Hypertrees

Nati Linial

Uri Feige 60, January 2020

Simplicial Complexes - A quick reminder/primer

A simplicial complex X on vertex set V is a collection of subsets of V that is closed down i.e.,

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

Simplicial Complexes - A quick reminder/primer

A simplicial complex X on vertex set V is a collection of subsets of V that is closed down i.e.,

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

A member $A \in X$ is called a simplex or a face of dimension $\dim(A) := |A| - 1$

Simplicial Complexes - A quick reminder/primer

A simplicial complex X on vertex set V is a collection of subsets of V that is closed down i.e.,

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

A member $A \in X$ is called a simplex or a face of dimension $\dim(A) := |A| - 1$ and $\dim(X) := \max\{\dim(A) | A \in X\}$

Familiar objects with new names

We see that a graph is synonymous with a one-dimensional simplicial complex. It has zero-dimensional faces (aka vertices)

Familiar objects with new names

We see that a graph is synonymous with a one-dimensional simplicial complex. It has zero-dimensional faces (aka vertices) and one dimensional faces = edges

What is expressible in this language?

How do you say a tree in the language of simplicial complexes?

What is expressible in this language?

How do you say a tree in the language of simplicial complexes?

Of course, a tree is a connected and acyclic graph

What is expressible in this language?

How do you say a tree in the language of simplicial complexes?

Of course, a tree is a connected and acyclic graph

Does the language of simplicial complexes provide analogous terms?

Gil Kalai 1983

ISRAEL JOURNAL OF MATHEMATICS, Vol 45, No. 4, 1983

ENUMERATION OF Q-ACYCLIC SIMPLICIAL COMPLEXES

BY

GIL KALAI

ABSTRACT

Let $\mathscr{C} = \mathscr{C}(n, k)$ be the class of all simplicial complexes C over a fixed set of n vertices $(2 \leqq k \leqq n)$ such that: (1) C has a complete $(k-1)$-skeleton, (2) C has precisely $\binom{n-1}{k}$ k-faces, (3) $H_k(C) = 0$. We prove that for $C \in \mathscr{C}$, $H_{k-1}(C)$ is a finite group, and our main result is:

$$\sum_{C} |H_{k-1}(C)|^2 = n^{\binom{n-2}{k}}.$$

Recall the incidence matrix of a graph

$V \times E$ Vertices vs. edges.

$$A_G = \begin{array}{c} \\ \vdots \\ i \\ \vdots \\ j \\ \vdots \end{array} \begin{array}{c} \begin{array}{ccccc} \dots & ij & \dots & \dots & \dots \end{array} \\ \begin{pmatrix} \dots & \dots & \dots & \dots & \dots \\ \dots & +1 & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots \\ \dots & -1 & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots \end{pmatrix} \end{array}$$

Incidence matrices tell many things

- G is connected iff A_G has a trivial left kernel.

Incidence matrices tell many things

- G is connected iff A_G has a trivial left kernel.
  - Because A_G's left kernel is the linear span of the indicator vectors of G's connected components.

Incidence matrices tell many things

- G is connected iff A_G has a trivial left kernel.
  - Because A_G's left kernel is the linear span of the indicator vectors of G's connected components.
- The cycle space of G is the right kernel of A_G.

Incidence matrices tell many things

- G is connected iff A_G has a trivial left kernel.
  - Because A_G's left kernel is the linear span of the indicator vectors of G's connected components.
- The cycle space of G is the right kernel of A_G.
  - Because A_G's right kernel is the linear span of the indicator vectors of G's cycle.

