The strange geometry of high-dimensional spanning forests

Tom Hutchcroft, UBC

Joint work with Yuval Peres

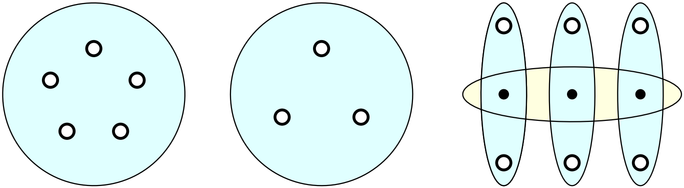

November 5, 2016

Uniform Spanning Trees

A **spanning tree** of a connected graph G is a connected subgraph of G that contains every vertex and no cycles. Given a finite connected graph G, the **uniform spanning tree** of G is simply a spanning tree of G chosen uniformly at random.

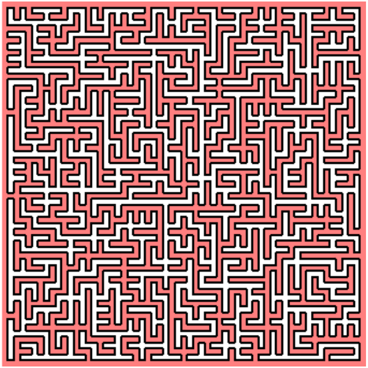

Figure: A uniform spanning tree (in white) of a 30x30 box in the square grid.

Connections

Uniform spanning trees are also closely related to many other interesting things in probability and elsewhere, for instance:

1. Loop-erased random walk.
2. Conformally invariant scaling limits.
3. The Abelian sandpile model.
4. Domino tiling.
5. Random cluster model.
6. Random interlacements.
7. Potential theory.
8. ℓ_2-Betti numbers and the fixed price problem of Gaboriau.
9. Rotor Routers.

Uniform Spanning Forests

The **uniform spanning forests** are extensions of the uniform spanning tree to infinite graphs, first studied by Pemantle ('91).

They are defined as limits of the uniform spanning trees on large finite induced subgraphs of an infinite graph G.

Uniform Spanning Forests

The **uniform spanning forests** are extensions of the uniform spanning tree to infinite graphs, first studied by Pemantle ('91).

They are defined as limits of the uniform spanning trees on large finite induced subgraphs of an infinite graph G.

Exhaustion of G: increasing sequence $\langle V_n \rangle_{n\geq 1}$ of finite sets $V_n \subset V$ with $\bigcup V_n = V$. Let G_n be the subgraph of G induced by V_n.

Uniform Spanning Forests

The **uniform spanning forests** are extensions of the uniform spanning tree to infinite graphs, first studied by Pemantle ('91).

They are defined as limits of the uniform spanning trees on large finite induced subgraphs of an infinite graph G.

Exhaustion of G: increasing sequence $\langle V_n \rangle_{n \geq 1}$ of finite sets $V_n \subset V$ with $\bigcup V_n = V$. Let G_n be the subgraph of G induced by V_n.

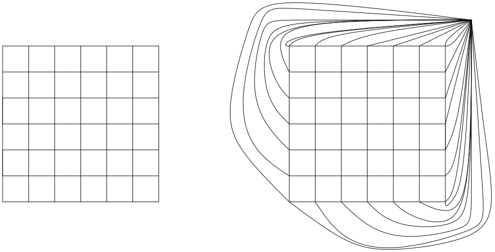

For each n, we also form a graph G_n^* from G by identifying (wiring) every vertex in $G \setminus G_n$ into a single vertex.

Uniform Spanning Forests

Free uniform spanning forest (FUSF): Distributional limit of the uniform spanning trees of G_n: for every finite $S \subseteq E$,

$$\mathrm{FUSF}_G(S \subseteq \mathfrak{F}) = \lim_{n\to\infty} \mathrm{UST}_{G_n}(S \subseteq T).$$

Uniform Spanning Forests

Free uniform spanning forest (FUSF): Distributional limit of the uniform spanning trees of G_n: for every finite $S \subseteq E$,

$$\mathsf{FUSF}_G(S \subseteq \mathfrak{F}) = \lim_{n \to \infty} \mathsf{UST}_{G_n}(S \subseteq T).$$

Wired uniform spanning forest (WUSF): Distributional limit

$$\mathsf{WUSF}_G(S \subseteq \mathfrak{F}) = \lim_{n \to \infty} \mathsf{UST}_{G_n^*}(S \subseteq T).$$

Both limits were shown to exist by Pemantle (and are independent of the choice of exhaustion).

Uniform Spanning Forests

Free uniform spanning forest (FUSF): Distributional limit of the uniform spanning trees of G_n: for every finite $S \subseteq E$,

$$\mathsf{FUSF}_G(S \subseteq \mathfrak{F}) = \lim_{n\to\infty} \mathsf{UST}_{G_n}(S \subseteq T).$$

Wired uniform spanning forest (WUSF): Distributional limit

$$\mathsf{WUSF}_G(S \subseteq \mathfrak{F}) = \lim_{n\to\infty} \mathsf{UST}_{G_n^*}(S \subseteq T).$$

Both limits were shown to exist by Pemantle (and are independent of the choice of exhaustion).

In $\mathbb{Z}^d$ (and other amenable transitive graphs) the two forests coincide – we'll call them simply the USF.

Connectivity

Although they are defined as limits of trees, the USFs need not be connected.

Connectivity

Although they are defined as limits of trees, the USFs need not be connected.

Pemantle ('91): The USF of $\mathbb{Z}^d$ is connected (i.e., a single tree) a.s. if and only if $d \leq 4$.

Connectivity

Although they are defined as limits of trees, the USFs need not be connected.

Pemantle ('91): The USF of $\mathbb{Z}^d$ is connected (i.e., a single tree) a.s. if and only if $d \leq 4$.

What's going on? Why is dimension four important?

Connectivity

Although they are defined as limits of trees, the USFs need not be connected.

Pemantle ('91): The USF of $\mathbb{Z}^d$ is connected (i.e., a single tree) a.s. if and only if $d \leq 4$.

What's going on? Why is dimension four important? The key to understanding this is **Wilson's algorithm**.

Connectivity

Although they are defined as limits of trees, the USFs need not be connected.

Pemantle ('91): The USF of $\mathbb{Z}^d$ is connected (i.e., a single tree) a.s. if and only if $d \leq 4$.

What's going on? Why is dimension four important? The key to understanding this is **Wilson's algorithm**.

In 1996, David Wilson introduced an algorithm that lets us build a uniform spanning tree of a finite graph G out of loop-erased random walks.

