cycles in G.

And what about the giant component?

A key idea in our work, and already in even more recent papers is that of a shadow.
As usual, we start with graphs. Let $G = (V, E)$ be a graph, and let e be an edge not in E. We say that e is in G's shadow if its addition to G creates new cycles in G. In other words, both vertices of e are in the same connected component of G.

Shadows in higher dimension

Let X be a d-complex and σ is a d-face that is not in X.

Shadows in higher dimension

Let X be a d-complex and σ is a d-face that is not in X. We say that σ is in X's shadow if its addition to X creates new cycles in X.

Back to random graphs

An easy but crucial observation:

Back to random graphs

An easy but crucial observation:
In the evolution of random graphs, around time $n/2$ the following things happen more-or-less simultaneously:

Back to random graphs

An easy but crucial observation:
In the evolution of random graphs, around time $n/2$ the following things happen more-or-less simultaneously:

- The graph almost surely has a cycle.

Back to random graphs

An easy but crucial observation:
In the evolution of random graphs, around time $n/2$ the following things happen more-or-less simultaneously:

- The graph almost surely has a cycle.
- The giant componnet of $\Omega(n)$ vertices emerges.

Back to random graphs

An easy but crucial observation:
In the evolution of random graphs, around time $n/2$ the following things happen more-or-less simultaneously:

- The graph almost surely has a cycle.
- The giant componnet of $\Omega(n)$ vertices emerges.
- The shadow attains its asymptotic full size of $\Theta(n^2)$ edges.

and in high dimensions

As part of our main theorem we prove that

and in high dimensions

As part of our main theorem we prove that
The threshold for cyclicity coincides with the emergence of the giant shadow of $\Theta(n^{d+1})$ d-faces.

and in high dimensions

As part of our main theorem we prove that
The threshold for cyclicity coincides with the emergence of the giant shadow of $\Theta(n^{d+1})$ d-faces.

But there is a difference

A view of phase transition in $G(n, p)$

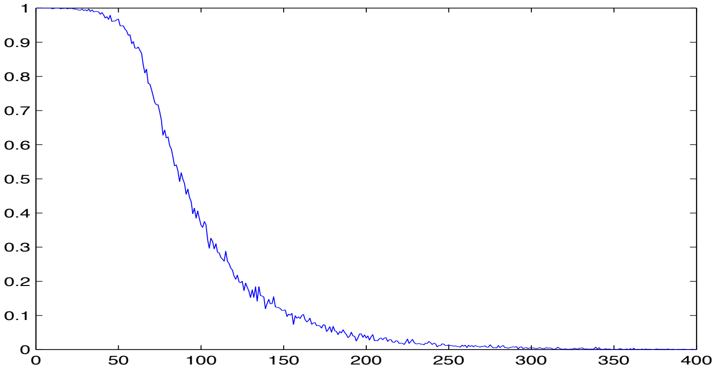

Phase transition in $X_2(n, p)$ complexes

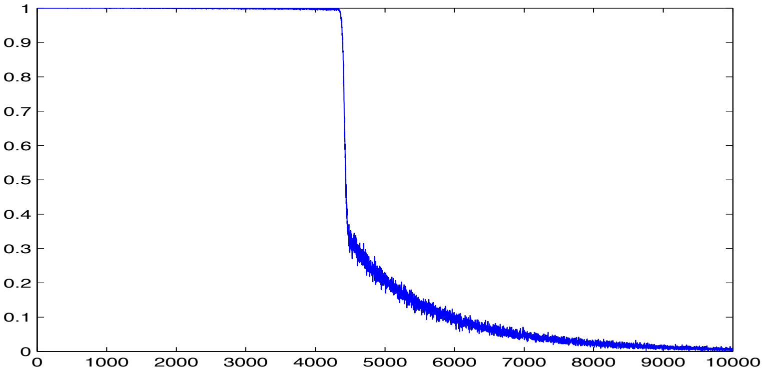

Phase transitions

We see a second order phase transition in $G(n, p)$ - the limit function is continuous, but its derivative has a discontinuity at the critical point.

Phase transitions

We see a second order phase transition in $G(n, p)$ - the limit function is continuous, but its derivative has a discontinuity at the critical point.
In contrast, at dimensions $d \geq 2$, the limit function is discontinuous at the critical point.

Phase transitions

We see a second order phase transition in $G(n, p)$ - the limit function is continuous, but its derivative has a discontinuity at the critical point.
In contrast, at dimensions $d \geq 2$, the limit function is discontinuous at the critical point.
We can now prove it.

Phase transitions

We see a second order phase transition in $G(n, p)$ - the limit function is continuous, but its derivative has a discontinuity at the critical point.
In contrast, at dimensions $d \geq 2$, the limit function is discontinuous at the critical point.
We can now prove it.
But don't worry, there are still many mysteries in this territory.

Mazal Tov Avi