Recall: Equivalent descriptions of trees

Theorem

If $G = (V, E)$ is a graph with n vertices and $n - 1$ edges, then TFAE

1. *G is connected.*
2. *G is acyclic.*

Recall: Equivalent descriptions of trees

Theorem

If $G = (V, E)$ is a graph with n vertices and $n - 1$ edges, then TFAE

1. *G is connected.*
2. *G is acyclic.*
3. *G is collapsible.*

Why G is connected iff it is acyclic

For every G $\text{rank}(A_G) \leq \text{rank}(A_{K_n}) = n - 1$

Why G is connected iff it is acyclic

For every G $\text{rank}(A_G) \leq \text{rank}(A_{K_n}) = n - 1$

The only linear dependence among the rows is

$$1A_{K_n} = 0.$$

Why G is connected iff it is acyclic

For every G $\text{rank}(A_G) \leq \text{rank}(A_{K_n}) = n - 1$

The only linear dependence among the rows is

$$1A_{K_n} = 0.$$

1. G is connected $\Leftrightarrow$

Why G is connected iff it is acyclic

For every G $\mathrm{rank}(A_G) \leq \mathrm{rank}(A_{K_n}) = n - 1$

The only linear dependence among the rows is

$$1A_{K_n} = 0.$$

1. G is connected $\Leftrightarrow$ A_G has a trivial left kernel.

Why G is connected iff it is acyclic

For every G $\text{rank}(A_G) \leq \text{rank}(A_{K_n}) = n - 1$

The only linear dependence among the rows is

$$1A_{K_n} = 0.$$

1. G is connected $\Leftrightarrow$ A_G has a trivial left kernel.
2. G is acyclic $\Leftrightarrow$

Why G is connected iff it is acyclic

For every G $\text{rank}(A_G) \leq \text{rank}(A_{K_n}) = n - 1$

The only linear dependence among the rows is

$$1A_{K_n} = 0.$$

1. G is connected $\Leftrightarrow$ A_G has a trivial left kernel.
2. G is acyclic $\Leftrightarrow$ A_G has a zero right kernel.

Why G is connected iff it is acyclic

For every G $\text{rank}(A_G) \leq \text{rank}(A_{K_n}) = n - 1$

The only linear dependence among the rows is

$$1A_{K_n} = 0.$$

1. G is connected $\Leftrightarrow$ A_G has a trivial left kernel.
2. G is acyclic $\Leftrightarrow$ A_G has a zero right kernel.
3. The $n-1$ columns of A_G are linearly independent.

Collapsibility

An elementary collapse is a step where you remove a vertex of degree one and the single edge that contains it.

Collapsibility

An elementary collapse is a step where you remove a vertex of degree one and the single edge that contains it.

A graph G is collapsible if by repeated application of elementary collapses you can eliminate all of the edges in G.

Collapsing - a linear algebra perspective

Let A_G be the incidence matrix of graph G. In an elementary collapse we erase row i and column e of A_G where the (i, e) entry is the only nonzero entry in the i-th row.

Collapsing - a linear algebra perspective

Let A_G be the incidence matrix of graph G. In an elementary collapse we erase row i and column e of A_G where the (i, e) entry is the only nonzero entry in the i-th row. Recall: e is the one and only edge incident with vertex i.

Collapsing - a linear algebra perspective

Let A_G be the incidence matrix of graph G. In an elementary collapse we erase row i and column e of A_G where the (i, e) entry is the only nonzero entry in the i-th row. Recall: e is the one and only edge incident with vertex i.

G is collapsible if it is possible to eliminate all its columns by a series of elementary collapses.

Collapsing - a linear algebra perspective

Let A_G be the incidence matrix of graph G. In an elementary collapse we erase row i and column e of A_G where the (i, e) entry is the only nonzero entry in the i-th row. Recall: e is the one and only edge incident with vertex i.

G is collapsible if it is possible to eliminate all its columns by a series of elementary collapses.

This implies that G is acyclic - Collapsing yields a proof that the right kernel is empty.

But note

As we saw, connectivity and acyclicity are linear algebraic. In contrast collapsibility is a purely combinatorial condition.

But note

As we saw, connectivity and acyclicity are linear algebraic. In contrast collapsibility is a purely combinatorial condition.

Indeed we will soon see that in higher dimensions collapsibility implies connectivity and acyclicity, but the reverse implication does not hold.

Setting the scene

We need a high-dimensional analog of the incidence matrix.

Boundary operators of simplicial cplexes

$(d-1)$-dimensional faces vs. d-dimensional faces.

$$\partial_2 = \begin{array}{c} \\ \vdots \\ ij \\ \vdots \\ ik \\ \vdots \\ \vdots \\ jk \\ \vdots \end{array} \begin{array}{c} \begin{array}{cccccccc} \dots & \dots & ijk & \dots & \dots & \dots & \dots & \dots \end{array} \\ \begin{pmatrix} \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & +1 & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & -1 & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & +1 & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \end{array}$$

Does this suggest what a hypertree is ?

We know where to start:

Does this suggest what a hypertree is ?

We know where to start:

Q: What is the rank of ∂_d ?

Does this suggest what a hypertree is ?

We know where to start:

Q: What is the rank of ∂_d ?

A: $\binom{n-1}{d}$,

Does this suggest what a hypertree is ?

We know where to start:

Q: What is the rank of ∂_d ?

A: $\binom{n-1}{d}$, let's prove it

That $\text{rank}(\partial_d) \leq \binom{n-1}{d}$ follows from $\partial_{d-1}\partial_d = 0$.

We will show that $\text{rank}(\partial_d) \geq \binom{n-1}{d}$ by exhibiting an explicit set of $\binom{n-1}{d}$ linearly independent columns, i.e., the set of d-faces of a d-dimensional hypertree.