This algorithm was extended to generate the WUSF of infinite, transient graphs by Benjamini, Lyons, Peres and Schramm (BLPS) in 2001.

Loop-erased Random Walk

The **loop-erasure** of a path is defined by deleting loops from the path chronologically as they are created.

Loop-erased Random Walk

The **loop-erasure** of a path is defined by deleting loops from the path chronologically as they are created.

Formally, given a path $\langle \gamma_i \rangle_i$, we define a sequence of times T_i by setting $T_0 = 0$ and, given T_i, setting T_{i+1} to be $1+$ the last time γ returns to γ_{T_i}, and define the loop-erasure of γ to be $\langle \gamma_{T_i} \rangle_i$.

Loop-erased Random Walk

The **loop-erasure** of a path is defined by deleting loops from the path chronologically as they are created.

Formally, given a path $\langle \gamma_i \rangle_i$, we define a sequence of times T_i by setting $T_0 = 0$ and, given T_i, setting T_{i+1} to be $1+$ the last time γ returns to γ_{T_i}, and define the loop-erasure of γ to be $\langle \gamma_{T_i} \rangle_i$.

The loop-erasure of a random walk is called **loop-erased random walk** (LERW) and was first studied by Lawler in 1980.

Wilson's Algorithm

Let G be an infinite transient graph. Let $\langle v_i \rangle_{i \geq 1}$ be an enumeration of the vertices of G. We define a sequence of forests F_i as follows:

- F_0 is the empty forest, that has no vertices or edges.

Wilson's Algorithm

Let G be an infinite transient graph. Let $\langle v_i \rangle_{i \geq 1}$ be an enumeration of the vertices of G. We define a sequence of forests F_i as follows:

- F_0 is the empty forest, that has no vertices or edges.
- Given F_i, start a random walk at v_{i+1}, independent of everything we have done so far. Stop this random walk if and when it hits the set of vertices already included in F_i. If this never happens, the walk runs forever.

Wilson's Algorithm

Let G be an infinite transient graph. Let $\langle v_i \rangle_{i \geq 1}$ be an enumeration of the vertices of G. We define a sequence of forests F_i as follows:

- F_0 is the empty forest, that has no vertices or edges.
- Given F_i, start a random walk at v_{i+1}, independent of everything we have done so far. Stop this random walk if and when it hits the set of vertices already included in F_i. If this never happens, the walk runs forever.
- Take the loop-erasure of this stopped random walk path, and let F_{i+1} be the union of F_i with this loop-erased path.

Wilson's Algorithm

Let G be an infinite transient graph. Let $\langle v_i \rangle_{i \geq 1}$ be an enumeration of the vertices of G. We define a sequence of forests F_i as follows:

- F_0 is the empty forest, that has no vertices or edges.
- Given F_i, start a random walk at v_{i+1}, independent of everything we have done so far. Stop this random walk if and when it hits the set of vertices already included in F_i. If this never happens, the walk runs forever.
- Take the loop-erasure of this stopped random walk path, and let F_{i+1} be the union of F_i with this loop-erased path.

Wilson '96, BLPS '01: The forest $F = \bigcup_{i \geq 0} F_i$ is the WUSF of G.

When we run Wilson's algorithm, a new component of the forest is created exactly at those steps when the walk does not hit the set of vertices already included in the forest.

When we run Wilson's algorithm, a new component of the forest is created exactly at those steps when the walk does not hit the set of vertices already included in the forest.

It follows that the WUSF of a transient graph is connected if and only if a random walk intersects an independent loop erased random walk almost surely.

Lyons, Peres, and Schramm '01: Let G be a transient graph and let u, v be vertices of G. Let X and Y be independent random walks started at u and v, and let L be the loop-erasure of Y. Then

$$\mathbb{P}(X \text{ and } L \text{ intersect}) \geq 2^{-8}\ \mathbb{P}(X \text{ and } Y \text{ intersect}).$$

Putting this together, we have a nice criterion for connectivity of the WUSF.

   

Benjamini, Lyons, Peres and Schramm (BLPS) ('01): the WUSF of a graph is almost surely connected if and only if two independent simple random walks on the graph intersect almost surely.

Note that if two independent random walks intersect almost surely, then in fact they intersect infinitely often almost surely.

It's not too hard to apply this to recover Pemantle's theorem.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

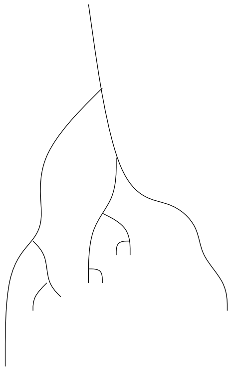

For $d \leq 4$, there is only one tree in the forest.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

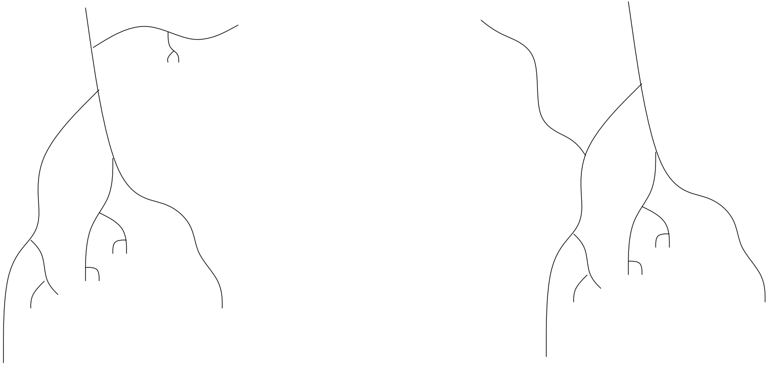

For $d \geq 5$, there are infinitely many trees in the forest.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

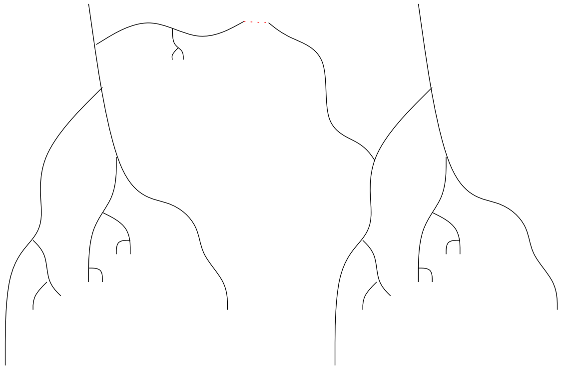

For $d \geq 5$, there are infinitely many trees in the forest. However, if $5 \leq d \leq 8$, then **every tree touches every other tree**.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

