Stochastics Meeting Lunteren 2012

Scaling limits for random trees and graphs

Christina Goldschmidt

Stochastics Meeting Lunteren 2012

Lecture 1: random trees

Principally based on work by

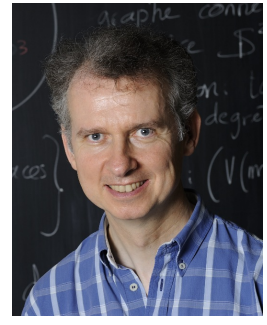

David Aldous
University of California, Berkeley

and

Jean-François Le Gall
Université Paris-Sud

Uniform random trees

We start with perhaps the simplest model of a random tree.

Uniform random trees

We start with perhaps the simplest model of a random tree.

Let $\mathbb{T}_n$ be the set of unordered trees on n vertices labelled by $[n] := \{1, 2, \ldots, n\}$.

Uniform random trees

We start with perhaps the simplest model of a random tree.

Let $\mathbb{T}_n$ be the set of unordered trees on n vertices labelled by $[n] := \{1, 2, \ldots, n\}$.

For example, $\mathbb{T}_3$ consists of the trees

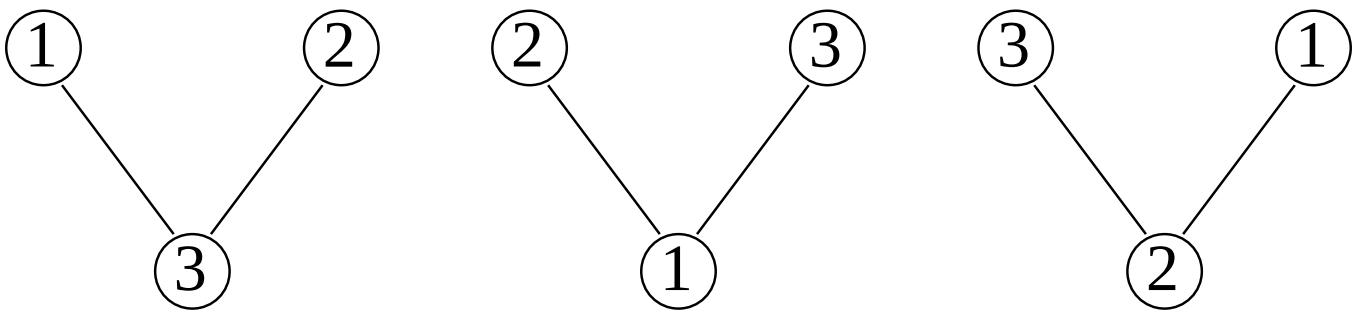

Uniform random trees

Unordered means that these trees are all the same:

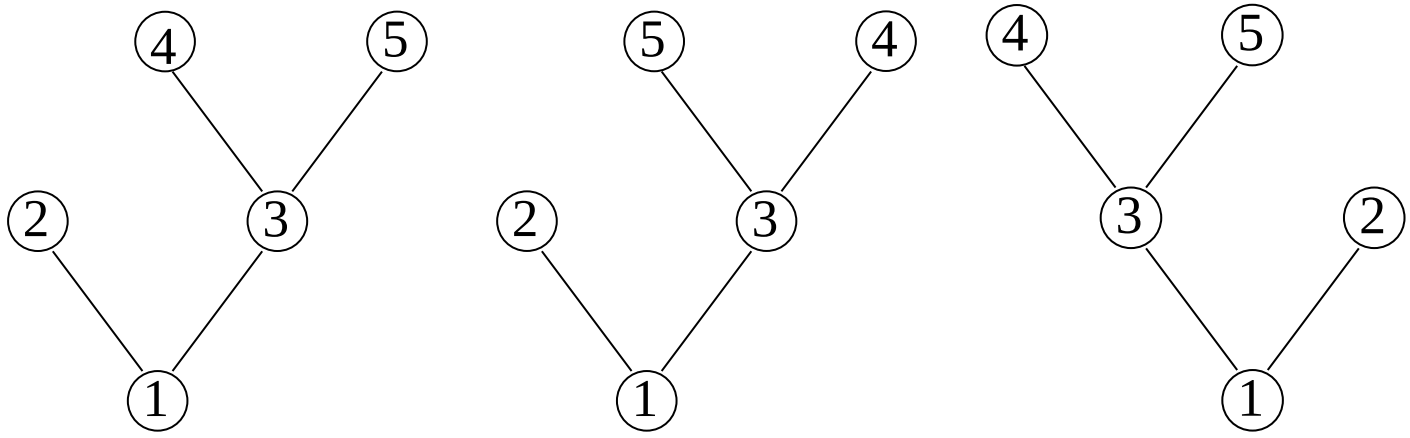

but this one is different:

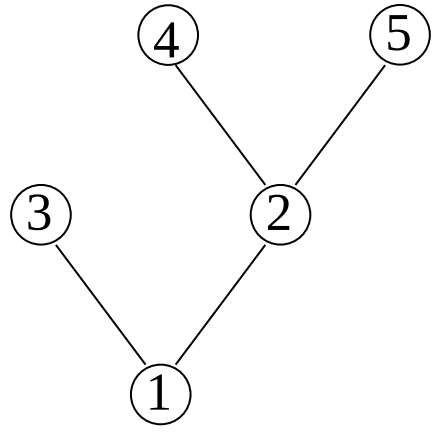

Uniform random trees

Cayley's formula: $|\mathbb{T}_n| = n^{n-2}$.

Uniform random trees

Cayley's formula: $|\mathbb{T}_n| = n^{n-2}$.

Write T_n for a tree chosen uniformly from $\mathbb{T}_n$.

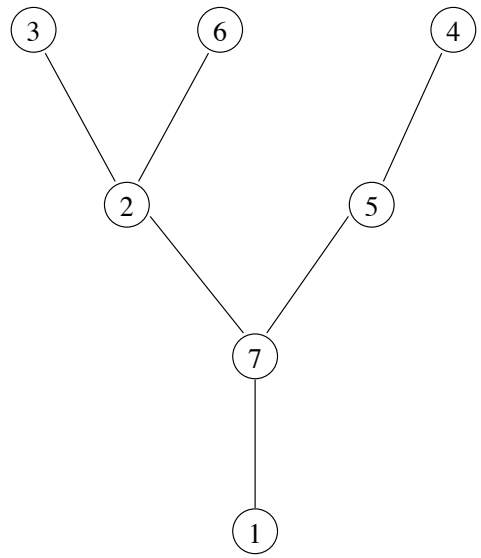

Uniform random trees

Cayley's formula: $|\mathbb{T}_n| = n^{n-2}$.

Write T_n for a tree chosen uniformly from $\mathbb{T}_n$.

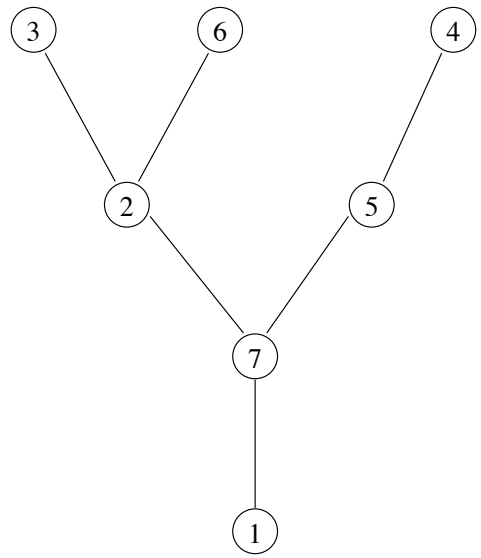

What happens as n grows?

Uniform random trees

Cayley's formula: $|\mathbb{T}_n| = n^{n-2}$.

Write T_n for a tree chosen uniformly from $\mathbb{T}_n$.

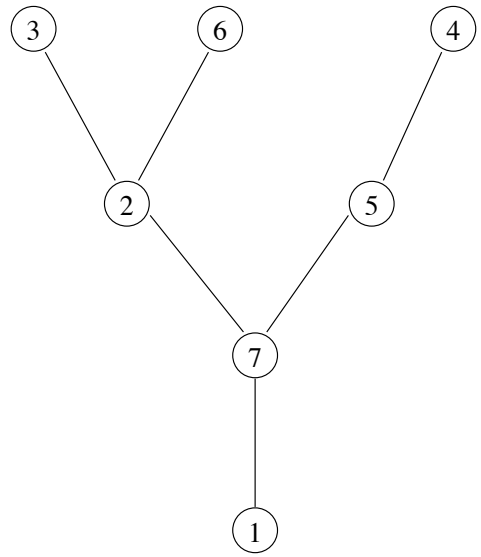

What happens as n grows?

It's useful to have an algorithm for building T_n.

The Aldous-Broder algorithm

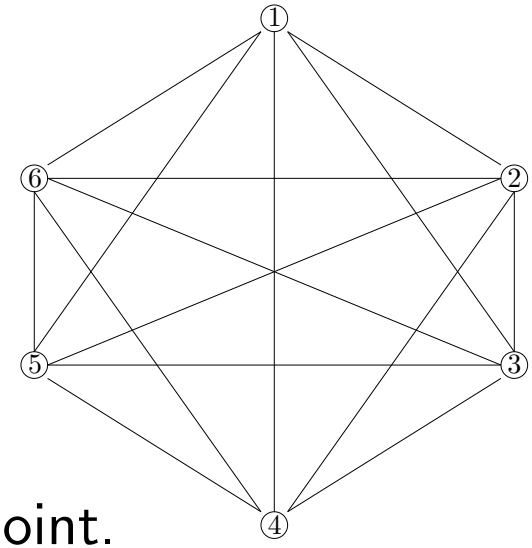

Take the complete graph on n vertices.

- Pick a uniform vertex to be the starting point.
- Run a simple random walk $(S_k)_{k\geq 0}$ on the graph (i.e. at each step, move to a neighbour chosen uniformly at random).
- Anytime the walk visits a new vertex, keep the edge along which it was reached.
- Stop when all vertices have been visited.

The resulting tree τ_0 is uniformly distributed on $\mathbb{T}_n$.