So, what is a d-dimensional hypertree?

It is a d-dimensional simplicial complex with

So, what is a d-dimensional hypertree?

It is a d-dimensional simplicial complex with

- A full $(d-1)$-dimensional skeleton.

So, what is a d-dimensional hypertree?

It is a d-dimensional simplicial complex with

- A full $(d-1)$-dimensional skeleton.
- It has $\binom{n-1}{d}$ d-dimensional faces.

So, what is a d-dimensional hypertree?

It is a d-dimensional simplicial complex with

- A full $(d-1)$-dimensional skeleton.
- It has $\binom{n-1}{d}$ d-dimensional faces.

Whose boundary operator ∂_d has

- a trivial left kernel.

So, what is a d-dimensional hypertree?

It is a d-dimensional simplicial complex with

- A full $(d-1)$-dimensional skeleton.
- It has $\binom{n-1}{d}$ d-dimensional faces.

Whose boundary operator ∂_d has

- a trivial left kernel.
- zero right kernel.

So, what is a d-dimensional hypertree?

It is a d-dimensional simplicial complex with

- A full $(d-1)$-dimensional skeleton.
- It has $\binom{n-1}{d}$ d-dimensional faces.

Whose boundary operator ∂_d has

- a trivial left kernel.
- zero right kernel.
- The $\binom{n-1}{d}$ columns of its ∂_d span the columns of the matrix of all $(d-1)$-faces vs. all d-faces.

Can you please show an example of hypertree?

Arguably the simplest one-dimensional (=graphic) tree is a star, i.e., all 1-dimensional faces (=edges) that contain, say, vertex n.

Can you please show an example of hypertree?

Arguably the simplest one-dimensional (=graphic) tree is a star, i.e., all 1-dimensional faces (=edges) that contain, say, vertex n.
The same works in every dimension: Take all d-faces (=sets of size $d+1$) which contain the vertex n.
Let's see how this works.

The d-dimensional hyperstar

Recall the boundary operator ∂_d of the full d-dimensional n-vertex full complex. It is a $\binom{n}{d} \times \binom{n}{d+1}$ matrix of $\{0, -1, 1\}$.

The d-dimensional hyperstar

Recall the boundary operator ∂_d of the full d-dimensional n-vertex full complex. It is a $\binom{n}{d} \times \binom{n}{d+1}$ matrix of $\{0, -1, 1\}$.
Rows in this matrix represent $(d-1)$-dimensional faces (=sets of size d). These sets fall in two categories:

The d-dimensional hyperstar (contd.)

I: Sets which do not contain the vertex n, there are $\binom{n-1}{d}$ of those.

The d-dimensional hyperstar (contd.)

I: Sets which do not contain the vertex n, there are $\binom{n-1}{d}$ of those.

II: Those which do contain the vertex n, their number is: $\binom{n-1}{d-1}$.

It is easily verified that the rows in category I linearly span those from category II. We can therefore eliminate the latter without losing in rank.

The d-dimensional hyperstar (contd.)

Its vertex set is $V = [n]$. Every subset of V of size $\leq d$ is a face (its $(d-1)$-skeleton is full). A set of size $d+1$ is a d-dimensional face iff it contains the vertex n.

The d-dimensional hyperstar (contd.)

Its vertex set is $V = [n]$. Every subset of V of size $\leq d$ is a face (its $(d-1)$-skeleton is full). A set of size $d+1$ is a d-dimensional face iff it contains the vertex n. Why is this a hypertree?

The d-dimensional hyperstar (contd.)

Its vertex set is $V = [n]$. Every subset of V of size $\leq d$ is a face (its $(d-1)$-skeleton is full). A set of size $d+1$ is a d-dimensional face iff it contains the vertex n. Why is this a hypertree?
Because after category II rows are eliminated, what remains from the matrix ∂_d is simply the identity matrix - clearly a full rank matrix. In the row corresponding to set S, the only nonzero entry is in the column corresponding to $S \dot{\cup} \{n\}$.

Collapsibility in higher dimensions

Let X be a d-dimensional complex.

Collapsibility in higher dimensions

Let X be a d-dimensional complex.
If some $(d-1)$-dimensional face τ is contained in a unique d-dimensional face σ, then the corresponding elementary collapse is to eliminate both τ and σ from X.

Collapsibility in higher dimensions

Let X be a d-dimensional complex.
If some $(d-1)$-dimensional face τ is contained in a unique d-dimensional face σ, then the corresponding elementary collapse is to eliminate both τ and σ from X.
X is d-collapsible if it is possible to eliminate all its d-faces by a series of elementary collapses.