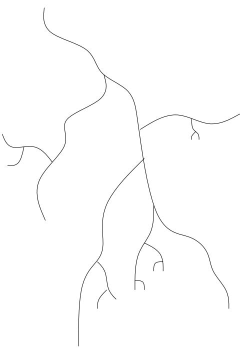

For $d \geq 9$, **there are trees in the forest that do not touch**.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

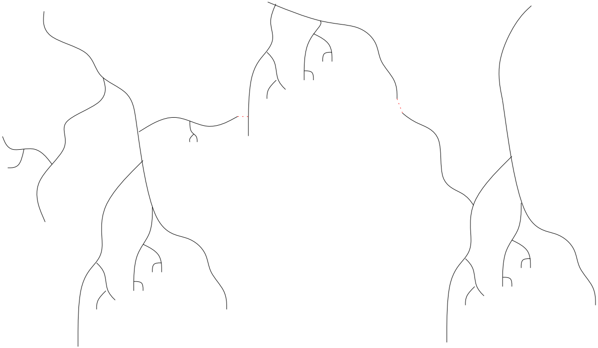

For $d \geq 9$, **there are trees in the forest that do not touch**. However, if $9 \leq d \leq 12$, then **for every two trees, there is a third 'intermediary' tree that touches both trees**.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

For $d \geq 13$, **there are trees in the forest for which we cannot find an intermediary tree**.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

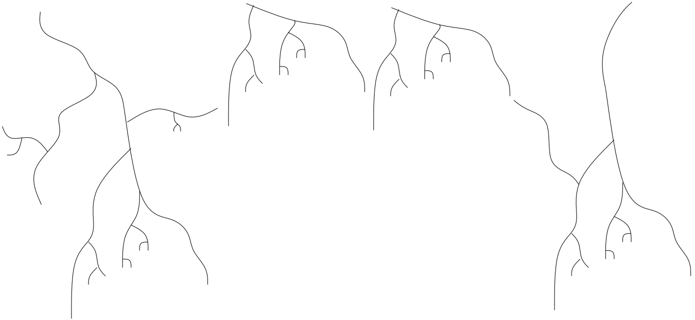

For $d \geq 13$, **there are trees in the forest for which we cannot find an intermediary tree**. However, if $13 \leq d \leq 16$, then **for every two trees, there is a pair of intermediary trees forming a path from one tree to the other**.

Transitions in dimensions $4, 8, 12, \ldots$

In 2004, Benjamini, Kesten, Peres, and Schramm discovered that the connectivity/disconnectivity transition in dimension four is merely the first of an infinite family of transitions, occuring every four dimensions.

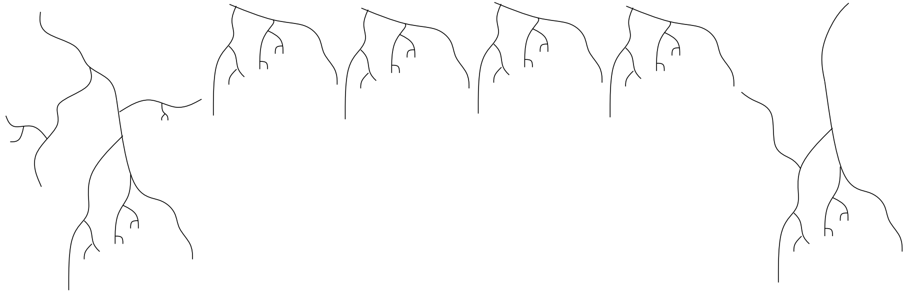

This pattern continues: **the number of intermediary trees required to connect two trees increases by one every time the dimension increases by four.**

Transitions every dimension!

Now for the new result:

Theorem (H. and Peres 2015)

BKPS is not the whole story. In fact, **the uniform spanning forest of $\mathbb{Z}^d$ undergoes qualitative changes to its connectivity properties every time the dimension increases (above four), not just every four dimensions.**

Transitions every dimension!

Now for the new result:

Theorem (H. and Peres 2015)

BKPS is not the whole story. In fact, **the uniform spanning forest of $\mathbb{Z}^d$ undergoes qualitative changes to its connectivity properties every time the dimension increases (above four), not just every four dimensions.**

The formal result will be phrased (at least for $d \geq 9$) in terms of *ubiquitous subgraphs* in the *component graph*.

Differences between $d = 9, 10, 11$

Suppose we take *three* trees.

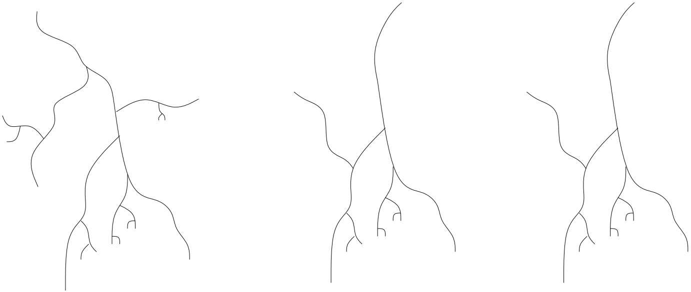

Differences between $d = 9, 10, 11$

Suppose we take *three* trees.

For $d \leq 10$, **we can always find a fourth tree that touches all three of the trees**.

Differences between $d = 9, 10, 11$

Suppose we take *three* trees.

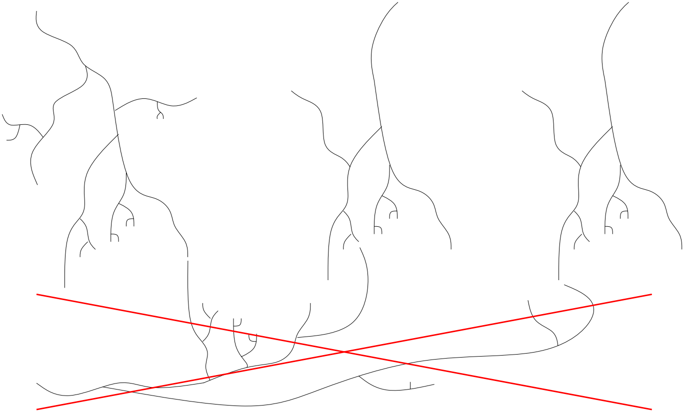

For $d \leq 10$, **we can always find a fourth tree that touches all three of the trees**. However, if $d \geq 11$, **then there are triples of trees for which such a fourth tree does not exist.**

Differences between $d = 9, 10, 11$

What if we take *four* trees?