The Aldous-Broder algorithm

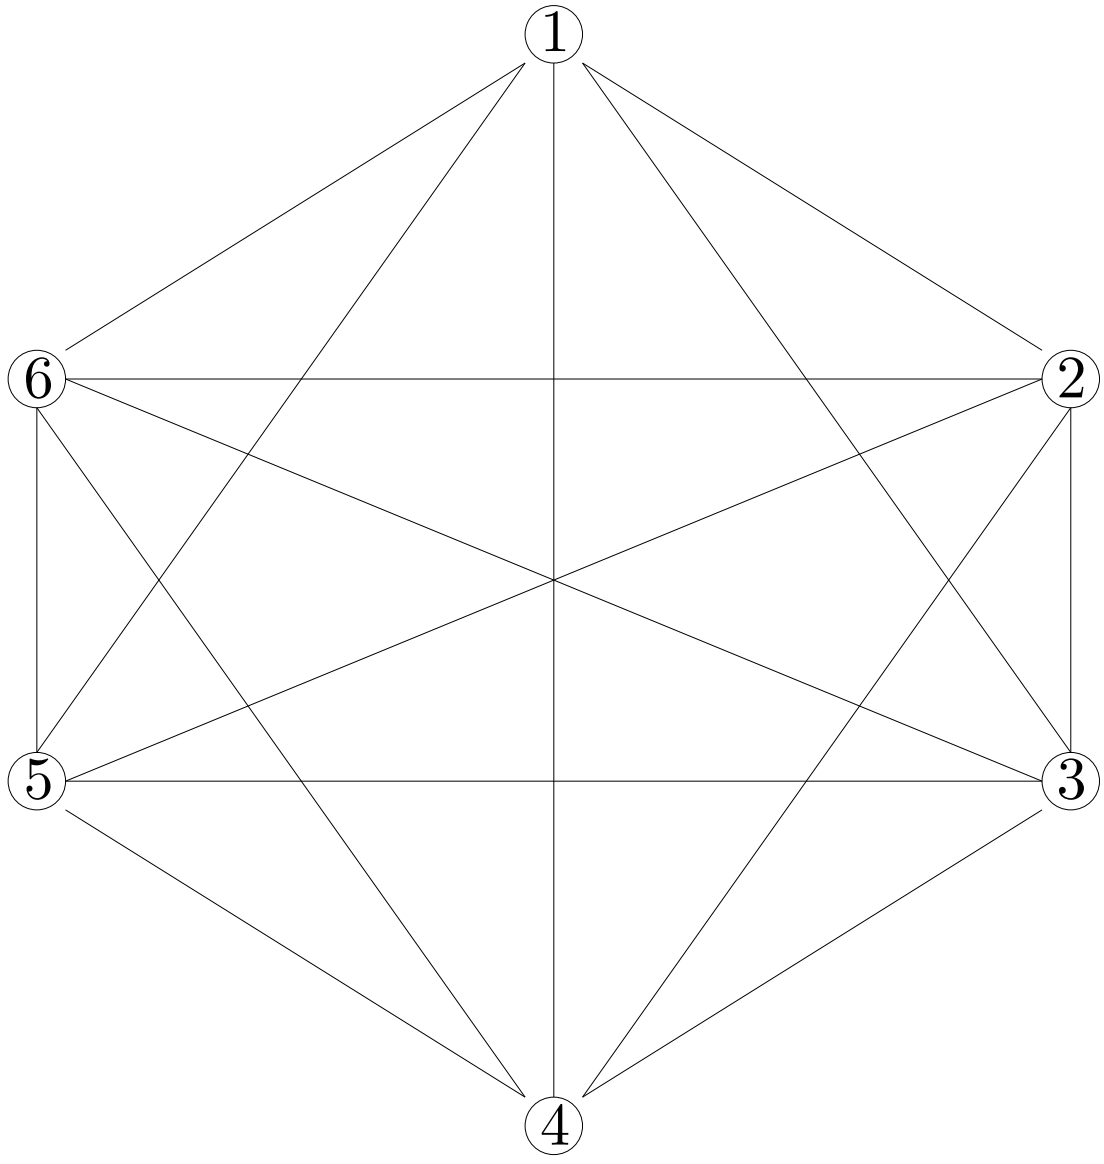

The Aldous-Broder algorithm

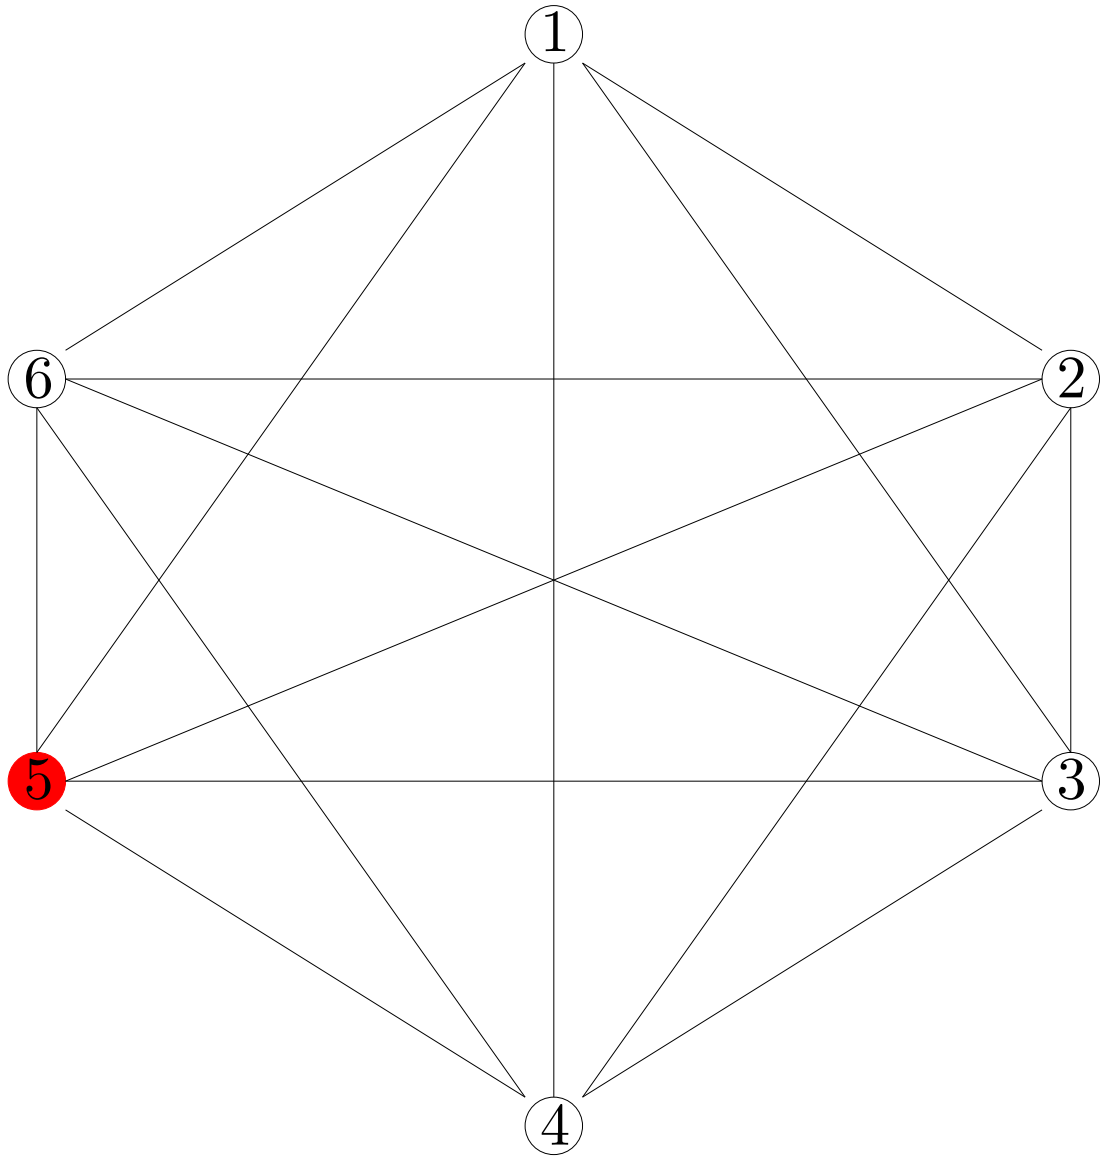

The Aldous-Broder algorithm

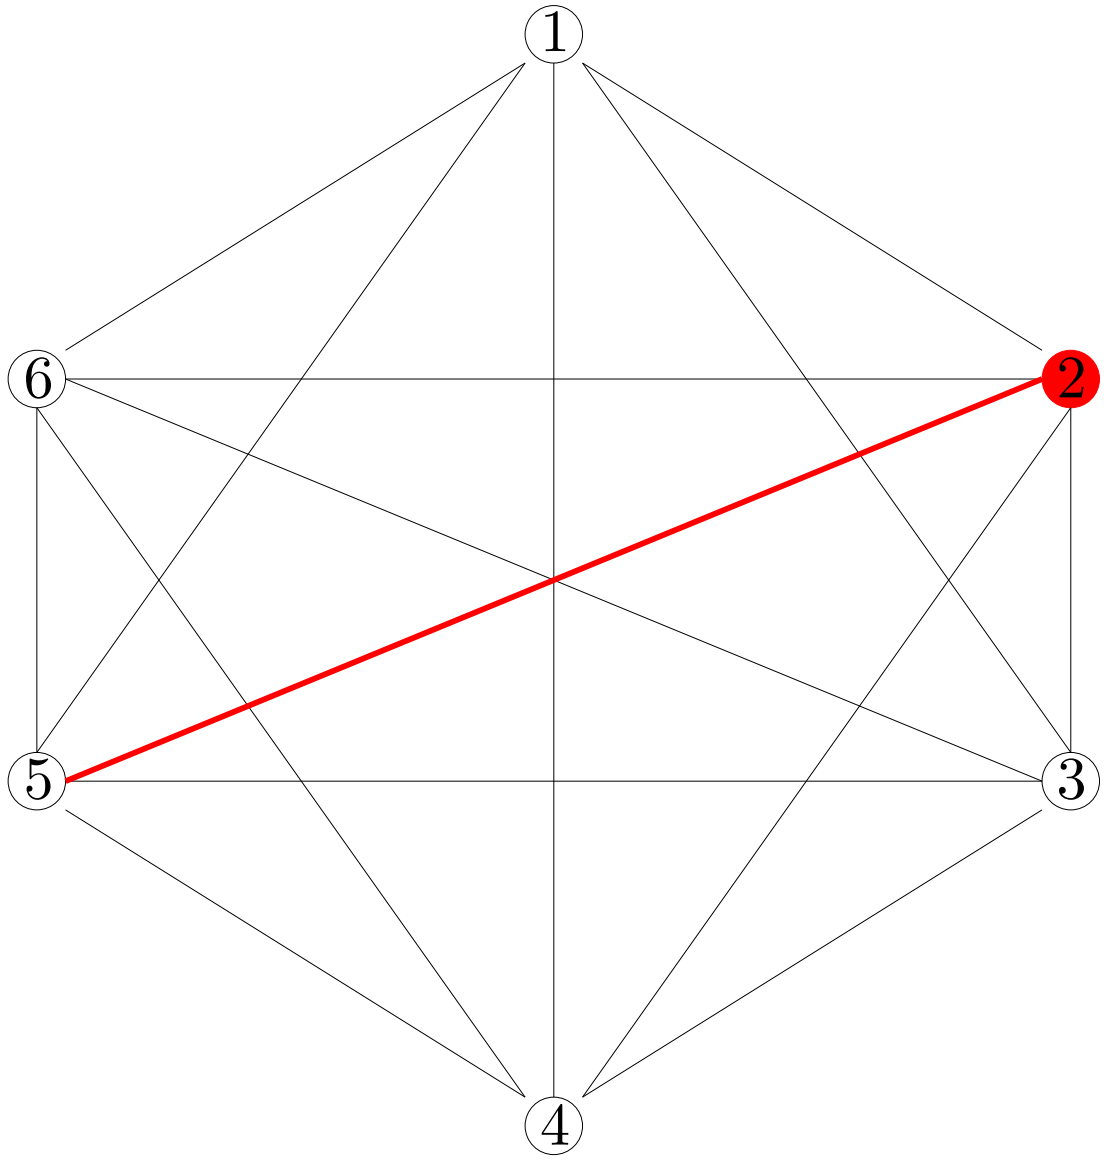

The Aldous-Broder algorithm

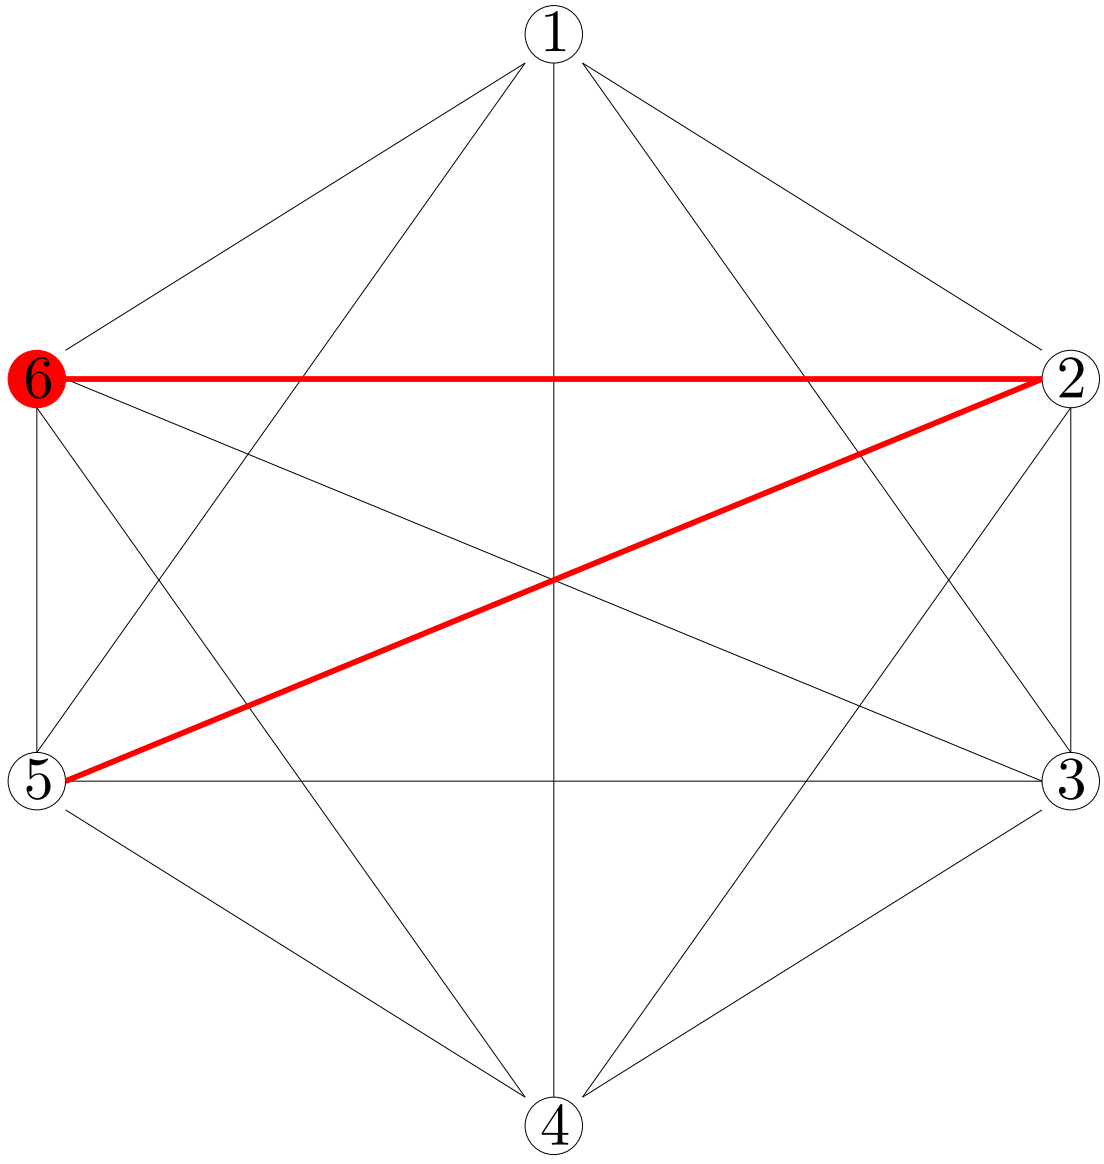

The Aldous-Broder algorithm

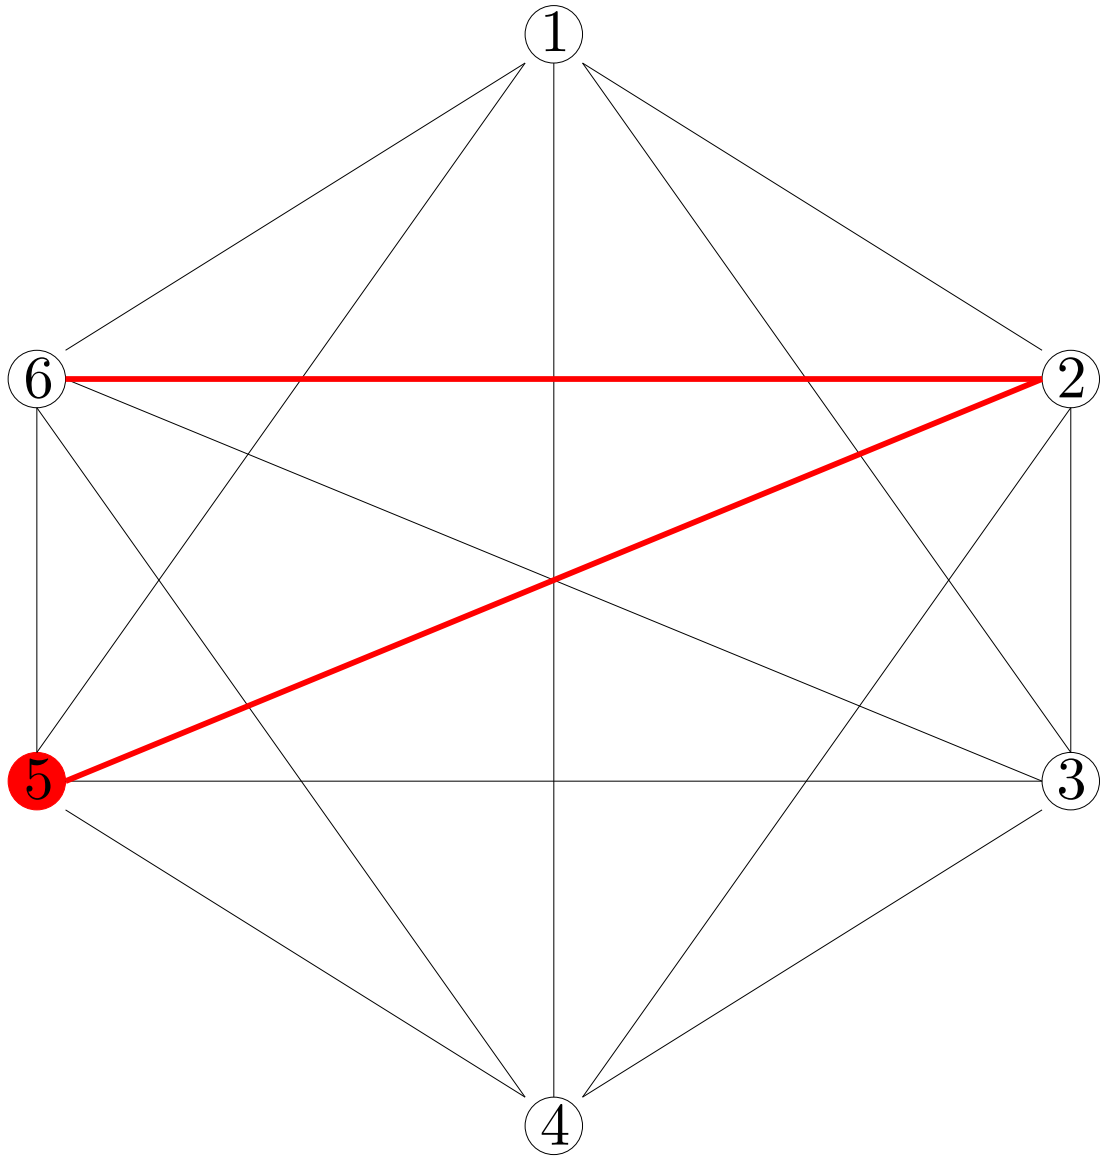

The Aldous-Broder algorithm

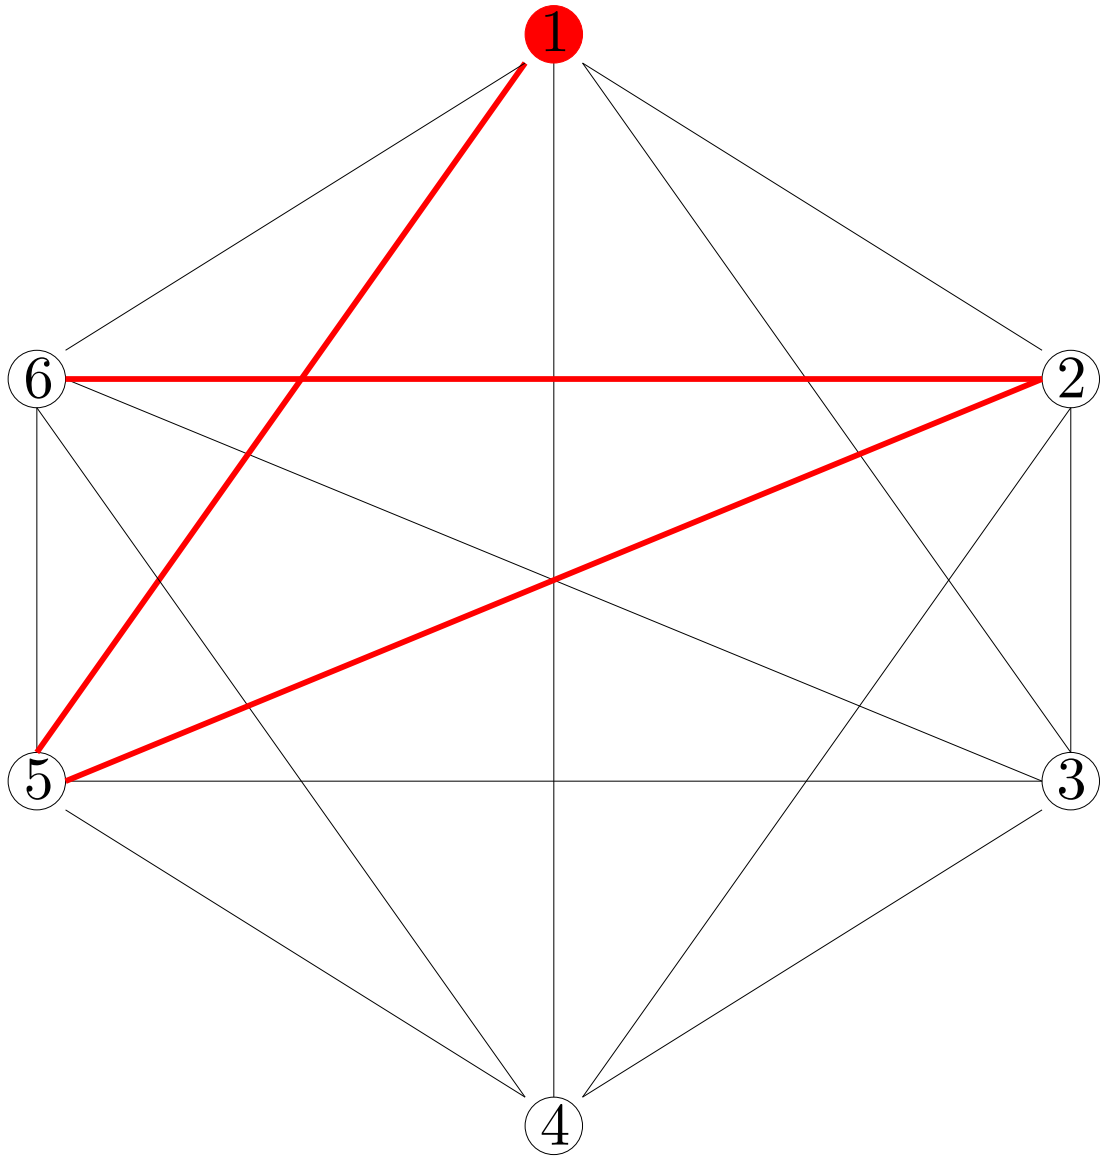

The Aldous-Broder algorithm

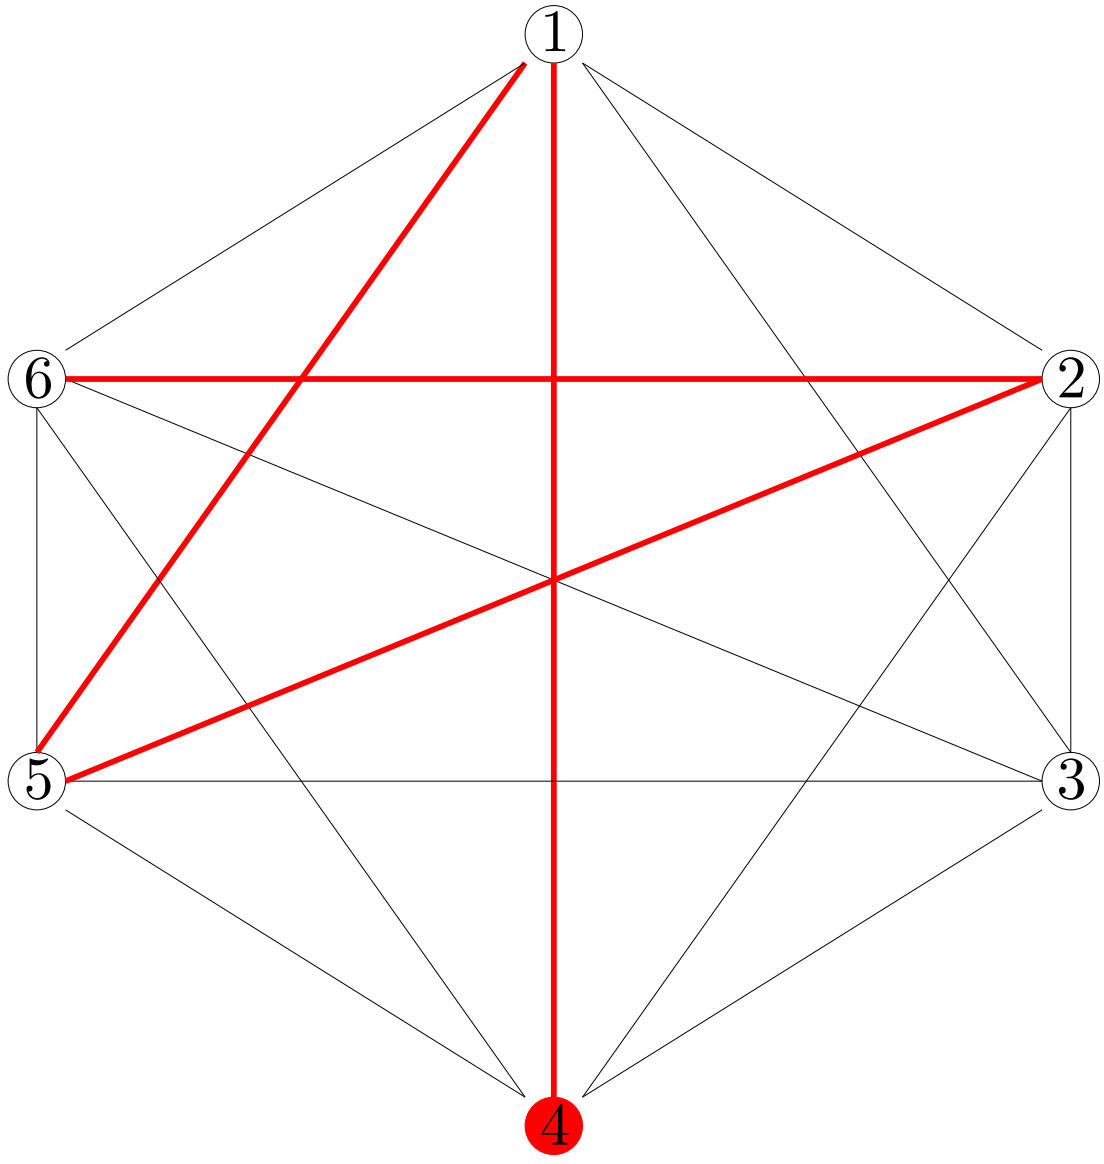

The Aldous-Broder algorithm

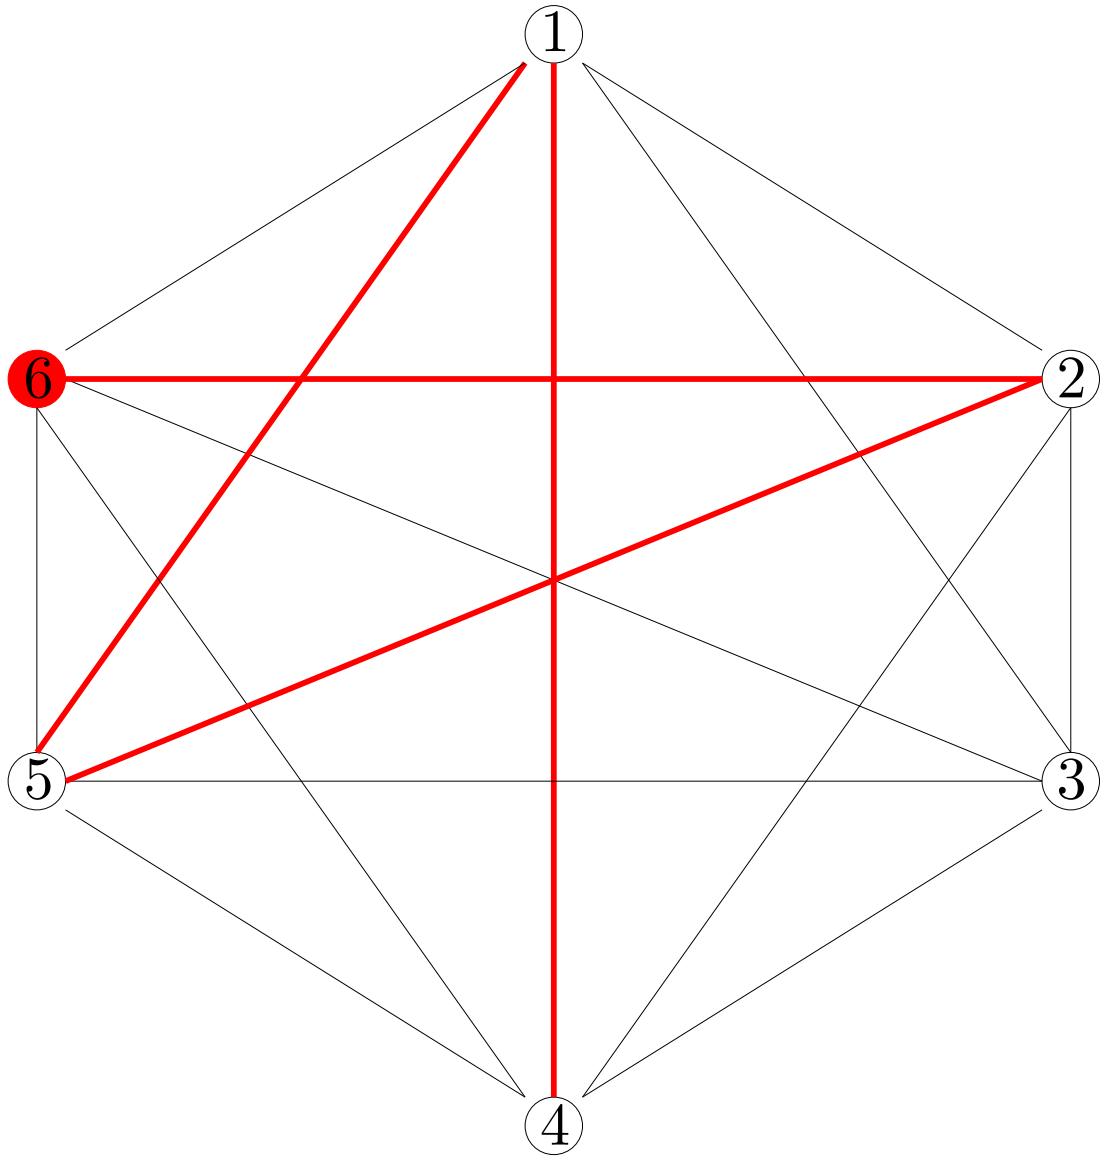

The Aldous-Broder algorithm

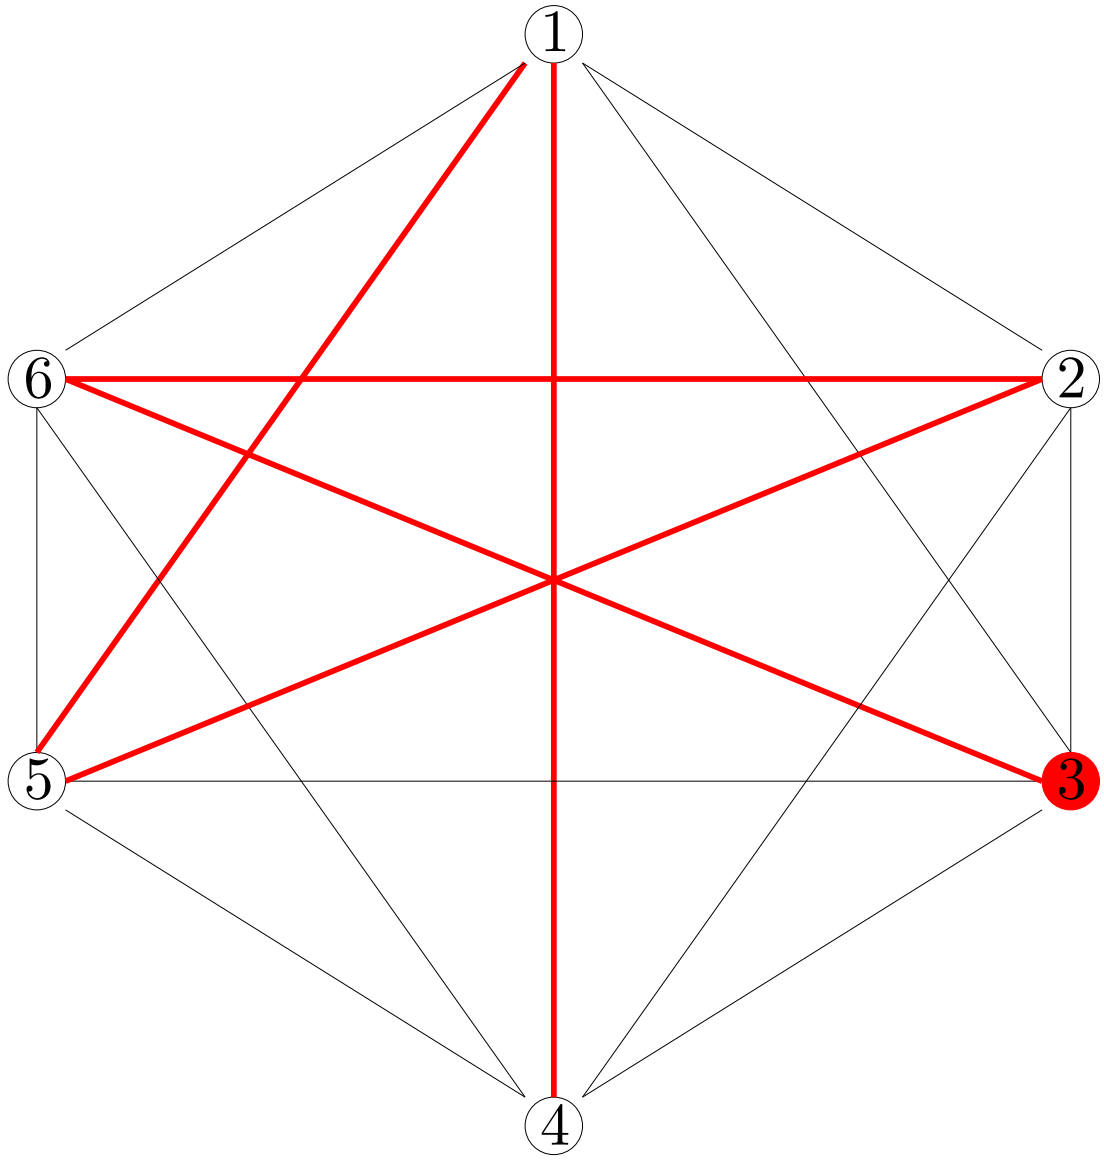

The Aldous-Broder algorithm

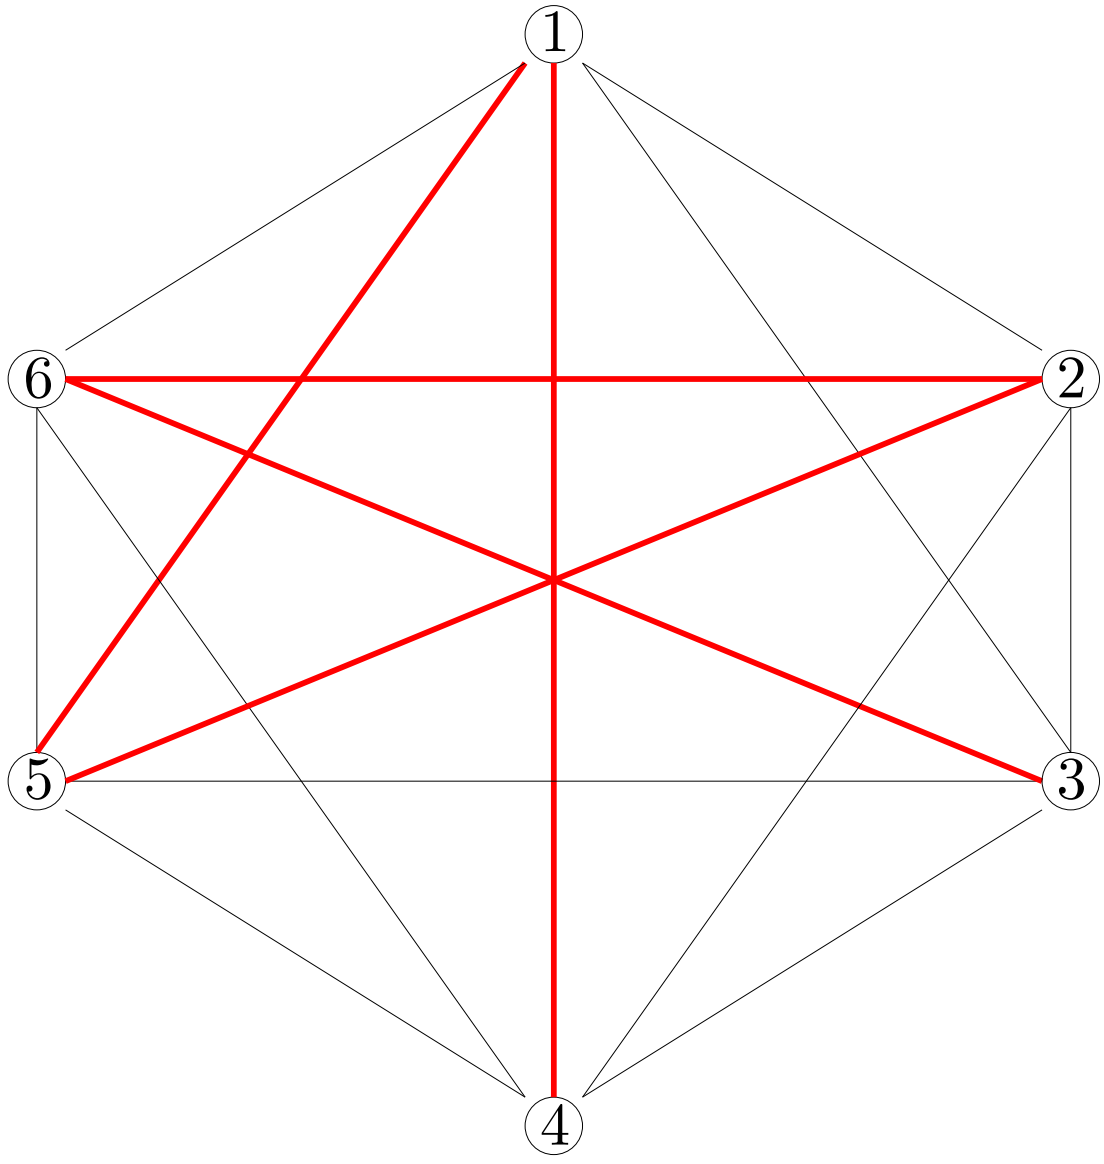

The Aldous-Broder algorithm

The random walk $(S_k)_{k\geq 0}$ has a uniform stationary distribution, and is reversible, so that it makes sense to talk about a stationary random walk $(S_k)_{k\in\mathbb{Z}}$.

The Aldous-Broder algorithm

The random walk $(S_k)_{k\geq 0}$ has a uniform stationary distribution, and is reversible, so that it makes sense to talk about a stationary random walk $(S_k)_{k\in\mathbb{Z}}$.

The dynamics of the random walk give rise to Markovian dynamics on $\mathbb{T}_n$.

The Aldous-Broder algorithm

The random walk $(S_k)_{k\geq 0}$ has a uniform stationary distribution, and is reversible, so that it makes sense to talk about a stationary random walk $(S_k)_{k\in\mathbb{Z}}$.

The dynamics of the random walk give rise to Markovian dynamics on $\mathbb{T}_n$.

Let τ_k be the tree constructed from the random walk started at time k. Since the random walk is stationary, the tree must be also.

The Aldous-Broder algorithm

It remains to show that the stationary distribution for $(\tau_k)_{k\in\mathbb{Z}}$ is uniform on $\mathbb{T}_n$. It turns out to be easier to think of τ_k as rooted at S_k, and also to think in reversed time.

The Aldous-Broder algorithm

It remains to show that the stationary distribution for $(\tau_k)_{k\in\mathbb{Z}}$ is uniform on $\mathbb{T}_n$. It turns out to be easier to think of τ_k as rooted at S_k, and also to think in reversed time.

Consider the transition probabilities $P(\tau, \tau')$ for the reversed-time chain (taking values in the set of rooted trees).

The Aldous-Broder algorithm

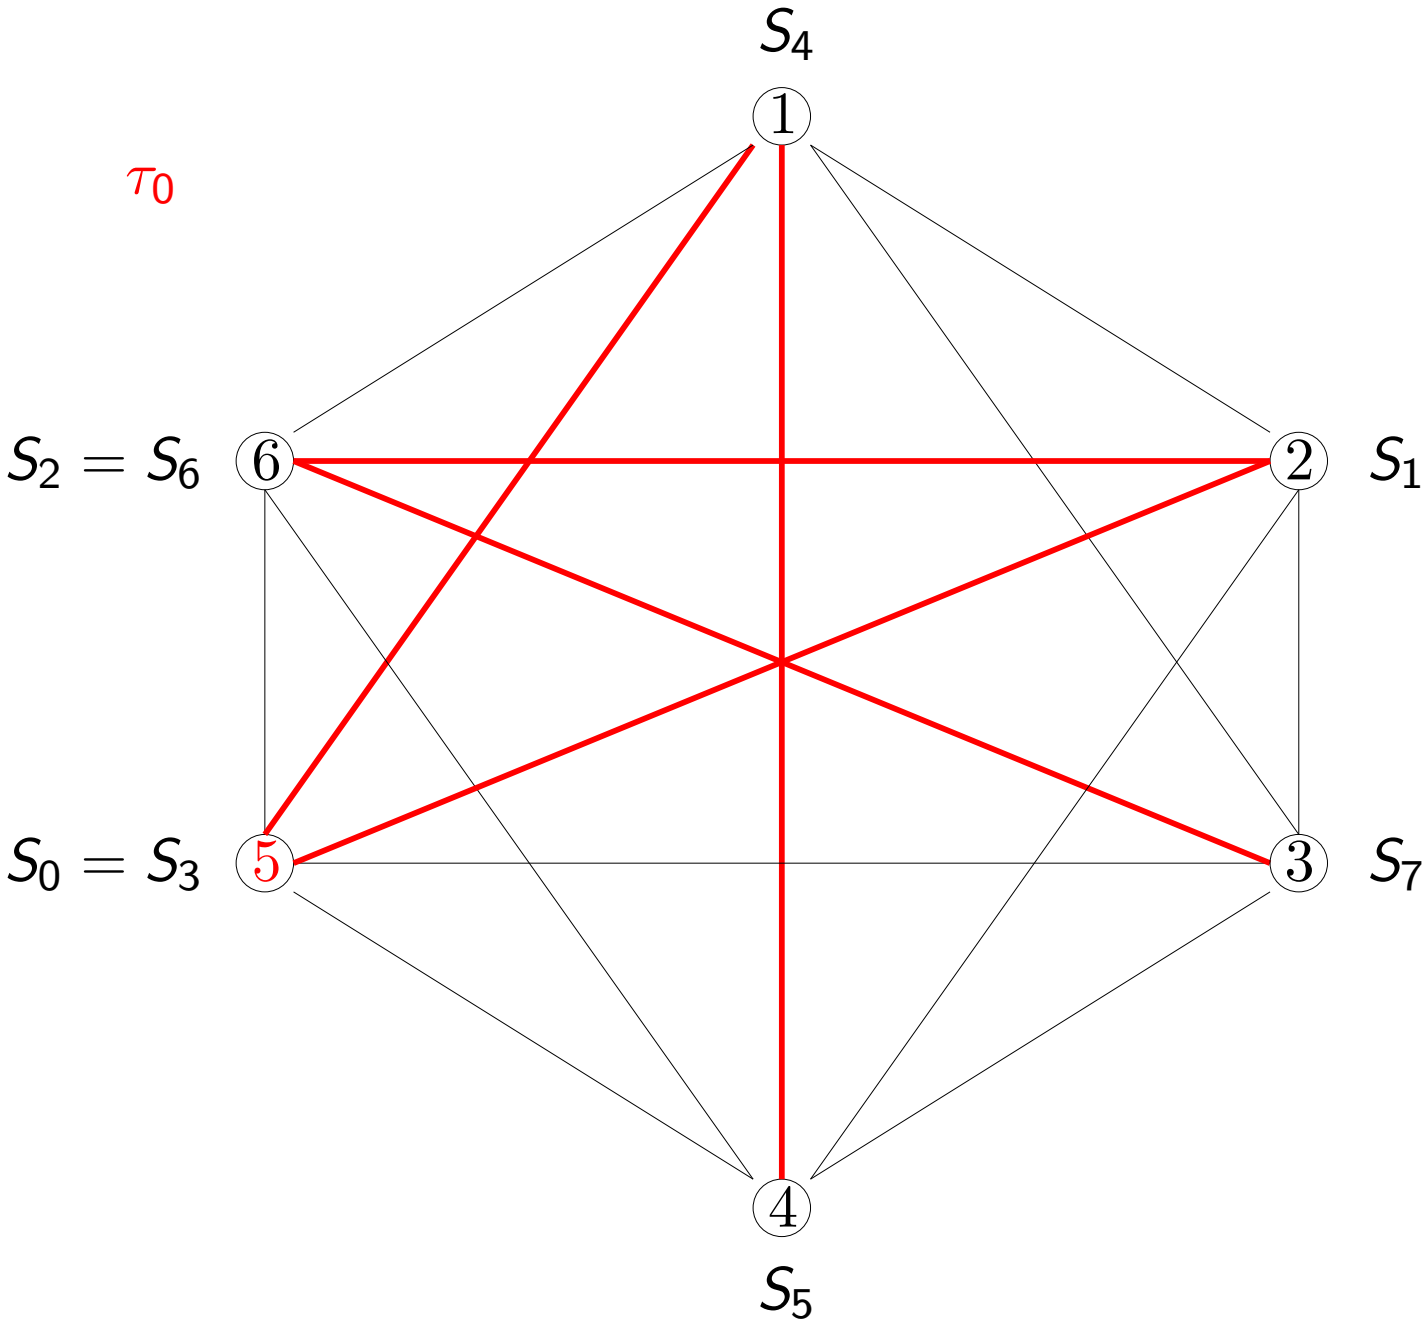

The Aldous-Broder algorithm

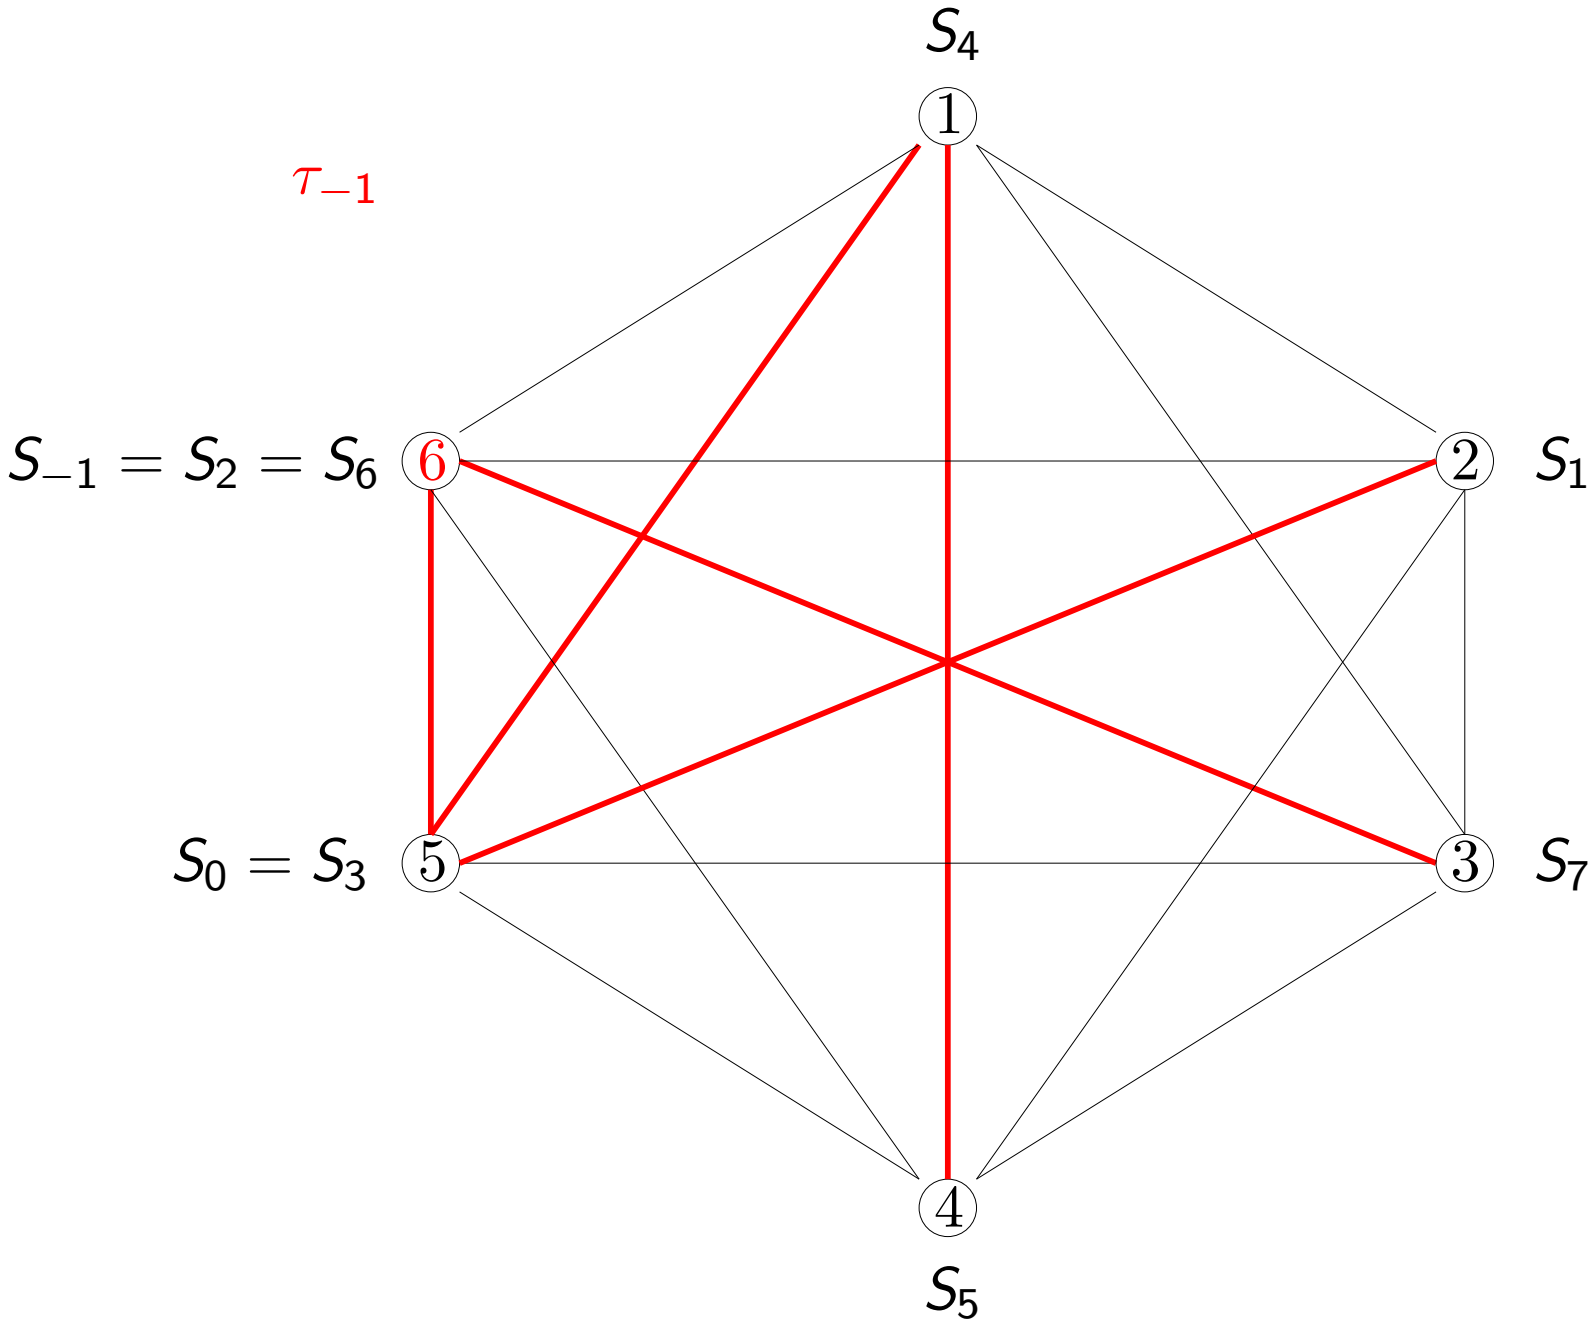

The Aldous-Broder algorithm

There are $n-1$ different places that the root might move to and so $n-1$ possible rooted trees we can reach going backwards in time, each equally likely.

So for fixed τ, $P(\tau, \tau') = 0$ or $1/(n-1)$.

The Aldous-Broder algorithm

There are $n-1$ different places that the root might move to and so $n-1$ possible rooted trees we can reach going backwards in time, each equally likely.

So for fixed τ, $P(\tau, \tau') = 0$ or $1/(n-1)$.

Similarly, for fixed τ', there are $n-1$ places the root might have jumped from, so that $P(\tau, \tau') = 0$ or $1/(n-1)$.

The Aldous-Broder algorithm

There are $n-1$ different places that the root might move to and so $n-1$ possible rooted trees we can reach going backwards in time, each equally likely.

So for fixed τ, $P(\tau, \tau') = 0$ or $1/(n-1)$.

Similarly, for fixed τ', there are $n-1$ places the root might have jumped from, so that $P(\tau, \tau') = 0$ or $1/(n-1)$.

So the matrix P is doubly stochastic and so has uniform stationary distribution on the set of rooted trees.

The Aldous-Broder algorithm

There are $n-1$ different places that the root might move to and so $n-1$ possible rooted trees we can reach going backwards in time, each equally likely.

So for fixed τ, $P(\tau, \tau') = 0$ or $1/(n-1)$.

Similarly, for fixed τ', there are $n-1$ places the root might have jumped from, so that $P(\tau, \tau') = 0$ or $1/(n-1)$.

So the matrix P is doubly stochastic and so has uniform stationary distribution on the set of rooted trees.

It's straightforward to show that the chain is irreducible and since the root is uniformly distributed, it follows that τ_0 is a uniform random tree.

A variant due to Aldous

"Do the labelling as we go, then relabel at the end."

1. Start from the vertex labelled 1.
2. For $2 \leq i \leq n$, connect vertex i to vertex V_i such that

$$V_i = \begin{cases} i-1 \text{ with probability } 1-(i-2)/n \\ \text{uniform on } \{1,2,\ldots,i-2\} \text{ otherwise.} \end{cases}$$