Collapsibility in higher dimensions

Let X be a d-dimensional complex.
If some $(d-1)$-dimensional face τ is contained in a unique d-dimensional face σ, then the corresponding elementary collapse is to eliminate both τ and σ from X.
X is d-collapsible if it is possible to eliminate all its d-faces by a series of elementary collapses.
The linear algebra perspective of collapsibility shows that it implies acyclicity.

Collapsibility in higher dimensions

Let X be a d-dimensional complex.
If some $(d-1)$-dimensional face τ is contained in a unique d-dimensional face σ, then the corresponding elementary collapse is to eliminate both τ and σ from X.
X is d-collapsible if it is possible to eliminate all its d-faces by a series of elementary collapses.
The linear algebra perspective of collapsibility shows that it implies acyclicity. But....

A little surprise

$\binom{6-1}{2} = 10$

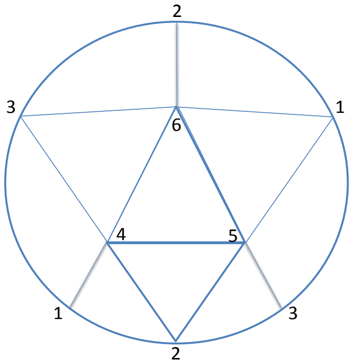

Figure: A triangulation of the projective plane

A little surprise

This example shows (at least) two things:

A little surprise

This example shows (at least) two things:
Unlike the 1-dimensional case of graphs, the definition of a d-dimensional hypertree depends on the underlying field.

A little surprise

This example shows (at least) two things:
Unlike the 1-dimensional case of graphs, the definition of a d-dimensional hypertree depends on the underlying field.

Indeed: The 6-point triangulation of the projective plane is a $\mathbb{Q}$-hypertree, but not a $\mathbb{F}_2$-hypertree.

In dimension $d \geq 2$ d-collapsibility is a stronger condition than being a hypertree.

In dimension $d \geq 2$ d-collapsibility is a stronger condition than being a hypertree.

In fact, already Kalai posed

In dimension $d \geq 2$ d-collapsibility is a stronger condition than being a hypertree.

In fact, already Kalai posed

Conjecture

For every $d \geq 2$ and for every field $\mathbb{F}$ and $n \to \infty$ almost none of the n-vertex d-dimensional $\mathbb{F}$-hypertrees are collapsible.

This remains open, and is supported by rigorous numerical experiments.

If so...

Q: Can you, at least, come up with more examples of non-collapsible hypertrees?

If so...

Q: Can you, at least, come up with more examples of non-collapsible hypertrees?

A construction: Let n be prime and $d \geq 2$. Fix a subset $A \subset \mathbb{Z}_n$ of cardinality $|A| = d + 1$. The sum complex X_A corresponding to A has a full $(d-1)$-dimensional skeleton and contains a d-face σ iff $\sum_{x \in \sigma} x \in A$.

If so...

Q: Can you, at least, come up with more examples of non-collapsible hypertrees?

A construction: Let n be prime and $d \geq 2$. Fix a subset $A \subset \mathbb{Z}_n$ of cardinality $|A| = d + 1$. The sum complex X_A corresponding to A has a full $(d-1)$-dimensional skeleton and contains a d-face σ iff $\sum_{x \in \sigma} x \in A$.

Theorem (L., Meshulam, Rosenthal)

The complex X_A is always a $\mathbb{Q}$-hypertree.

If so...

Q: Can you, at least, come up with more examples of non-collapsible hypertrees?

A construction: Let n be prime and $d \geq 2$. Fix a subset $A \subset \mathbb{Z}_n$ of cardinality $|A| = d + 1$. The sum complex X_A corresponding to A has a full $(d-1)$-dimensional skeleton and contains a d-face σ iff $\sum_{x \in \sigma} x \in A$.

Theorem (L., Meshulam, Rosenthal)

The complex X_A is always a $\mathbb{Q}$-hypertree. It is collapsible iff A forms an arithmetic progression.

Cayley-Kalai Formula

Theorem (Cayley's Formula, Borchardt 1860)

The number of trees with vertex set $[n]$ is n^{n-2}.

Theorem (Kalai 1983)

$$\sum_{T} |H_{d-1}(T)|^2 = n^{\binom{n-2}{d}}$$

where the sum is over all n-vertex d-dimensional hypertrees T.

But how many d-hypertrees are there?

Open Problem

For $d \geq 2$ and large n, find (at least approximately) the number of d-dimensional n-vertex $\mathbb{Q}$-hypertrees.

Kalai's Formula yields estimates, but falls short of an asymptotic formula. In joint work with Y. Peled these estimates were significantly improved, though a full answer still eludes us.