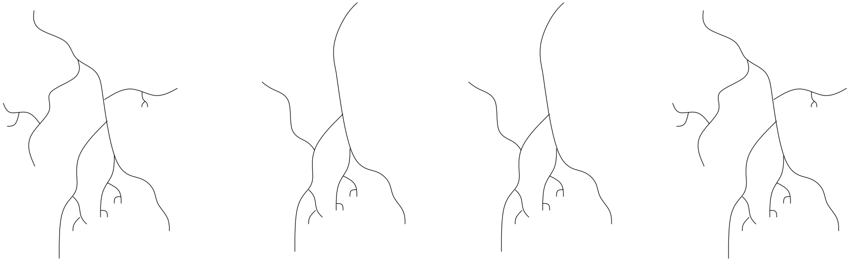

Differences between $d = 9, 10, 11$

What if we take *four* trees?

We can always find a fifth tree that touches all four of the trees if and only if $d \leq 9$.

Differences between $d = 9, 10, 11$

For *five* trees, we can still find a tree touching all of them in nine dimensions.

This is the most we can do when $d = 9$: There exist collections of six trees in nine dimensions that do not have a common neighbour.

The component graph

Let's first phrase the BKPS result in a different way.

We define the **component graph**, $\mathcal{C}(F)$, of the USF F of $\mathbb{Z}^d$ to be the graph that the trees of the USF as vertices, and has an edge between two vertices if and only if they touch.

The component graph

Let's first phrase the BKPS result in a different way.

We define the **component graph**, $\mathcal{C}(F)$, of the USF F of $\mathbb{Z}^d$ to be the graph that the trees of the USF as vertices, and has an edge between two vertices if and only if they touch.

More generally, we can consider the range r component graph, $\mathcal{C}_r(F)$, where two trees share an edge if they come within distance r of each other in the graph.

The component graph

Let's first phrase the BKPS result in a different way.

We define the **component graph**, $\mathcal{C}(F)$, of the USF F of $\mathbb{Z}^d$ to be the graph that the trees of the USF as vertices, and has an edge between two vertices if and only if they touch.

More generally, we can consider the range r component graph, $\mathcal{C}_r(F)$, where two trees share an edge if they come within distance r of each other in the graph.

BKPS: $C_r(F)$ has diameter $\lfloor (d-1)/4 \rfloor$

In particular, $C_r(F)$ is a single vertex for $1 \leq d \leq 4$, and is a complete graph with infinitely many vertices when $5 \leq d \leq 8$.

The component graph

Let's first phrase the BKPS result in a different way.

We define the **component graph**, $\mathcal{C}(F)$, of the USF F of $\mathbb{Z}^d$ to be the graph that the trees of the USF as vertices, and has an edge between two vertices if and only if they touch.

More generally, we can consider the range r component graph, $\mathcal{C}_r(F)$, where two trees share an edge if they come within distance r of each other in the graph.

BKPS: $C_r(F)$ has diameter $\lfloor (d-1)/4 \rfloor$

In particular, $C_r(F)$ is a single vertex for $1 \leq d \leq 4$, and is a complete graph with infinitely many vertices when $5 \leq d \leq 8$.

Question: Beyond BKPS, is there anything we can say about the component graph in higher dimensions?

A **graph with boundary** is a graph $H = (\partial V, V_\circ, E)$ whose vertices are partitioned into **boundary vertices** ∂V and **interior vertices** $V_\circ$.

A **graph with boundary** is a graph $H = (\partial V, V_\circ, E)$ whose vertices are partitioned into **boundary vertices** ∂V and **interior vertices** $V_\circ$.

Suppose G is a graph, H is a graph with boundary, and $x = \{x_v : v \in \partial V\}$ is a collection of vertices of G indexed by the boundary vertices of H.

A **graph with boundary** is a graph $H = (\partial V, V_\circ, E)$ whose vertices are partitioned into **boundary vertices** ∂V and **interior vertices** $V_\circ$.

Suppose G is a graph, H is a graph with boundary, and $x = \{x_v : v \in \partial V\}$ is a collection of vertices of G indexed by the boundary vertices of H.

We say that H is **present** at x if there exist a collection $\{x_u : u \in V_\circ\}$ of vertices of G indexed by the interior vertices of H such that x_u and x_v are adjacent in G whenever u and v are adjacent in H.

We say that H is **faithfully present** at x if the vertices x_u can be chosen to be distinct.

A **graph with boundary** is a graph $H = (\partial V, V_\circ, E)$ whose vertices are partitioned into **boundary vertices** ∂V and **interior vertices** $V_\circ$.

Suppose G is a graph, H is a graph with boundary, and $x = \{x_v : v \in \partial V\}$ is a collection of vertices of G indexed by the boundary vertices of H.

We say that H is **present** at x if there exist a collection $\{x_u : u \in V_\circ\}$ of vertices of G indexed by the interior vertices of H such that x_u and x_v are adjacent in G whenever u and v are adjacent in H.

We say that H is **faithfully present** at x if the vertices x_u can be chosen to be distinct.

We say that H is **ubiquitous** in G if it is present at every x, and **faithfully ubiquitous** if it is faithfully present at every x.

Examples

A path of length n, with boundary given by the endpoints, is ubiquitous in G if and only if G has diameter at most n.

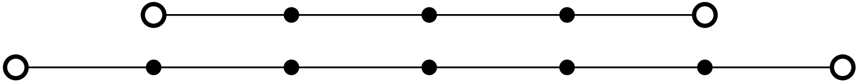

It is faithfully ubiquitous if every two vertices in the graph can be connected by a simple path of length *exactly* n.

Examples

An n-star, with its leaves as its boundary, is ubiquitous if and only if every n vertices share a common neighbour.

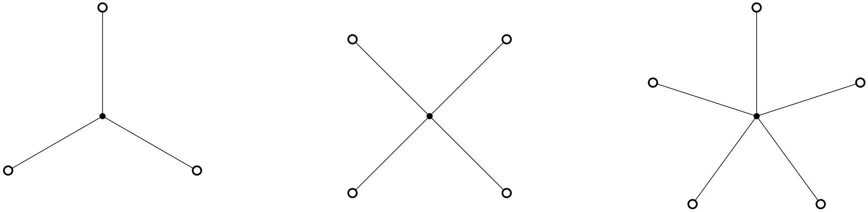

Theorem (H. and Peres 2016)

Let $d_1 > d_2 \geq 9$.

- *The choice of distance r in the component graph does not affect which finite graphs with boundary are ubiquitous. (Actually we replace $\mathbb{Z}^d$ with another transitive d-dimensional graph too.)*

Theorem (H. and Peres 2016)

Let $d_1 > d_2 \geq 9$.