3. Take a uniform random permutation of the labels.

The Aldous-Broder algorithm

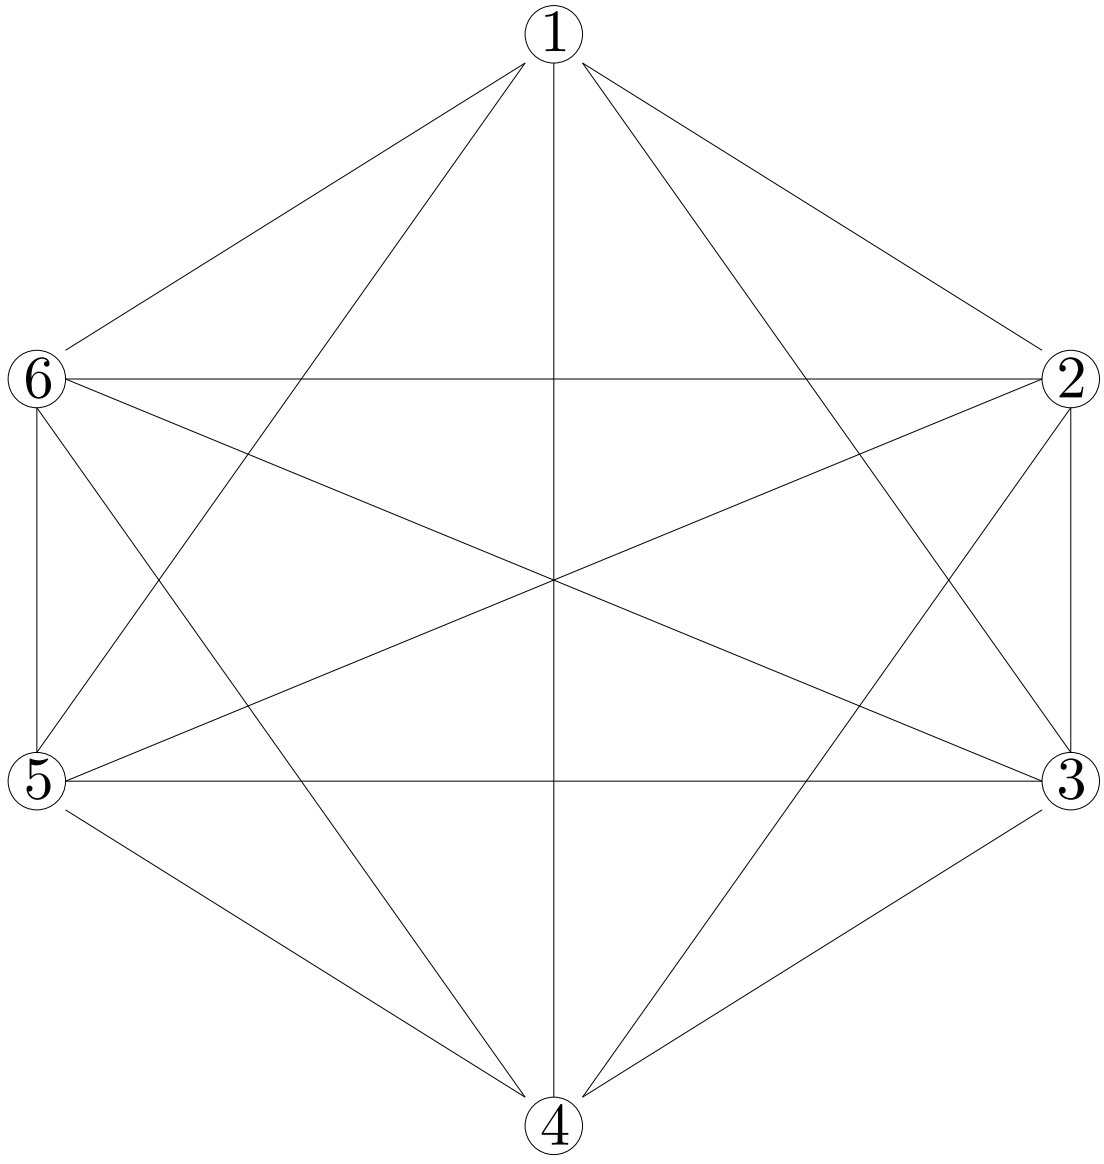

The Aldous-Broder algorithm

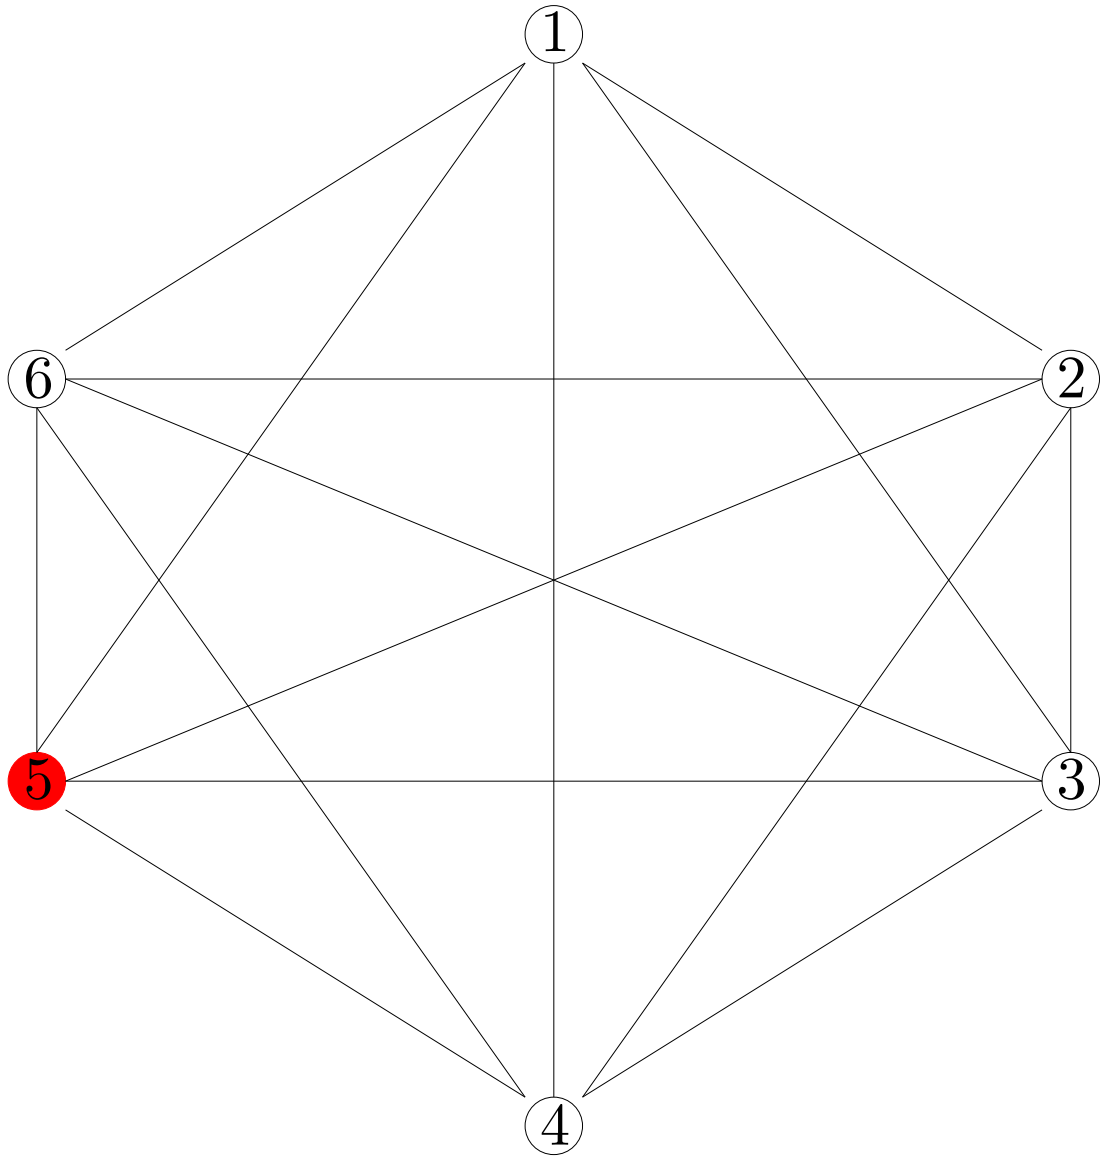

The Aldous-Broder algorithm

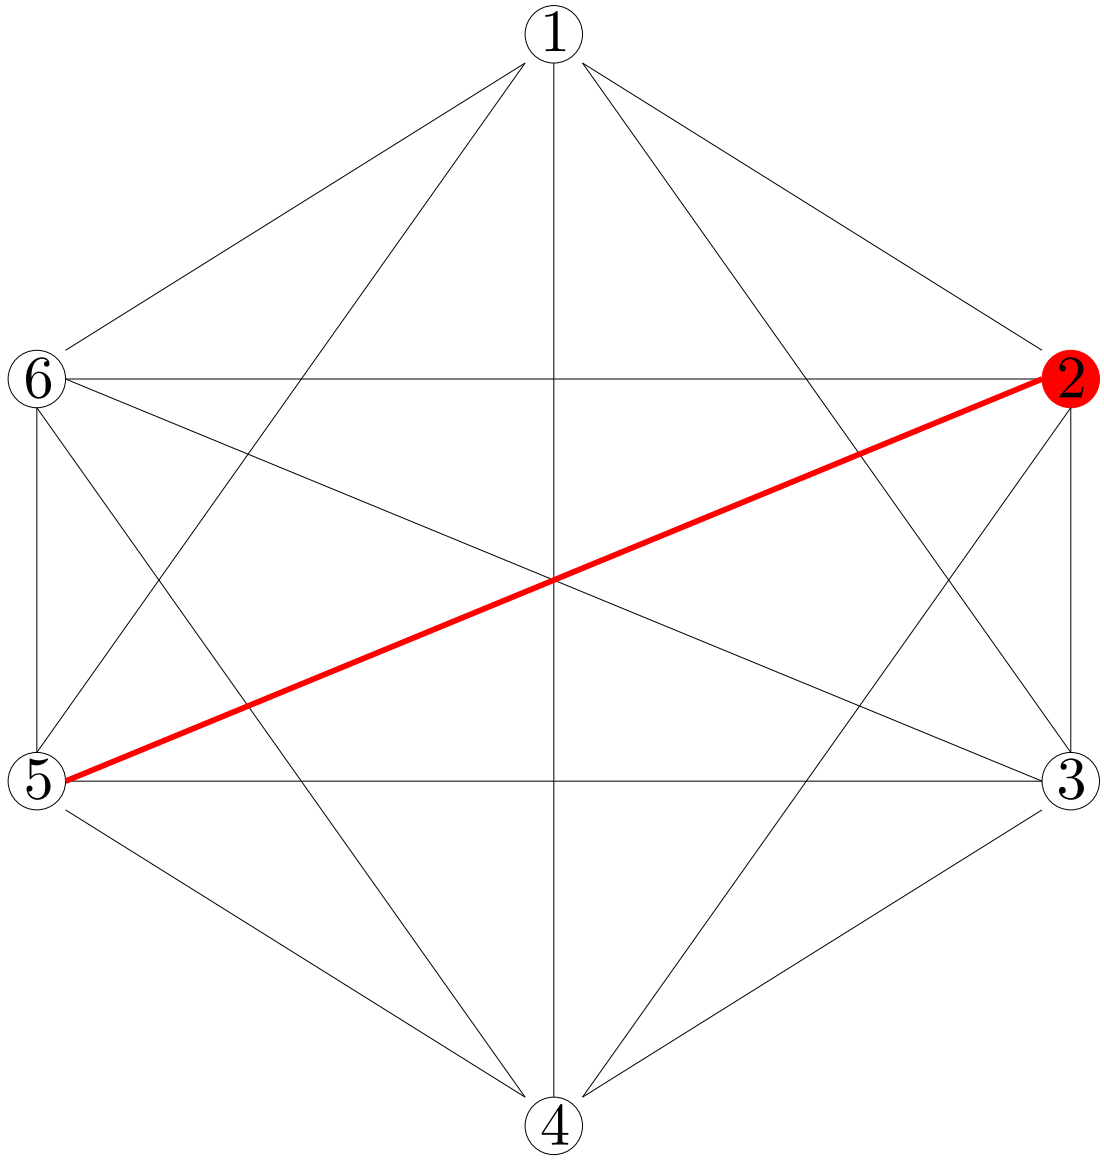

The Aldous-Broder algorithm

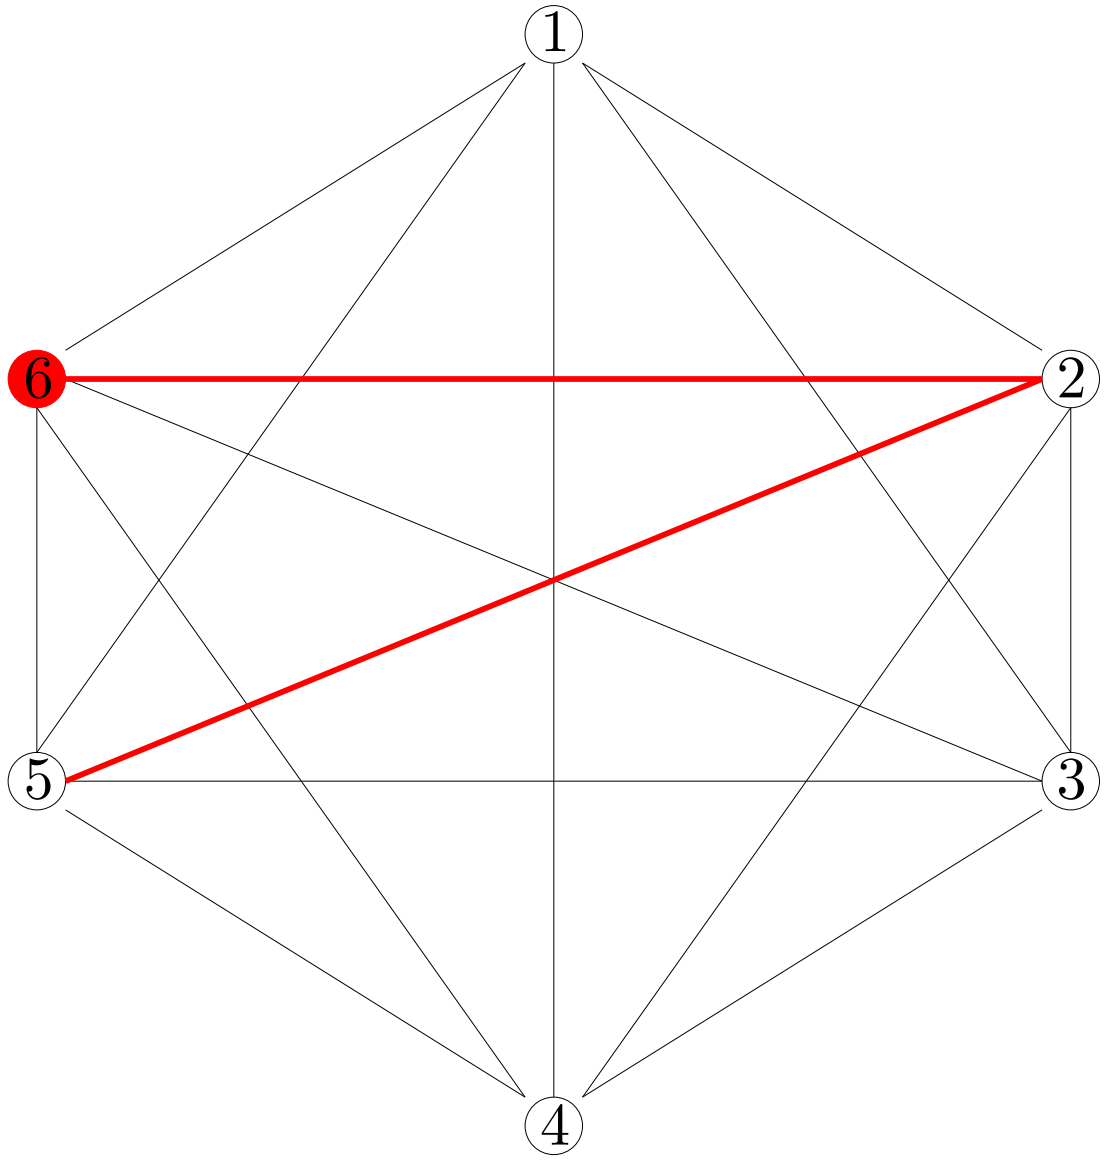

The Aldous-Broder algorithm

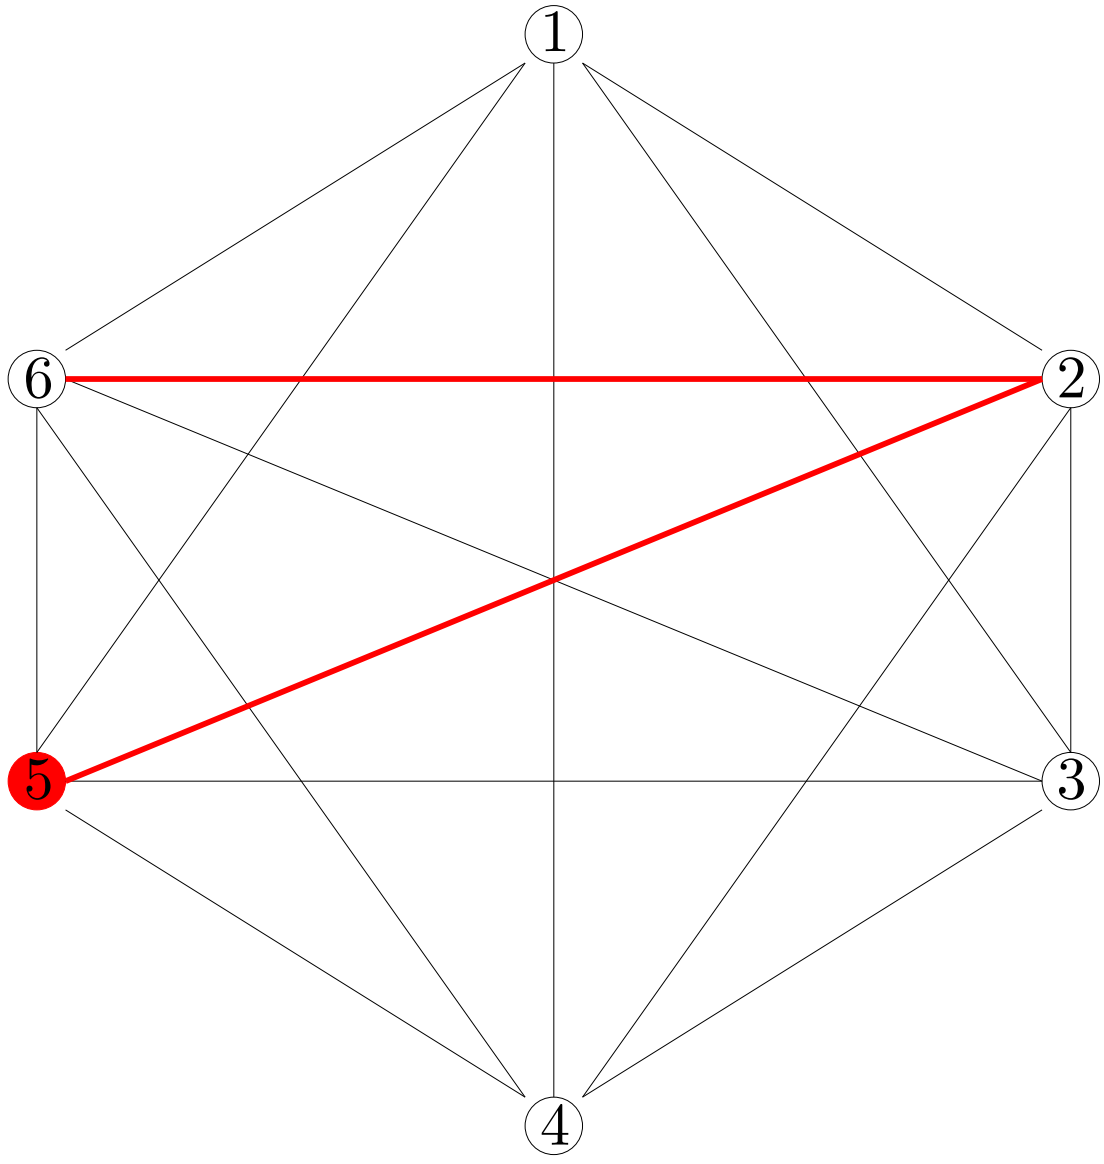

The Aldous-Broder algorithm

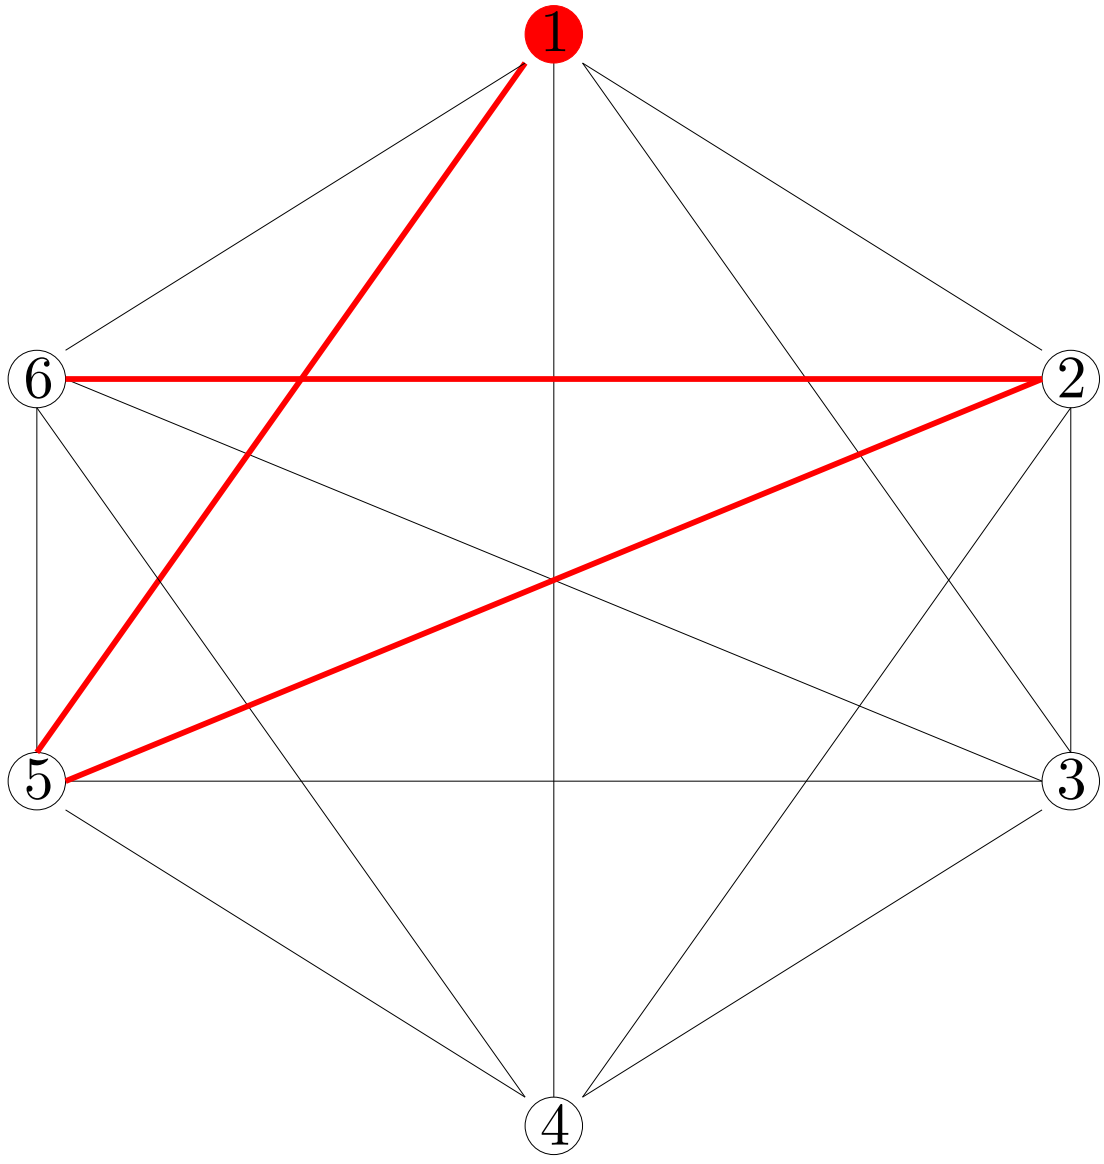

The Aldous-Broder algorithm

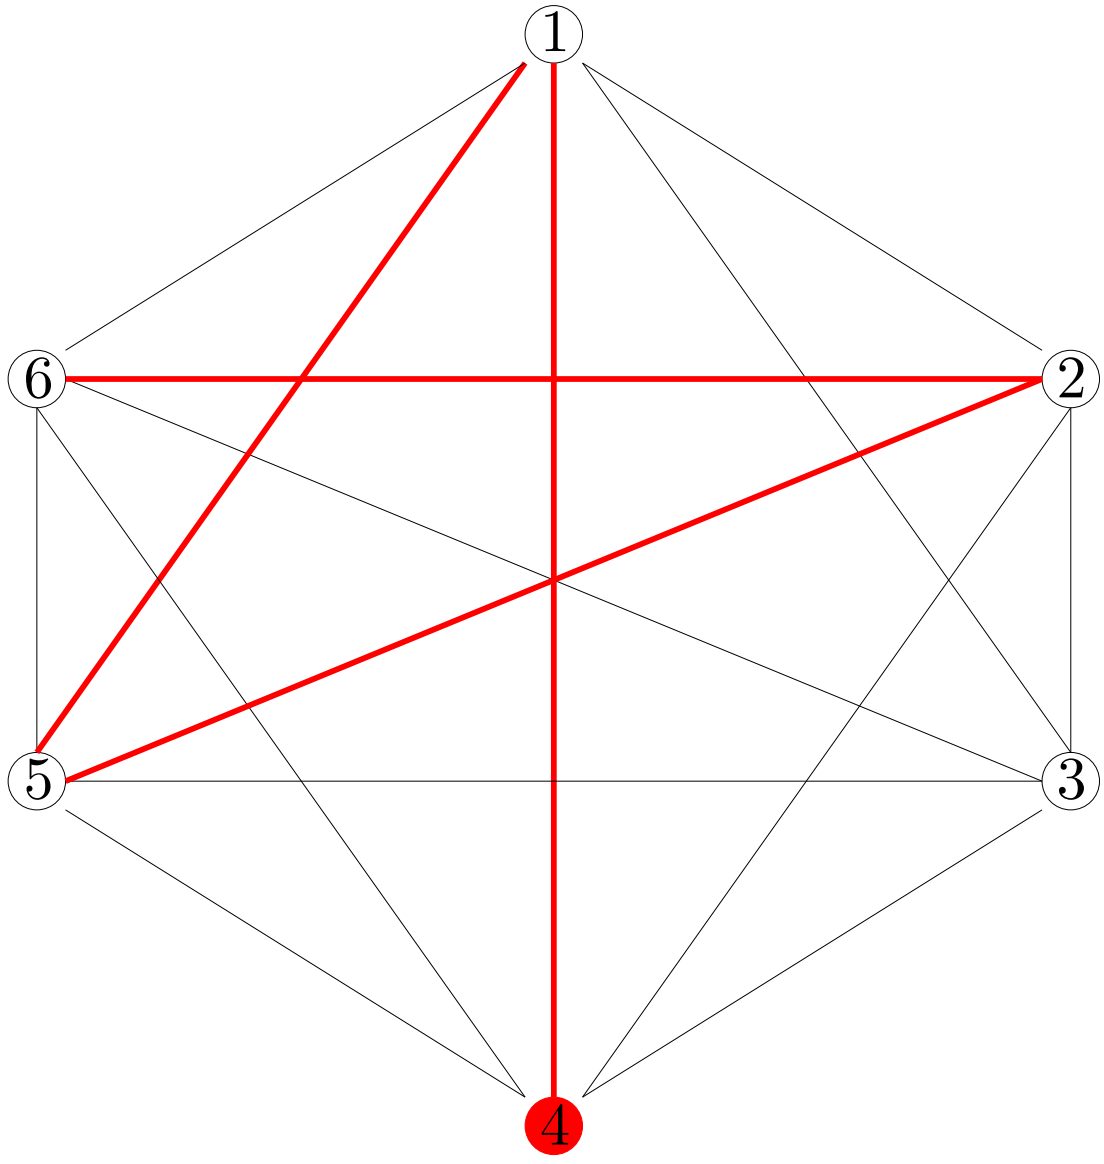

The Aldous-Broder algorithm

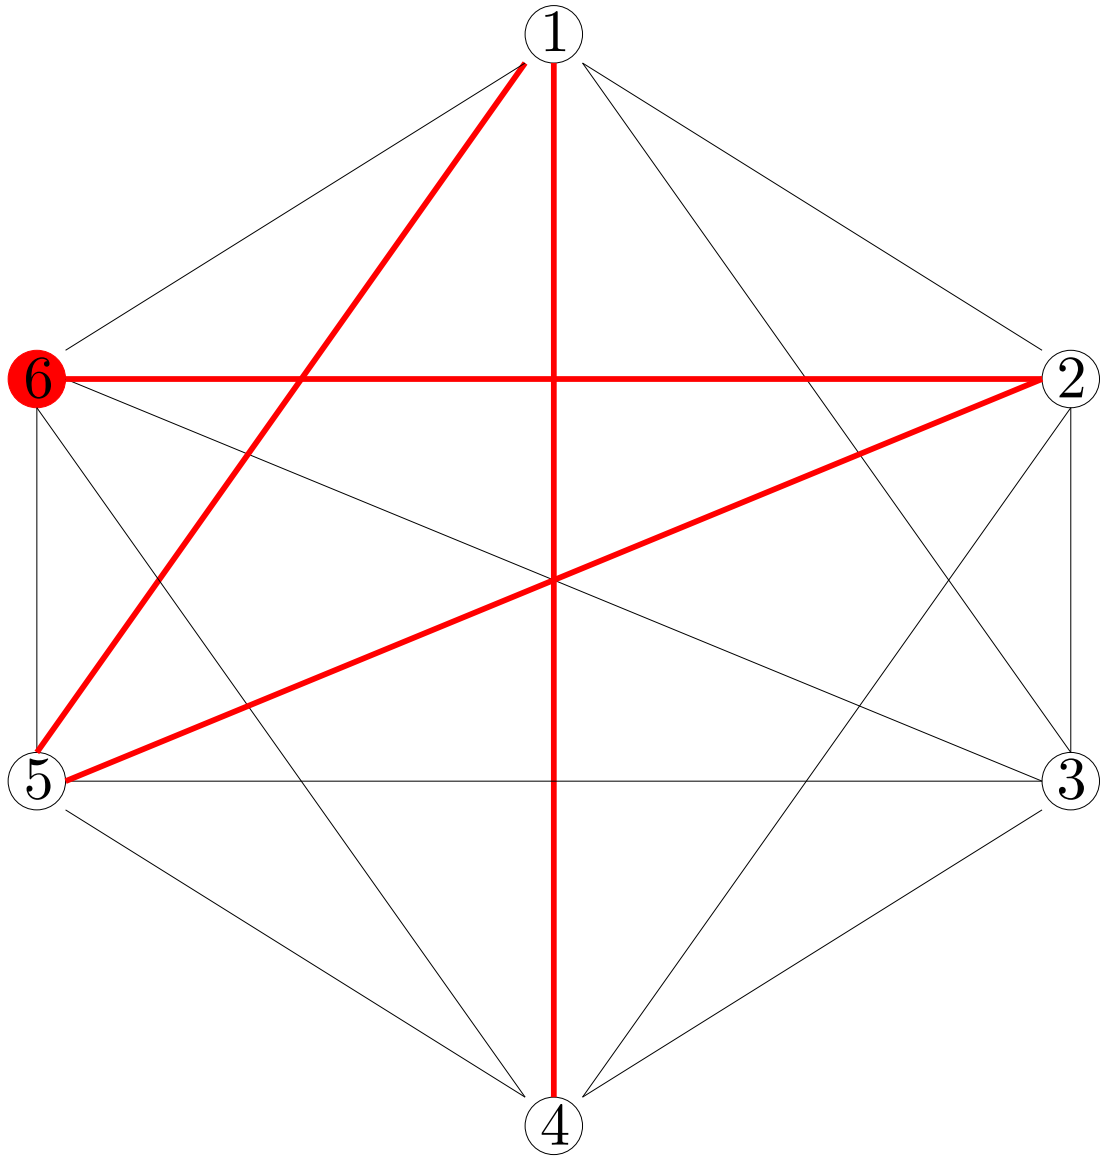

The Aldous-Broder algorithm

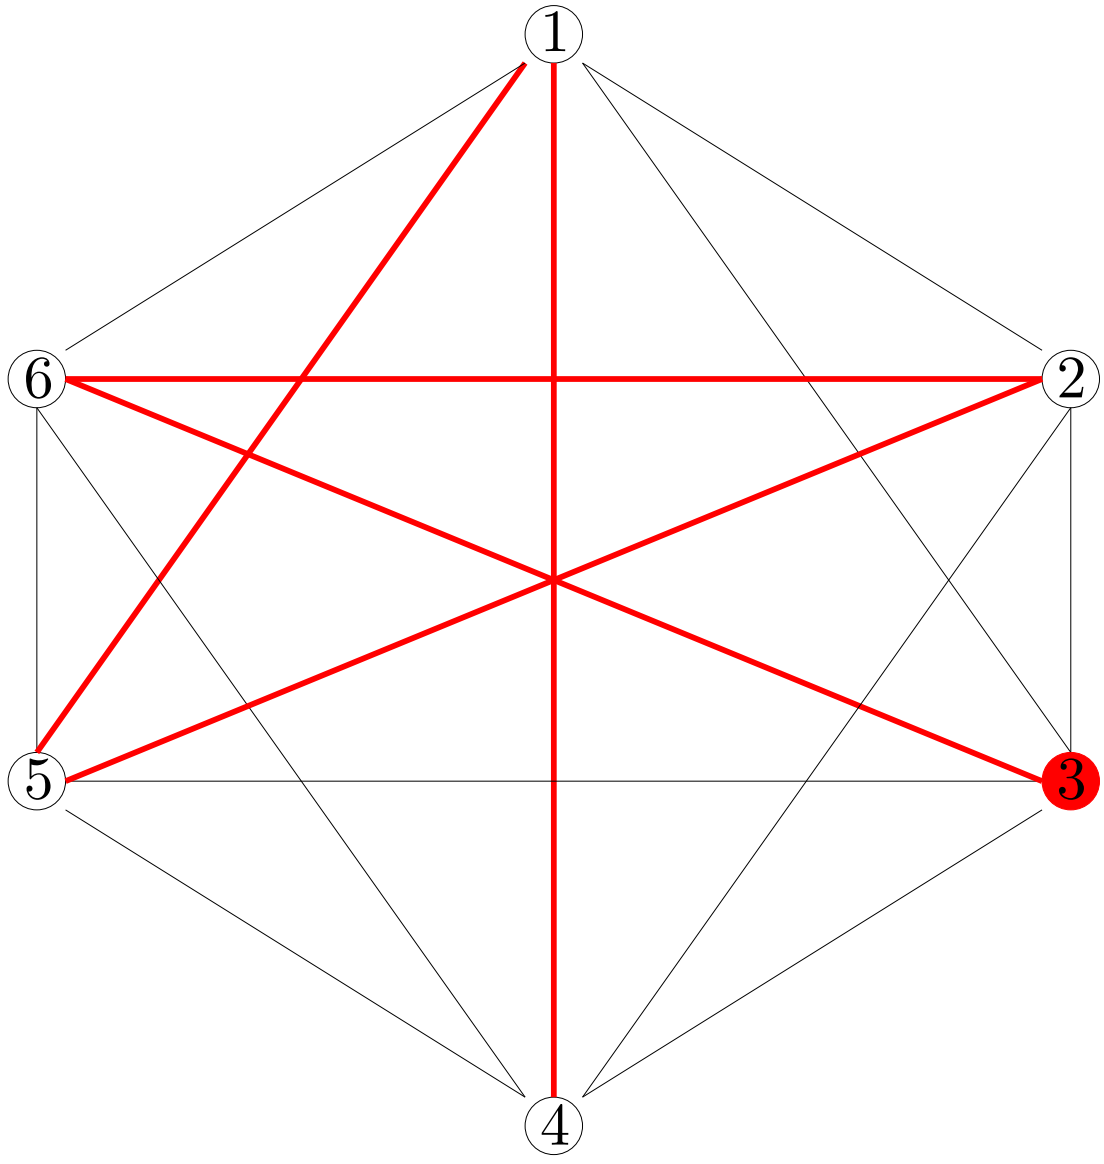

The Aldous-Broder algorithm

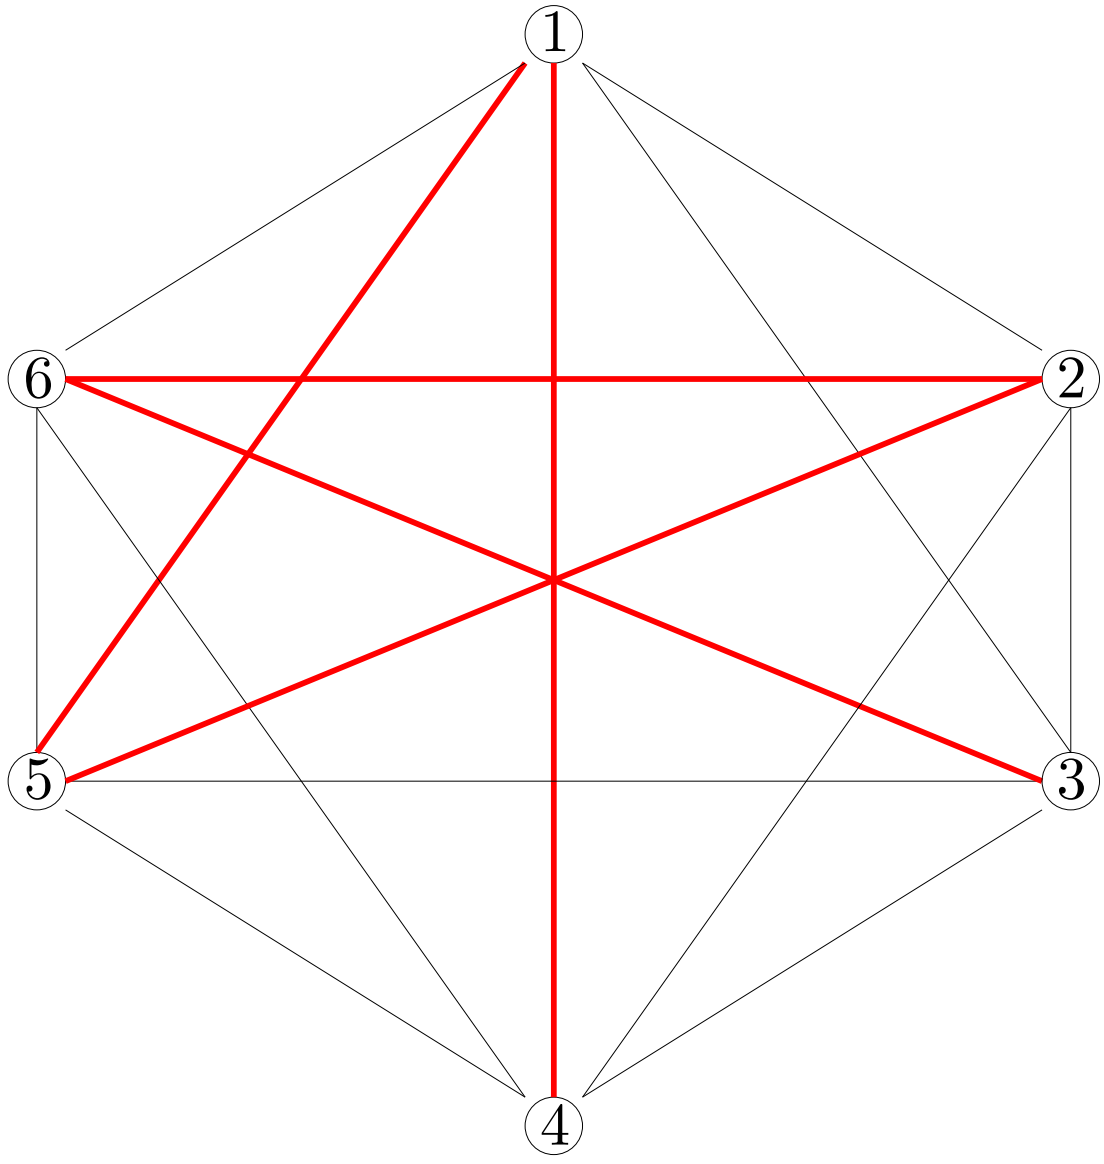

Aldous' algorithm

Consider $n = 10$.

Aldous' algorithm

$V_2 = 1$ with probability 1

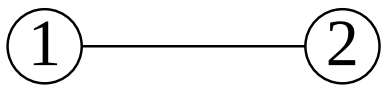

Aldous' algorithm

$$V_3 = \begin{cases} 1 & \text{with probability } 1/10 \\ 2 & \text{with probability } 9/10 \end{cases}$$

Aldous' algorithm

$$V_4 = \begin{cases} j & \text{with probability } 1/10,\ 1 \le j \le 2 \\ 3 & \text{with probability } 8/10 \end{cases}$$

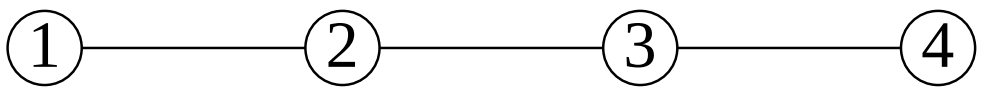

Aldous' algorithm

$$V_5 = \begin{cases} j & \text{with probability } 1/10,\ 1 \leq j \leq 3 \\ 4 & \text{with probability } 7/10 \end{cases}$$

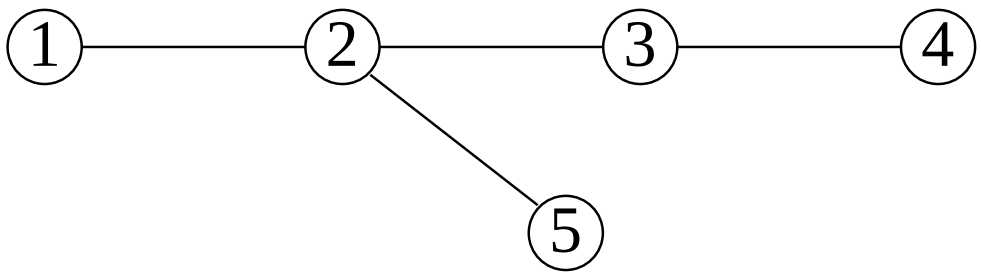

Aldous' algorithm

$$V_6 = \begin{cases} j & \text{with probability } 1/10,\ 1 \le j \le 4 \\ 5 & \text{with probability } 6/10 \end{cases}$$

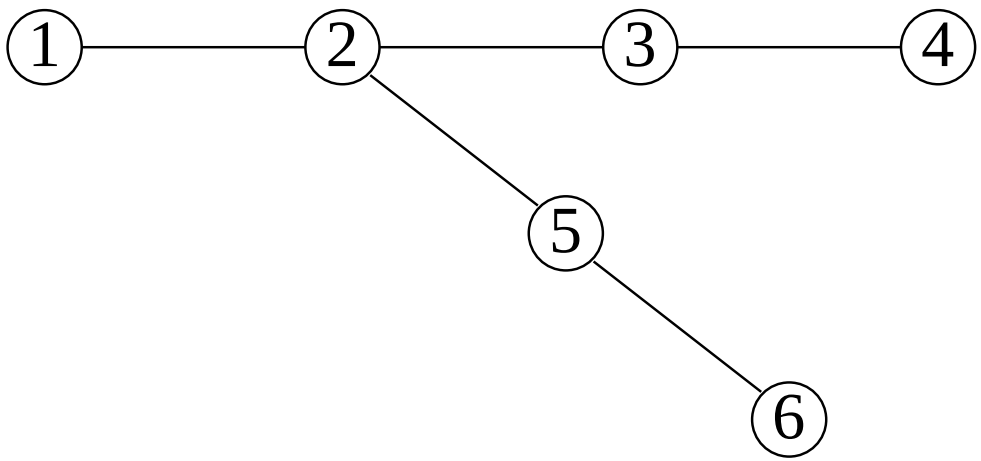

Aldous' algorithm

$$V_7 = \begin{cases} j & \text{with probability } 1/10,\ 1 \le j \le 5 \\ 6 & \text{with probability } 5/10 \end{cases}$$

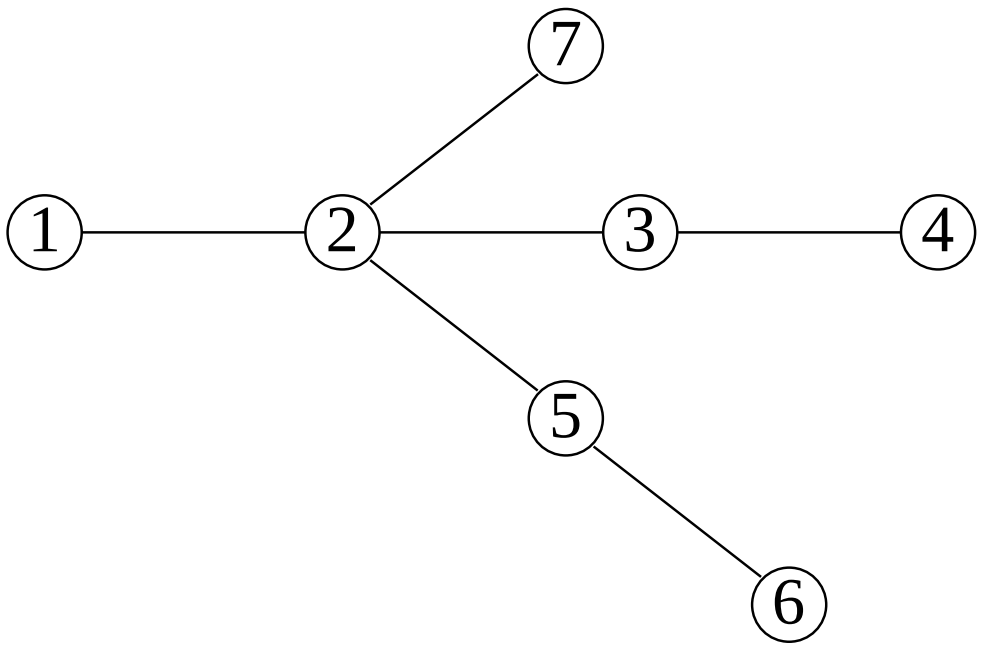

Aldous' algorithm

$$V_8 = \begin{cases} j & \text{with probability } 1/10,\ 1 \le j \le 6 \\ 7 & \text{with probability } 4/10 \end{cases}$$

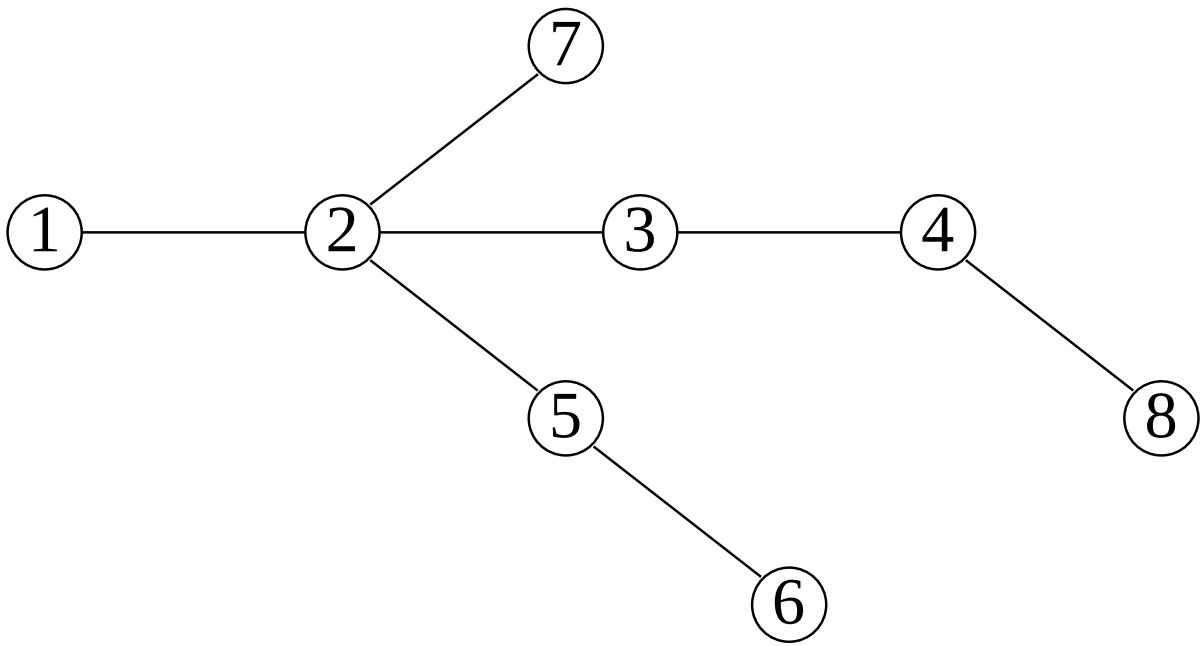

Aldous' algorithm

$$V_9 = \begin{cases} j & \text{with probability } 1/10,\ 1 \le j \le 7 \\ 8 & \text{with probability } 3/10 \end{cases}$$

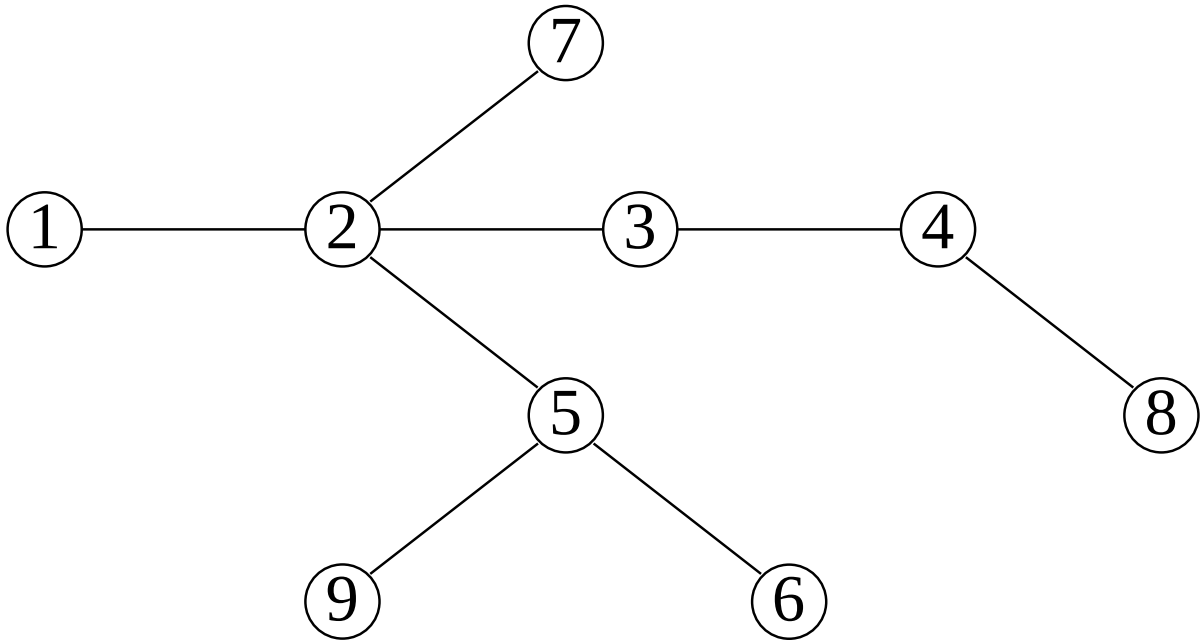

Aldous' algorithm

$$V_{10} = \begin{cases} j & \text{with probability } 1/10,\ 1 \le j \le 8 \\ 9 & \text{with probability } 2/10 \end{cases}$$

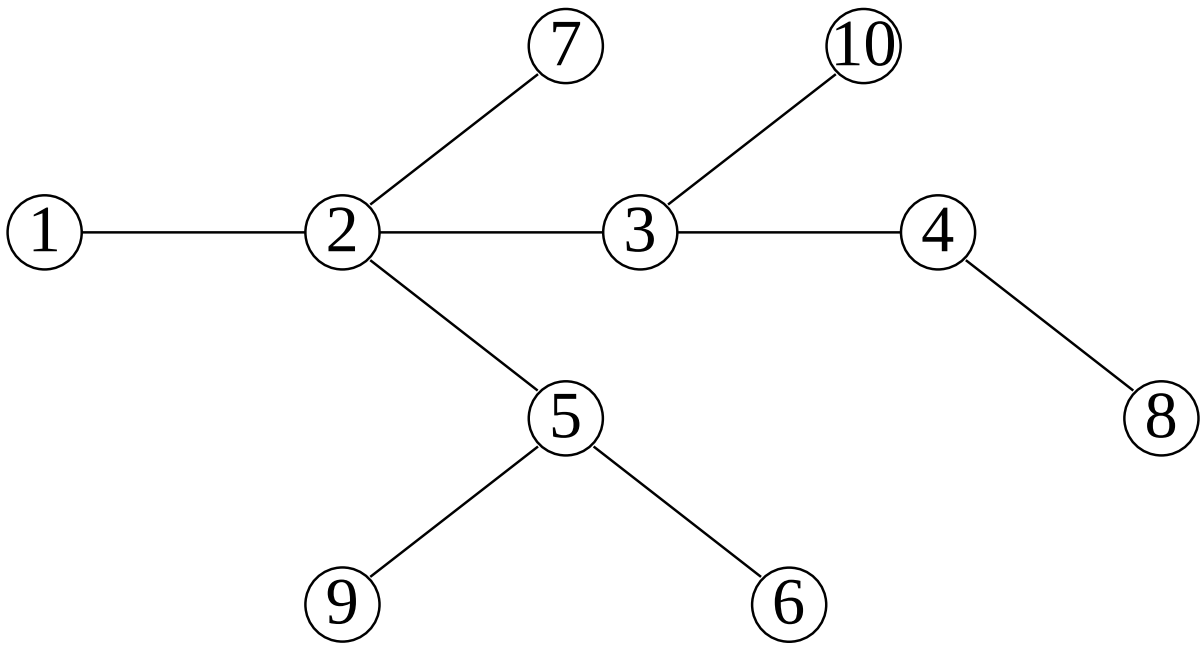

Aldous' algorithm

Permute.

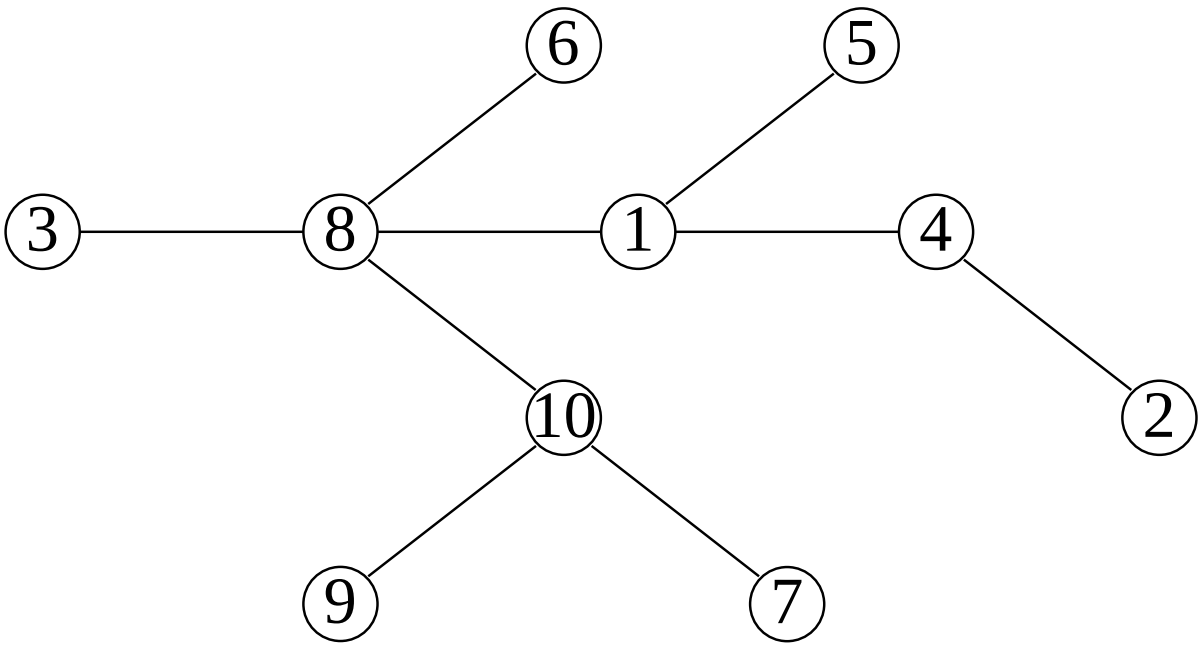

Typical distances

Consider the tree before we permute. Let

$$J_n = \inf\{i \geq 2 : V_i \neq i - 1\}.$$

We can use J_n to give us an idea of typical distances in the tree.

In our example, $J_{10} = 4$:

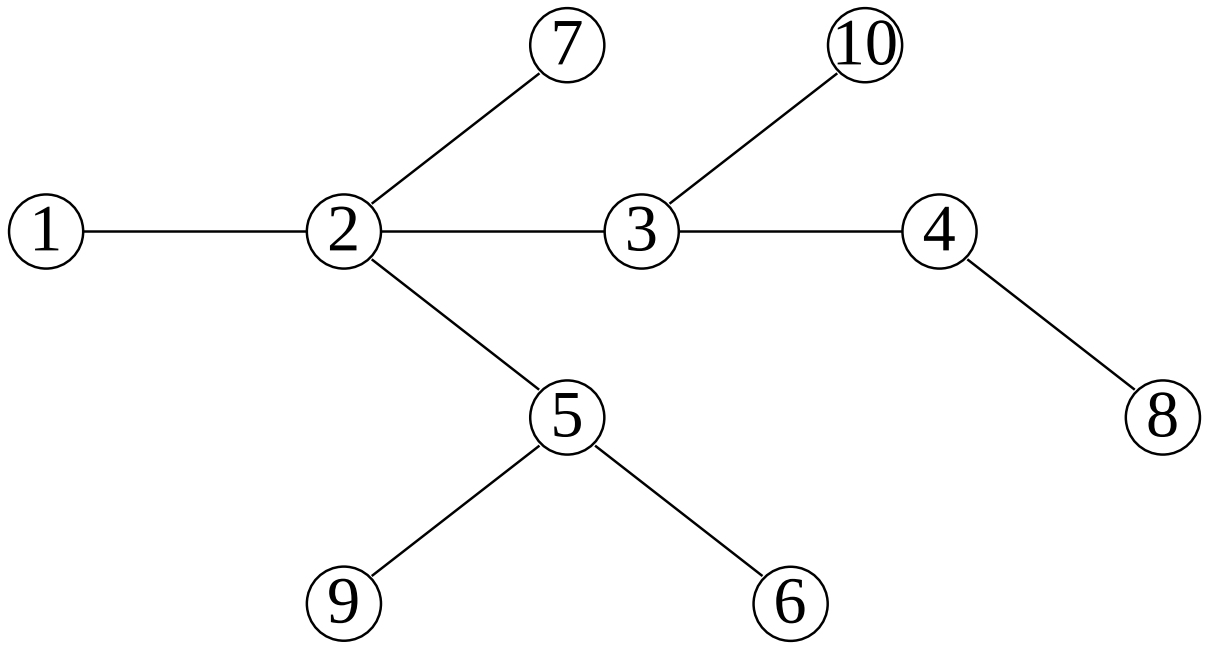

Typical distances

For $2 \leq i \leq n$, connect vertex i to vertex V_i such that

$$V_i = \begin{cases} i-1 \text{ with probability } 1-(i-2)/n \\ \text{uniform on } \{1,2,\ldots,i-2\} \text{ otherwise.} \end{cases}$$

$J_n = \inf\{i \geq 2 : V_i \neq i-1\}$

Proposition

$n^{-1/2} J_n$ converges in distribution as $n \to \infty$.

Typical distances

Once we have built this first stick of consecutive labels, we pick a uniform starting point along that stick and attach a new stick with a random length, and so on.

Typical distances

Once we have built this first stick of consecutive labels, we pick a uniform starting point along that stick and attach a new stick with a random length, and so on.

Imagine now that edges in the tree have length 1. The proposition suggests that rescaling edge-lengths by $n^{-1/2}$ will give some sort of limit for the whole tree. The limiting version of the algorithm is as follows.

Stick-breaking procedure

Take an inhomogeneous Poisson process on $\mathbb{R}^+$ of intensity t at t.

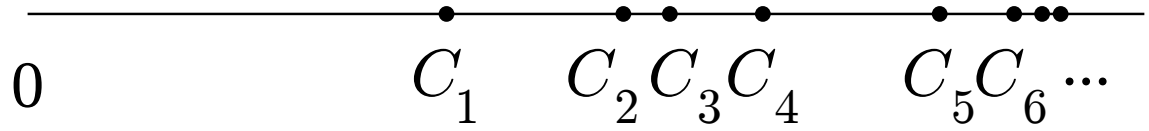

Stick-breaking procedure

Take an inhomogeneous Poisson process on $\mathbb{R}^+$ of intensity t at t.

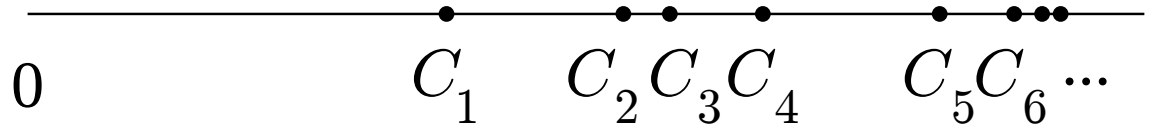

$\mathbb{P}(C_1 > x) = \exp\left(-\int_0^x t dt\right) = \exp(-x^2/2).$

Stick-breaking procedure

Take an inhomogeneous Poisson process on $\mathbb{R}^+$ of intensity t at t.

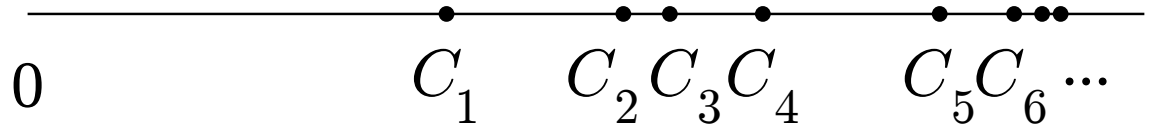

$\mathbb{P}(C_1 > x) = \exp\left(-\int_0^x t dt\right) = \exp(-x^2/2)$.

Consider the line-segments $[0, C_1)$, $[C_1, C_2), \ldots$.

Stick-breaking procedure

Take an inhomogeneous Poisson process on $\mathbb{R}^+$ of intensity t at t.

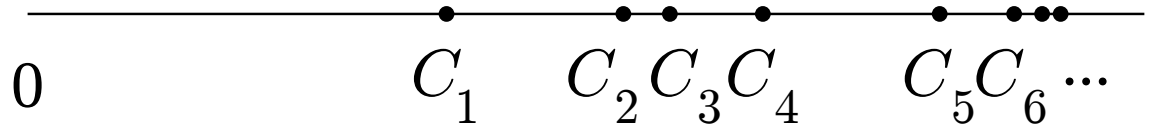

$\mathbb{P}(C_1 > x) = \exp\left(-\int_0^x t dt\right) = \exp(-x^2/2)$.