But how many d-hypertrees are there?

Open Problem

For $d \geq 2$ and large n, find (at least approximately) the number of d-dimensional n-vertex $\mathbb{Q}$-hypertrees.

Kalai's Formula yields estimates, but falls short of an asymptotic formula. In joint work with Y. Peled these estimates were significantly improved, though a full answer still eludes us.
Here is the strategy that we used to derive our bounds.

A random process

We consider the random process that starts with a full $(d-1)$-dimensional skeleton. At each step pick a random d-dimensional face $\sigma \notin X$. If possible, we add σ to X. Otherwise, we discard σ.

A random process

We consider the random process that starts with a full $(d-1)$-dimensional skeleton. At each step pick a random d-dimensional face $\sigma \notin X$. If possible, we add σ to X. Otherwise, we discard σ.

We cannot add σ to X iff this creates a new cycle. In this case we say that σ is in the shade of X.

To wit: At each step we add to the current complex a random d-face σ whose addition creates no new cycle ("σ is not in the shade of X").

Living in the shades

Let $G = (V, E)$ be a disconnected graph, and let $ij \notin E$. We say that ij is in G's shadow if i and j belong to the same connected component of G.

Living in the shades

Let $G = (V, E)$ be a disconnected graph, and let $ij \notin E$. We say that ij is in G's shadow if i and j belong to the same connected component of G. In other words ij is in G's shadow iff the column corresponding to the edge ij is in the linear span of the columns of A_G.

Outside the shadow of an evolving graph

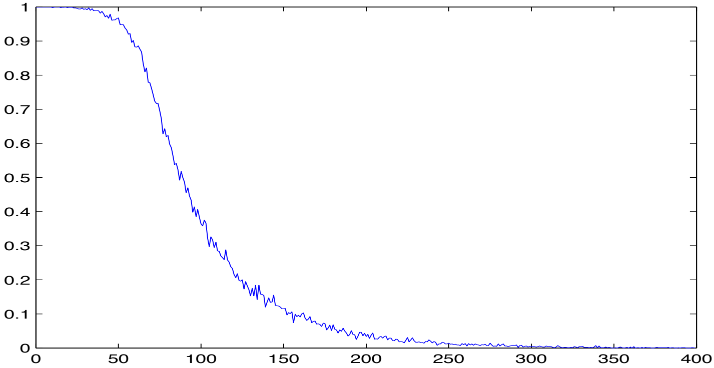

Shadow of an evolving 2-complex

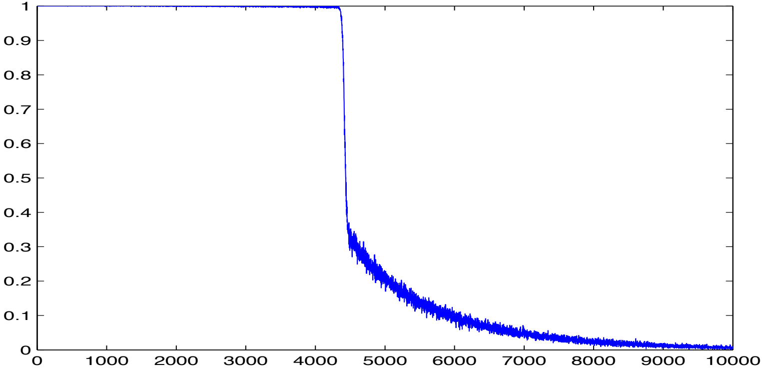

A lower bound on the number of hypertrees: A taste of the proof

Claim: If X is an acyclic d-complex, then the number of d-faces in its shadow Y is $\leq \frac{n \cdot |X|}{d+1}$

A lower bound on the number of hypertrees: A taste of the proof

Claim: If X is an acyclic d-complex, then the number of d-faces in its shadow Y is $\leq \frac{n \cdot |X|}{d+1}$
There are exactly $(d+1) \cdot |Y|$ pairs (v, σ) with v a vertex in σ, a d-face in Y. Let W be the set of d-faces in Y that contain v.

A lower bound on the number of hypertrees: A taste of the proof

Claim: If X is an acyclic d-complex, then the number of d-faces in its shadow Y is $\leq \frac{n \cdot |X|}{d+1}$
There are exactly $(d+1) \cdot |Y|$ pairs (v, σ) with v a vertex in σ, a d-face in Y. Let W be the set of d-faces in Y that contain v.
W is an acyclic complex, being part of v's hyperstar. So, the columns corresponding to W are linearly independent and spanned by X. Therefore $|W| \leq |X|$, which proves our claim.