- *The choice of distance r in the component graph does not affect which finite graphs with boundary are ubiquitous. (Actually we replace $\mathbb{Z}^d$ with another transitive d-dimensional graph too.)*
- *Every finite graph with boundary that is ubiquitous a.s. in the component graph of the USF of $\mathbb{Z}^{d_1}$ is also ubiquitous a.s. in the component graph of the USF of $\mathbb{Z}^{d_2}$.*

Theorem (H. and Peres 2016)

Let $d_1 > d_2 \geq 9$.

- *The choice of distance r in the component graph does not affect which finite graphs with boundary are ubiquitous. (Actually we replace $\mathbb{Z}^d$ with another transitive d-dimensional graph too.)*
- *Every finite graph with boundary that is ubiquitous a.s. in the component graph of the USF of $\mathbb{Z}^{d_1}$ is also ubiquitous a.s. in the component graph of the USF of $\mathbb{Z}^{d_2}$.*
- *There exists a finite graph with boundary H such that H is ubiquitous in the component graph of the USF $\mathbb{Z}^{d_2}$ but not of $\mathbb{Z}^{d_1}$.*
- *For every finite graph with boundary H, we can compute the dimensions at which H is or isn't ubiquitous in the component graph of the USF.*

Overview

- To show non-ubiquity in above the critical dimension: upper bound the first moment of the number of **witnesses** for the presence of H, i.e., collections of 'connection points'.

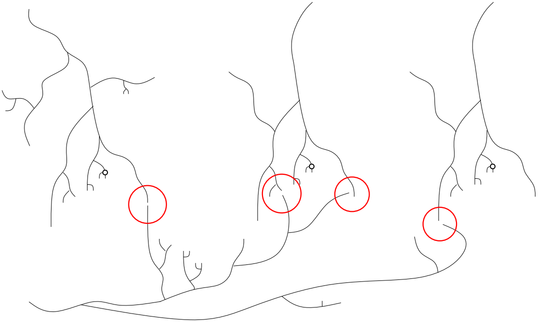

Overview

- To show non-ubiquity in above the critical dimension: upper bound the first moment of the number of **witnesses** for the presence of H, i.e., collections of ‘connection points’.

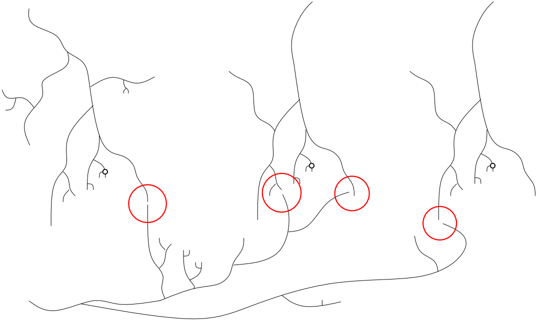

- To show ubiquity at and below the critical dimension: Start with first and second moment bounds on the number of witnesses, to show that the number of distinct witnesses is infinite with positive probability.

Overview

- To show non-ubiquity in above the critical dimension: upper bound the first moment of the number of **witnesses** for the presence of H, i.e., collections of 'connection points'.

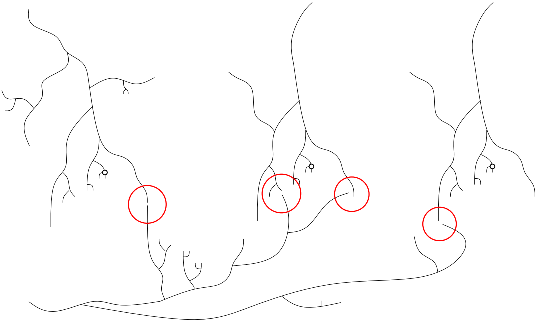

- To show ubiquity at and below the critical dimension: Start with first and second moment bounds on the number of witnesses, to show that the number of distinct witnesses is infinite with positive probability.
  - Improve this to an almost sure statement using *indistinguishability*.

The spread

We'll write $\asymp$, $\preceq$, and $\succeq$ for equalities and inequalities that hold up to a positive multiplicative constant.

The spread

We'll write $\asymp$, $\preceq$, and $\succeq$ for equalities and inequalities that hold up to a positive multiplicative constant.

Given a set of points W in $\mathbb{Z}^d$, we write $\langle W \rangle$ for the **spread** of W, defined to be

$$\langle W \rangle = \inf \prod_{\{u,v\} \in T} (1 + d(u,v)) \,:\, T \text{ is an abstract tree with vertex set } W$$

The spread

We'll write $\asymp$, $\preceq$, and $\succeq$ for equalities and inequalities that hold up to a positive multiplicative constant.

Given a set of points W in $\mathbb{Z}^d$, we write $\langle W \rangle$ for the **spread** of W, defined to be

$$\langle W \rangle = \inf \prod_{\{u,v\} \in T} (1 + d(u,v)) \ : \ T \text{ is an abstract tree with vertex set } W$$

In practice we don't use this definition but rather pick an enumeration $u_1, \ldots, u_n$ of W, and note that

$$\langle W \rangle \asymp \prod_{i=2}^{n} (1 + \min\{d(u_i, u_j) : j < i\})$$

The main estimates

The starting point for our calculations are the following estimates of BKPS. If W is a finite set then

$$\mathbb{P}(W \text{ all in same tree}) \asymp \langle W \rangle^{-(d-4)}$$

The main estimates

The starting point for our calculations are the following estimates of BKPS. If W is a finite set then

$$\mathbb{P}(W \text{ all in same tree}) \asymp \langle W \rangle^{-(d-4)}$$

and, if W_i are a collection of finite sets, then

$$\mathbb{P}(W_i \text{ all in the same tree for each } i\text{, different } i\text{s in different trees}) \preceq \prod_i \langle W_i \rangle^{-(d-4)}.$$

The main estimates

The starting point for our calculations are the following estimates of BKPS. If W is a finite set then

$$\mathbb{P}(W \text{ all in same tree}) \asymp \langle W \rangle^{-(d-4)}$$

and, if W_i are a collection of finite sets, then

$$\mathbb{P}(W_i \text{ all in the same tree for each } i\text{, different } i\text{s in different trees}) \preceq \prod_i \langle W_i \rangle^{-(d-4)}.$$

If the sets are well separated enough, we have a lower bound of the same order - I'll pretend for the talk that we always have this lower bound.

These estimates are provun using Wilson's alrogithm.