Consider the line-segments $[0, C_1)$, $[C_1, C_2), \ldots$.

Start from $[0, C_1)$ and proceed inductively.

Stick-breaking procedure

Take an inhomogeneous Poisson process on $\mathbb{R}^+$ of intensity t at t.

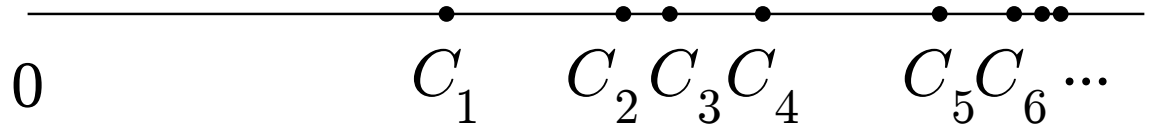

$\mathbb{P}(C_1 > x) = \exp\left(-\int_0^x t dt\right) = \exp(-x^2/2)$.

Consider the line-segments $[0, C_1)$, $[C_1, C_2), \ldots$.

Start from $[0, C_1)$ and proceed inductively.

For $i \geq 2$, attach $[C_{i-1}, C_i)$ at a random point chosen uniformly over the existing tree.

Stick-breaking procedure

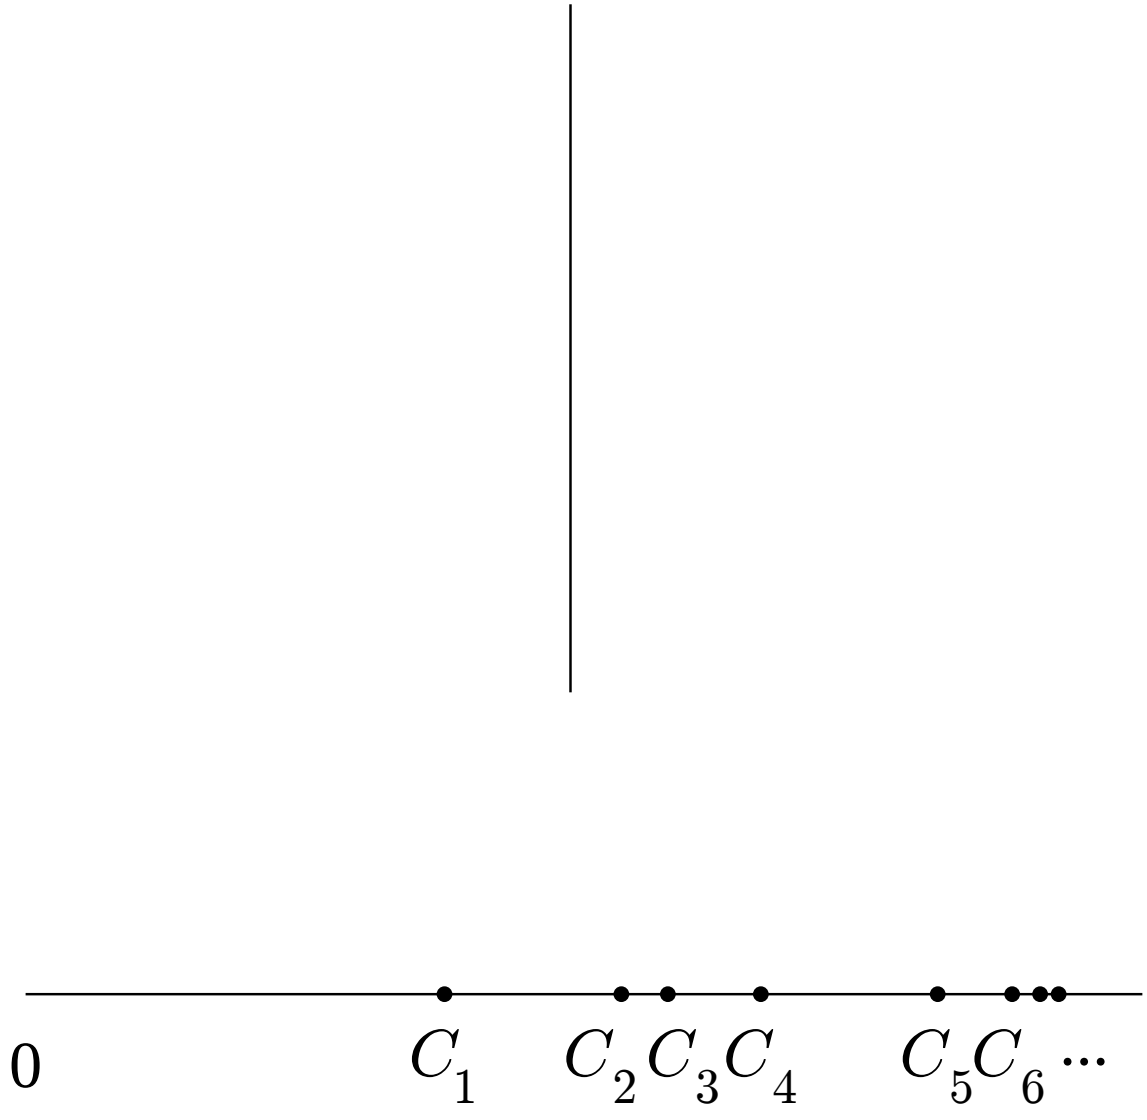

Stick-breaking procedure

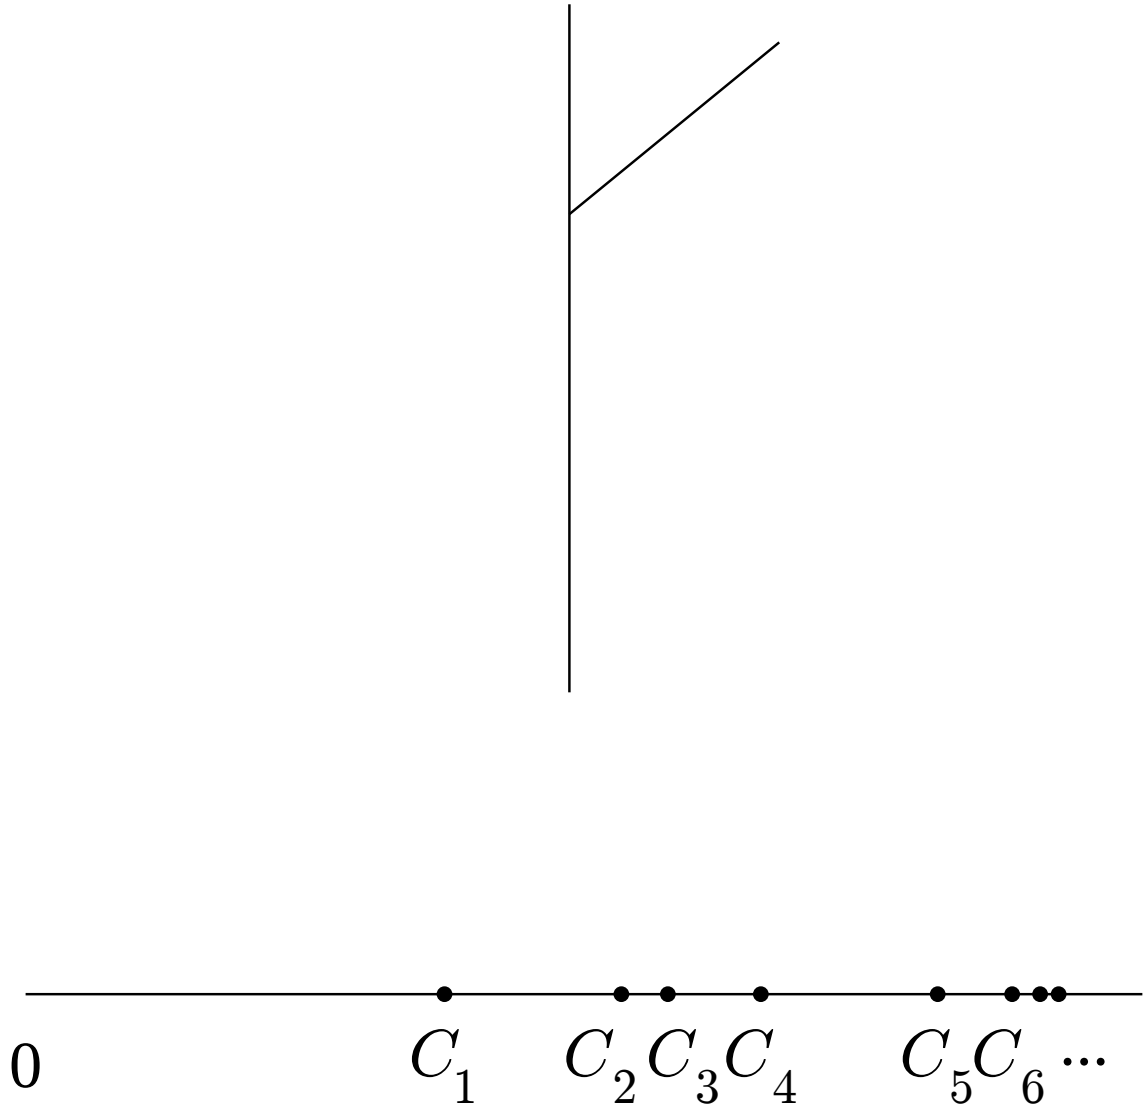

Stick-breaking procedure

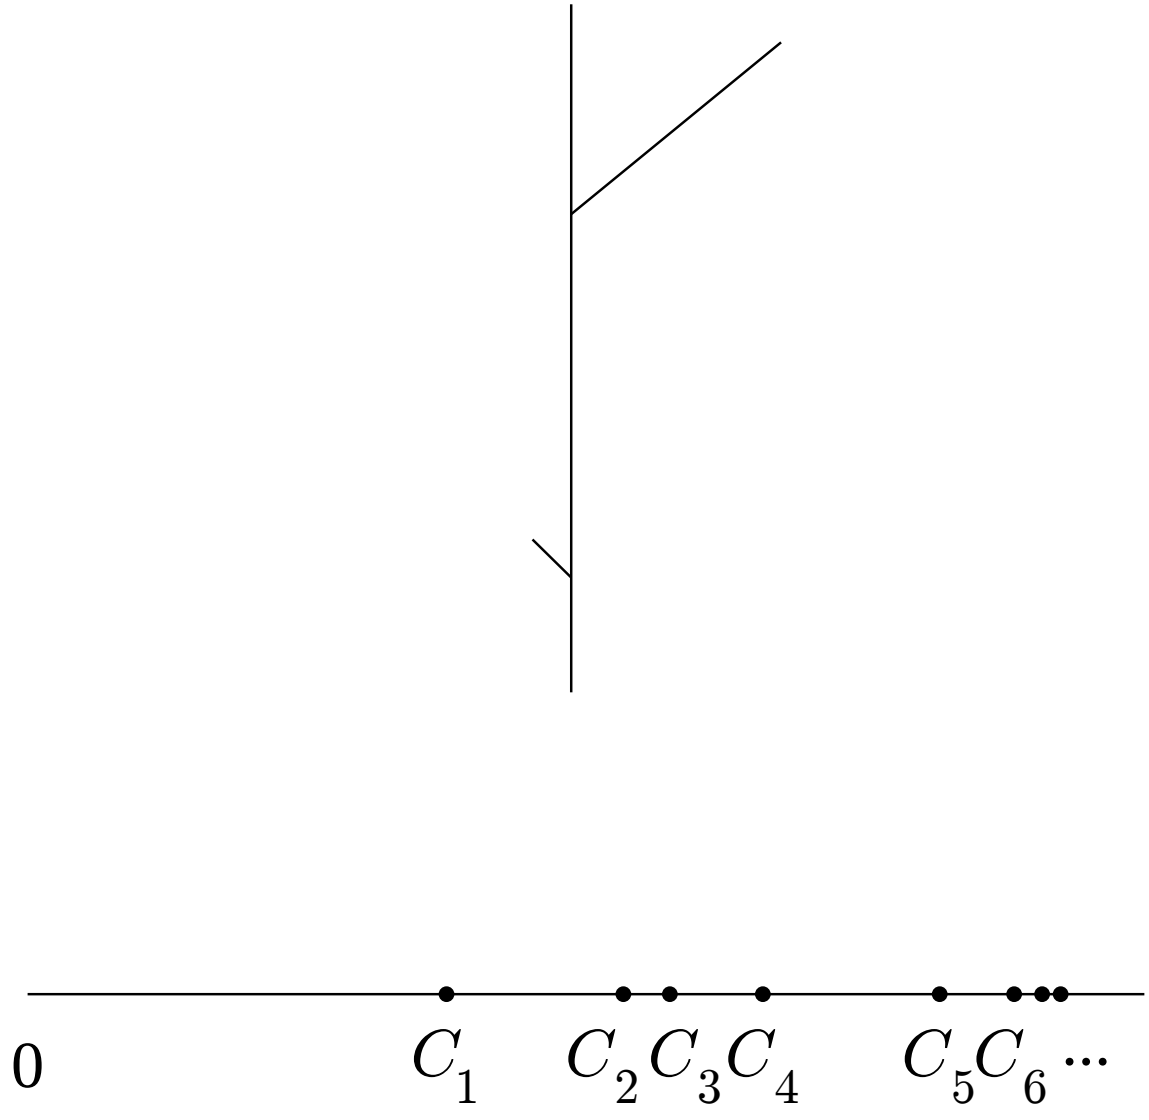

Stick-breaking procedure

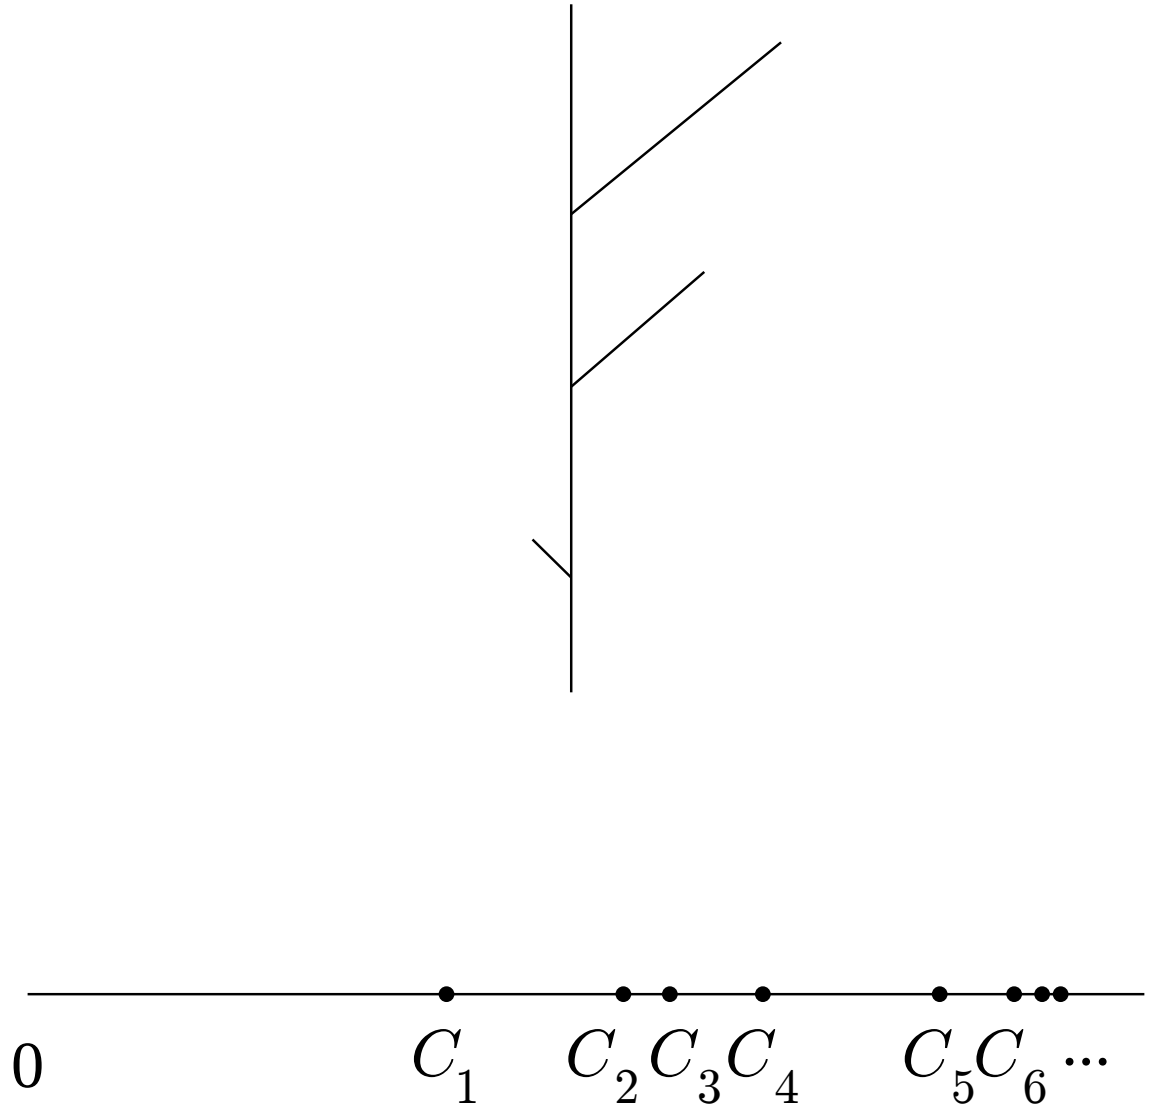

Stick-breaking procedure

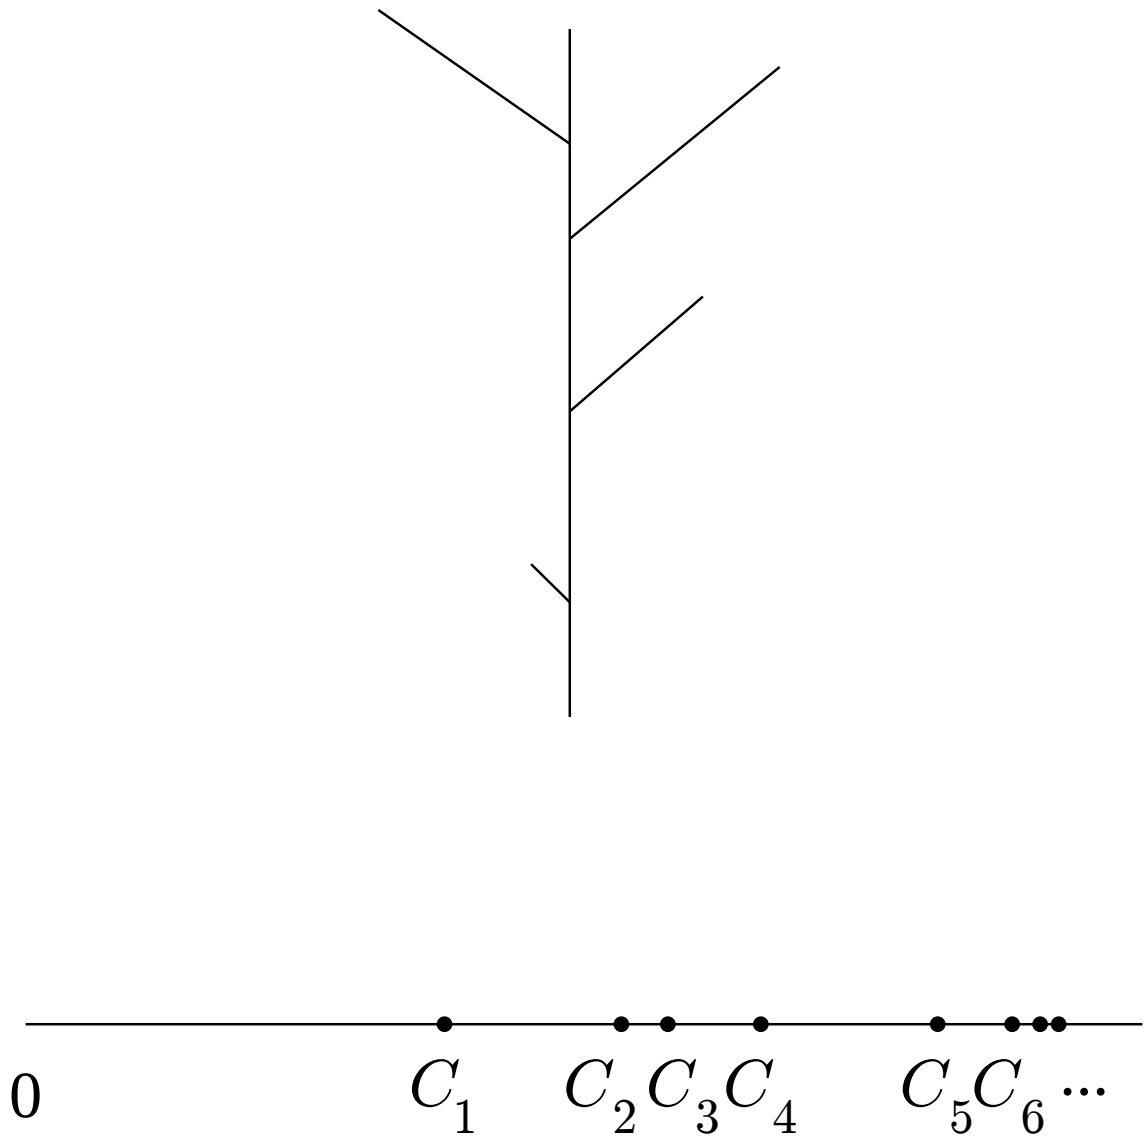

Stick-breaking procedure

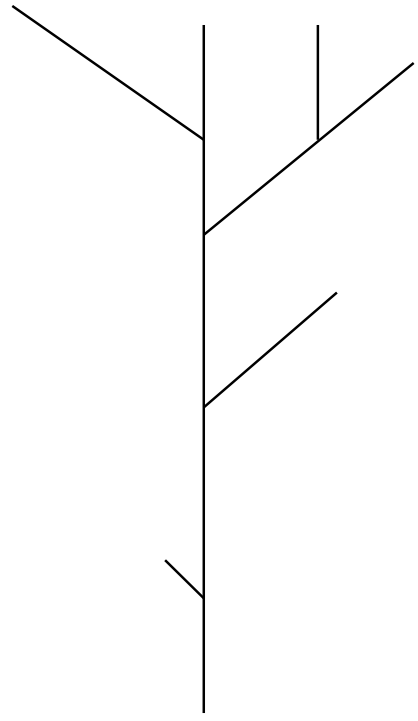

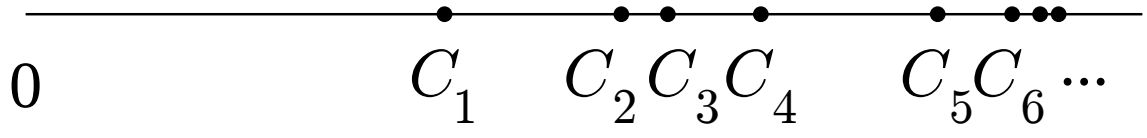

Stick-breaking procedure

Start from $[0, C_1)$ and proceed inductively.

For $i \geq 2$, attach $[C_{i-1}, C_i)$ at a random point chosen uniformly over the existing tree.

Take the union of all the branches, and then take its closure.

Stick-breaking procedure

Start from $[0, C_1)$ and proceed inductively.

For $i \geq 2$, attach $[C_{i-1}, C_i)$ at a random point chosen uniformly over the existing tree.

Take the union of all the branches, and then take its closure.

This procedure gives (a rather informally expressed!) definition of Aldous' Brownian continuum random tree (CRT) which will be the key object in this lecture.

The Brownian continuum random tree

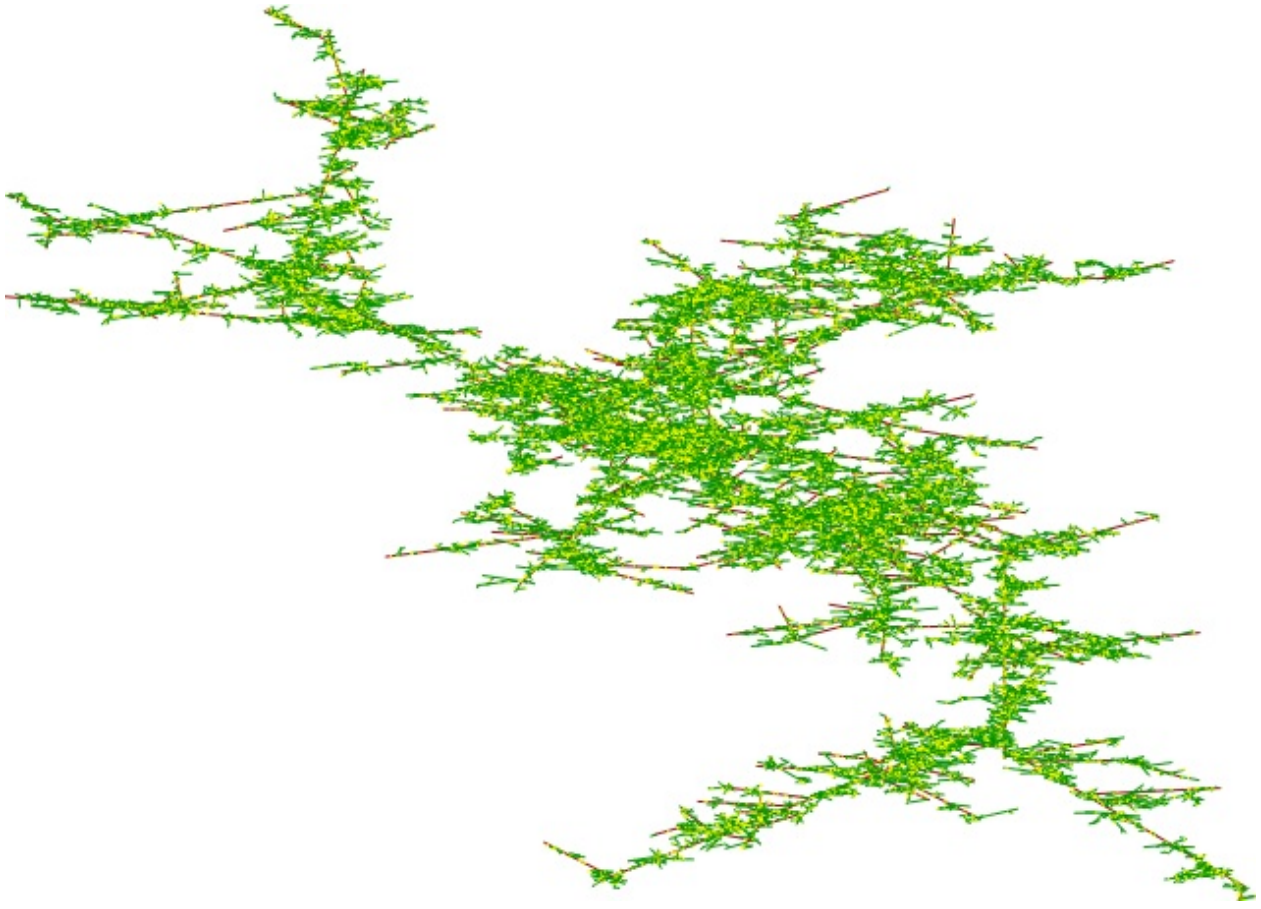

[Picture by Grégory Miermont]

Scaling limits of random trees

The purpose of today's lecture is

- to show that the Brownian CRT is indeed the scaling limit of the uniform random tree
- to explain in what sense this is true
- and also to show that the limit is universal for a whole family of random trees.

An aside: a prototypical scaling limit

Suppose that $Z_1, Z_2, \ldots$ are independent and identically distributed random variables with mean 0 and variance 1. Let $X(0) = 0$ and $X(k) = \sum_{i=1}^{k} Z_i$. Then $(X(k), k \geq 0)$ is a random walk.

An aside: a prototypical scaling limit

Suppose that $Z_1, Z_2, \ldots$ are independent and identically distributed random variables with mean 0 and variance 1. Let $X(0) = 0$ and $X(k) = \sum_{i=1}^{k} Z_i$. Then $(X(k), k \geq 0)$ is a random walk.

Donsker's theorem. As $n \to \infty$,

$$\frac{1}{\sqrt{n}}(X(\lfloor nt \rfloor), t \geq 0) \xrightarrow{d} (W(t), t \geq 0),$$

where $(W(t), t \geq 0)$ is a standard Brownian motion.

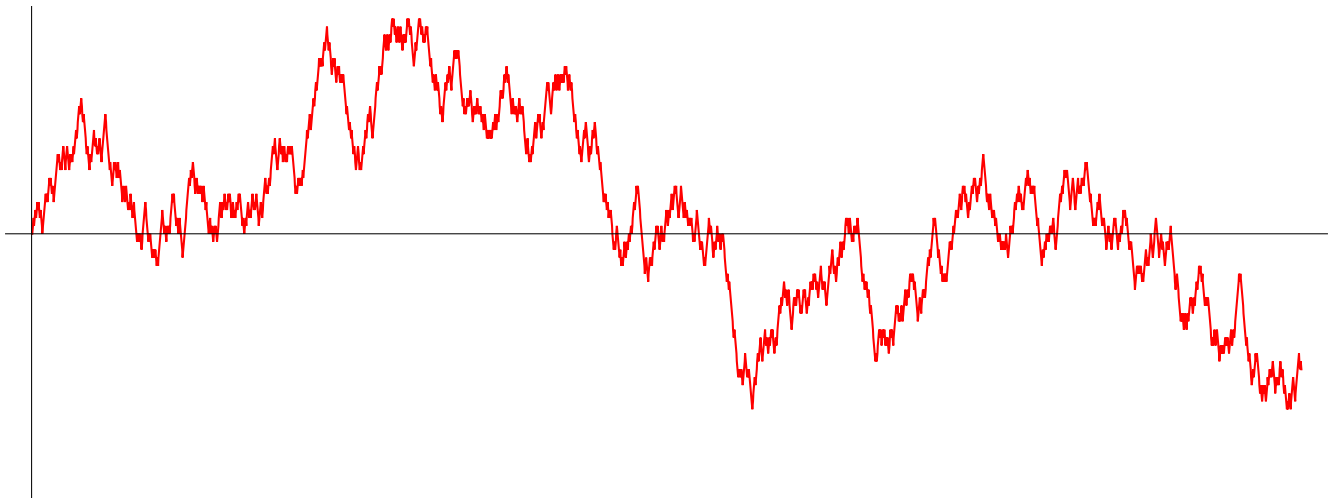

Specifying distributions

We want to deal with random objects which are not real-valued. We will typically specify their distributions via expectations of real-valued test-functions.

Specifying distributions

We want to deal with random objects which are not real-valued. We will typically specify their distributions via expectations of real-valued test-functions.

To take Brownian motion as an example: the random object takes values in $C(\mathbb{R}_+, \mathbb{R})$ and its distribution is specified by the values of

$$\mathbb{E}\left[f(W(t), t \geq 0)\right]$$

for functions $f : C(\mathbb{R}_+, \mathbb{R}) \to \mathbb{R}$ in a suitable class.

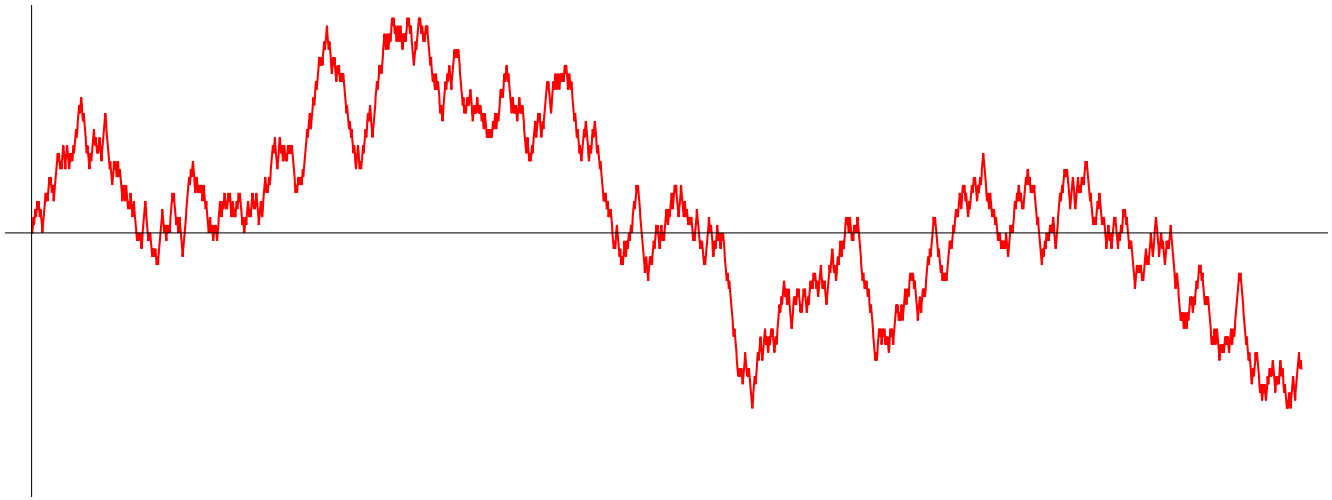

Convergence in distribution

The convergence in distribution in Donsker's theorem means that for all bounded continuous real-valued functions f,

$$\mathbb{E}\left[f\left(\frac{1}{\sqrt{n}}(X(\lfloor nt \rfloor), t \geq 0)\right)\right] \to \mathbb{E}\left[f(W(t), t \geq 0)\right].$$

Convergence in distribution

The convergence in distribution in Donsker's theorem means that for all bounded continuous real-valued functions f,

$$\mathbb{E}\left[f\left(\frac{1}{\sqrt{n}}(X(\lfloor nt\rfloor), t \geq 0)\right)\right] \to \mathbb{E}\left[f(W(t), t \geq 0)\right].$$

(Here, f is continuous for the metric on $D(\mathbb{R}_+, \mathbb{R})$ given by

$$d(x, y) = \sum_{k=1}^{\infty} 2^{-k} \left(\sup_{t\in[0,k]} |x(t) - y(t)| \wedge 1\right),$$

yielding uniform convergence on compact time-intervals).

The scaling limit of the uniform random tree

Theorem. (Aldous (1993); Le Gall (2006)) As $n \to \infty$,

$$\frac{1}{\sqrt{n}} T_n \xrightarrow{d} \mathcal{T},$$

where $\mathcal{T}$ is the Brownian CRT.

Convergence in distribution

What is the the sense of the convergence in distribution

$$\frac{1}{\sqrt{n}} T_n \xrightarrow{d} \mathcal{T} \quad \text{as } n \to \infty?$$

We don't really care about vertex labels (except that they give a combinatorial weight to our trees): the important things are the shape of the tree and the distances between its vertices.

So we are going to think of $(T_n, n \geq 1)$ and $\mathcal{T}$ as metric spaces.

Trees as metric spaces

The vertices of T_n come equipped with a natural metric: the graph distance.

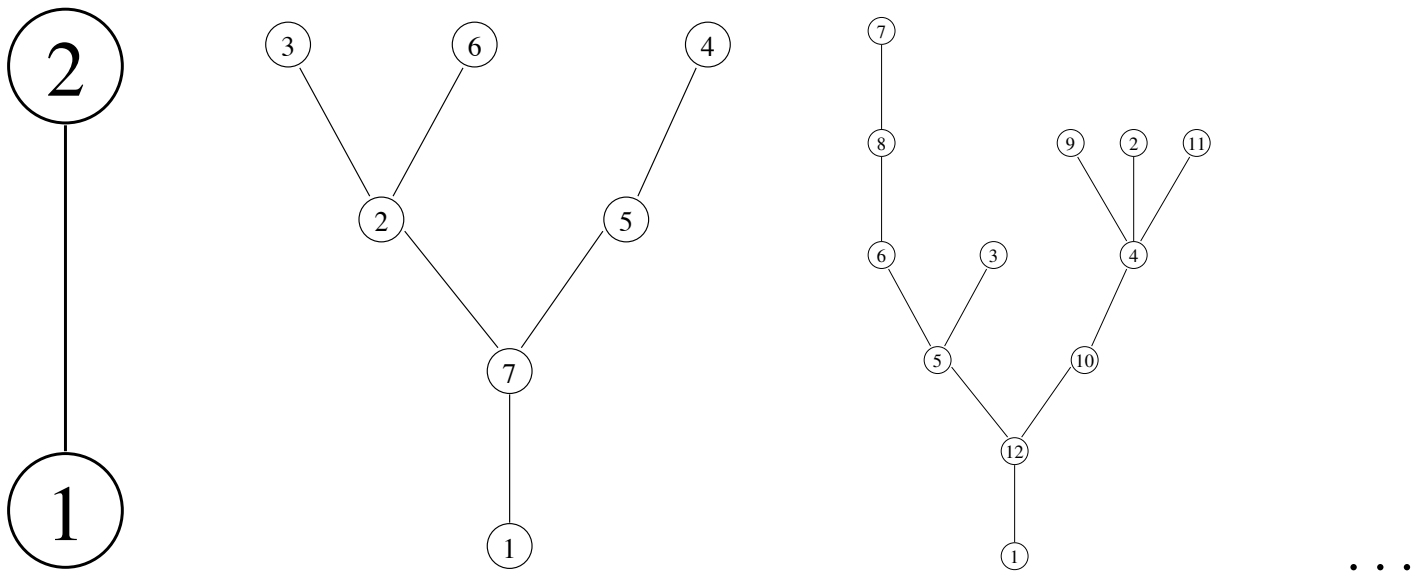

$\ldots$

Write $\frac{1}{\sqrt{n}} T_n$ for the metric space given by the vertices of T_n with the graph distance divided by $\sqrt{n}$.

We will see later how to formulate $\mathcal{T}$ properly as a metric space.

Convergence in distribution

Consider the space, $\mathbb{M}$, of compact metric spaces. We'll define a distance $\mathrm{d_{GH}}$ on $\mathbb{M}$ in a moment. Then

$$\frac{1}{\sqrt{n}} T_n \xrightarrow{d} \mathcal{T} \quad \text{as } n \to \infty$$

will mean that for any bounded function $f : \mathbb{M} \to \mathbb{R}$ which is continuous with respect to $\mathrm{d_{GH}}$, we have

$$\mathbb{E}\left[f\left(\frac{1}{\sqrt{n}} T_n\right)\right] \to \mathbb{E}\left[f\left(\mathcal{T}\right)\right] \quad \text{as } n \to \infty.$$

Measuring the distance between metric spaces

Suppose that (X, d) and (X', d') are compact metric spaces.

Measuring the distance between metric spaces

Suppose that (X, d) and (X', d') are compact metric spaces.

A correspondence R is a subset of $X \times X'$ such that for every $x \in X$, there exists $x' \in X'$ with $(x, x') \in R$ and vice versa.

Measuring the distance between metric spaces

Suppose that (X, d) and (X', d') are compact metric spaces.

A correspondence R is a subset of $X \times X'$ such that for every $x \in X$, there exists $x' \in X'$ with $(x, x') \in R$ and vice versa.

The distortion of R is

$$\mathrm{dis}(R) = \sup\{|d(x, y) - d'(x', y')| : (x, x'), (y, y') \in R\}.$$

Measuring the distance between metric spaces

Suppose that (X, d) and (X', d') are compact metric spaces.

A correspondence R is a subset of $X \times X'$ such that for every $x \in X$, there exists $x' \in X'$ with $(x, x') \in R$ and vice versa.

The distortion of R is

$$\mathrm{dis}(R) = \sup\{|d(x, y) - d'(x', y')| : (x, x'), (y, y') \in R\}.$$

(X, d) and (X', d') are at Gromov-Hausdorff distance less than $\epsilon > 0$ if there exists a correspondence R between X and X' such that $\mathrm{dis}(R) < 2\epsilon$. Write

$$\mathrm{d_{GH}}((X, d), (X', d')) < \epsilon.$$

Coding discrete trees by functions

It turns out to be useful to encode the shapes of our trees by more familiar objects: functions.

Coding discrete trees by functions

It turns out to be useful to encode the shapes of our trees by more familiar objects: functions.

We will do this in two different ways: the contour function and the depth-first walk.

Contour function

The contour function is obtained by simply tracing the "contour" of the tree from left to right at speed 1, so that we pass along each edge twice. Record the distance from the root at each time to get $(C(t), 0 \leq t \leq 2(n-1))$.