A little context - $G(n, p)$

This is the grandfather of all models of random graphs. Introduced by Erdős and Rényi in the 60's, a mainstay of modern combinatorics and still an important source of ideas and inspiration.

A little context - $G(n, p)$

This is the grandfather of all models of random graphs. Introduced by Erdős and Rényi in the 60's, a mainstay of modern combinatorics and still an important source of ideas and inspiration.
Start with n vertices.

A little context - $G(n, p)$

This is the grandfather of all models of random graphs. Introduced by Erdős and Rényi in the 60's, a mainstay of modern combinatorics and still an important source of ideas and inspiration.
Start with n vertices. For each pair of vertices $\{x, y\}$ independently with probability p include $e = xy$ in the random graph that you generate.

A little context - $G(n, p)$

This is the grandfather of all models of random graphs. Introduced by Erdős and Rényi in the 60's, a mainstay of modern combinatorics and still an important source of ideas and inspiration.

Start with n vertices. For each pair of vertices $\{x, y\}$ independently with probability p include $e = xy$ in the random graph that you generate.

Closely related model: the evolution of random graphs starts with n vertices and no edges. At each step add a random edge to the evolving graph.

A d-dimensional analog of $G(n, p)$

About 15 years ago, Roy Meshulam and I introduced a model of a random d-dimensional n-vertex complex $X_d(n, p)$. In dimension $d = 1$ the $X_1(n, p)$ model coincides with $G(n, p)$.

A d-dimensional analog of $G(n, p)$

About 15 years ago, Roy Meshulam and I introduced a model of a random d-dimensional n-vertex complex $X_d(n, p)$. In dimension $d = 1$ the $X_1(n, p)$ model coincides with $G(n, p)$.
Start with a full $(d-1)$-dimensional skeleon. (In the case of graphs - start with n vertices.)

A d-dimensional analog of $G(n,p)$

About 15 years ago, Roy Meshulam and I introduced a model of a random d-dimensional n-vertex complex $X_d(n,p)$. In dimension $d = 1$ the $X_1(n,p)$ model coincides with $G(n,p)$.
Start with a full $(d-1)$-dimensional skeleon. (In the case of graphs - start with n vertices.)
For each d-dimensional face σ, independently and with probability p, decide whether $\sigma \in X$. (For graphs - same with every edge).

Some basic facts in $G(n, p)$ theory

Theorem (Erdős and Rényi '60)

The threshold for graph connectivity in $G(n, p)$ is

$$p = \frac{\ln n}{n}$$

Some basic facts in $G(n,p)$ theory

Theorem (Erdős and Rényi '60)

The threshold for graph connectivity in $G(n,p)$ is

$$p = \frac{\ln n}{n}$$

E.g., if $p \leq (1-\epsilon)\frac{\ln n}{n}$, then whp a graph in $G(n,p)$ is disconnected.

Some basic facts in $G(n, p)$ theory

Theorem (Erdős and Rényi '60)

The threshold for graph connectivity in $G(n, p)$ is

$$p = \frac{\ln n}{n}$$

E.g., if $p \leq (1 - \epsilon)\frac{\ln n}{n}$, then whp a graph in $G(n, p)$ is disconnected.
Whereas if $p \geq (1 + \epsilon)\frac{\ln n}{n}$, whp a graph in $G(n, p)$ is connected.

Spelling "Connectivity" in high dimension

The boundary operator ∂_d has a trivial left kernel.

Theorem (L. - Meshulam - Wallach)

The threshold for connectivity of $X_d(n, p)$ is

$$p = \frac{d \ln n}{n}.$$

Spelling "Connectivity" in high dimension

The boundary operator ∂_d has a trivial left kernel.

Theorem (L. - Meshulam - Wallach)

The threshold for connectivity of $X_d(n, p)$ is

$$p = \frac{d \ln n}{n}.$$

Specifically, whp, left kernel$(\partial_d(X))$ is

Spelling "Connectivity" in high dimension

The boundary operator ∂_d has a trivial left kernel.

Theorem (L. - Meshulam - Wallach)

The threshold for connectivity of $X_d(n, p)$ is

$$p = \frac{d \ln n}{n}.$$

Specifically, whp, left kernel$(\partial_d(X))$ is

- *nontrivial for $p < (1 - \epsilon)\frac{d \ln n}{n}$, and*
- *trivial for $p > (1 + \epsilon)\frac{d \ln n}{n}$.*

Phase transition in $G(n, p)$ theory

Erdős and Rényi's most dramatic discovery is the phase transition in the evolution of random graphs.

Phase transition in $G(n, p)$ theory

Erdős and Rényi's most dramatic discovery is the phase transition in the evolution of random graphs.