Suppose H is a finite graph with boundary, and $x = \{x_v : v \in \partial V\}$ are points in $\mathbb{Z}^d$ indexed by ∂V. Let's suppose that all the points x are in a ball of radius 2^n about the origin.

Define S_k to be the number of collections ξ_e in the dyadic shell of edges at distance between 2^k and 2^{k+1} from the origin that witness the faithfull presence of H at x in the following sense:

- For each $v \in \partial V$, the tree containing x_v also contains one of the endpoints of each of the edges ξ_e where e is incident to v in H.

Suppose H is a finite graph with boundary, and $x = \{x_v : v \in \partial V\}$ are points in $\mathbb{Z}^d$ indexed by ∂V. Let's suppose that all the points x are in a ball of radius 2^n about the origin.

Define S_k to be the number of collections ξ_e in the dyadic shell of edges at distance between 2^k and 2^{k+1} from the origin that witness the faithfull presence of H at x in the following sense:

- For each $v \in \partial V$, the tree containing x_v also contains one of the endpoints of each of the edges ξ_e where e is incident to v in H.
- For each $v \in V_\circ$, there is a tree that contains one of the endpoints of each of the edges ξ_e where e is incident to v in H.

Suppose H is a finite graph with boundary, and $x = \{x_v : v \in \partial V\}$ are points in $\mathbb{Z}^d$ indexed by ∂V. Let's suppose that all the points x are in a ball of radius 2^n about the origin.

Define S_k to be the number of collections ξ_e in the dyadic shell of edges at distance between 2^k and 2^{k+1} from the origin that witness the faithfull presence of H at x in the following sense:

- For each $v \in \partial V$, the tree containing x_v also contains one of the endpoints of each of the edges ξ_e where e is incident to v in H.
- For each $v \in V_\circ$, there is a tree that contains one of the endpoints of each of the edges ξ_e where e is incident to v in H.
- The trees corresponding to each of the vertices of H are different.

In other words, we have a tree for each vertex of H, and they touch each other in the way they are supposed to at the edges ξ_e.

Examples of witnesses ξ: we have a tree for each vertex of H, and they touch each other in the way they are supposed to at the edges ξ_e.

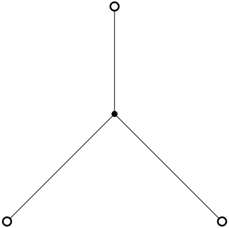

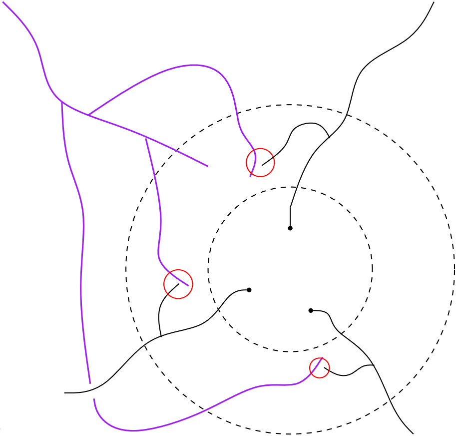

Examples of witnesses ξ: we have a tree for each vertex of H, and they touch each other in the way they are supposed to at the edges ξ_e.

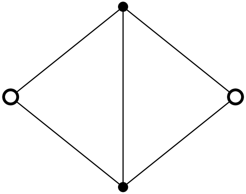

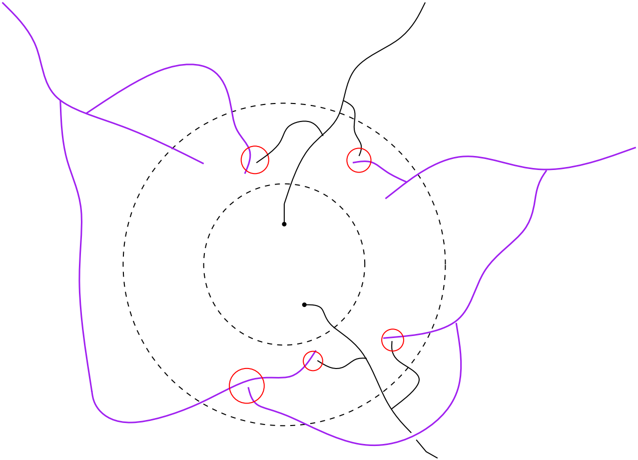

First moment

We can write down the probability that a collection in the shell is a witness – it's approximately

$$W(\xi) = \prod_{v \in \partial V} \langle \{x_v\} \cup \{\xi_e : e \sim v\} \rangle^{-(d-4)} \prod_{v \in V_\circ} \langle \{\xi_e : e \sim v\} \rangle^{-(d-4)}.$$

To calculate the expected number of witnesses in the shell, we need to sum this over all collections ξ in the shell.

Our first guess might be that the main contribution to this sum comes from ξs that are 'spread out', so that all the distances are on the order 2^k.

Our first guess might be that the main contribution to this sum comes from ξs that are 'spread out', so that all the distances are on the order 2^k.

There are about $2^{d|E|k}$ such configurations, while the value of $W(\xi)$ for each such configuration is about $2^{-(d-4)(2|E|-|V_\circ|)k}$.

Our first guess might be that the main contribution to this sum comes from ξs that are 'spread out', so that all the distances are on the order 2^k.

There are about $2^{d|E|k}$ such configurations, while the value of $W(\xi)$ for each such configuration is about $2^{-(d-4)(2|E|-|V_\circ|)k}$.

So the contribution to the first moment of these kind of configurations is about

$$(2^k)^{-(d-8)|E|+(d-4)|V_\circ|}$$

Our first guess might be that the main contribution to this sum comes from ξs that are 'spread out', so that all the distances are on the order 2^k.

There are about $2^{d|E|k}$ such configurations, while the value of $W(\xi)$ for each such configuration is about $2^{-(d-4)(2|E|-|V_\circ|)k}$.

So the contribution to the first moment of these kind of configurations is about

$$(2^k)^{-(d-8)|E|+(d-4)|V_\circ|}$$

Our first guess might therefore be that H is faithfully ubiquitous if and only if $-(d-8)|E|+(d-4)|V_\circ| \geq 0$. This gives the right answer for a path by BKPS!

Our first guess might be that the main contribution to this sum comes from ξs that are 'spread out', so that all the distances are on the order 2^k.

There are about $2^{d|E|k}$ such configurations, while the value of $W(\xi)$ for each such configuration is about $2^{-(d-4)(2|E|-|V_\circ|)k}$.