Contour function

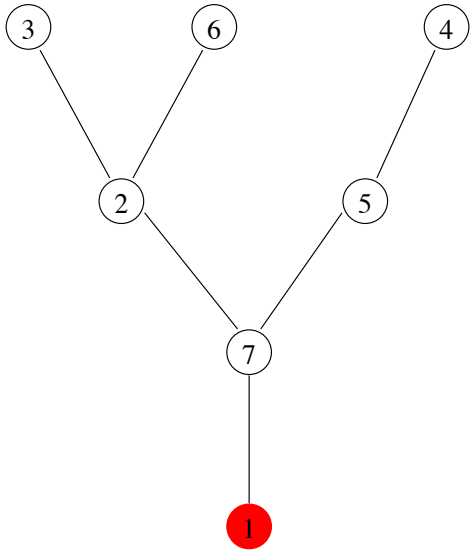

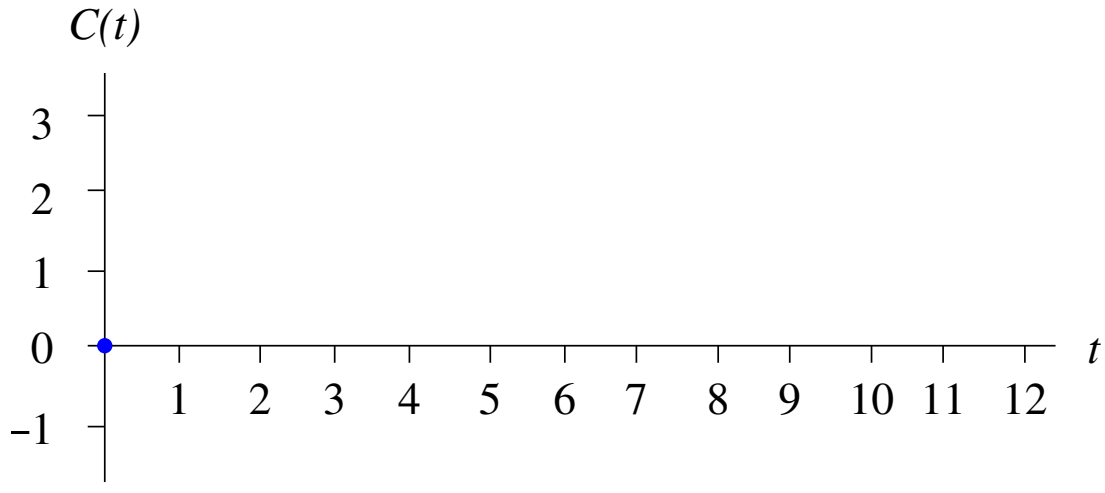

Contour function

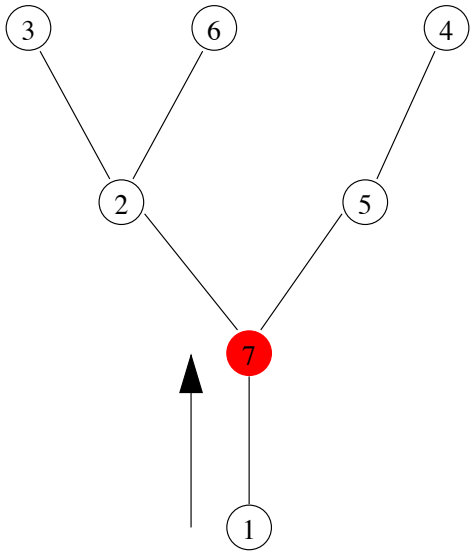

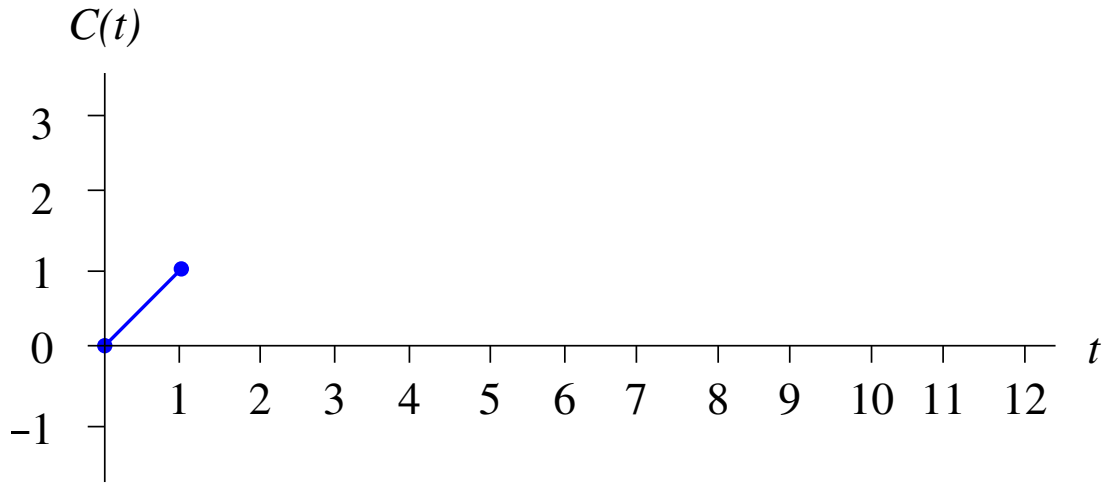

Contour function

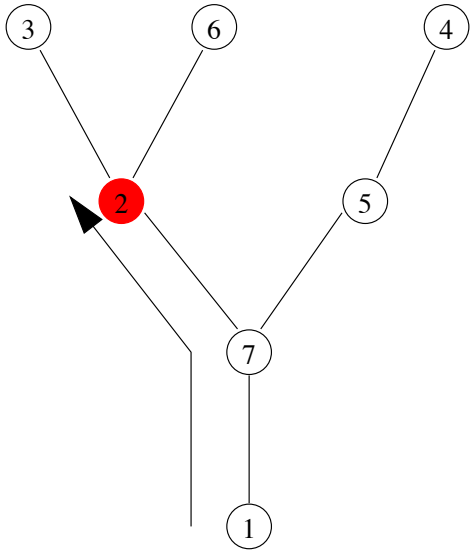

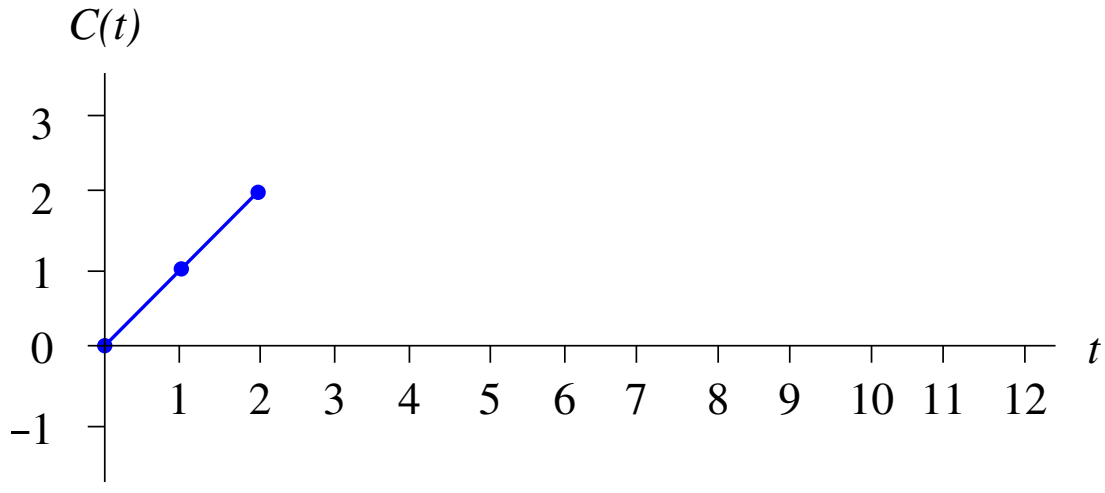

Contour function

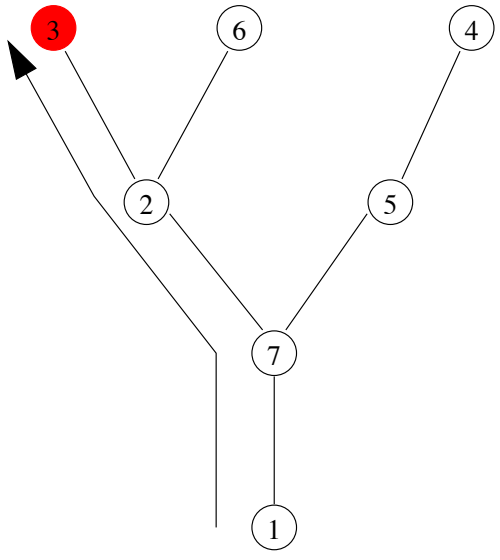

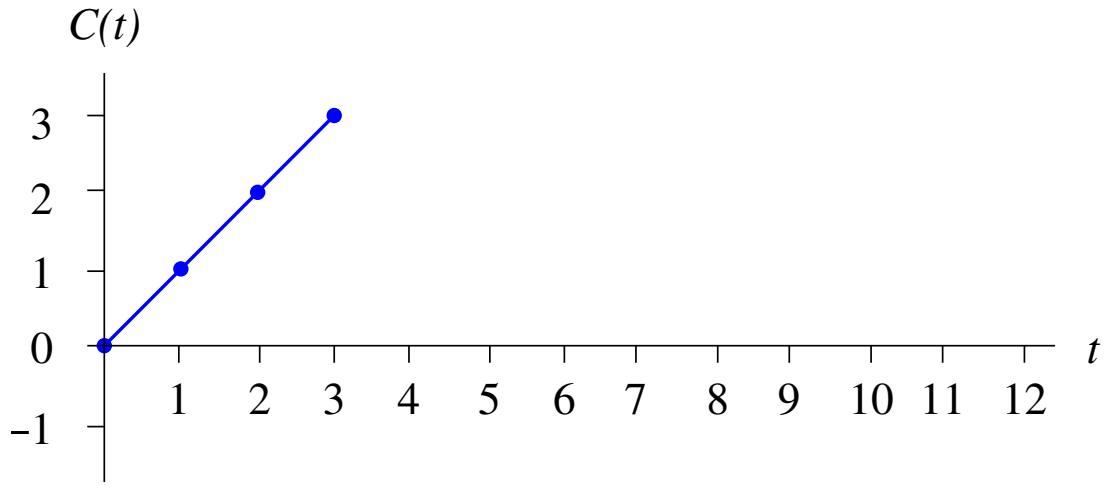

Contour function

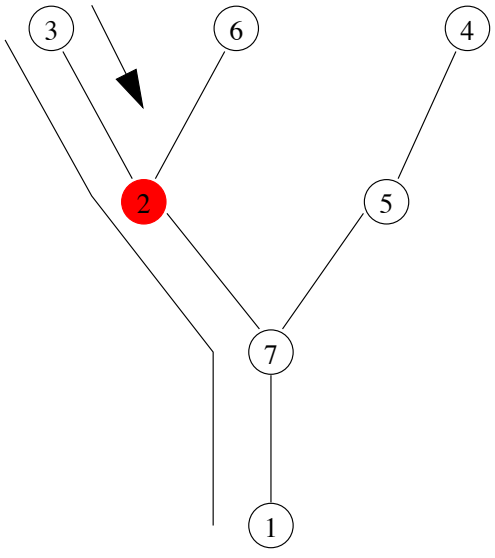

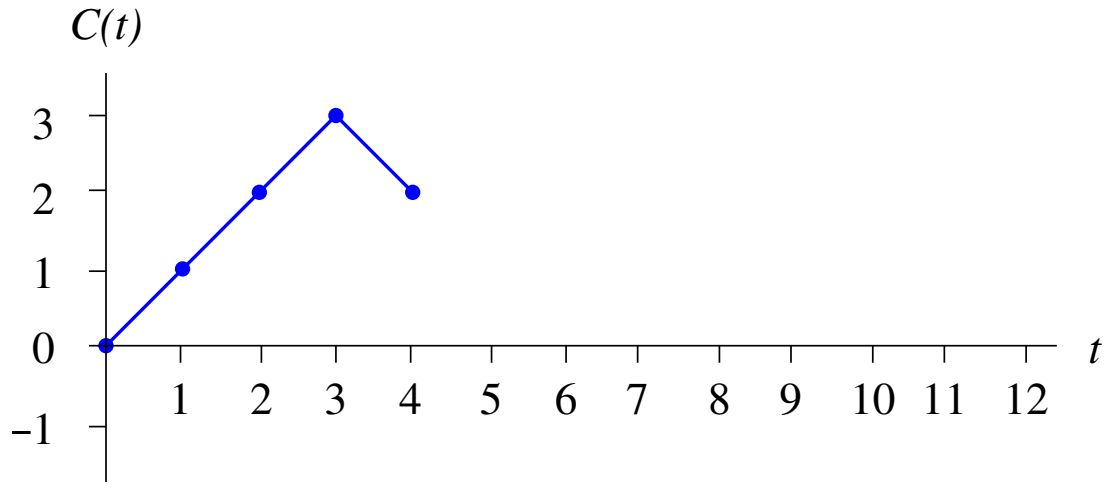

Contour function

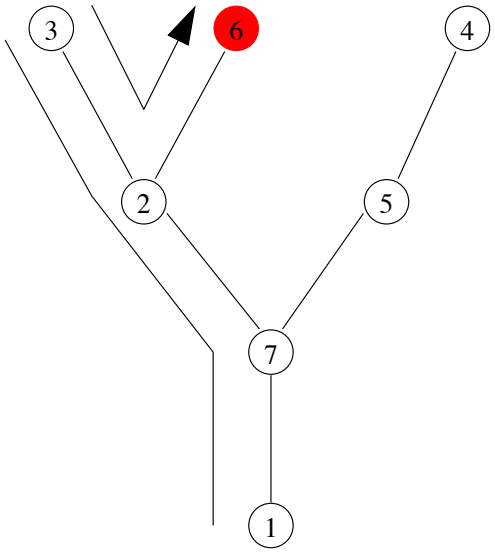

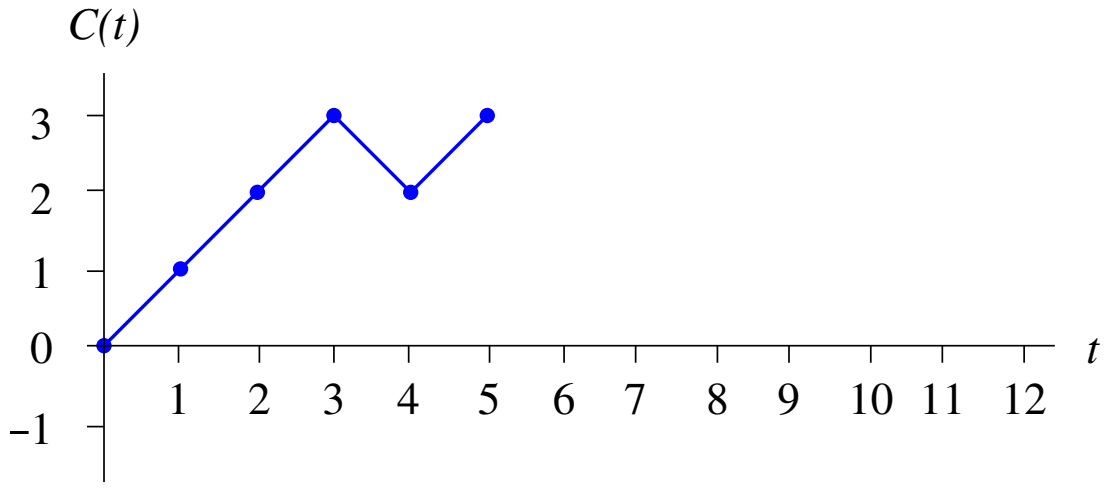

Contour function

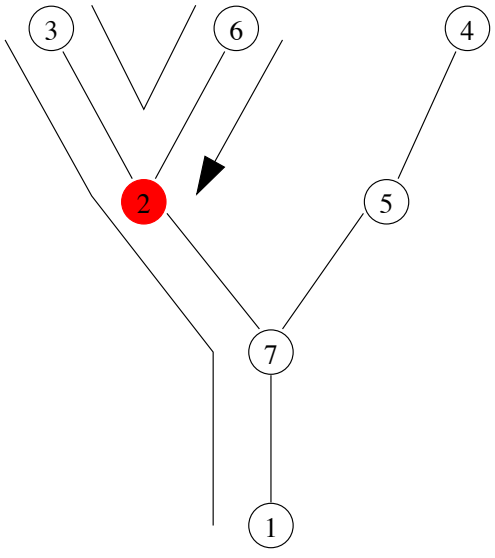

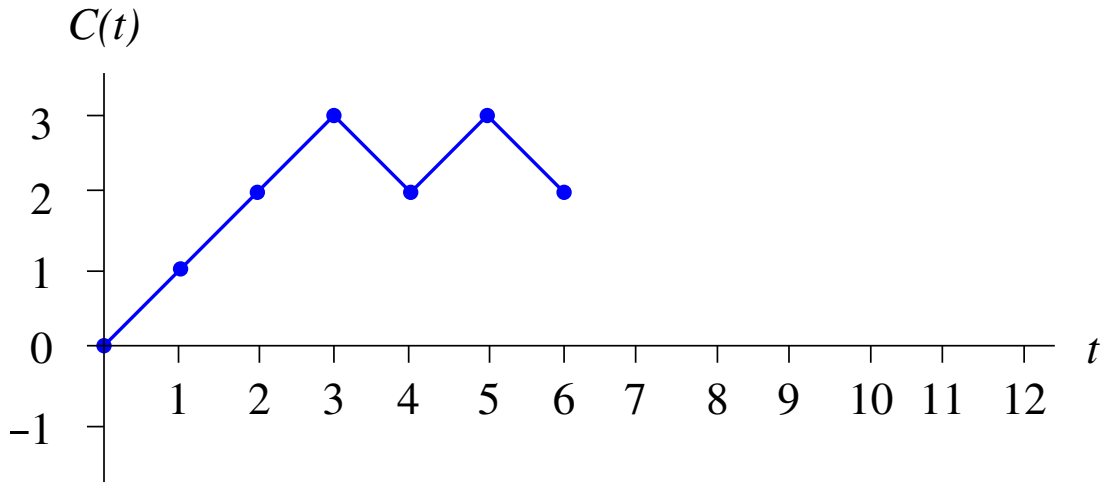

Contour function

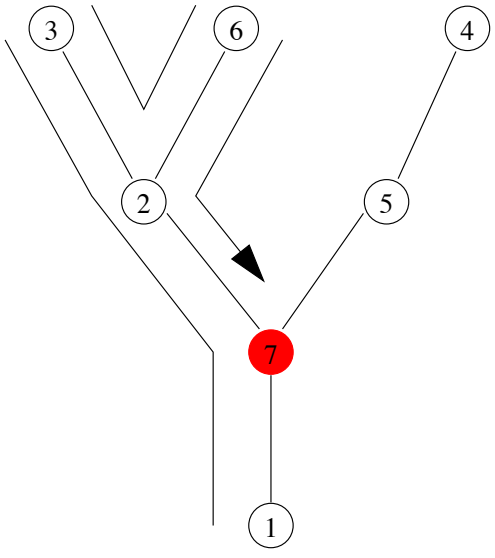

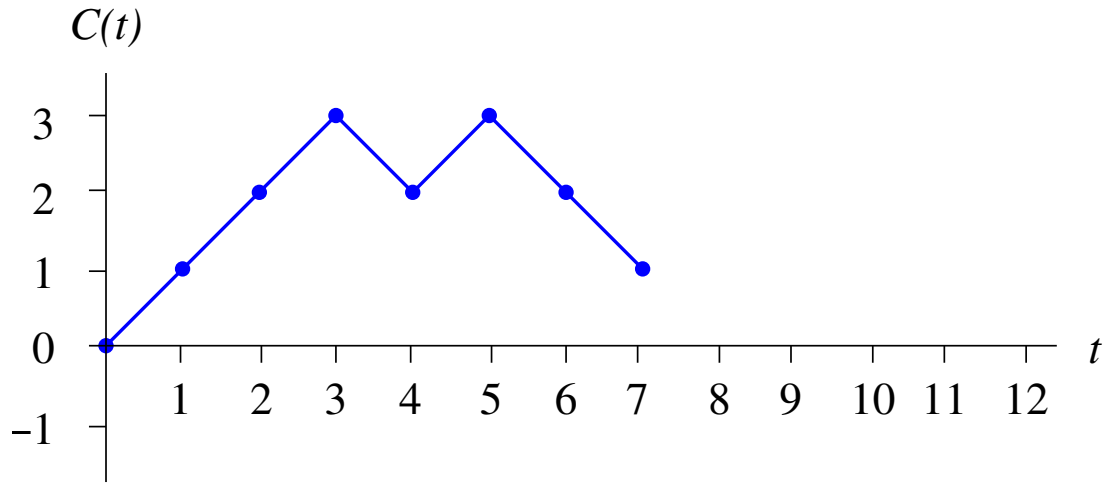

Contour function

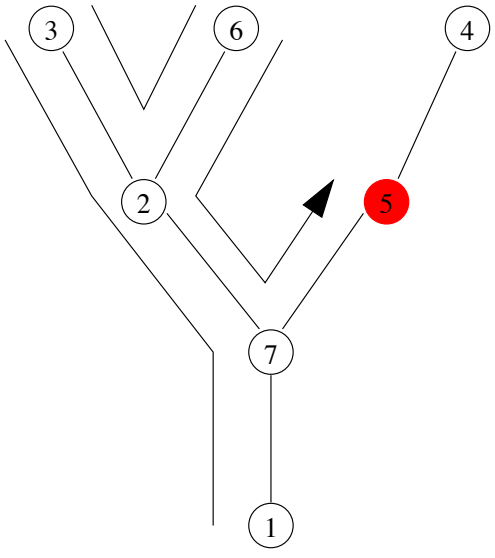

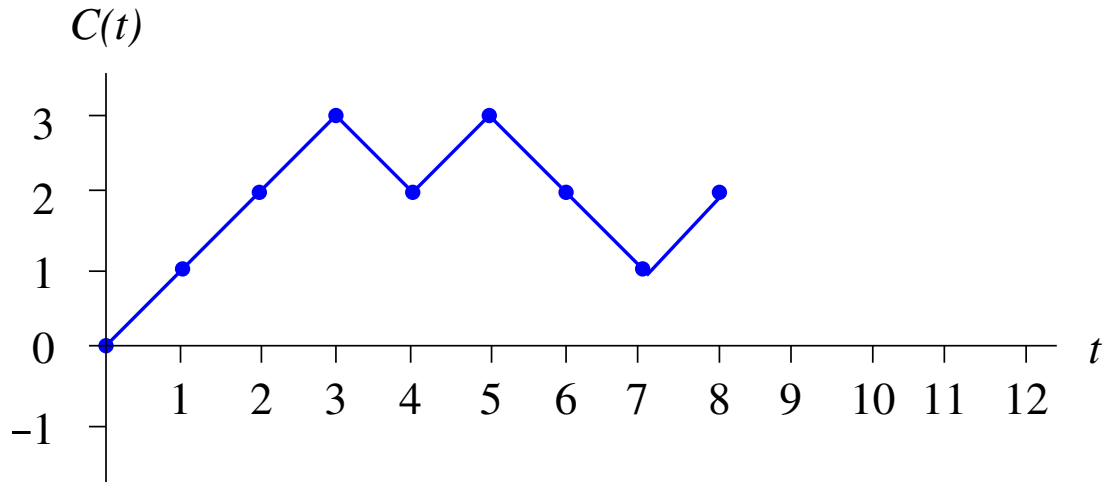

Contour function

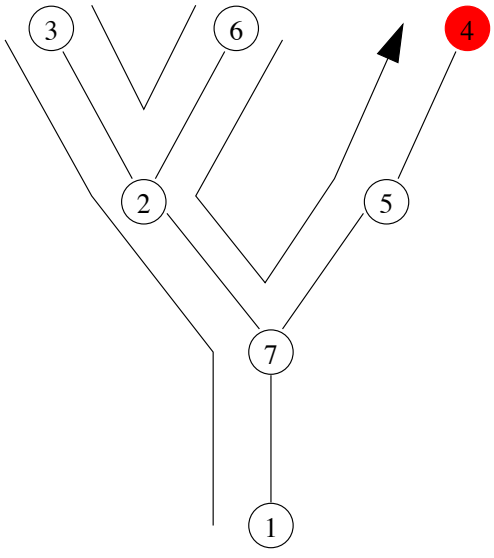

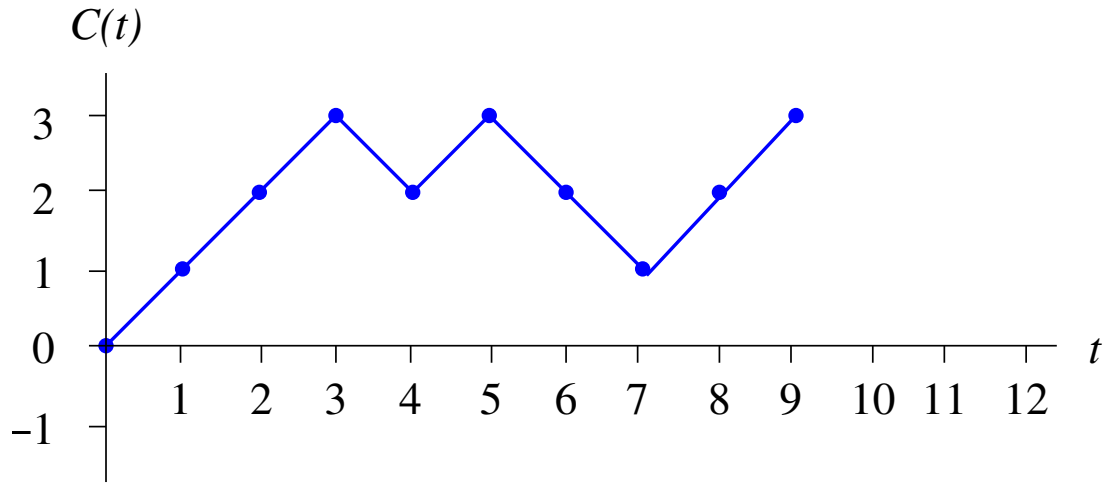

Contour function

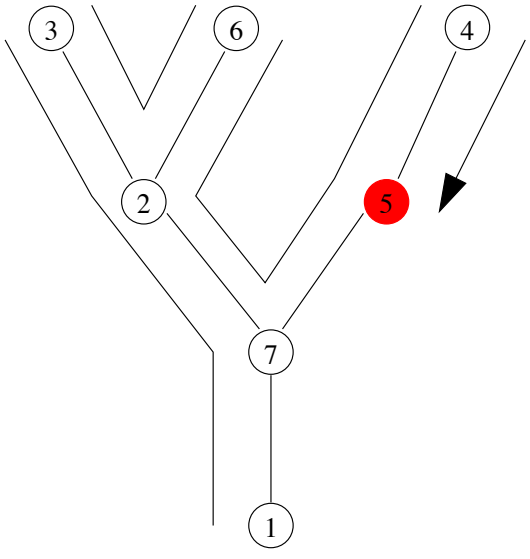

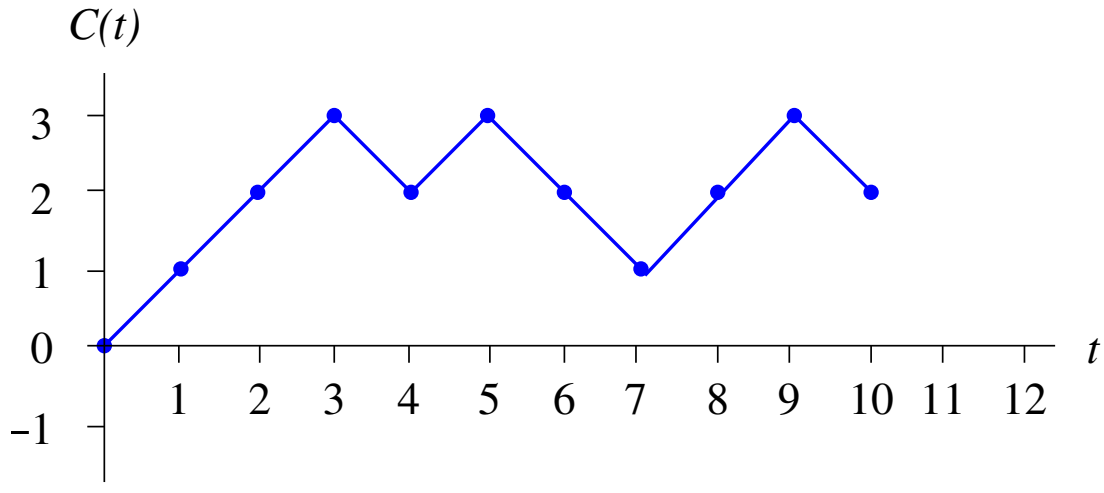

Contour function

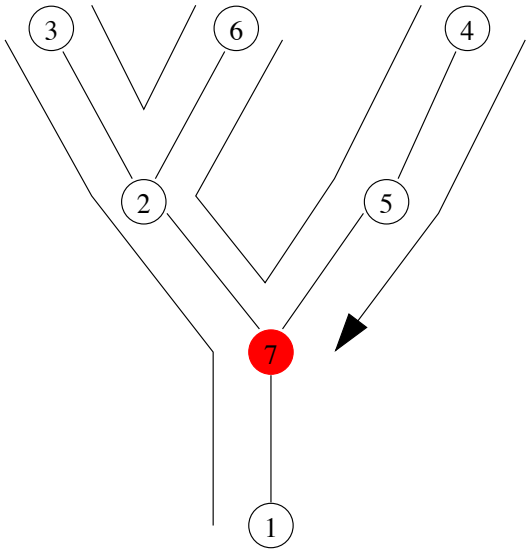

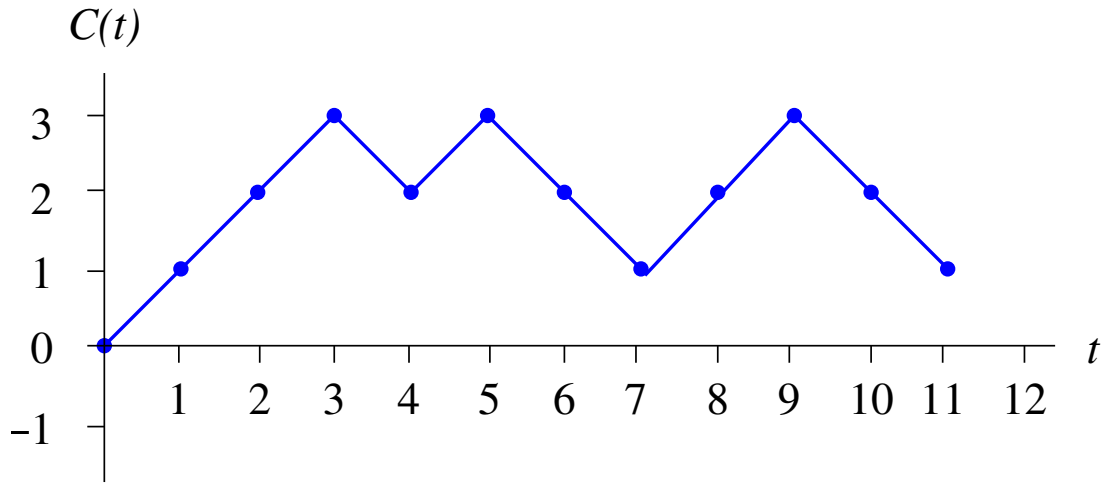

Contour function

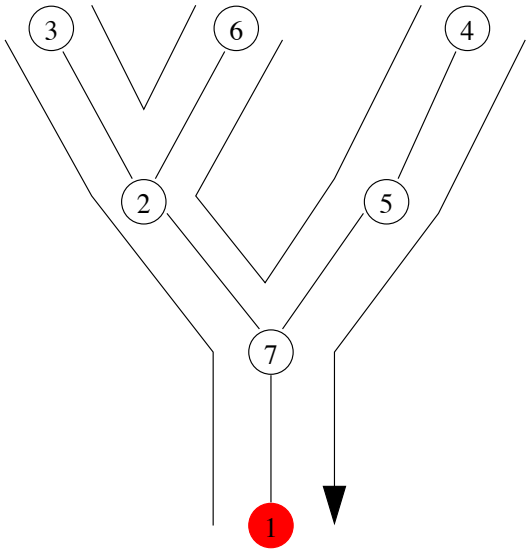

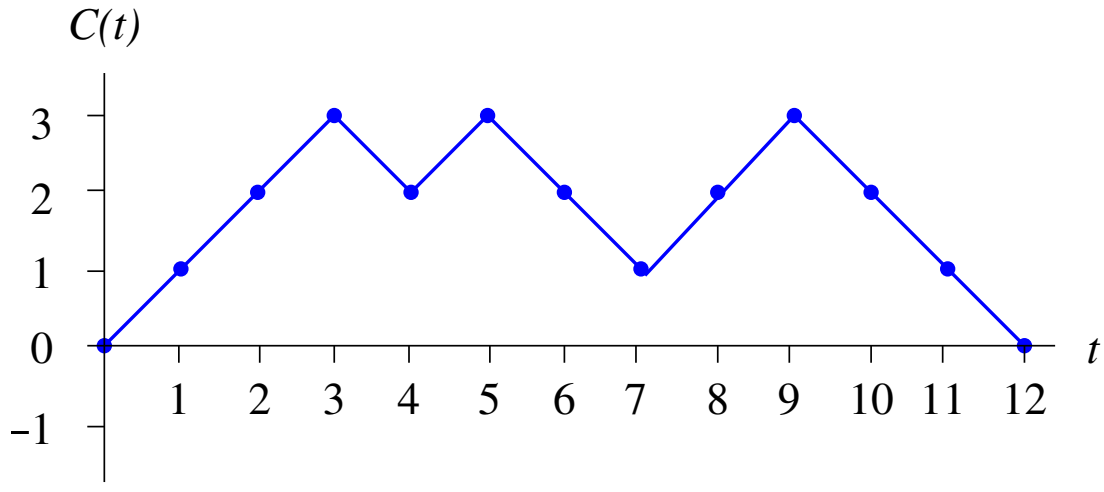

Contour function

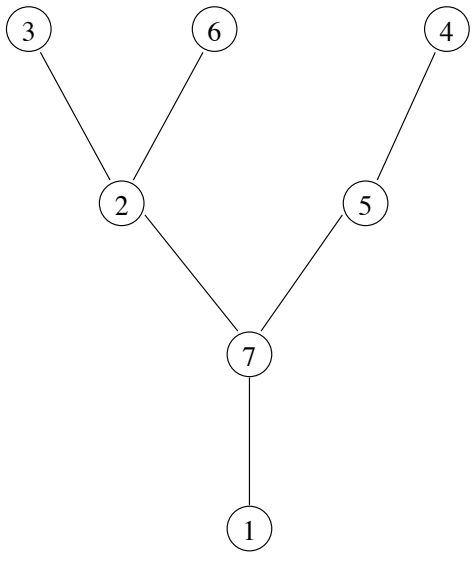

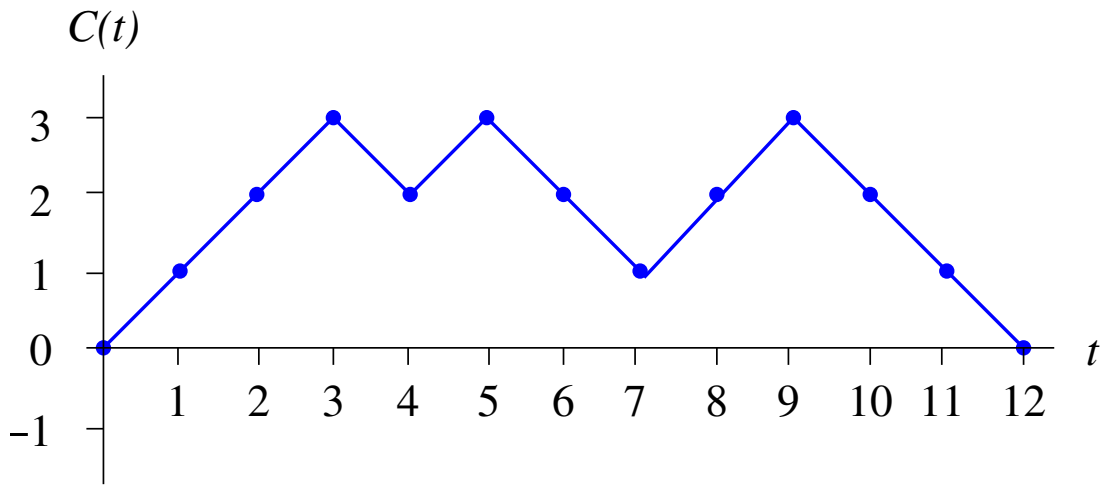

The contour function is a sort of "expanded" version of the tree.

Depth-first walk

The depth-first walk is constructed using a depth-first search of the tree. At each step, the vertices may be in one of four states: unexplored, current, alive or dead.

Depth-first walk

The depth-first walk is constructed using a depth-first search of the tree. At each step, the vertices may be in one of four states: unexplored, current, alive or dead.

Let $X(k)$ be the number of vertices alive at step k.

Depth-first walk

Step 0: initialisation

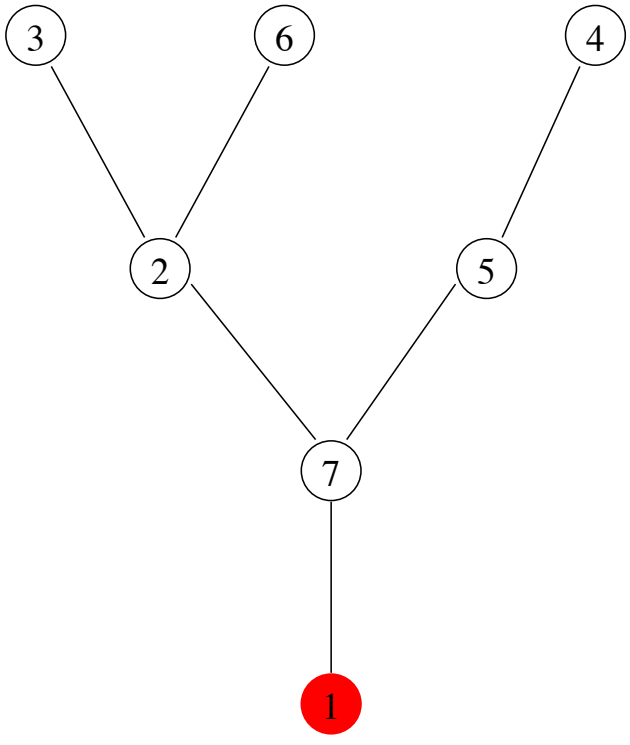

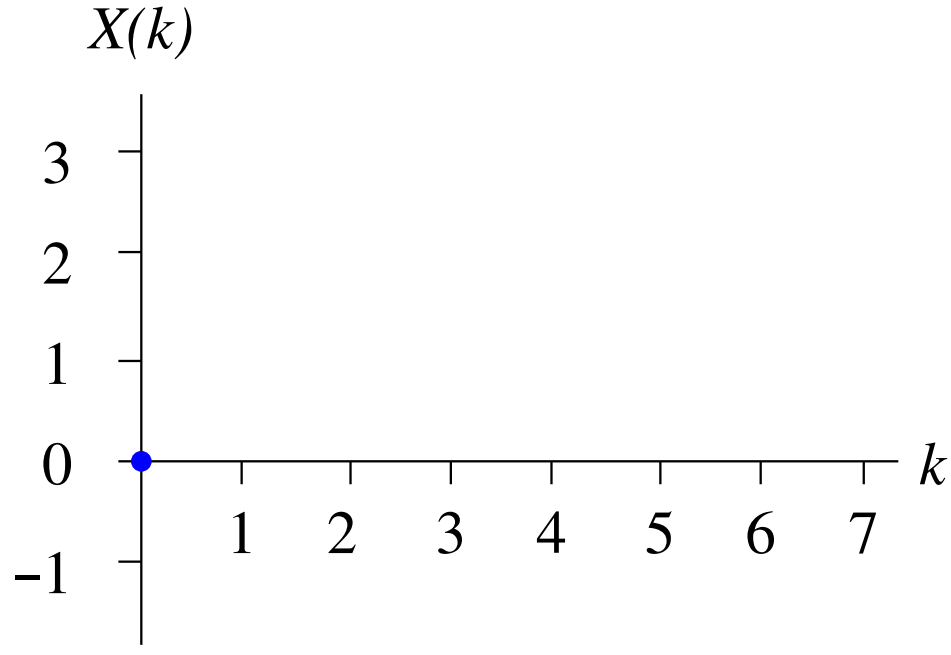

Current: 1 Alive: none Dead: none $X(0) = 0$.

Depth-first walk

Step 1

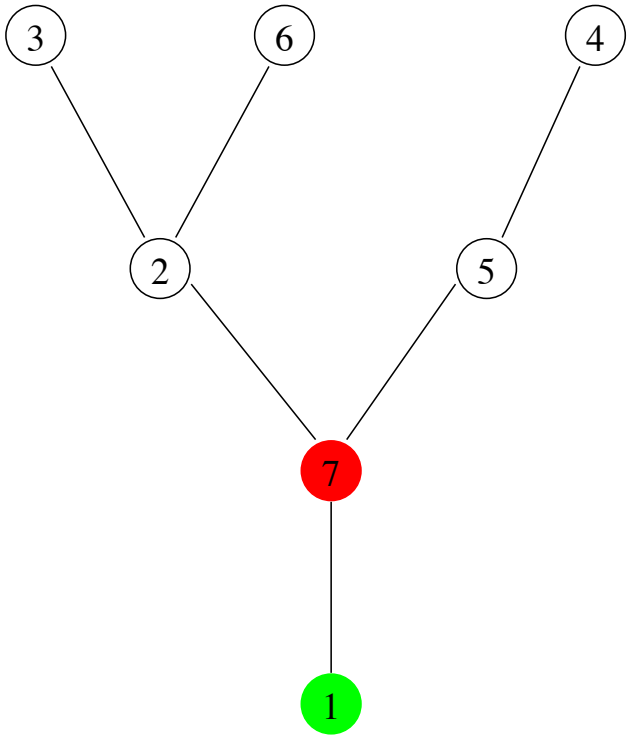

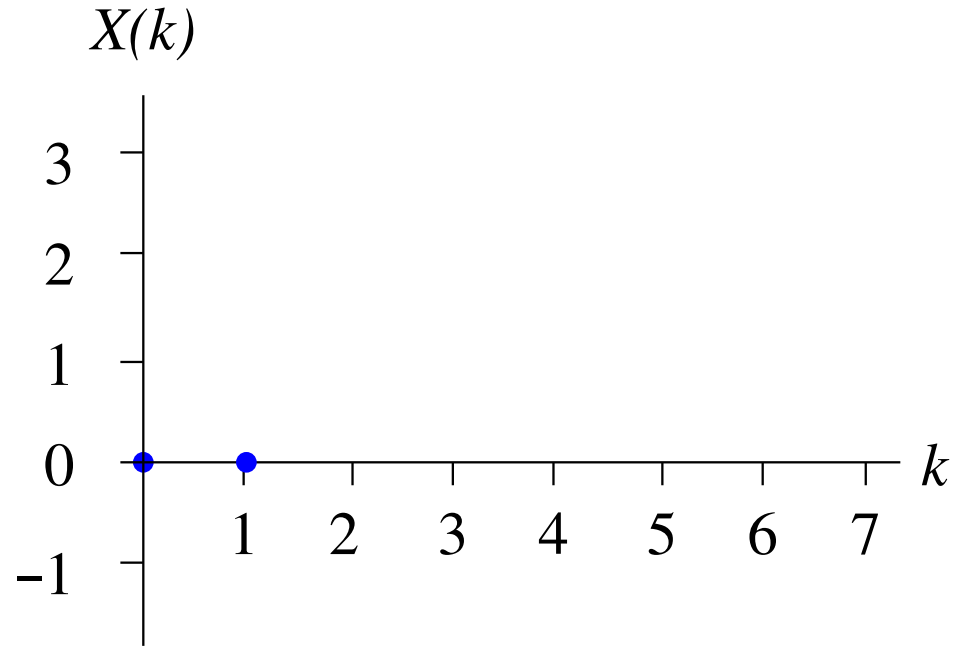

Current: 7 Alive: none Dead: 1 $X(1) = 0$.

Depth-first walk

Step 2

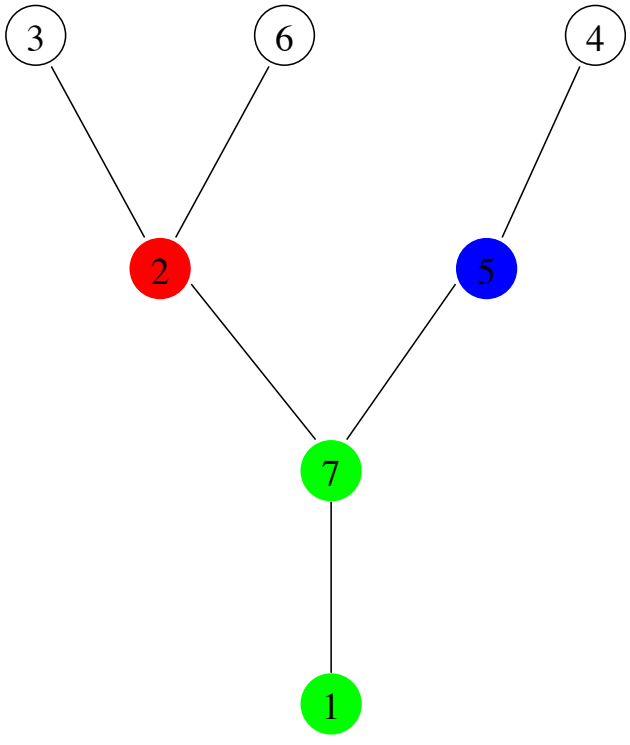

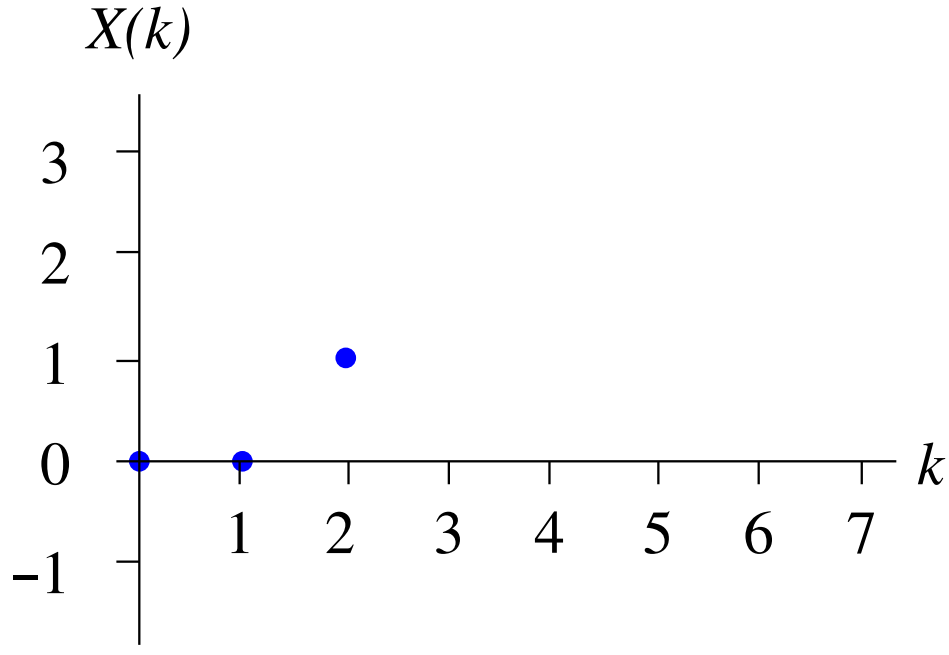

Current: 2 Alive: 5 Dead: 1, 7 $X(2) = 1$.