Start with n isolated vertices and sequentially add a new random edge, one at a time.

Phase transition in $G(n,p)$ theory

Erdős and Rényi's most dramatic discovery is the phase transition in the evolution of random graphs.

Start with n isolated vertices and sequentially add a new random edge, one at a time. Initially every edge is isolated. Later, small and simple connected components appear.

Phase transition in $G(n, p)$ theory

Erdős and Rényi's most dramatic discovery is the phase transition in the evolution of random graphs.

Start with n isolated vertices and sequentially add a new random edge, one at a time. Initially every edge is isolated. Later, small and simple connected components appear.

- small = cardinality $O(\log n)$.

Phase transition in $G(n,p)$ theory

Erdős and Rényi's most dramatic discovery is the phase transition in the evolution of random graphs.

Start with n isolated vertices and sequentially add a new random edge, one at a time. Initially every edge is isolated. Later, small and simple connected components appear.

- small = cardinality $O(\log n)$.
- simple = a tree.

Phase transition in $G(n, p)$ theory

Erdős and Rényi's most dramatic discovery is the phase transition in the evolution of random graphs.

Start with n isolated vertices and sequentially add a new random edge, one at a time. Initially every edge is isolated. Later, small and simple connected components appear.

- small = cardinality $O(\log n)$.
- simple = a tree.
- Plus a Poisson number of unicylic graphs with $O(\log n)$ vertices.

Crescendo - The phase transition

Around step $\frac{n}{2}$ and over a very short period of time

Crescendo - The phase transition

Around step $\frac{n}{2}$ and over a very short period of time a giant component emerges.

Crescendo - The phase transition

Around step $\frac{n}{2}$ and over a very short period of time a giant component emerges.

giant$=$ cardinality $\Omega(n)$.

Crescendo - The phase transition

Around step $\frac{n}{2}$ and over a very short period of time a giant component emerges.

giant$=$ cardinality $\Omega(n)$.

Note: Time $\frac{n}{2}$ corresponds to $p = \frac{1}{n}$.

In the wake of the revolution

Around step $\frac{n}{2}$ many sudden changes occur.

In the wake of the revolution

Around step $\frac{n}{2}$ many sudden changes occur. E.g., for $p < \frac{1-\epsilon}{n}$, the probability that the graph is acyclic is bounded away from both zero and one.

In the wake of the revolution

Around step $\frac{n}{2}$ many sudden changes occur. E.g., for $p < \frac{1-\epsilon}{n}$, the probability that the graph is acyclic is bounded away from both zero and one.

But if $p > \frac{1+\epsilon}{n}$, G almost surely contains a cycle.

Meanwhile in high dimensions.....

There are at least two high-dimensional analogs of the forest/non-forest dichotomy in graphs.

Meanwhile in high dimensions.....

There are at least two high-dimensional analogs of the forest/non-forest dichotomy in graphs.

- Collapsible/non-collapsible complex.

Meanwhile in high dimensions.....

There are at least two high-dimensional analogs of the forest/non-forest dichotomy in graphs.

- Collapsible/non-collapsible complex.
- Acyclic/cyclic: right kernel$(\partial_d) \stackrel{?}{=} 0$.

Meanwhile in high dimensions.....

There are at least two high-dimensional analogs of the forest/non-forest dichotomy in graphs.

- Collapsible/non-collapsible complex.
- Acyclic/cyclic: right kernel$(\partial_d) \stackrel{?}{=} 0$.

Recall: collapsible complexes are acyclic, so clearly

$$p_{collapse} \leq p_{acyclic}$$

Meanwhile in high dimensions.....

There are at least two high-dimensional analogs of the forest/non-forest dichotomy in graphs.

- Collapsible/non-collapsible complex.
- Acyclic/cyclic: right kernel(∂_d) $\stackrel{?}{=} 0$.

Recall: collapsible complexes are acyclic, so clearly

$$p_{collapse} \leq p_{acyclic}$$

This inequality is strict.

Concretely

Theorem (Lior Aronshtam, L., Tomasz Łuczak, Roy Meshulam, Yuval Peled)

- *The collapsibility threshold in* $X_d(n, p)$ *is*

$$(1 + o_d(1))\frac{\log d}{n}.$$

Concretely

Theorem (Lior Aronshtam, L., Tomasz Łuczak, Roy Meshulam, Yuval Peled)

- *The collapsibility threshold in $X_d(n, p)$ is*

$$(1 + o_d(1))\frac{\log d}{n}.$$

- *The threshold for having a cycle whp is*

$$\frac{d + 1 - o_d(1)}{n}.$$

Resolving the remaining major difficulty

We have no notion of a high-dimensional connected component

Theorem (L., Y. Peled)

Exactly at the same $p = \frac{c}{n}$ *where* $X_d(n,p)$ *almost surely acquires a cycle, the shadow of the complex becomes gigantic* ($=\Omega(n^{d+1})$ *faces).*

Resolving the remaining major difficulty

We have no notion of a high-dimensional connected component

Theorem (L., Y. Peled)

Exactly at the same $p = \frac{c}{n}$ where $X_d(n, p)$ almost surely acquires a cycle, the shadow of the complex becomes gigantic ($=\Omega(n^{d+1})$ faces).