So the contribution to the first moment of these kind of configurations is about

$$(2^k)^{-(d-8)|E|+(d-4)|V_\circ|}$$

Our first guess might therefore be that H is faithfully ubiquitous if and only if $-(d-8)|E|+(d-4)|V_\circ| \geq 0$. This gives the right answer for a path by BKPS!

In fact, however, this doesn't even give us the right order for the first moment in general (and there will be more issues with the second moment later).

In fact, however, this doesn't even give us the right order for the first moment in general (and there will be more issues with the second moment later).

What *is* true is that the main contribution to the first moment comes from configurations that consist of a few 'clumps' of points at $O(1)$ distance, with different clumps at around the maximal distance 2^k,

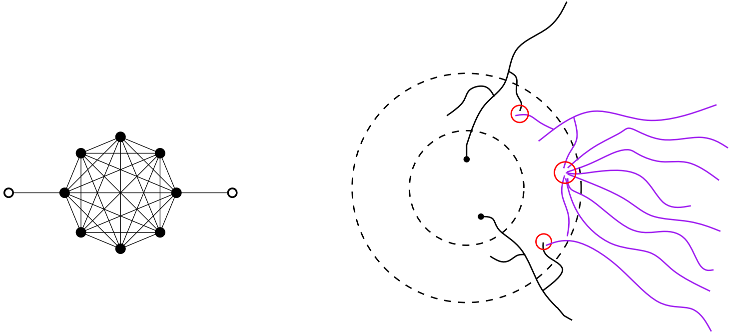

It turns out there is a nice way to interpret this clumping in terms of *hypergraphs*.

Recall that a **hypergraph** is just like a graph, except that an 'edge' can connect arbitrarily many vertices. We can also define hypergraphs with boundary $H = (\partial V, V_\circ, E)$.

Recall that a **hypergraph** is just like a graph, except that an 'edge' can connect arbitrarily many vertices. We can also define hypergraphs with boundary $H = (\partial V, V_\circ, E)$.

Given a hypergraph (with boundary) H and an equivalence relation on its edge set, we define the **coarsening** by taking the unions of the edges in each equivalence class to form a new, bigger edge.

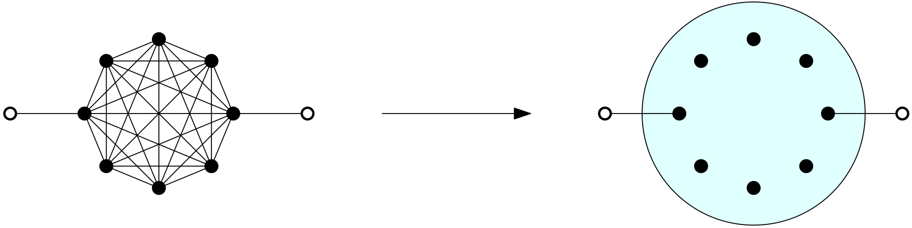

If we have a configuration ξ in which some of the points are clumped and the clumps are well separated, we can use this clumping structure to define a coarsening H' of H.

If we have a configuration ξ in which some of the points are clumped and the clumps are well separated, we can use this clumping structure to define a coarsening H' of H.

The contribution of configurations clumped according to the combinatorics of H' can be computed to be

$$\exp_2 \left(d|E(H')|\, k + (d-4)|V_\circ(H')|\, k - (d-4)\Delta(H')\, k \right),$$

where

$$\Delta(H') = \sum_v \deg(v) = \sum_e \deg(e).$$

So, to find the main contribution to the first moment, we need to optimize this expression over all coarsenings H' of H. (Trees are always optimal - no need to coarsen!)

If we have a configuration ξ in which some of the points are clumped and the clumps are well separated, we can use this clumping structure to define a coarsening H' of H.

The contribution of configurations clumped according to the combinatorics of H' can be computed to be

$$\exp_2\left(d|E(H')|\,k + (d-4)|V_\circ(H')|\,k - (d-4)\Delta(H')\,k\right),$$

where

$$\Delta(H') = \sum_v \deg(v) = \sum_e \deg(e).$$

So, to find the main contribution to the first moment, we need to optimize this expression over all coarsenings H' of H. (Trees are always optimal - no need to coarsen!)

Let's define

$$\eta_d(H) = d|E| + (d-4)|V_\circ| - (d-4)\Delta$$

and define

$$\bar{\eta}_d(H) = \max\{\eta_d(H') : \ H' \text{ a coarsening of } H\},$$

so that the expected number of witnesses in the shell is about

$$2^{\bar{\eta}_d(H)\,k}.$$

Let's define

$$\eta_d(H) = d|E| + (d-4)|V_\circ| - (d-4)\Delta$$

and define

$$\bar{\eta}_d(H) = \max\{\eta_d(H') : \ H' \text{ a coarsening of } H\},$$

so that the expected number of witnesses in the shell is about

$$2^{\bar{\eta}_d(H)\, k + O(\log k)}.$$

In fact this is really true up to a polynomial correction!

Depending on whether $\bar{\eta}_d(H)$ is positive, negative, or zero, we either get that the expected number of witnesses in the dyadic shell grows exponentially, decays exponentially, or neither.

In particular, it's very plausible that we will have non-ubiquity in the case that the above max is negative. This is true – to prove it we need to control the number of witnesses that do not appear inside a single dyadic shell.

What if $\bar{\eta}_d(H) \geq 0$?

If we knew that the *second moments* of the number of witnesses in an annulus were of the same order as the first moment squared, we could deduce that the probability of having a witness inside a given dyadic shell did not decay as the shell got larger, and it would follow that there were infinitely many witnesses with positive probability.

What if $\bar{\eta}_d(H) \geq 0$?

If we knew that the *second moments* of the number of witnesses in an annulus were of the same order as the first moment squared, we could deduce that the probability of having a witness inside a given dyadic shell did not decay as the shell got larger, and it would follow that there were infinitely many witnesses with positive probability.

It turns out that there is an easy criterion to see if this happens, coming from an obvious obstruction to the first moment giving the right answer.

If a graph is faithfully ubiquitous then so are all of its subgraphs.

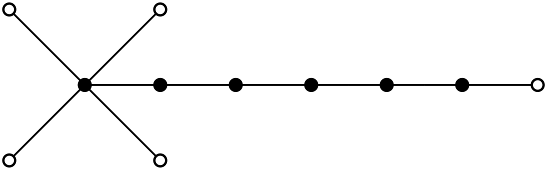

If we attach some ‘easy to find’ graph onto a ‘hard to find graph’, we can make $\bar{\eta}_d(H)$ large even though H clearly isn’t ubiquitous due to having a non-ubiquitous subgraph.