Depth-first walk

Step 3

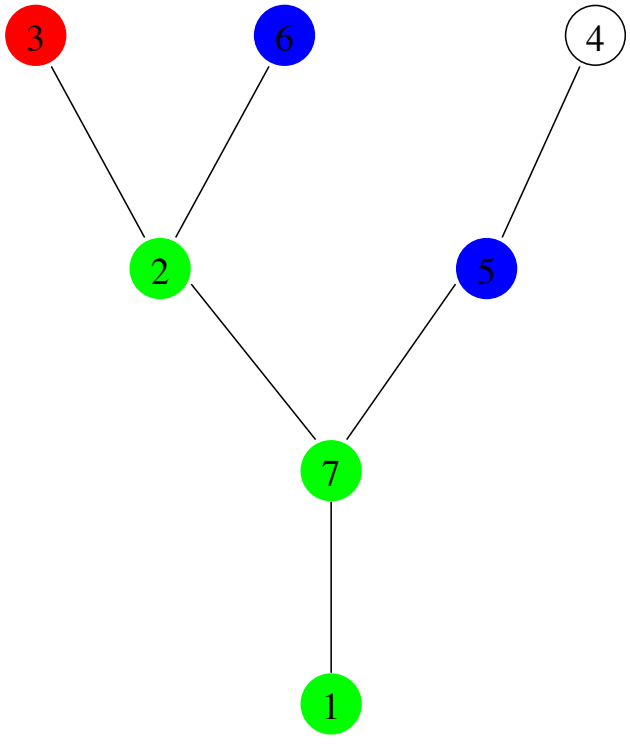

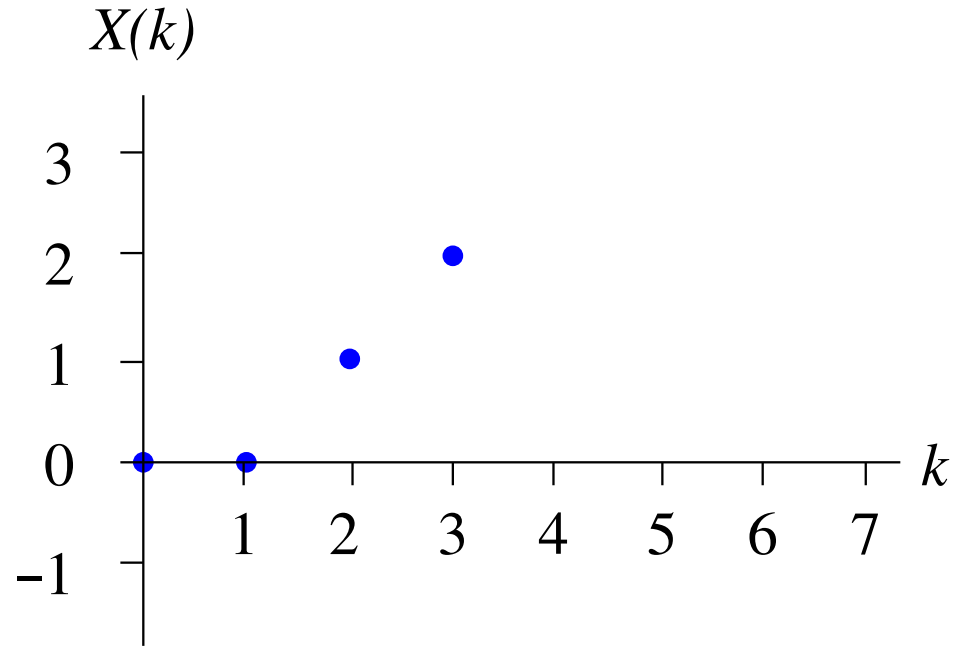

Current: 3 Alive: 6, 5 Dead: 1, 7, 2 $X(3) = 2$.

Depth-first walk

Step 4

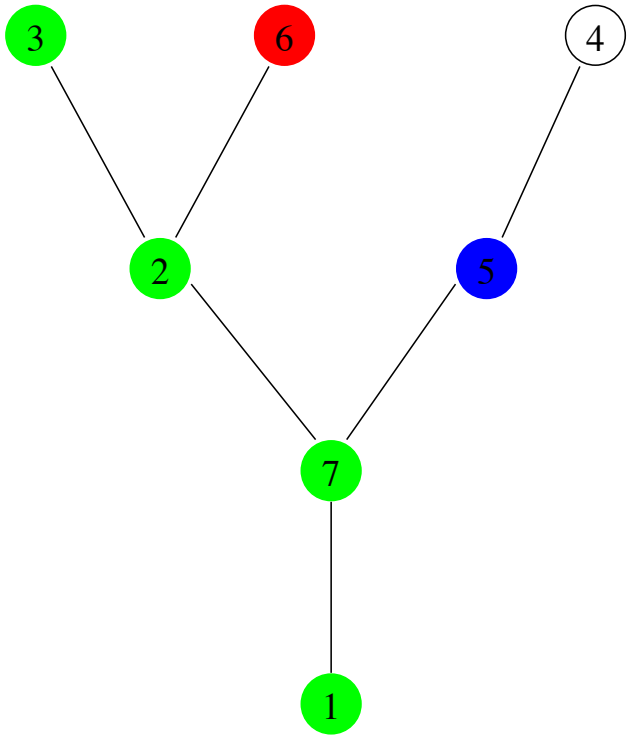

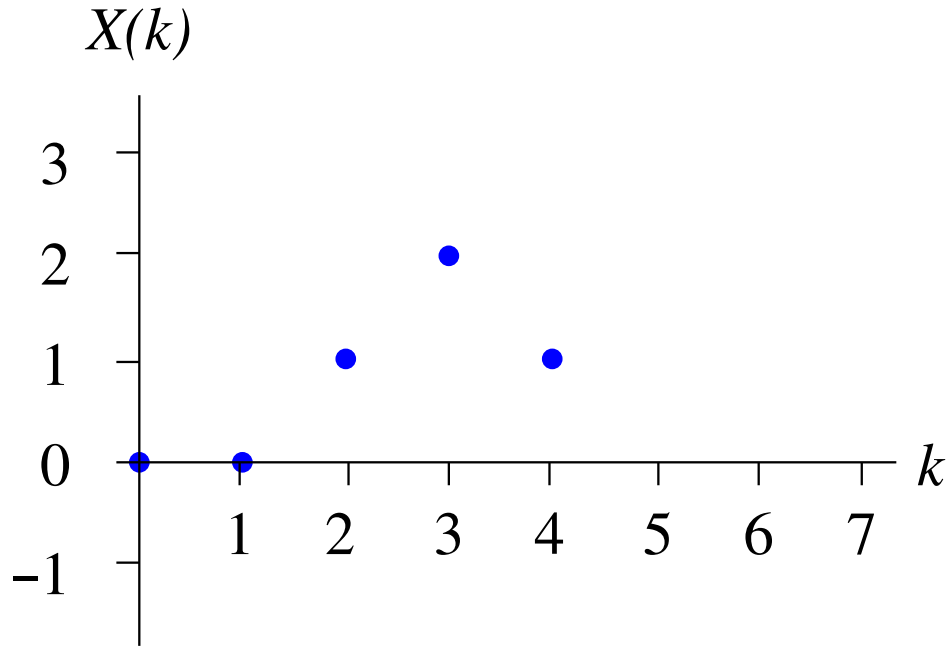

Current: 6 Alive: 5 Dead: 1, 7, 2, 3 $X(4) = 1$.

Depth-first walk

Step 5

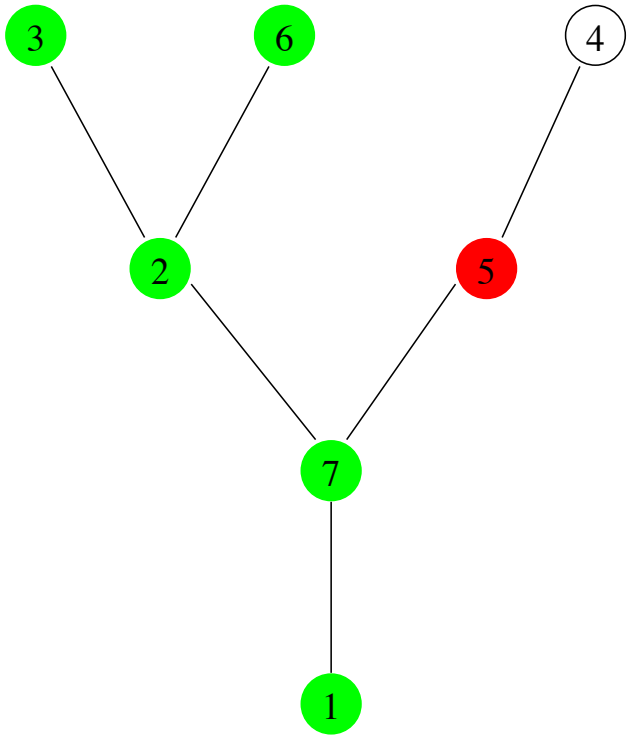

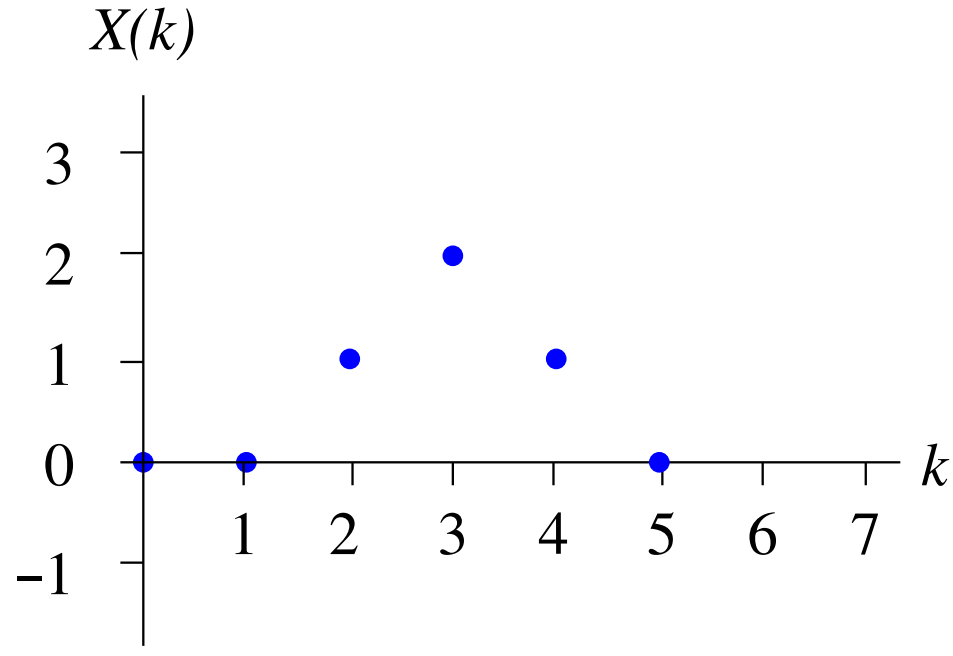

Current: 5 Alive: none Dead: 1, 7, 2, 3, 6 $X(5) = 0$.

Depth-first walk

Step 6

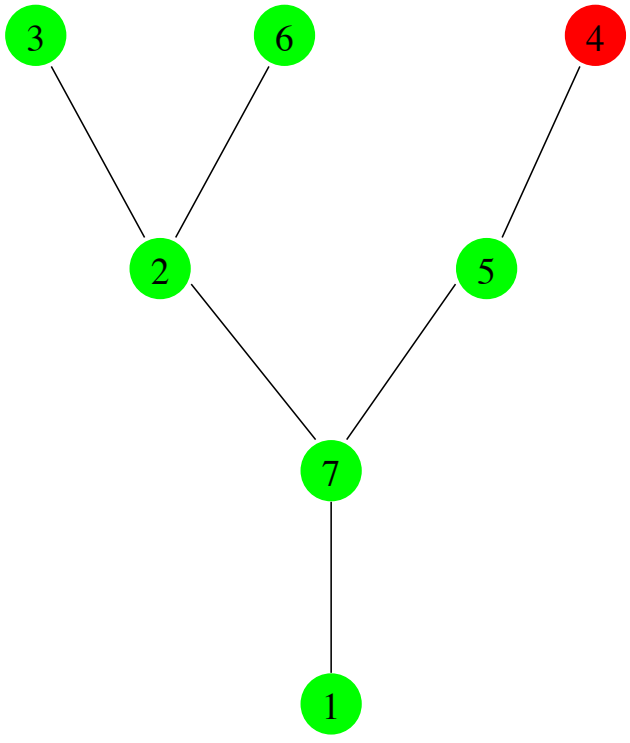

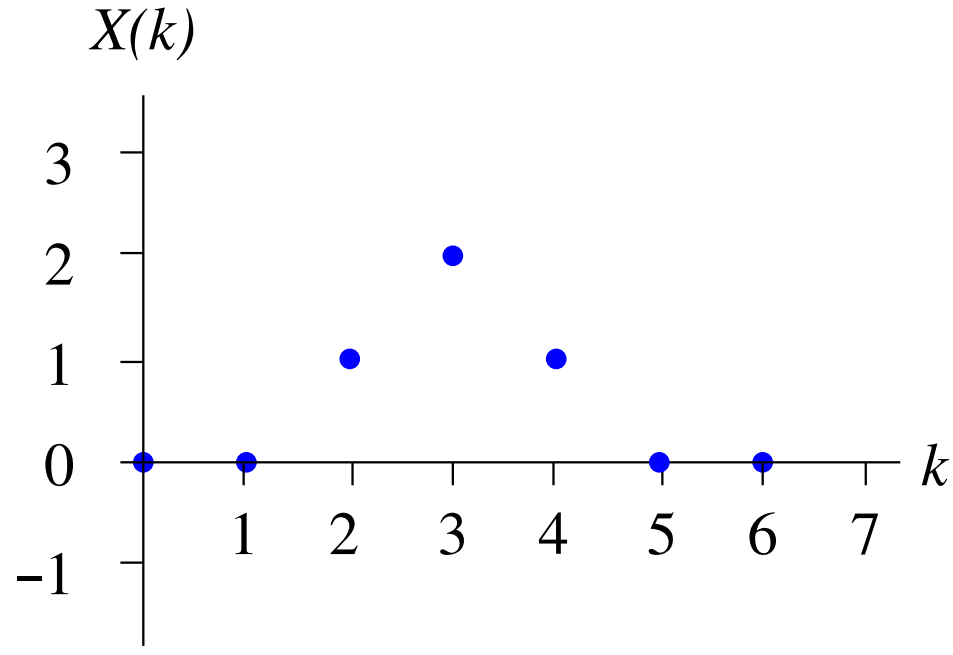

Current: 4 Alive: none Dead: 1, 7, 2, 3, 6, 5 $X(6) = 0$.

Depth-first walk

Step 7

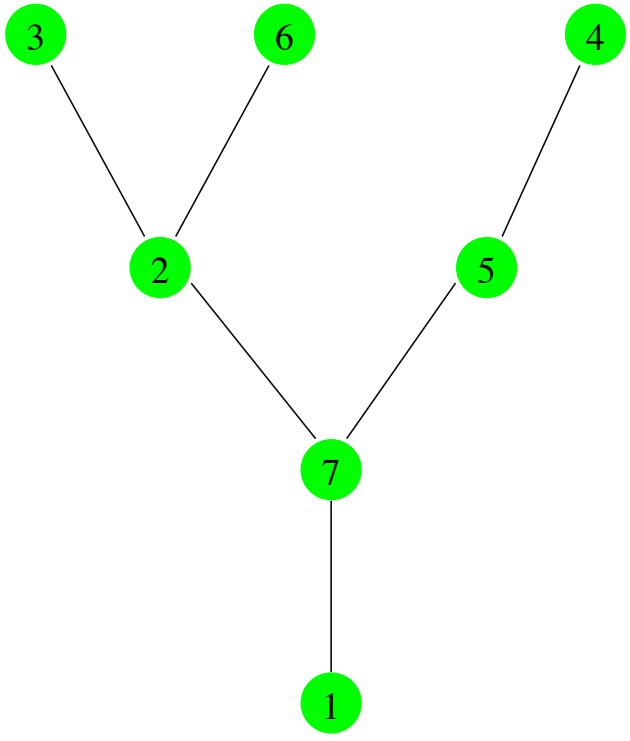

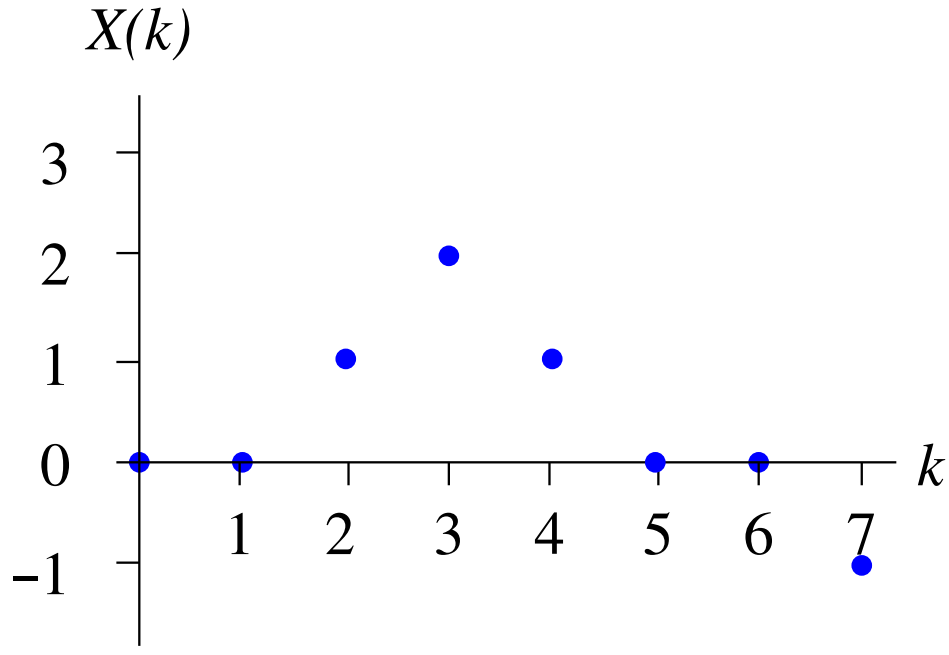

Dead: 1, 7, 2, 3, 6, 5, 4

Depth-first walk

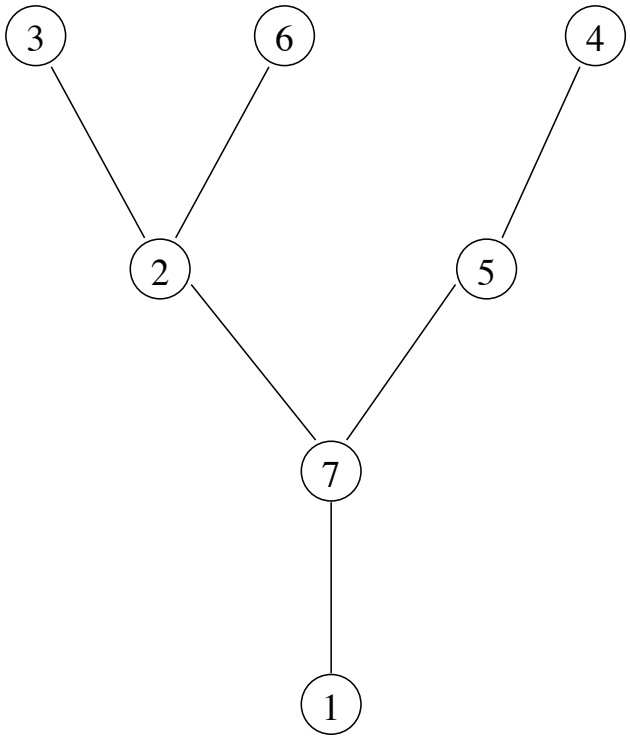

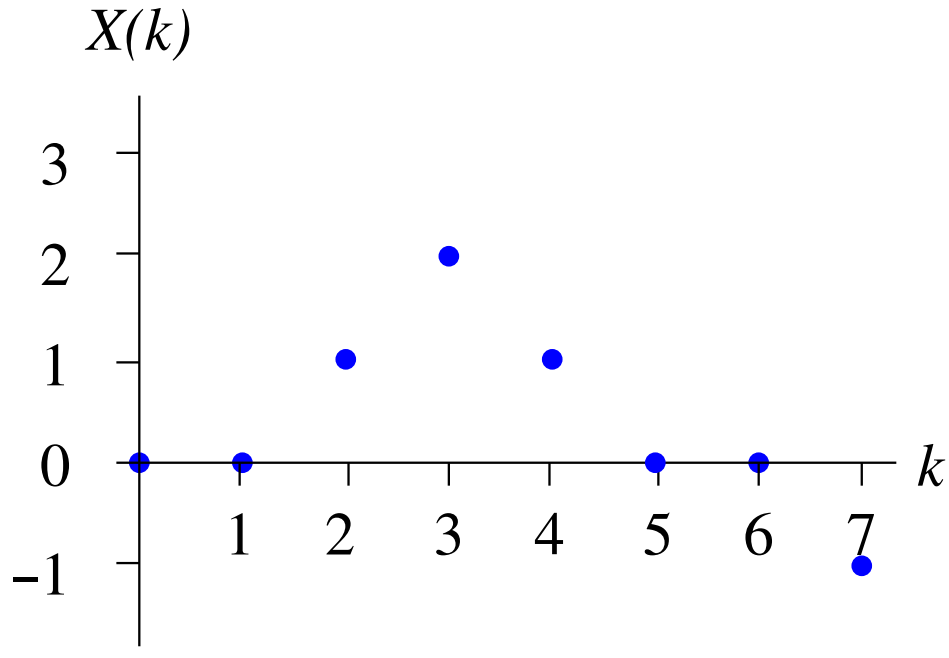

Depth-first walk

Why does $(X(k), 0 \leq k \leq n)$ encode the shape of the tree? At each step, we had

$$X(k) = X(k-1) + \#\text{children of the } k\text{th vertex visited} - 1.$$

So we can recover the sequence of numbers of children, which is enough to get to the shape.

Depth-first walk

Why does $(X(k), 0 \leq k \leq n)$ encode the shape of the tree? At each step, we had

$$X(k) = X(k-1) + \#\text{children of the } k\text{th vertex visited} - 1.$$

So we can recover the sequence of numbers of children, which is enough to get to the shape.

$X(n) = -1$ precisely at the end of the tree. (This is simply because we added up the number of children and subtracted the number of vertices.)

Coding a uniform random tree by functions

For a uniform random tree, the contour function and depth-first walk are random. We would like to understand their distributions.

Coding a uniform random tree by functions

For a uniform random tree, the contour function and depth-first walk are random. We would like to understand their distributions.

Although it's a little easier to see why the contour function encodes the shape of the tree, the depth-first walk is easier to work with. To see why, it's actually helpful to generalize our tree model.

Reminder: Galton-Watson process

A Galton-Watson branching process $(B_k)_{k\geq 0}$ describes the size of a population which evolves as follows:

Reminder: Galton-Watson process

A Galton-Watson branching process $(B_k)_{k\geq 0}$ describes the size of a population which evolves as follows:

- Start with a single individual.

Reminder: Galton-Watson process

A Galton-Watson branching process $(B_k)_{k\geq 0}$ describes the size of a population which evolves as follows:

- Start with a single individual.
- This individual has a number of children distributed according to the offspring distribution μ, where $\mu(i)$ gives the probability of i children, $i \geq 0$.

Reminder: Galton-Watson process

A Galton-Watson branching process $(B_k)_{k\geq 0}$ describes the size of a population which evolves as follows:

- Start with a single individual.
- This individual has a number of children distributed according to the offspring distribution μ, where $\mu(i)$ gives the probability of i children, $i \geq 0$.
- Each child reproduces as an independent copy of the original individual.

Reminder: Galton-Watson process

A Galton-Watson branching process $(B_k)_{k\geq 0}$ describes the size of a population which evolves as follows:

- Start with a single individual.
- This individual has a number of children distributed according to the offspring distribution μ, where $\mu(i)$ gives the probability of i children, $i \geq 0$.
- Each child reproduces as an independent copy of the original individual.

B_k gives the number of individuals in generation k (in particular, $B_0 = 1$).

Galton-Watson trees

A Galton-Watson tree is the family tree arising from a Galton-Watson branching process.

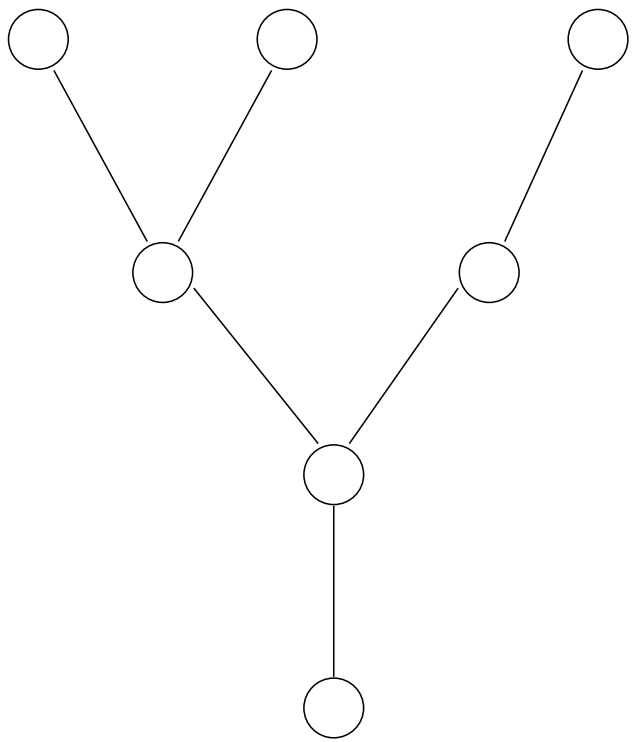

$B_4 = 0$

$B_3 = 3$

$B_2 = 2$

$B_1 = 1$

$B_0 = 1$

Galton-Watson trees

We will only consider the case where the offspring distribution μ is critical and has finite offspring variance i.e.

$$\sum_{i=1}^{\infty} i\mu(i) = 1, \qquad \sum_{i=1}^{\infty} i^2\mu(i) < \infty.$$

We also impose that $\mu(1) < 1$. These conditions ensure that the Galton-Watson process becomes extinct with probability 1, which entails that the tree is finite.

The uniform random tree is a Galton-Watson tree

Proposition. The uniform random tree T_n has the same distribution as a tree generated as follows:

- Take a Galton-Watson tree with Poisson(1) offspring distribution;
- Condition it to have total progeny precisely n;
- Assign the vertices random labels chosen from $\{1, 2, \ldots, n\}$.

The depth-first walk of a Galton-Watson tree

Let T_n be a Galton-Watson tree with critical offspring distribution μ conditioned to have total progeny n (and given a uniform random labelling). What does this mean for the depth-first walk $(X^n(k), 0 \leq k \leq n)$ which encodes it?

The depth-first walk of a Galton-Watson tree

Let T_n be a Galton-Watson tree with critical offspring distribution μ conditioned to have total progeny n (and given a uniform random labelling). What does this mean for the depth-first walk $(X^n(k), 0 \leq k \leq n)$ which encodes it?

The numbers of children of different vertices in the branching process are i.i.d. random variables with distribution μ.

The depth-first walk of a Galton-Watson tree

Let T_n be a Galton-Watson tree with critical offspring distribution μ conditioned to have total progeny n (and given a uniform random labelling). What does this mean for the depth-first walk $(X^n(k), 0 \leq k \leq n)$ which encodes it?

The numbers of children of different vertices in the branching process are i.i.d. random variables with distribution μ.

Recall that

$$X^n(k) = X^n(k-1) + \#\text{children of the } k\text{th vertex visited} - 1.$$

So $(X^n(k), 0 \leq k \leq n)$ is like a random walk with step sizes having distribution $\mu(\cdot) - 1$, which has mean 0 and finite variance σ^2.

The depth-first walk of a Galton-Watson tree

Let T_n be a Galton-Watson tree with critical offspring distribution μ conditioned to have total progeny n (and given a uniform random labelling). What does this mean for the depth-first walk $(X^n(k), 0 \leq k \leq n)$ which encodes it?

The numbers of children of different vertices in the branching process are i.i.d. random variables with distribution μ.

Recall that

$$X^n(k) = X^n(k-1) + \#\text{children of the } k\text{th vertex visited} - 1.$$

So $(X^n(k), 0 \leq k \leq n)$ is like a random walk with step sizes having distribution $\mu(\cdot) - 1$, which has mean 0 and finite variance σ^2.

The effect of the conditioning is to force $X^n(k) > -1$ for $1 \leq k \leq n-1$ and $X^n(n) = -1$.

Convergence to a Brownian motion

Let's ignore the conditioning for the moment. Then we have a random walk $(X(k), k \geq 0)$ with step sizes having mean 0 and variance σ^2, and so Donsker's theorem tells us that

$$\frac{1}{\sigma\sqrt{m}}(X(\lfloor mt \rfloor), t \geq 0) \xrightarrow{d} (W(t), t \geq 0),$$

as $m \to \infty$, where $(W(t), t \geq 0)$ is a Brownian motion.

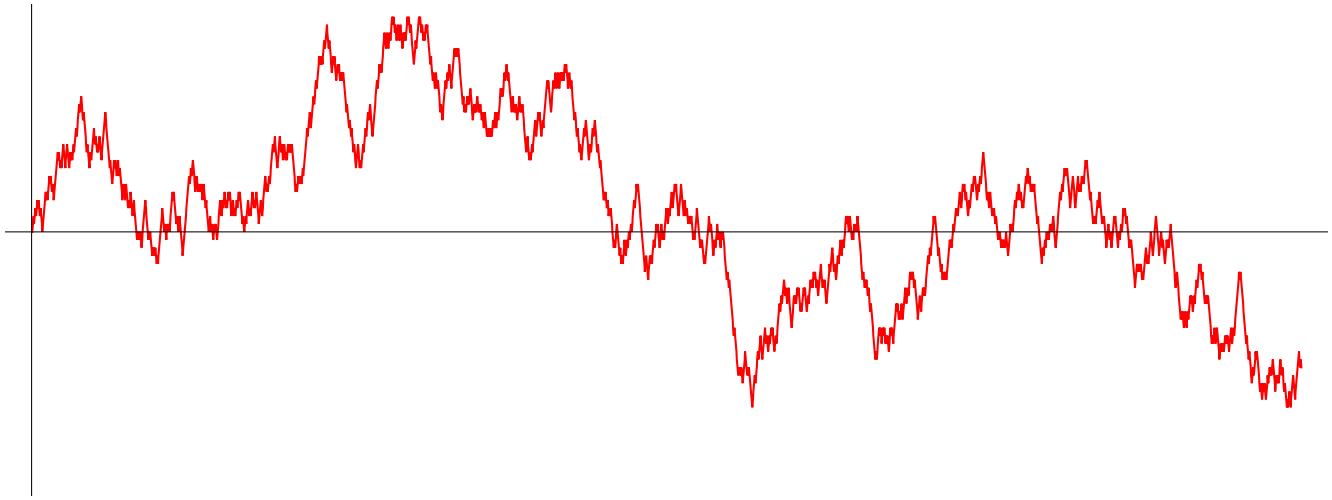

Convergence to a Brownian excursion

We want a version of this result where the time-interval is restricted to be $[0, n]$ and the process is conditioned to stay above -1 until time n when it hits -1.

Since we're going to rescale the process by $1/\sqrt{n}$ and let n go to ∞, this is essentially the same as conditioning the process to stay above 0 until time n.

Convergence to a Brownian excursion

We get

$$\frac{1}{\sigma\sqrt{n}}(X^n(\lfloor nt \rfloor), 0 \leq t \leq 1) \xrightarrow{d} (e(t), 0 \leq t \leq 1),$$

as $n \to \infty$, where $(e(t), 0 \leq t \leq 1)$ is a standard Brownian excursion. (Think of this as a Brownian motion conditioned to stay positive until time 1 when it returns to 0.)

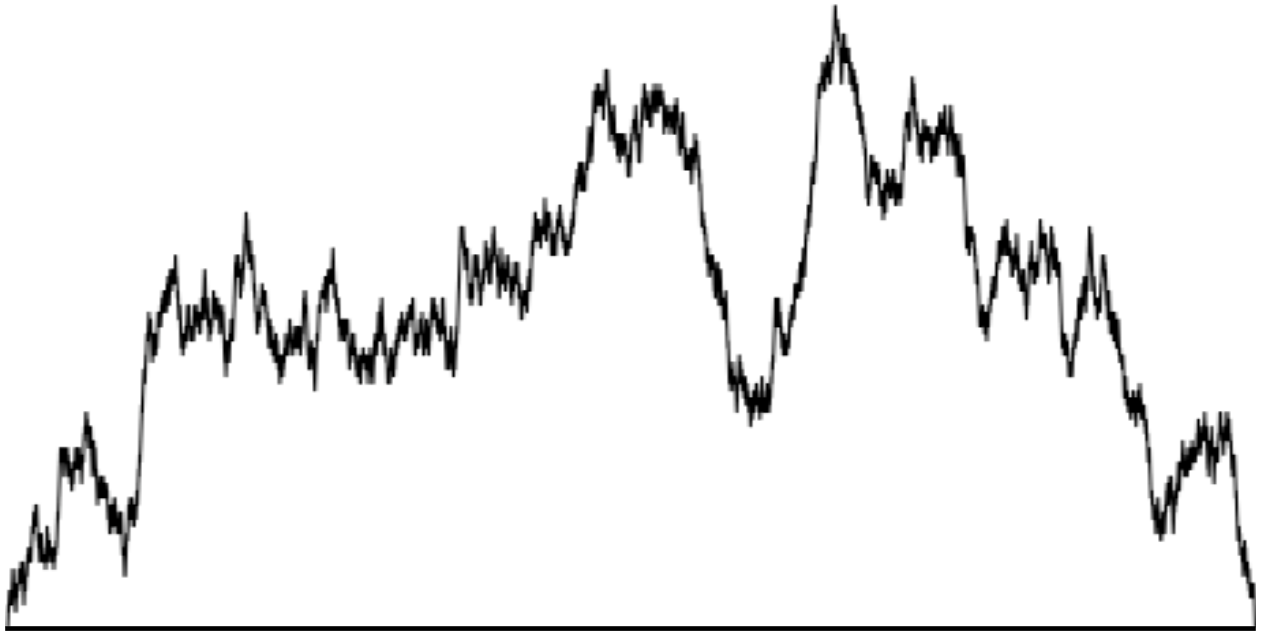

Convergence to a Brownian excursion

It turns out that more-or-less the same limit holds for the contour function:

$$\frac{\sigma}{\sqrt{n}}(C^n(2nt), 0 \leq t \leq 1) \xrightarrow{d} (2e(t), 0 \leq t \leq 1),$$

as $m \to \infty$, where $(e(t), 0 \leq t \leq 1)$ is a standard Brownian excursion.

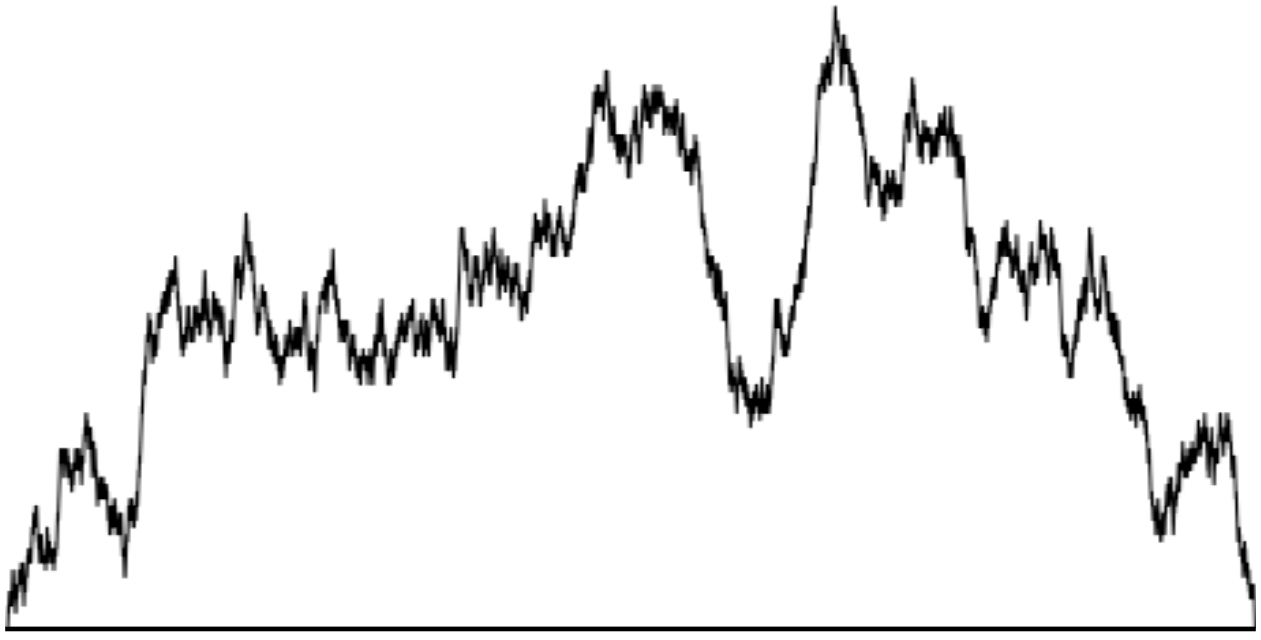

From excursions back to trees

This result certainly suggests that there should be some sort of limiting tree encoded by the Brownian excursion.

From excursions back to trees

This result certainly suggests that there should be some sort of limiting tree encoded by the Brownian excursion.

Indeed, the Brownian excursion acts as a sort of contour function for the limit tree.

From excursions back to trees

This result certainly suggests that there should be some sort of limiting tree encoded by the Brownian excursion.

Indeed, the Brownian excursion acts as a sort of contour function for the limit tree.

The procedure to get the Brownian continuum random tree from a Brownian excursion generalises to other similar functions (deterministic or random), so I'll present it in general.

From excursions back to trees

Let $h : [0,1] \to \mathbb{R}^+$ be an excursion, that is a continuous function such that $h(0) = h(1) = 0$ and $h(x) > 0$ for $x \in (0,1)$.