This statement applies in all dimensions. However, when $d = 1$ the limit distribution is continuous but not smooth, while for $d \geq 2$ the limit distribution is discontinuous.

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

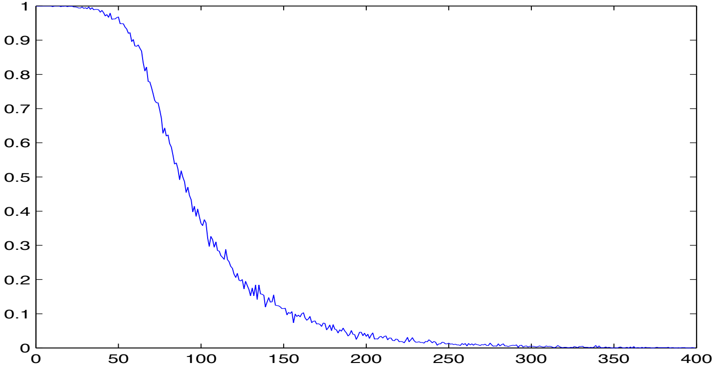

Phase transition in $X_2(n, p)$ complexes

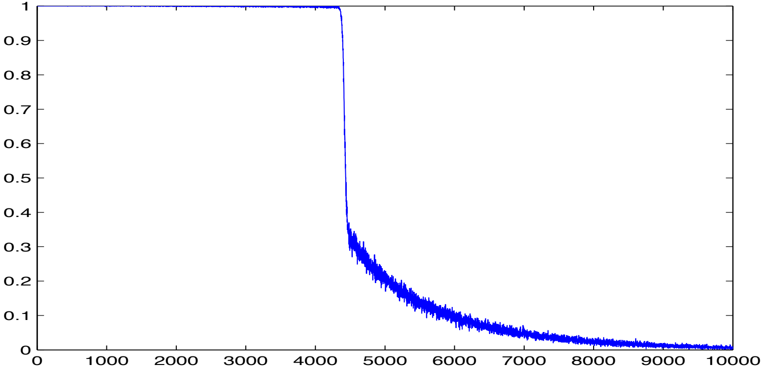

More surprises in the shadows

Easy Observation

Let G be an "almost tree", i.e., an n vertex forest with $n-2$ edges (and hence with two connected components). Then at least $(1-o(1))\frac{n^2}{4}$, i.e., at least half of the remaining edges, are in G's shadow.

Surprises in the shadows (contd.)

Construction: Let X be a 2-dimensional n-vertex complex with a full 1-dimensional skeleton. The 2-faces of X are the arithmetic triples of difference $\neq 1$. Easy fact: The number of 2-faces in X is $\binom{n-1}{2} - 1$ (one less than a 2-dimensional hypertree).

[1] It actually suffices to assume the weaker Artin's conjecture

Surprises in the shadows (contd.)

Construction: Let X be a 2-dimensional n-vertex complex with a full 1-dimensional skeleton. The 2-faces of X are the arithmetic triples of difference $\neq 1$. Easy fact: The number of 2-faces in X is $\binom{n-1}{2} - 1$ (one less than a 2-dimensional hypertree).

Theorem (L., Newman, Peled, Rabinovich)

The complex X is $\mathbb{Q}$-acyclic. Assuming the Riemann hypothesis[1]*, there are infinitely many primes n for which X has an empty shadow.*

[1] It actually suffices to assume the weaker Artin's conjecture

Sampling hypertrees

Recent work of Anari, Liu, Oveis Gharan and Vinzant shows that a most natural algorithm for sampling hypertrees converges in polynomial time.

Hyperpaths?

A path is a tree where every vertex is in two or fewer edges.

Can we construct d-dimensional hypertree where every $(d-1)$-dimensional face is contained in $d+1$ or fewer d-dimensional face?

Hyperpaths (contd.)

If so, how many are they? In dimension one: If P_n denotes the number of n-vertex paths and T_n is the number of n-vertex trees, then

$$\frac{P_n}{T_n} = (\frac{1}{e} + o(1))^n$$

and in higher dimension?

MAZALTOV URI