However, if this does not happen, and $\bar{\eta}_d(H') \geq 0$ for every subgraph H' of H, we can show that the second moment of the number of witnesses in a shell is comparable to the first moment squared, as we wanted.

Now, we'd like to go from knowing that there are infinitely many witnesses with positive probability to knowing this almost surley.

We need a zero-one law!

We need indistinguishability of tuples of trees in the USF by tail properties!

Indistinguishability of trees

Theorem (H. and Nachmias 2015)

Let G be a unimodular transitive graph (e.g. $\mathbb{Z}^d$), and let F be either the FUSF of WUSF of G. Then for every measurable, automorphism invariant set $\mathscr{A}$ of subgraphs of G, either every component of F is in $\mathscr{A}$ or none of the components of F are in $\mathscr{A}$.

Timar (2015) proved this theorem for the FUSF in the case that it is different from the WUSF in independent work.

We need a version of indistinguishability for *tuples* of components.

Let $G = (V, E)$ be a graph. A measurable subset of $V^k \times \{0,1\}^E$ is a k-**multicomponent property** if it is invariant to changing the vertices within their component, i.e.,

$$(v_1, \ldots, v_k, \omega) \in \mathscr{A} \Rightarrow (u_1, \ldots, u_k, \omega) \in \mathscr{A}$$
$$\text{if } v_i \text{ is connected to } u_i \text{ in } \omega \text{ for each } i = 1, \ldots, k.$$

Let $G = (V, E)$ be a graph. A measurable subset of $V^k \times \{0,1\}^E$ is a k-**multicomponent property** if it is invariant to changing the vertices within their component, i.e.,

$$(v_1, \ldots, v_k, \omega) \in \mathscr{A} \Rightarrow (u_1, \ldots, u_k, \omega) \in \mathscr{A}$$
$$\text{if } v_i \text{ is connected to } u_i \text{ in } \omega \text{ for each } i = 1, \ldots, k.$$

A multicomponent property is *tail* if it is stable if we replace the configuration ω by ω' such that both

- ω and ω' have finite symmetric difference, and
- the components of v_i in ω and ω' have finite symmetric difference for each $i = 1, \ldots, n$.

e.g. $\mathscr{A}$ = the event that there exist infinitely many witnesses for the presence of some finite graph with boundary in the component graph at the trees containing $v_1, \ldots, v_k$.

Theorem (H. and Peres 2016)

k-tuples of distinct components of the USF of $\mathbb{Z}^d$ are indistinguishable by tail multicomponent properties: If $\mathscr{A}$ is a tail k-multicomponent property, then either every k-tuple of distinct trees in the USF has property $\mathscr{A}$ or none of them do.

Proven using techniques from H. and Nachmias (2015), BKPS.

Theorem (H. and Peres 2016)

k-tuples of distinct components of the USF of $\mathbb{Z}^d$ are indistinguishable by tail multicomponent properties: If $\mathscr{A}$ is a tail k-multicomponent property, then either every k-tuple of distinct trees in the USF has property $\mathscr{A}$ or none of them do.

Proven using techniques from H. and Nachmias (2015), BKPS.

(In fact the result holds for the WUSF of every Liouville graph, even without any transitivity.)

Putting everything together...

Theorem (H. and Peres 2016)

Let H be a finite graph with boundary. Then H is faithfully ubiquitous in the component graph of the USF of $\mathbb{Z}^d$ if and only if

$$\min\left\{\max\left\{\eta_d(H'') : H'' \text{ is a coarsening of } H'\right\} H' \text{ is a subgraph of } H\right\}$$

is non-negative. Moreover, H is ubiquitous if and only if it has a quotient that if faithfully ubiquitous.

Putting everything together...

Theorem (H. and Peres 2016)

Let H be a finite graph with boundary. Then H is faithfully ubiquitous in the component graph of the USF of $\mathbb{Z}^d$ if and only if

$$\min\Big\{\max\left\{\eta_d(H'') : H'' \text{ is a coarsening of } H'\right\} H' \text{ is a subgraph of } H\Big\}$$

is non-negative. Moreover, H is ubiquitous if and only if it has a quotient that if faithfully ubiquitous.

Given the theorem, it's easy to find families of graphs with boundary that separate all dimensions ≥ 9, e.g.

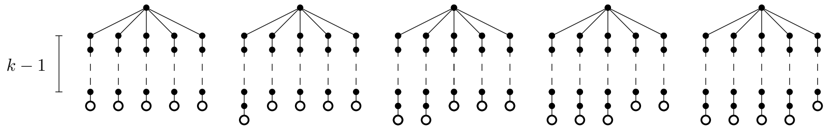

Separating $d = 5, 6, 7, 8$

Given $r \geq 1$, define the **component hypergraph** by letting a set of trees form an edge if there is a ball of radius r that intersects every tree in the set. We can define (faithful) ubiquity of hypergraphs with boundary in the component hypergraph exactly as we did with graphs.

Theorem (H. and Peres 2016)

Let H be a finite hypergraph with boundary. Then H is faithfully ubiquitous in the component hypergraph (for sufficiently large r) of the USF of $\mathbb{Z}^d$ if and only if

$$\min\left\{\max\left\{\eta_d(H'') : H'' \textit{ is a coarsening of } H'\right\} H' \textit{ is a subhypergraph of } H\right\}$$

is non-negative. Moreover, H is ubiquitous if and only if it has a quotient that if faithfully ubiquitous.

Separating $d = 5, 6, 7, 8$.

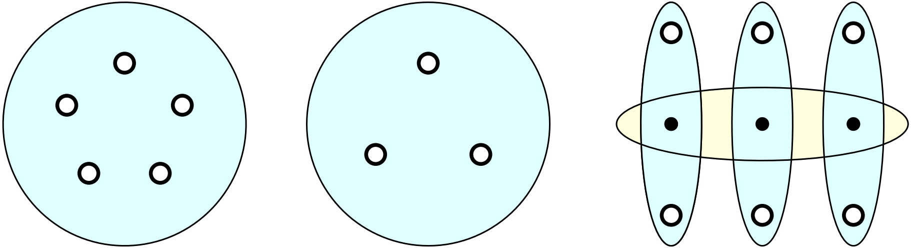

In five dimensions, for every five trees in the USF, there is a ball of constant radius that intersects all five trees. This is not true with six or more trees.

In six dimensions, for every three trees in the USF, there is a ball of constant radius that intersects all three trees. This is not true with four or more trees, nor is it true for $d \geq 7$.

To separate 7 and 8, we can use the hypergraph on the right...

Thank you!