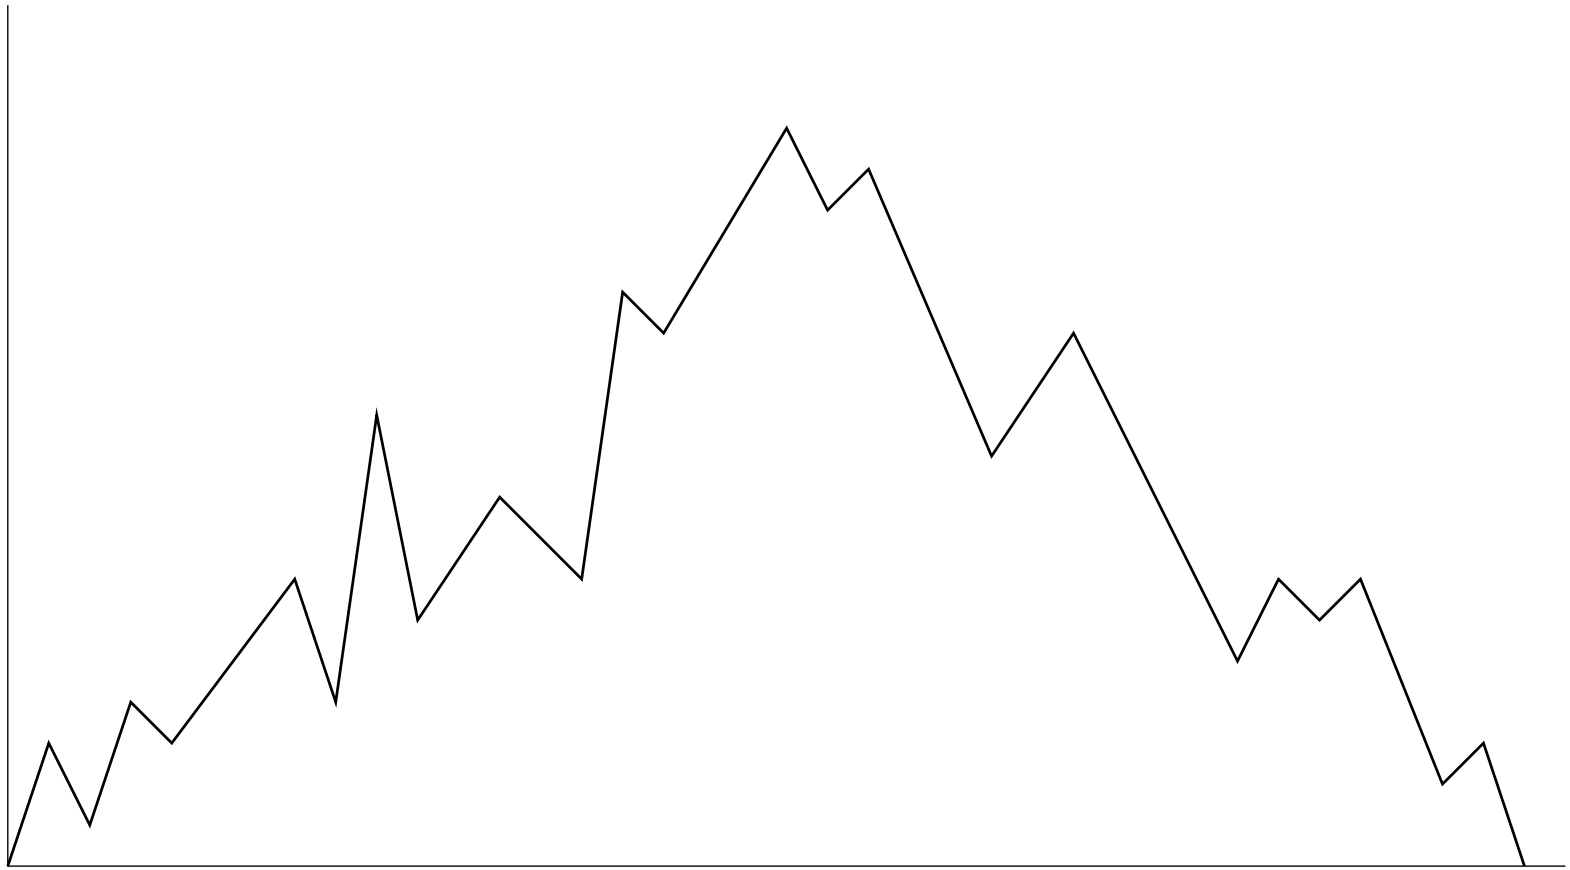

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

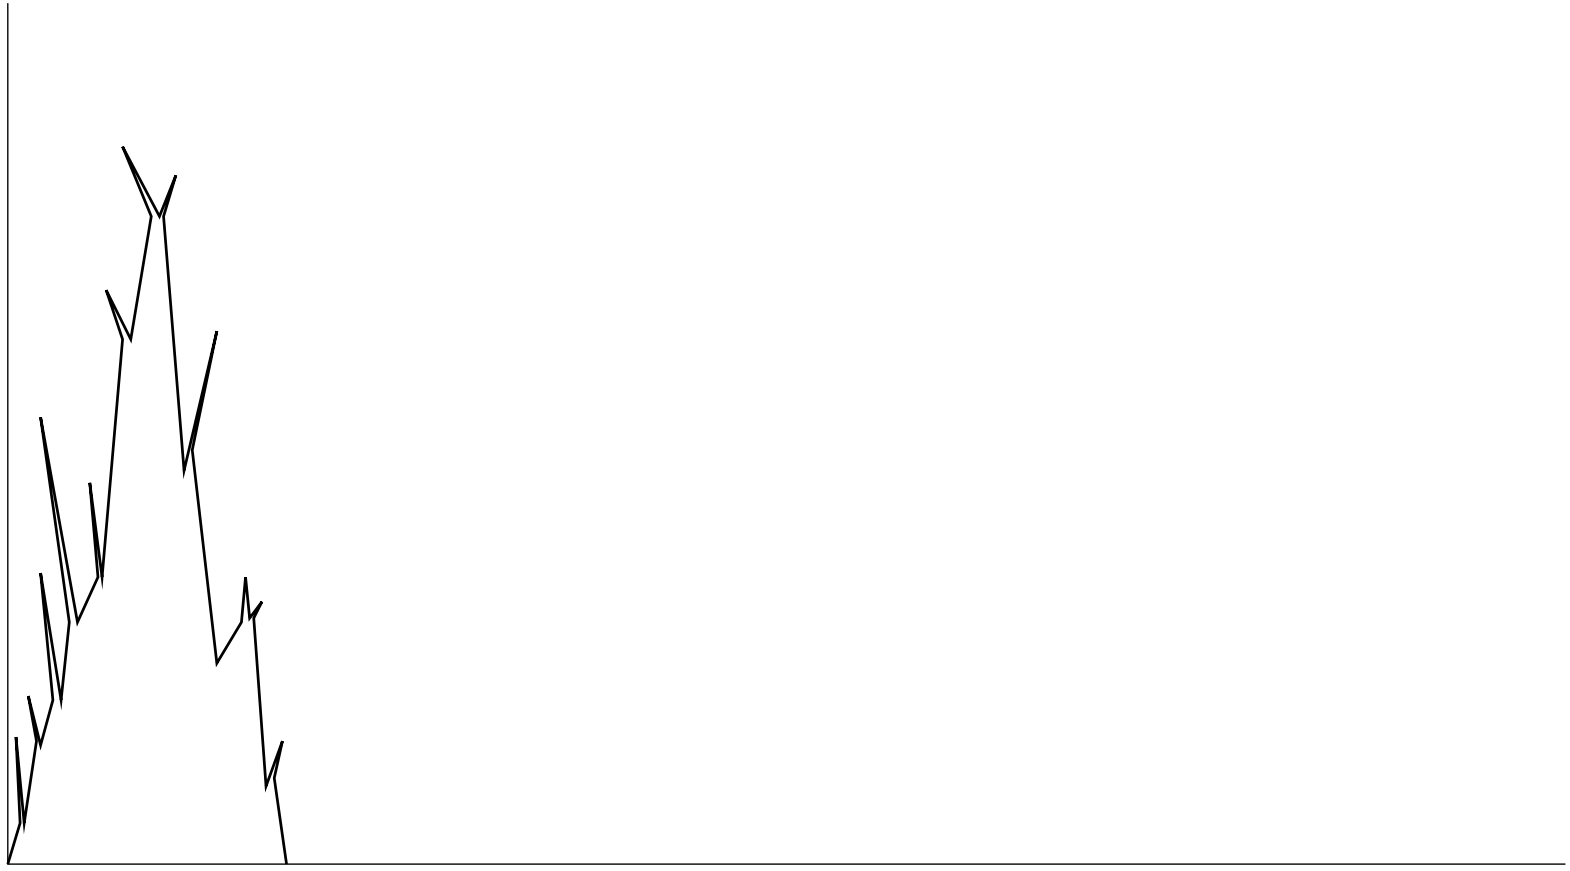

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

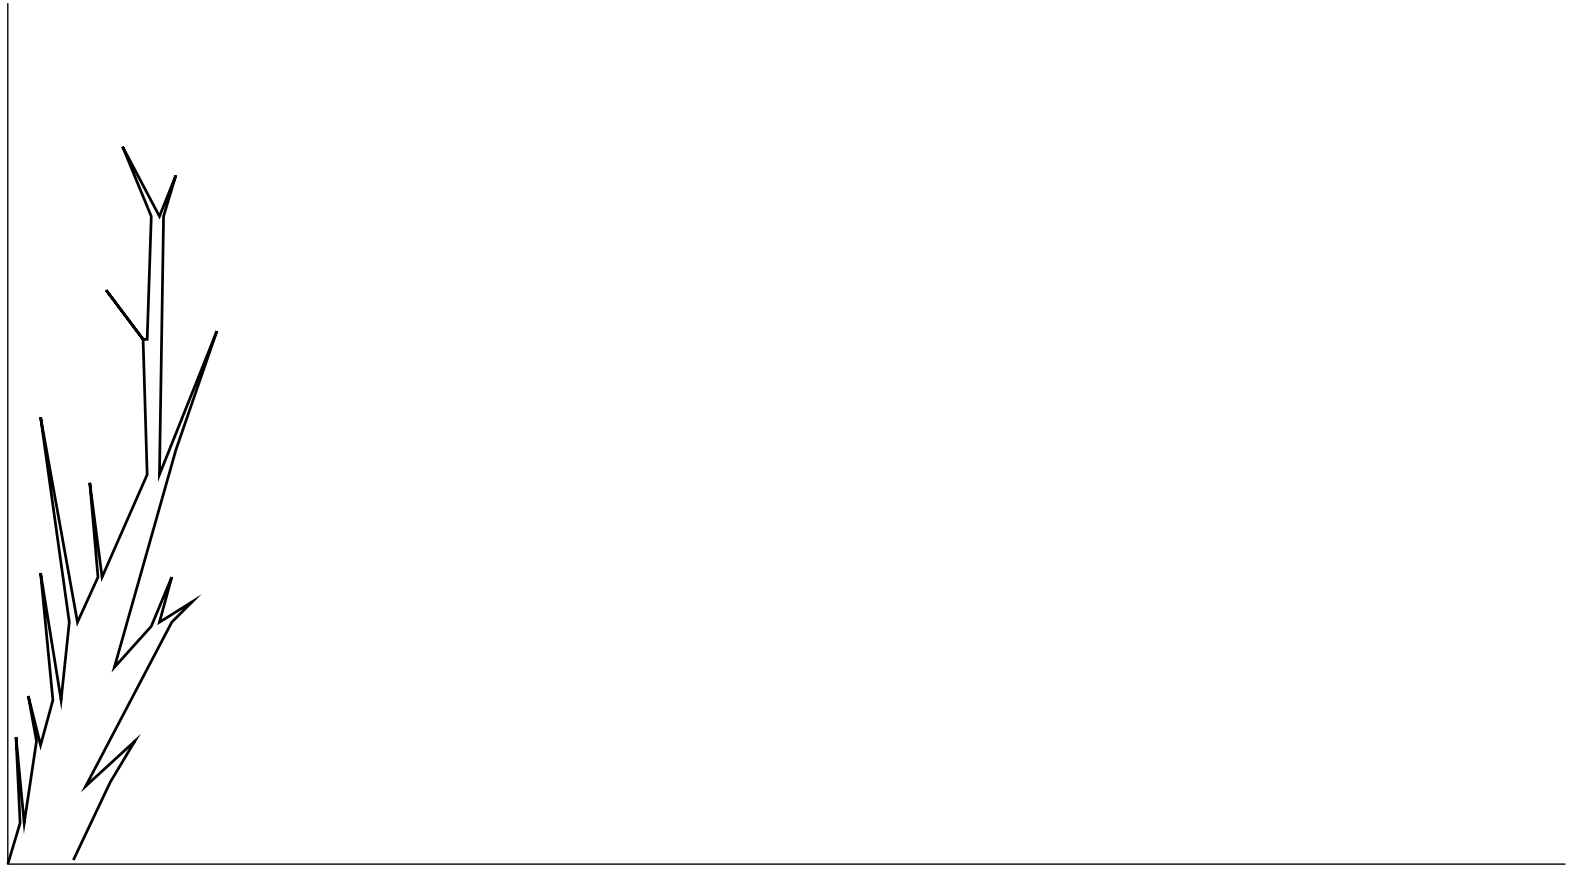

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together...

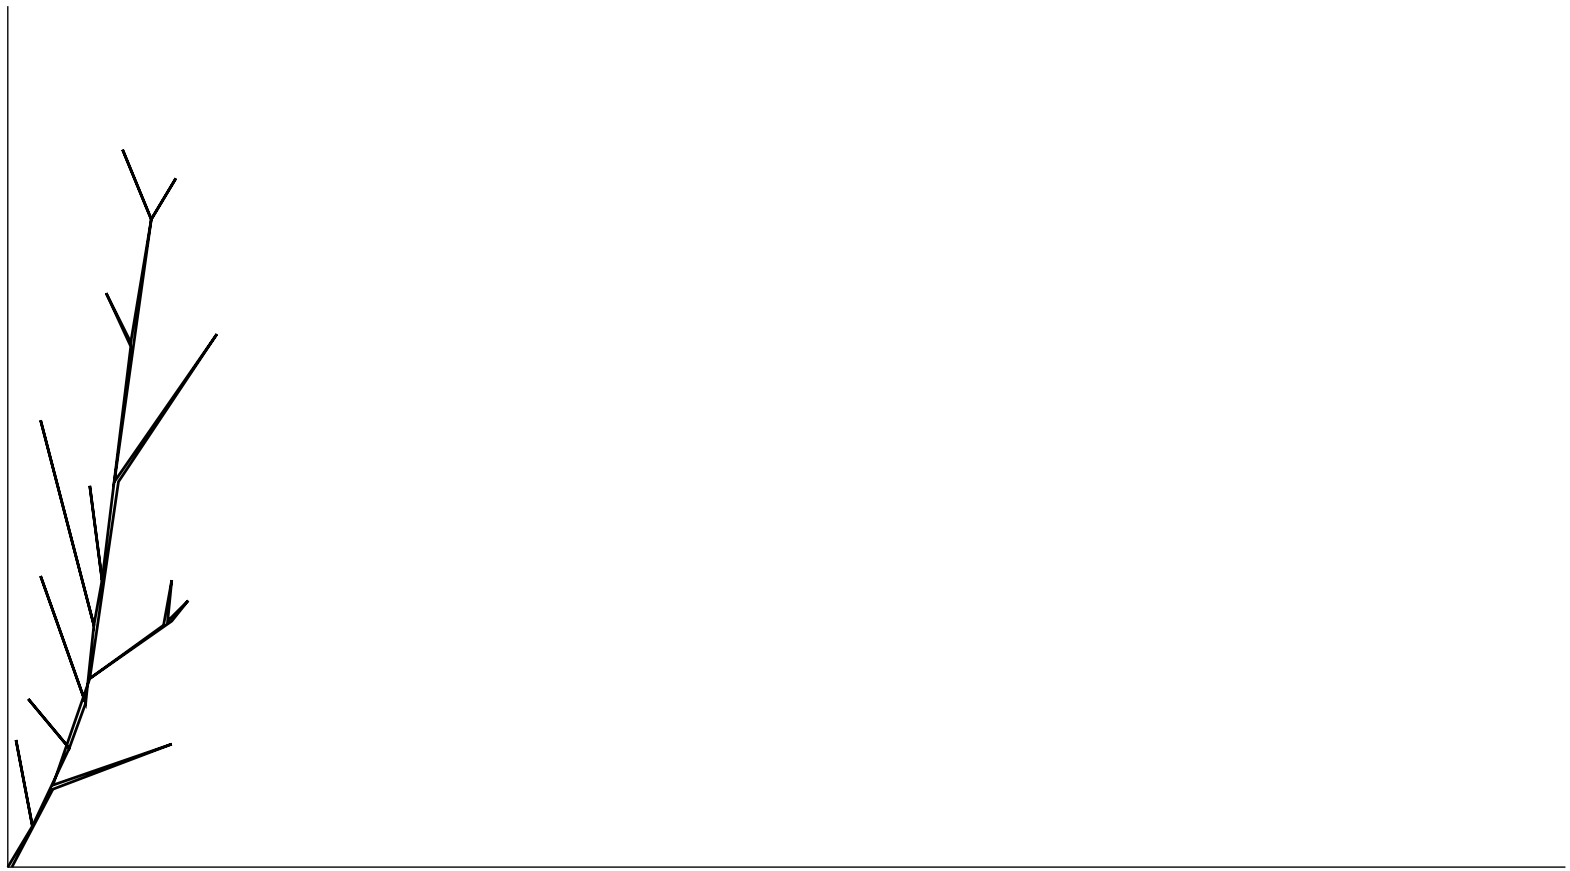

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together to get a tree.

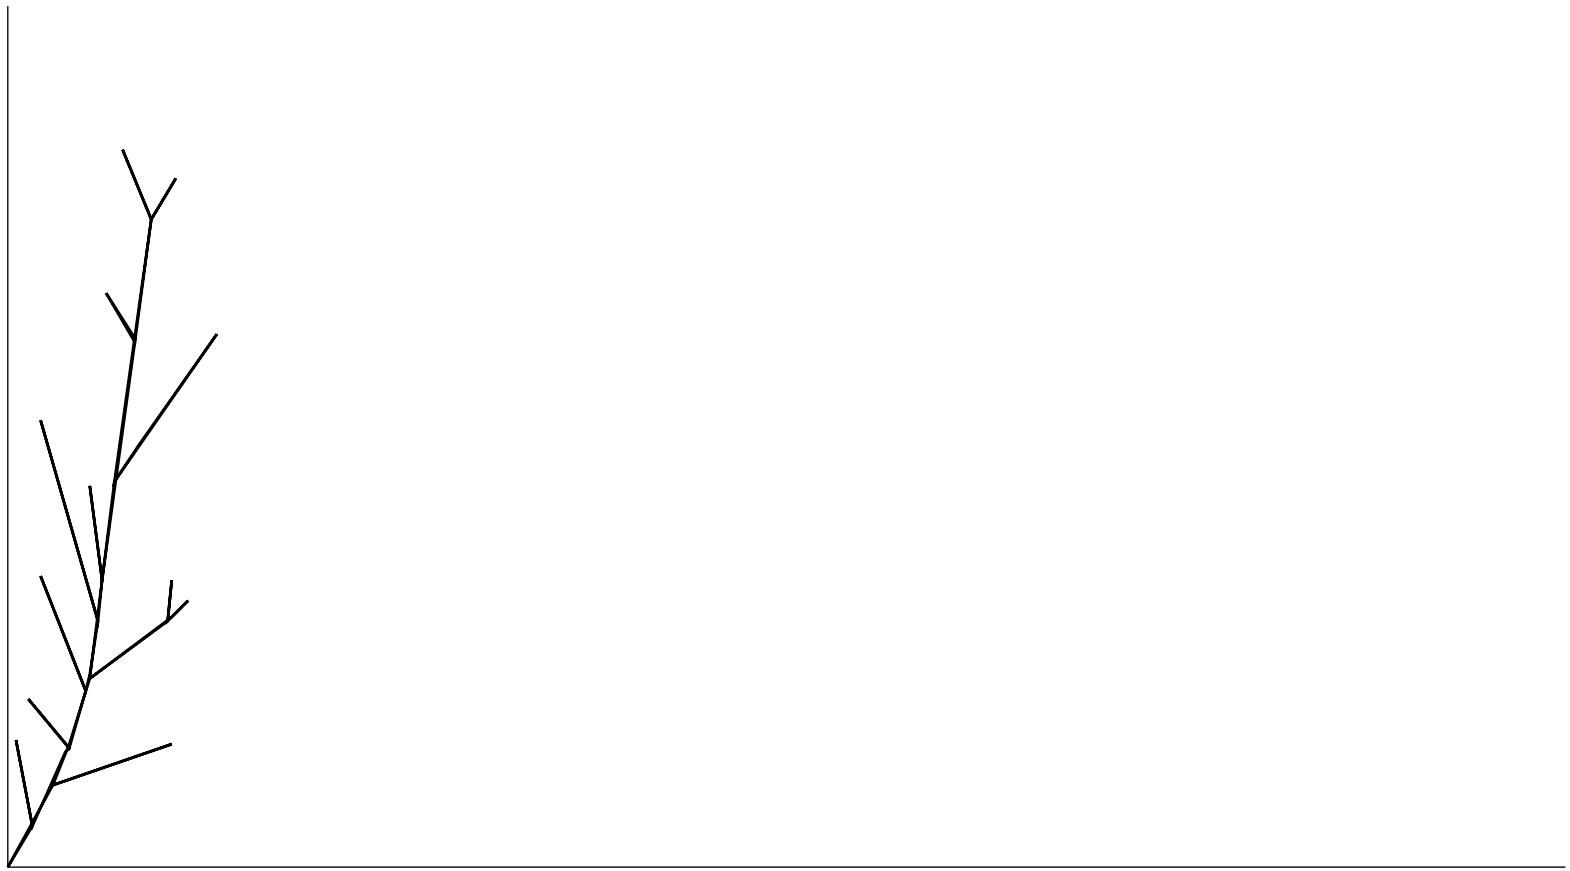

From excursions back to trees

Now put glue on the underside of the excursion and push the two sides together to get a real tree.

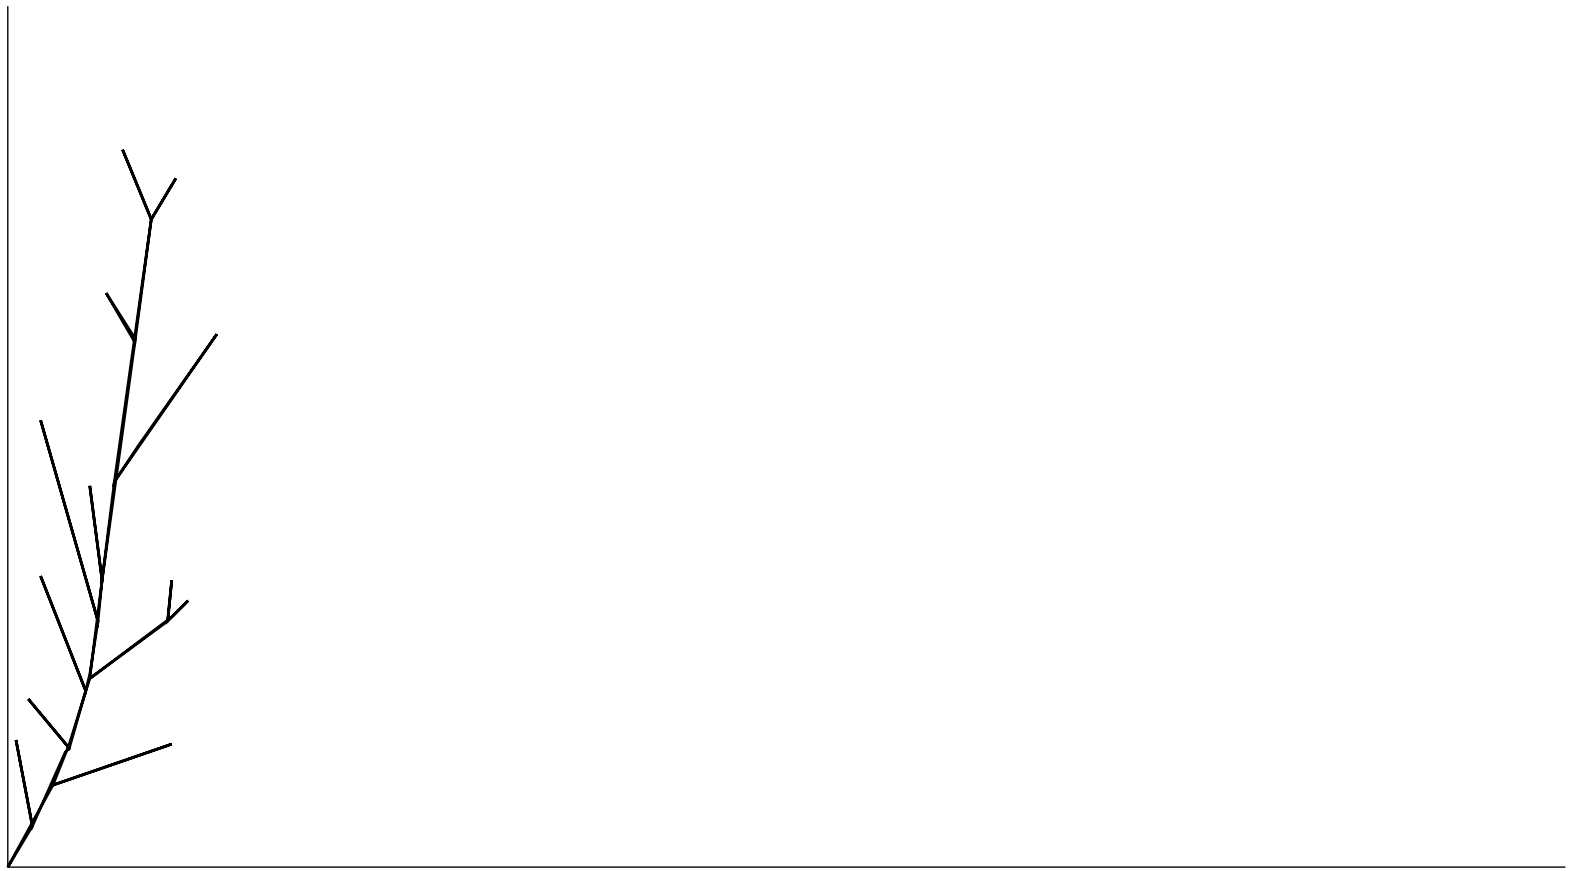

From excursions back to trees

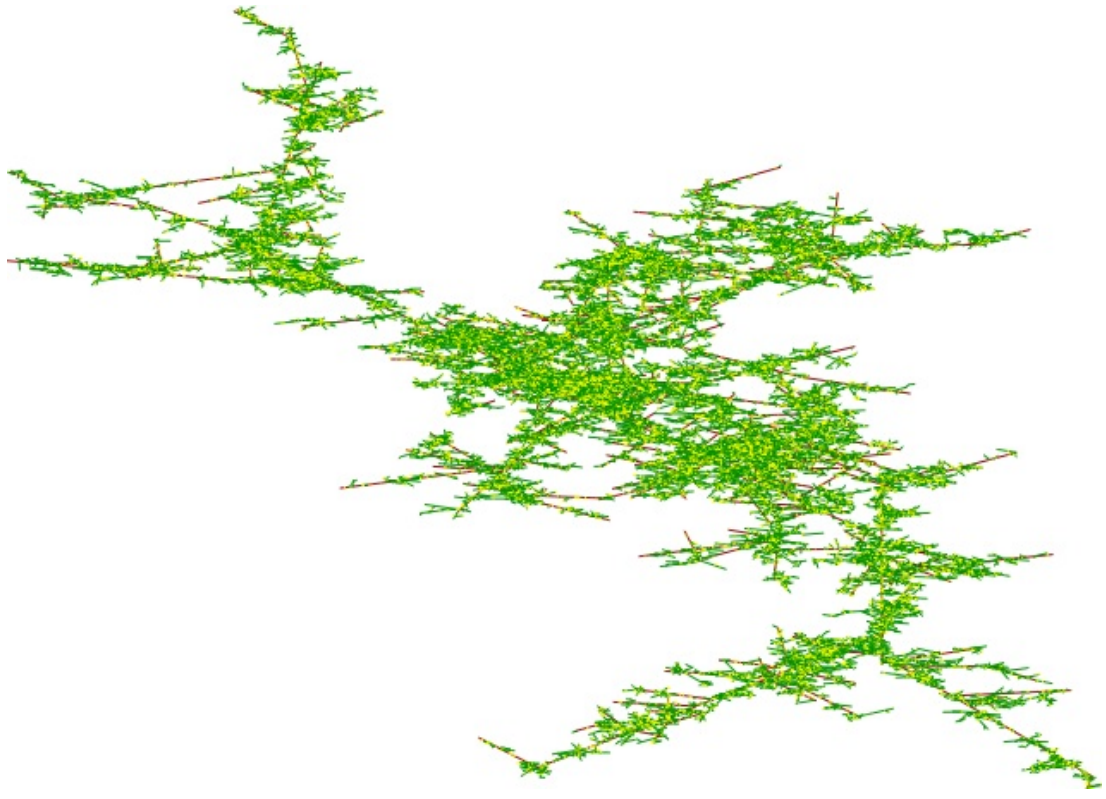

The Brownian continuum random tree is the real tree we obtain by doing this gluing procedure to the function $(2e(t), 0 \leq t \leq 1)$.

From excursions back to trees

If we do the same gluing operation to the contour function of a discrete tree, we get a real tree which is at Gromov-Hausdorff distance 1 from the discrete tree.

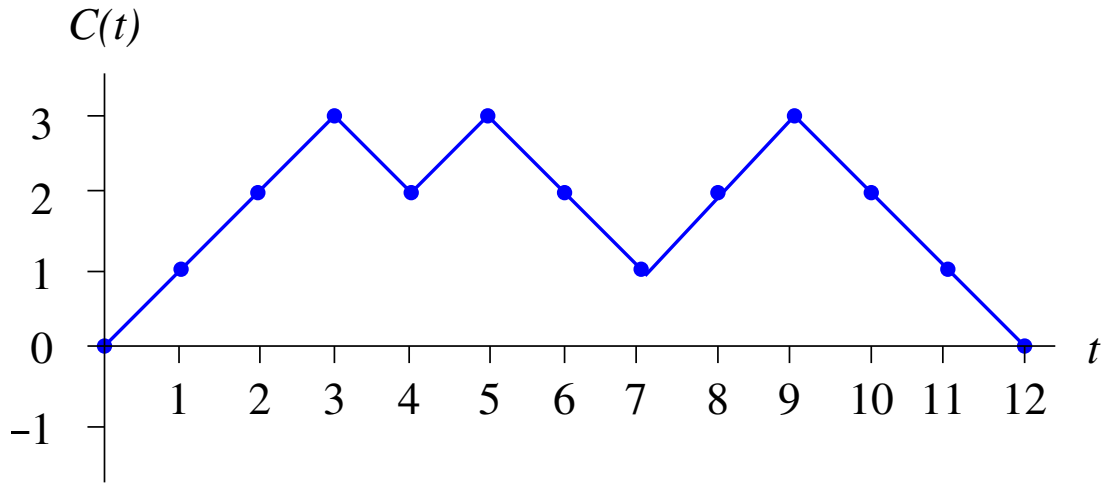

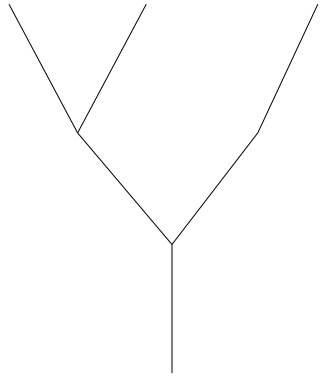

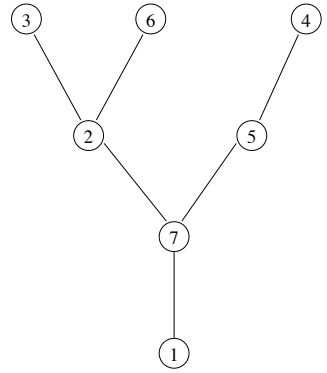

The convergence result

Recall that $(C^n(t), 0 \leq t \leq 2n)$ is the contour function of a critical Galton-Watson tree conditioned to have n vertices and that

$$\frac{\sigma}{\sqrt{n}}(C^n(2nt), 0 \leq t \leq 1) \xrightarrow{d} (2e(t), 0 \leq t \leq 1).$$

Write $\mathcal{T}_n$ for the real tree version of $\frac{\sigma}{\sqrt{n}} T_n$. We have

$$\mathrm{d_{GH}}\left(\frac{\sigma}{\sqrt{n}} T_n, \mathcal{T}_n\right) \leq \frac{\sigma}{\sqrt{n}}.$$

The convergence result

$$\frac{\sigma}{\sqrt{n}}(C^n(2nt), 0 \leq t \leq 1) \xrightarrow{d} (2e(t), 0 \leq t \leq 1).$$

These contour functions give us a natural correspondence between $\mathcal{T}_n$ and $\mathcal{T}$: just match up the positions parametrized by t in the two trees.

Since the contour function converges to $2e$ in distribution, the distortion of this correspondence goes to 0 and so $\frac{\sigma}{\sqrt{n}} T_n$ converges to $\mathcal{T}$ in the Gromov-Hausdorff distance.

Random trees: a summary

- T_n is a uniform random tree on vertices labelled by $\{1, 2, \ldots, n\}$.
- We think of T_n as a (random) metric space by using the graph distance.
- Write $\frac{1}{\sqrt{n}} T_n$ for the same metric space with all distances divided by $\sqrt{n}$.
- $\frac{1}{\sqrt{n}} T_n \xrightarrow{d} \mathcal{T}$ as $n \to \infty$ in the sense of the Gromov-Hausdorff distance.
- $\mathcal{T}$ is the Brownian continuum random tree, the tree encoded by $(2e(t), 0 \leq t \leq 1)$ (or constructed via the stick-breaking procedure).
- The same limit holds for any Galton-Watson tree T_n with critical offspring distribution of finite variance (up to a scaling constant depending on the variance).

Scaling limits for other random trees

The Brownian CRT is a universal scaling limit for a whole class of trees including critical Galton-Watson trees with finite offspring variance, as well as unordered binary trees (Marckert & Miermont), uniform unordered trees (Haas & Miermont) and critical multitype Galton-Watson trees (Miermont).

Scaling limits for other random trees

The Brownian CRT is a universal scaling limit for a whole class of trees including critical Galton-Watson trees with finite offspring variance, as well as unordered binary trees (Marckert & Miermont), uniform unordered trees (Haas & Miermont) and critical multitype Galton-Watson trees (Miermont).

If the offspring distribution is critical but in the domain of attraction of a stable law of index in $(1, 2)$ then there is convergence to the so-called stable tree which is encoded by an excursion of a stable Lévy process (Duquesne). More generally, there is a whole family of Lévy trees (Duquesne & Le Gall).

Bibliography: random trees

The continuum random tree I
D. Aldous, *Annals of Probability* **19** (1991) pp.1-28.

The continuum random tree II. An overview
D. Aldous, in *Stochastic analysis (Durham 1990)*, vol. 167 of London Mathematical Society Lecture Note Series (1991) pp.23-70.

The continuum random tree III
D. Aldous, *Annals of Probability* **21** (1993) pp.248-289.

Random trees and applications
J.-F. Le Gall, *Probability Surveys* **2** (2005) pp.245-311.

Random trees, Lévy processes and spatial branching processes
T. Duquesne and J.-F. Le Gall, *Asterisque* **281** (2002).

Stochastics Meeting Lunteren 2012

Lecture 2: random graphs

Joint work with

Louigi Addario-Berry and **Nicolas Broutin**

McGill University

INRIA Rocquencourt, Paris

Random trees: a summary from Lecture 1

- T_m is a uniform random tree on vertices labelled by $\{1, 2, \ldots, m\}$.
- We think of T_m as a (random) metric space by using the graph distance.
- Write $\frac{1}{\sqrt{m}} T_m$ for the same metric space with all distances divided by $\sqrt{m}$.
- $\frac{1}{\sqrt{m}} T_m \xrightarrow{d} \mathcal{T}$ as $m \to \infty$ in the sense of the Gromov-Hausdorff distance.
- $\mathcal{T}$ is the Brownian continuum random tree, the tree encoded by $(2e(t), 0 \leq t \leq 1)$ (or constructed via the stick-breaking procedure).

The Erdős-Rényi random graph

Take n vertices labelled by $[n] := \{1, 2, \ldots, n\}$ and put an edge between any pair independently with probability p. Call the resulting model $G(n, p)$.

Example: $n = 10$, $p = 0.4$.

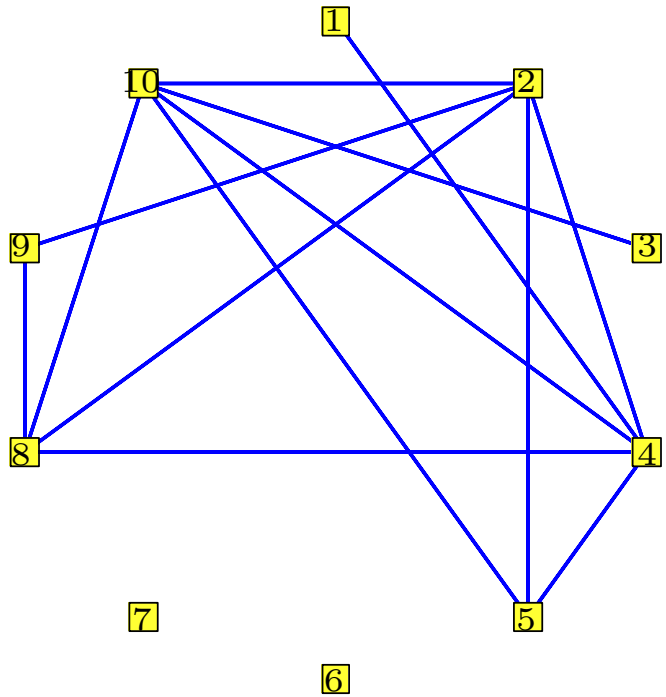

Connected components

We're going to be interested in the connected components of these graphs.

Below, there are three of them.

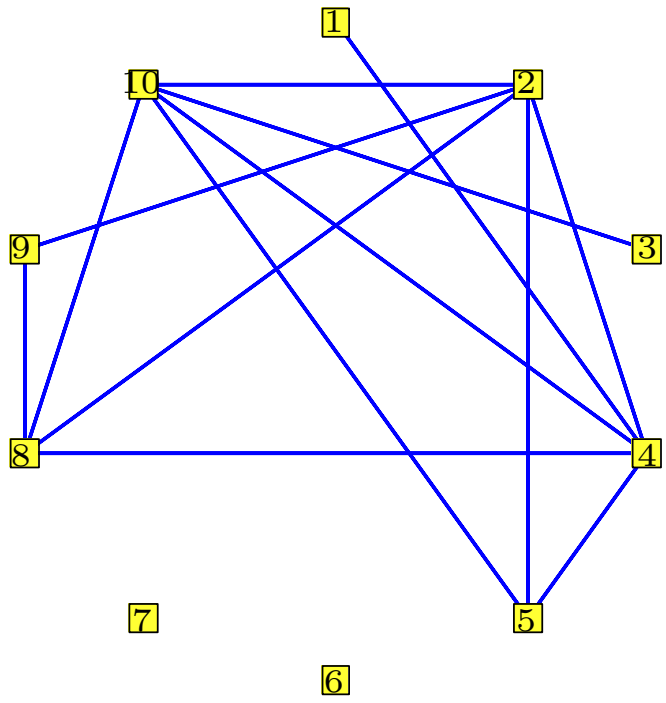

The phase transition

Let $p = c/n$ and consider the largest component (vertices in green, edges in red).

$n = 200,\ c = 0.4$

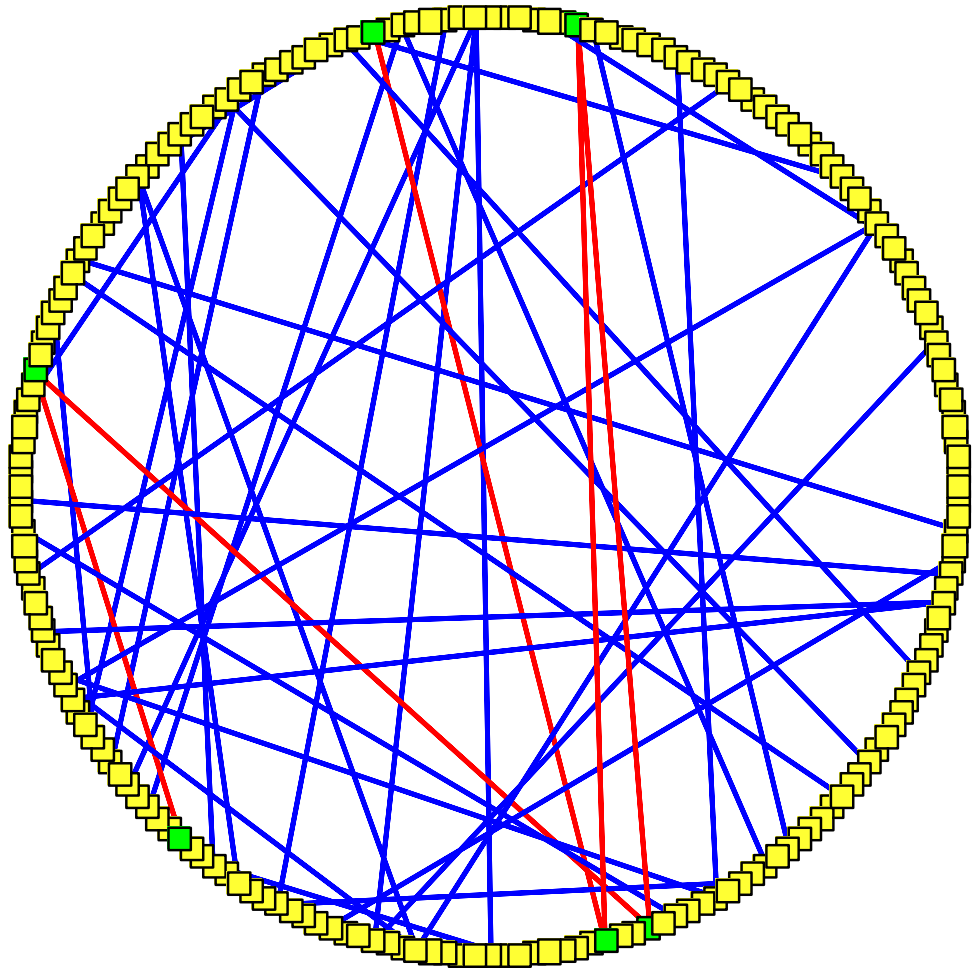

The phase transition

Let $p = c/n$ and consider the largest component (vertices in green, edges in red).

$n = 200,\ c = 0.8$

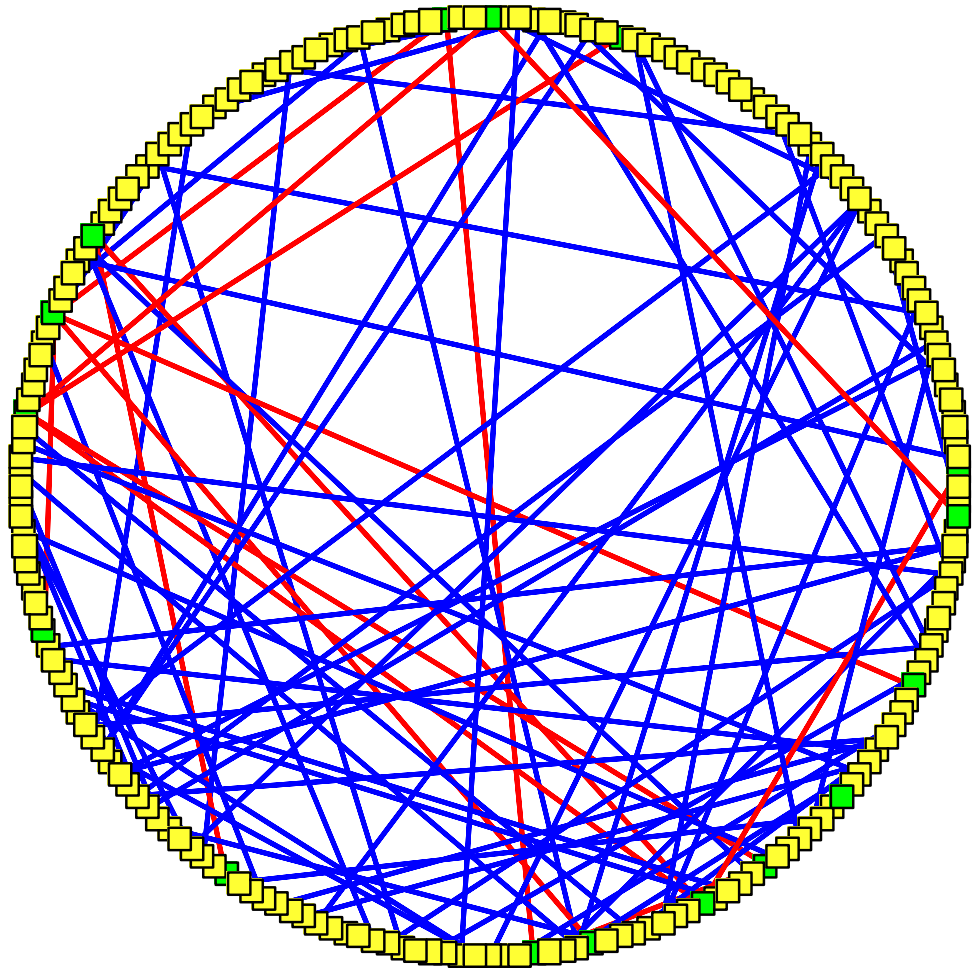

The phase transition

Let $p = c/n$ and consider the largest component (vertices in green, edges in red).

$n = 200,\ c = 1.2$

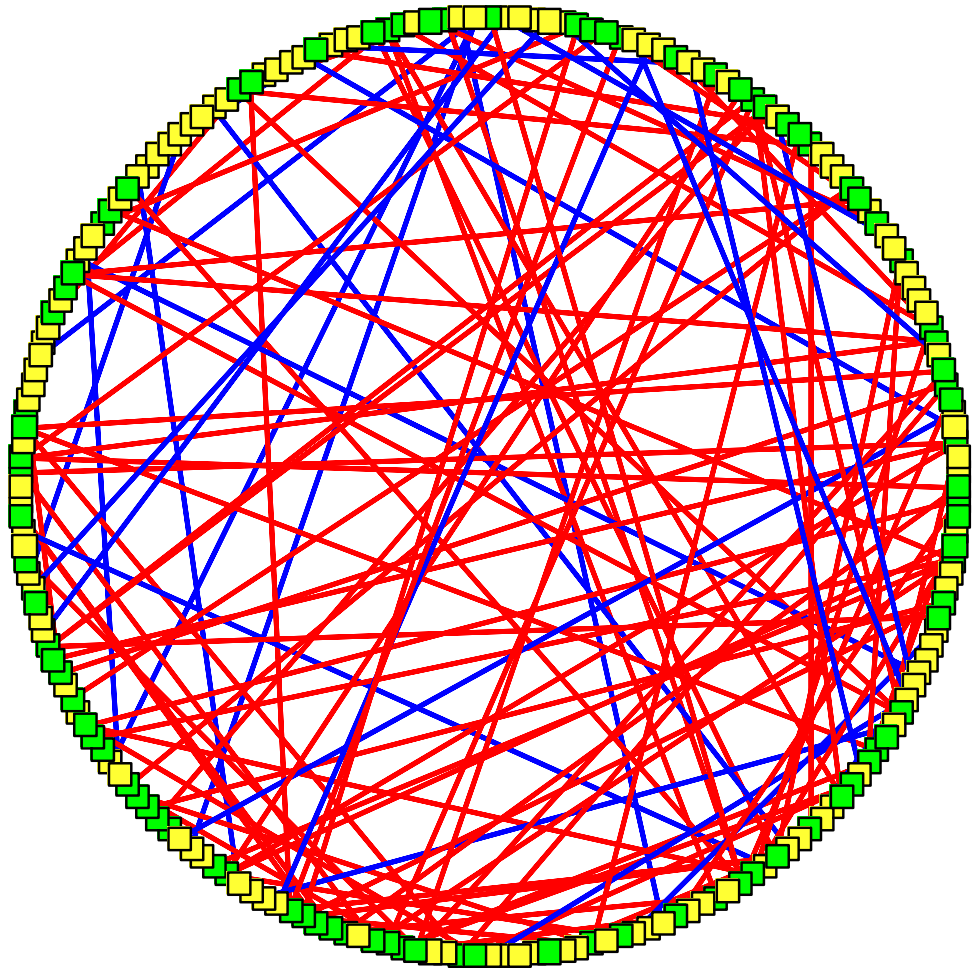

The phase transition (Erdős and Rényi (1960))

By the size of a component, we mean its number of vertices.

The phase transition (Erdős and Rényi (1960))

By the size of a component, we mean its number of vertices.

Consider $p = c/n$.

- For $c < 1$, the largest connected component has size $O(\log n)$;
- for $c > 1$, the largest connected component has size $\Theta(n)$ (and the others are all $O(\log n)$).

(These statements hold with probability tending to 1 as $n \to \infty$.)

Heuristic picture of the phase transition

Vertex 1 has a Binomial$(n-1, c/n) \approx$ Poisson(c) number of neighbours, N.

Heuristic picture of the phase transition

Vertex 1 has a Binomial$(n-1, c/n) \approx$ Poisson(c) number of neighbours, N.

Consider now one of those neighbours. It has a Binomial$(n-N-1, c/n)$ number of neighbours we haven't seen before, which is still well-approximated by Poisson(c) as long as $N = o(n)$.

Heuristic picture of the phase transition

Vertex 1 has a Binomial$(n-1, c/n) \approx$ Poisson(c) number of neighbours, N.

Consider now one of those neighbours. It has a Binomial$(n-N-1, c/n)$ number of neighbours we haven't seen before, which is still well-approximated by Poisson(c) as long as $N = o(n)$.

Continuing in this way, we see that we can approximate the size of the component containing vertex 1 by the total progeny in a branching process with Poisson(c) offspring distribution (as long as the population doesn't get too large...).

Heuristic picture of the phase transition

If $c \leq 1$, this branching process dies out with probability 1, which corresponds to getting only a small component containing vertex 1. A similar argument works for subsequent components.

Heuristic picture of the phase transition

If $c \leq 1$, this branching process dies out with probability 1, which corresponds to getting only a small component containing vertex 1. A similar argument works for subsequent components.

If, on the other hand, $c > 1$, there is a positive probability that the branching process will survive. The branching process approximation holds good until we explore the first component which does not die out; this component ends up being the giant.

The critical point of the phase transition

Recall: $p = c/n$.

- For $c < 1$, the largest connected component has size $O(\log n)$;
- for $c > 1$, the largest connected component has size $\Theta(n)$ (and the others are all $O(\log n)$).

The critical point of the phase transition

Recall: $p = c/n$.

- For $c < 1$, the largest connected component has size $O(\log n)$;
- for $c > 1$, the largest connected component has size $\Theta(n)$ (and the others are all $O(\log n)$).

If $c = 1$, the largest component has size $\Theta(n^{2/3})$ and, indeed, there is a whole sequence of components of this order.

The critical random graph

The critical window: $p = \frac{1}{n} + \frac{\lambda}{n^{4/3}}$, where $\lambda \in \mathbb{R}$. For such p, the largest components have size $\Theta(n^{2/3})$.

The critical random graph

The critical window: $p = \frac{1}{n} + \frac{\lambda}{n^{4/3}}$, where $\lambda \in \mathbb{R}$. For such p, the largest components have size $\Theta(n^{2/3})$.

We will also be interested in the surplus of a component, the number of edges more than a tree that it has.

A component with surplus 3:

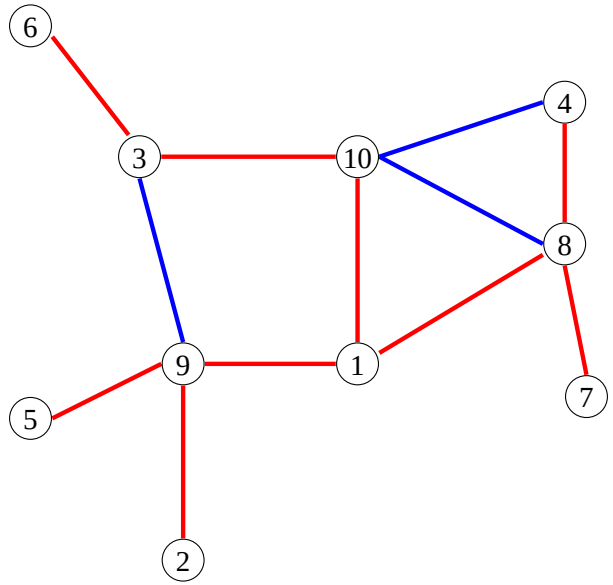

Convergence of the sizes and surpluses

Fix λ and let $C_1^n, C_2^n, \ldots$ be the sequence of component sizes of $G\left(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}}\right)$ in decreasing order, and let $S_1^n, S_2^n, \ldots$ be the corresponding surpluses.

Write $\mathbf{C}^n = (C_1^n, C_2^n, \ldots)$ and $\mathbf{S}^n = (S_1^n, S_2^n, \ldots)$.

Convergence of the sizes and surpluses

Fix λ and let $C_1^n, C_2^n, \ldots$ be the sequence of component sizes of $G\left(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}}\right)$ in decreasing order, and let $S_1^n, S_2^n, \ldots$ be the corresponding surpluses.

Write $\mathbf{C}^n = (C_1^n, C_2^n, \ldots)$ and $\mathbf{S}^n = (S_1^n, S_2^n, \ldots)$.

Theorem. (Aldous (1997)) As $n \to \infty$,

$$(n^{-2/3}\mathbf{C}^n, \mathbf{S}^n) \xrightarrow{d} (\mathbf{C}, \mathbf{S}).$$

Convergence of the sizes and surpluses

Theorem. (Aldous (1997)) As $n \to \infty$,

$$(n^{-2/3}\mathbf{C}^n, \mathbf{S}^n) \xrightarrow{d} (\mathbf{C}, \mathbf{S}) \quad \text{as } n \to \infty.$$

Convergence for the first co-ordinate takes place in

$$\ell^2_{\searrow} := \left\{ \mathbf{x} = (x_1, x_2, \ldots) : x_1 \geq x_2 \geq \ldots \geq 0, \sum_{i=1}^{\infty} x_i^2 < \infty \right\}$$

with the usual ℓ^2-distance $\|\mathbf{x} - \mathbf{y}\|_2 = \sqrt{\sum_{i=1}^{\infty}(x_i - y_i)^2}$. For the second co-ordinate, convergence is in the distance

$$d(\mathbf{u}, \mathbf{v}) = 2^{-\inf\{i \geq 1 : u_i \neq v_i\}}$$

between integer sequences $\mathbf{u}$ and $\mathbf{v}$.

Limiting sizes and surpluses

Let $W_\lambda(t) = W(t) + \lambda t - \frac{t^2}{2}$, $t \geq 0$, where $(W(t), t \geq 0)$ is a standard Brownian motion.

Limiting sizes and surpluses

Let $W_\lambda(t) = W(t) + \lambda t - \frac{t^2}{2}$, $t \geq 0$, where $(W(t), t \geq 0)$ is a standard Brownian motion.

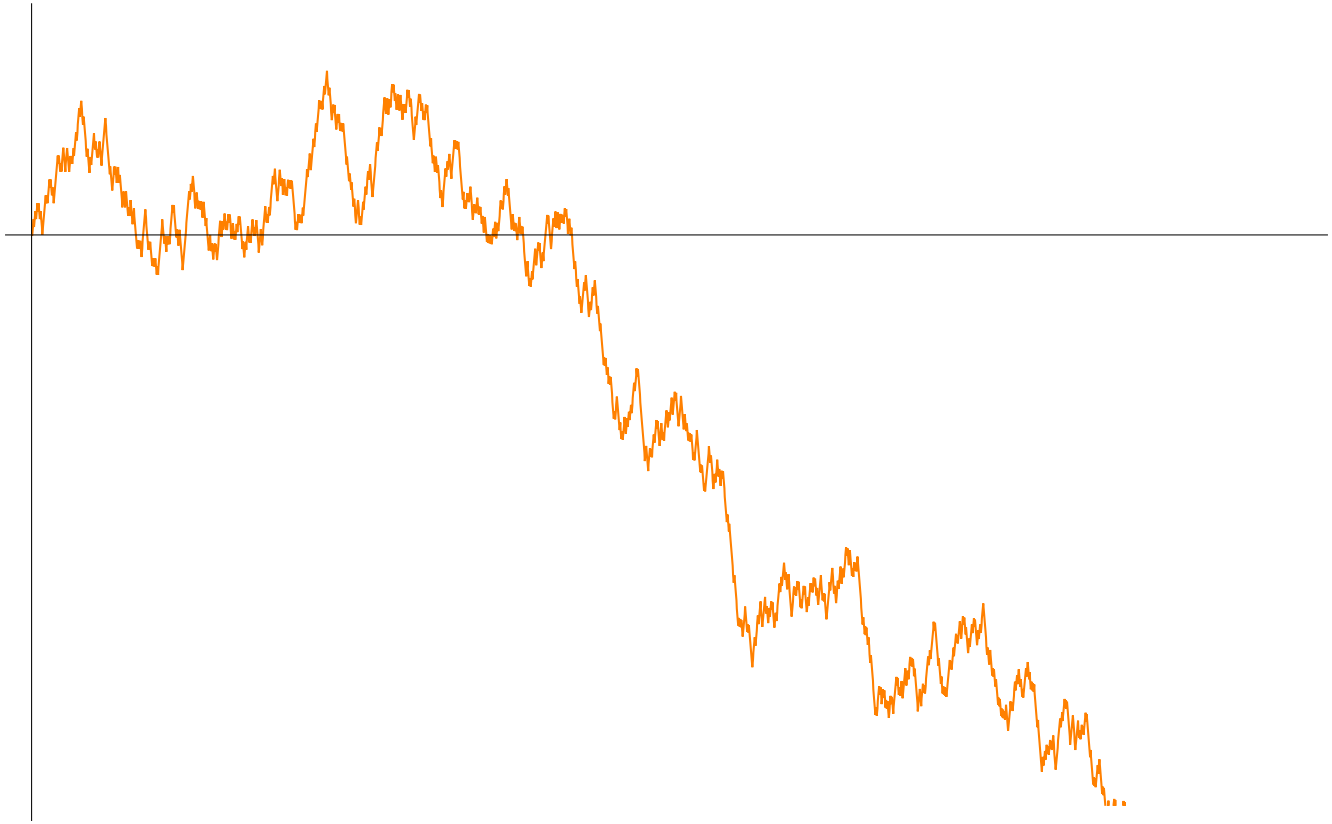

[Pictures by Louigi Addario-Berry]

Limiting sizes and surpluses

Let $W_\lambda(t) = W(t) + \lambda t - \frac{t^2}{2}$, $t \geq 0$, where $(W(t), t \geq 0)$ is a standard Brownian motion.

Let $B_\lambda(t) = W_\lambda(t) - \min_{0 \leq s \leq t} W_\lambda(s)$ be the process reflected at its minimum.

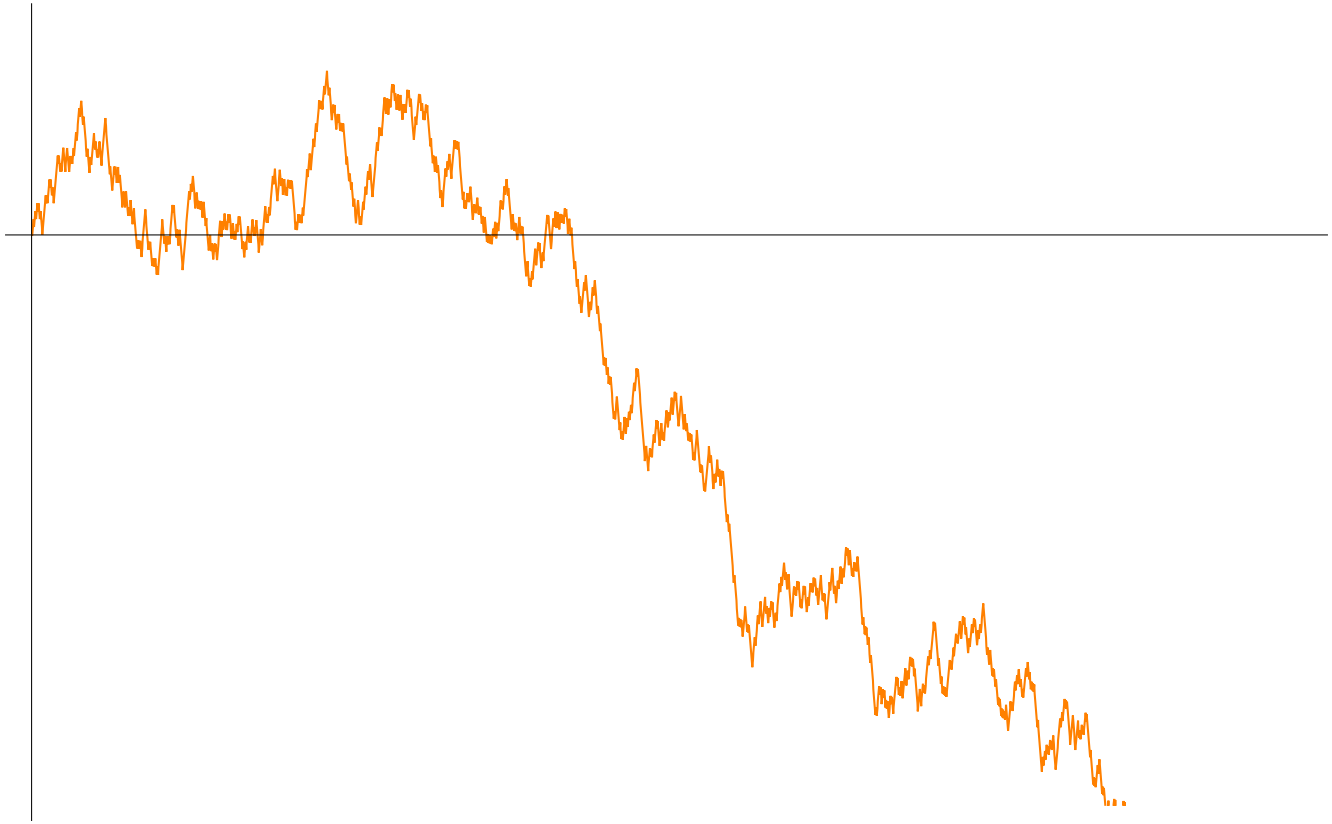

[Pictures by Louigi Addario-Berry]

Limiting sizes and surpluses

Let $W_\lambda(t) = W(t) + \lambda t - \frac{t^2}{2}$, $t \geq 0$, where $(W(t), t \geq 0)$ is a standard Brownian motion.

Let $B_\lambda(t) = W_\lambda(t) - \min_{0 \leq s \leq t} W_\lambda(s)$ be the process reflected at its minimum.

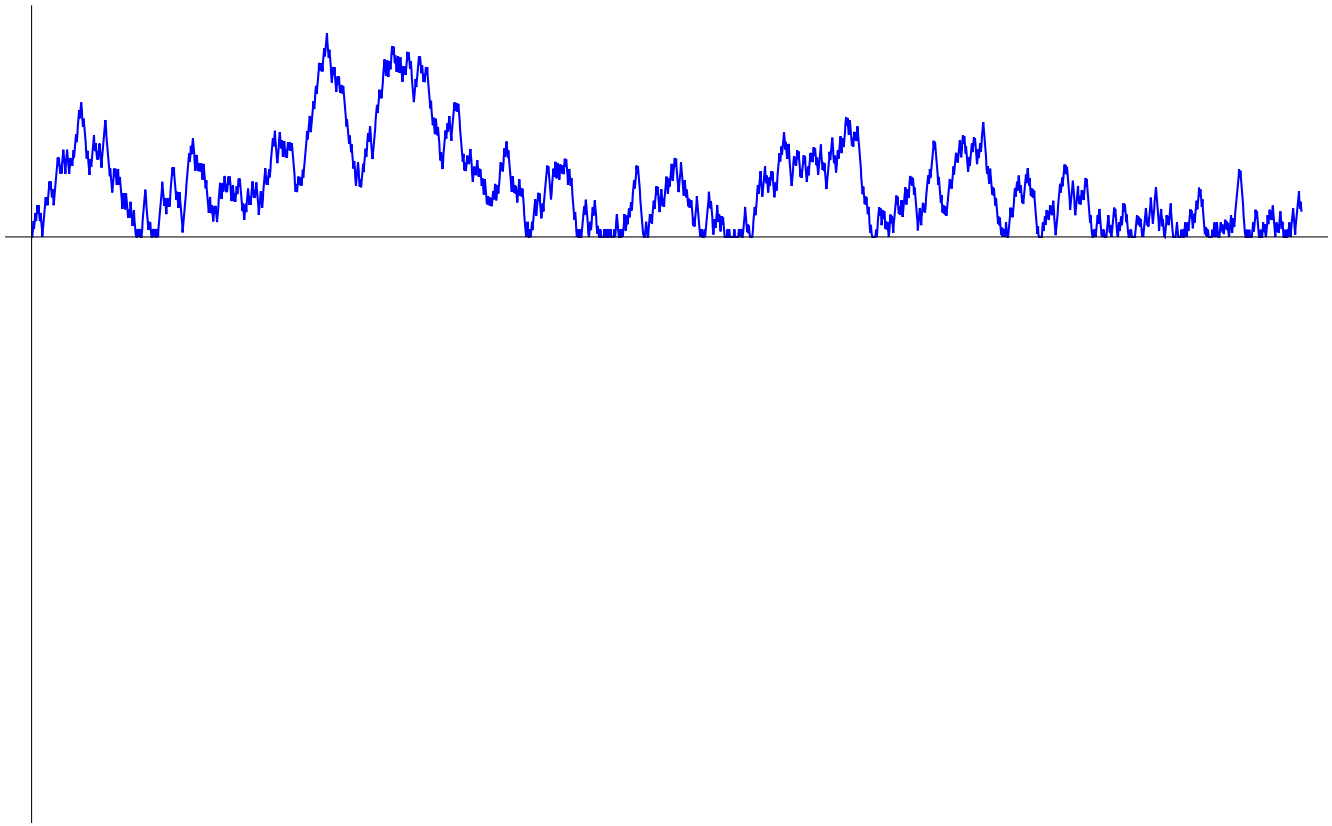

[Pictures by Louigi Addario-Berry]

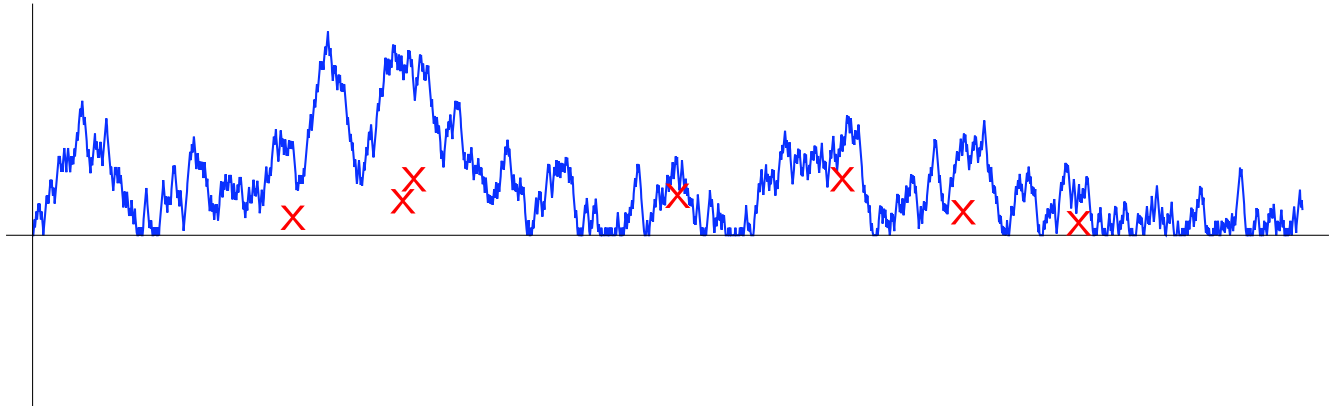

Decorate the picture with the points of a rate one Poisson process in the plane which fall above the x-axis and below the graph.

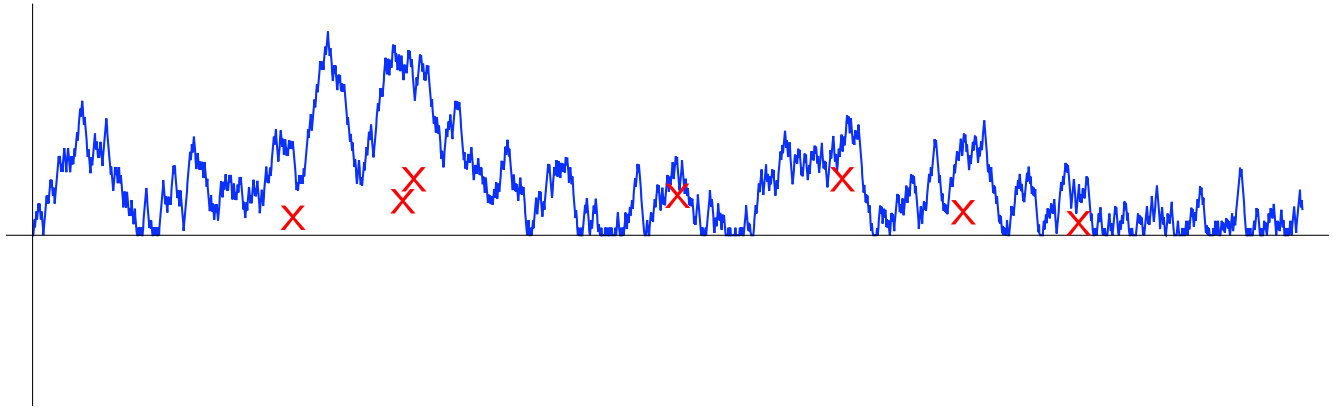

Decorate the picture with the points of a rate one Poisson process in the plane which fall above the x-axis and below the graph.

The path of B_λ can be split up into excursions above 0.

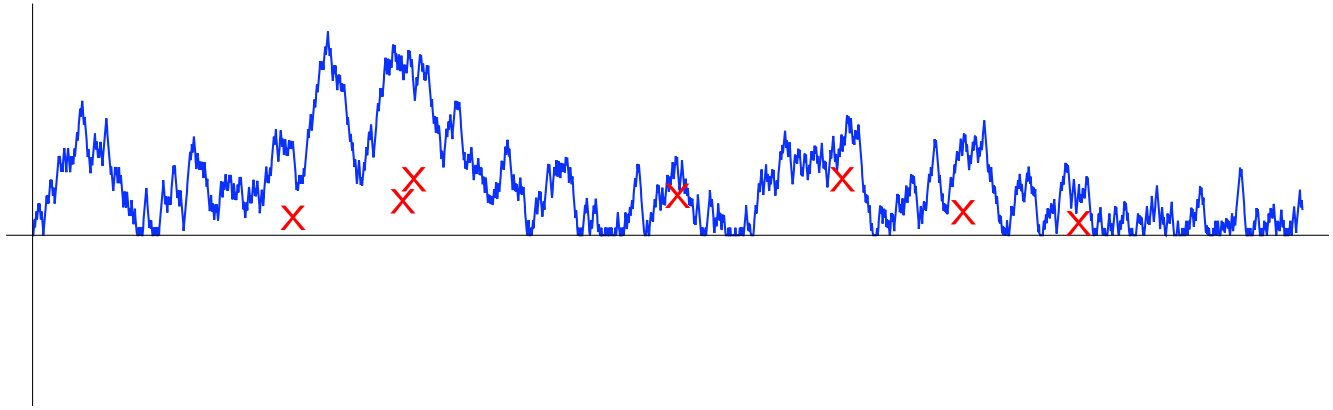

Decorate the picture with the points of a rate one Poisson process in the plane which fall above the x-axis and below the graph.

The path of B_λ can be split up into excursions above 0.

C is the sequence of lengths of these excursions, in decreasing order.

S is the sequence of numbers of points falling under those excursions.

Proof technique: depth-first exploration

As for the random trees yesterday, a key tool is a depth-first exploration. For a tree, we defined the depth-first walk by $X(0) = 0$ and, for $1 \leq k \leq m$,

$$X(k) = \sum_{i=0}^{k-1} (k(v_i) - 1),$$

where $k(v)$ is the number of children of vertex v and $v_0, v_1, \ldots, v_{m-1}$ are the vertices in depth-first order.

The problem with this is that the components of a random graph are not (in general) trees. This is resolved by simply ignoring the extra edges!

Depth-first exploration

Recall that vertices can have four states: unexplored, current, alive or dead. There's no need to stop when we hit the end of the first component. $X(k)$ will be the number of vertices alive at step k minus the number of components already fully explored.

Depth-first exploration: an example

Step 0

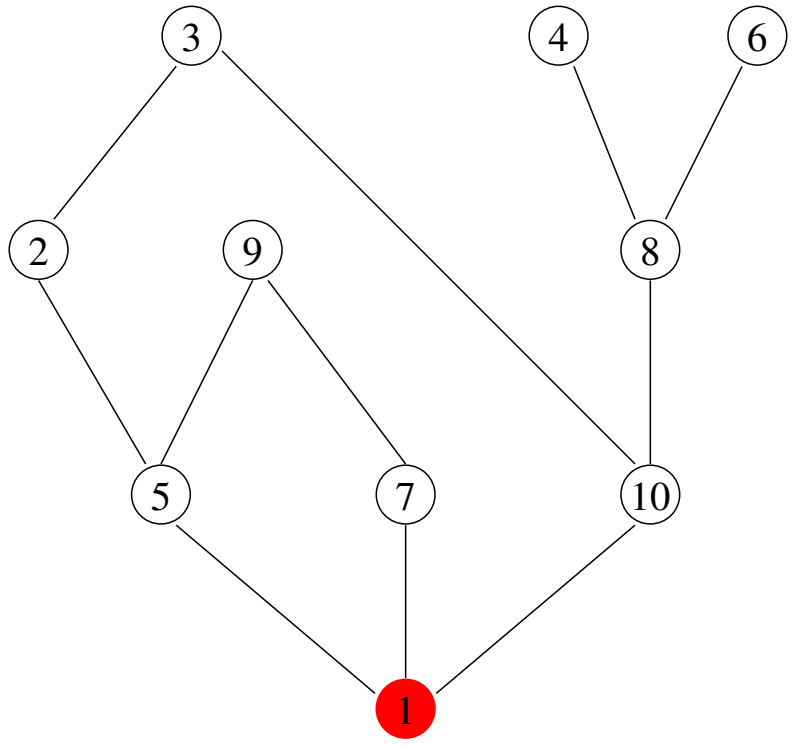

Current: 1 Alive: none Dead: none $X(0) = 0$.

Depth-first exploration: an example

Step 1

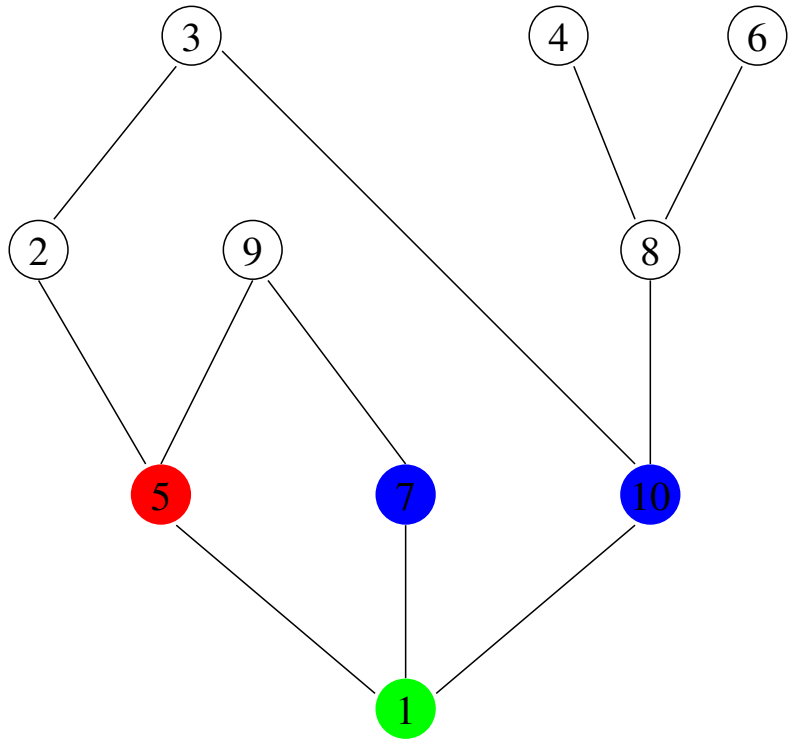

Current: 5 Alive: 7, 10 Dead: 1 $X(1) = 2$.

Depth-first exploration: an example

Step 2

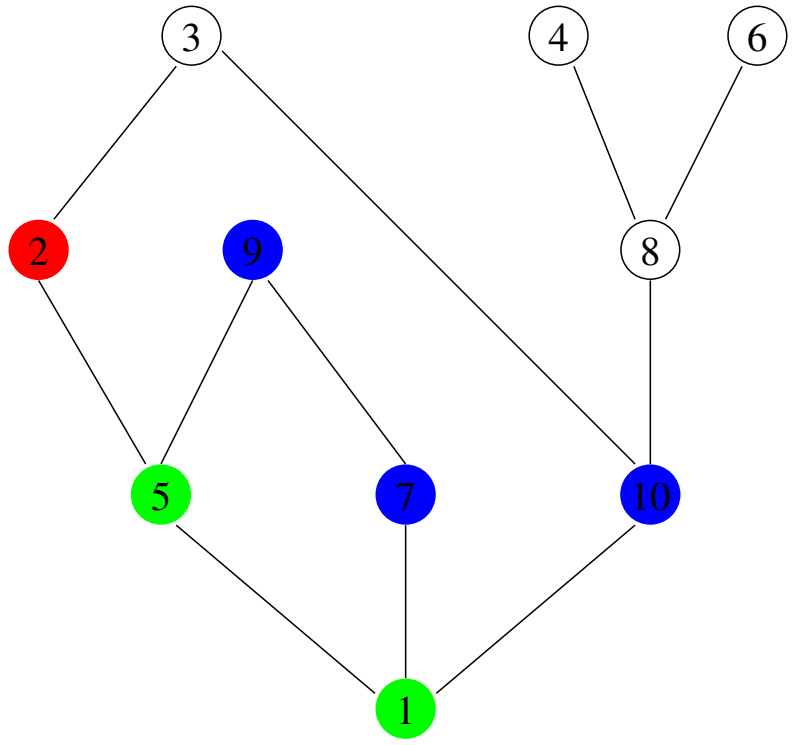

Current: 2 Alive: 9, 7, 10 Dead: 1, 5 $X(2) = 3$.

Depth-first exploration: an example

Step 3

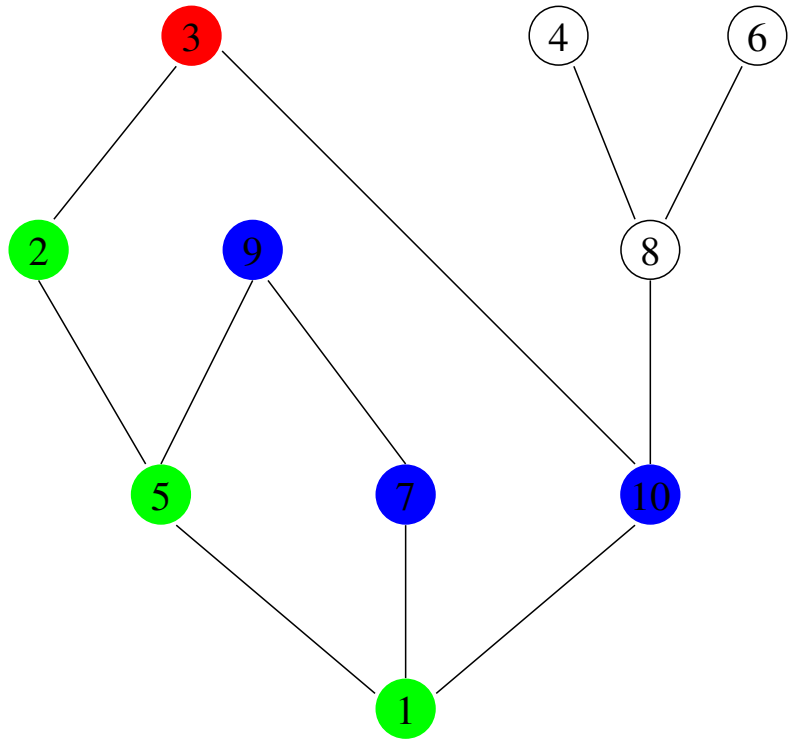

Current: 3 Alive: 9, 7, 10 Dead: 1, 5, 2 $X(3) = 3$.

Depth-first exploration: an example

Step 4

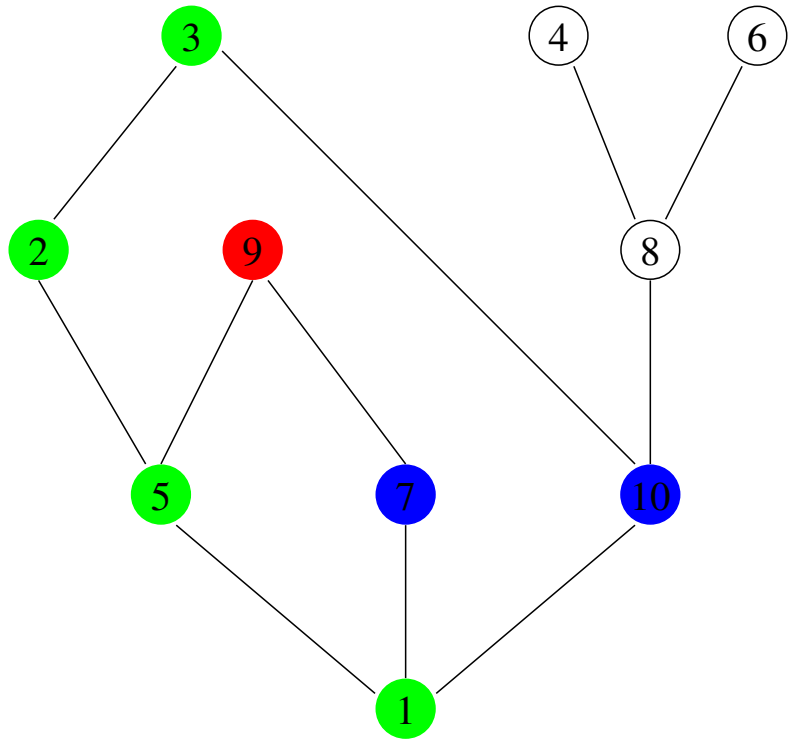

Current: 9 Alive: 7, 10 Dead: 1, 5, 2, 3 $X(4) = 2$.

Depth-first exploration: an example

Step 5

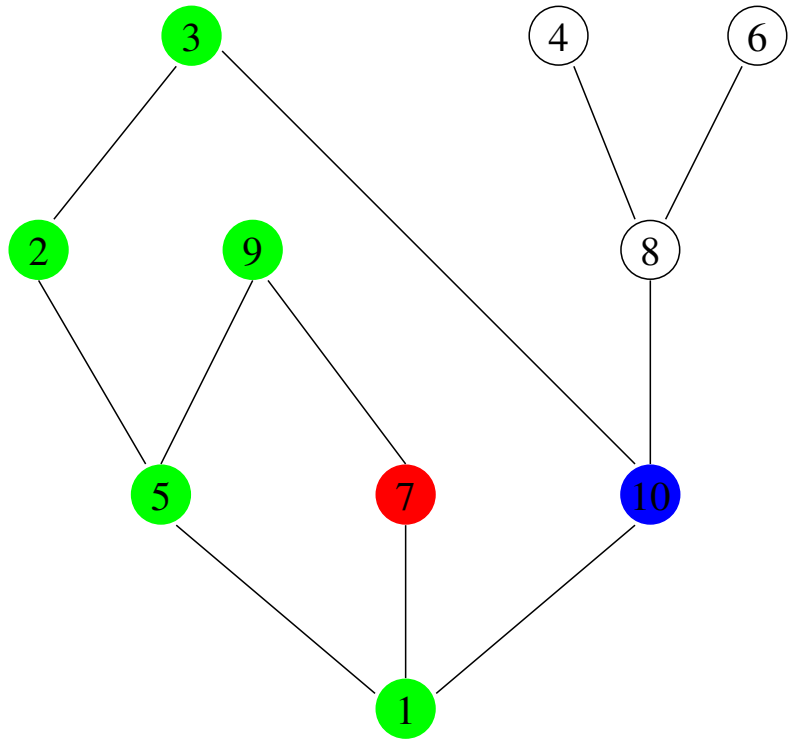

Current: 7 Alive: 10 Dead: 1, 5, 2, 3, 9 $X(5) = 1$.

Depth-first exploration: an example

Step 6

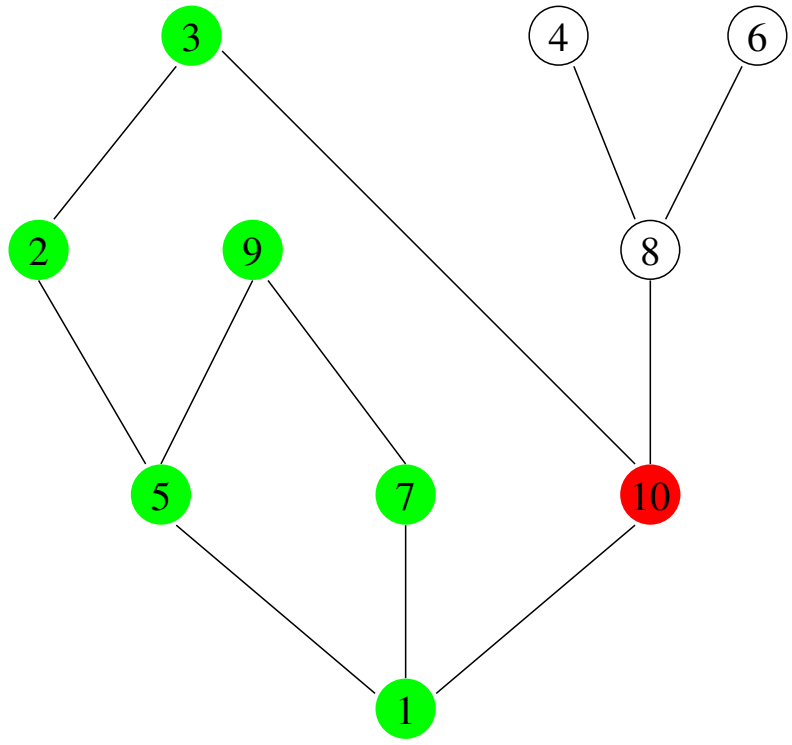

Current: 10 Alive: none Dead: 1, 5, 2, 3, 9, 7 $X(6) = 0$.

Depth-first exploration: an example

Step 7

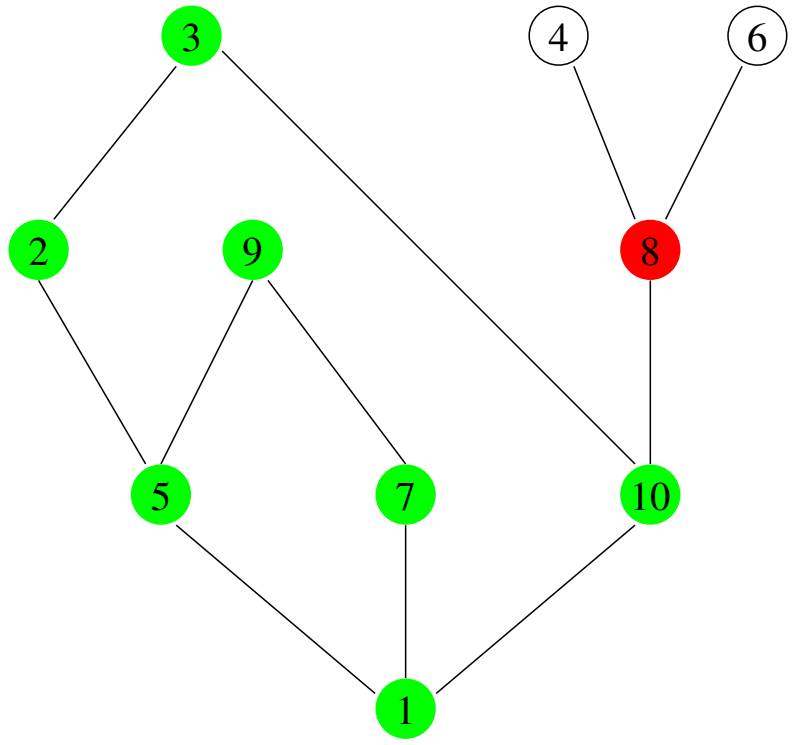

Current: 8 Alive: none Dead: 1, 5, 2, 3, 9, 7, 10 $X(7) = 0$.

Depth-first exploration: an example

Step 8

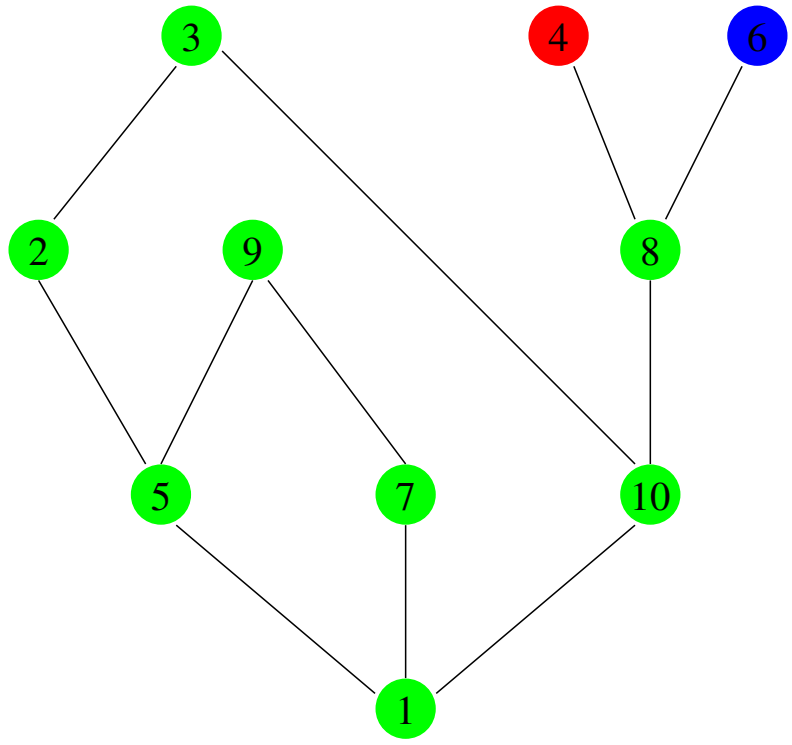

Current: 4 Alive: 6 Dead: 1, 5, 2, 3, 9, 7, 10, 8 $X(8) = 1$.

Depth-first exploration: an example

Step 9

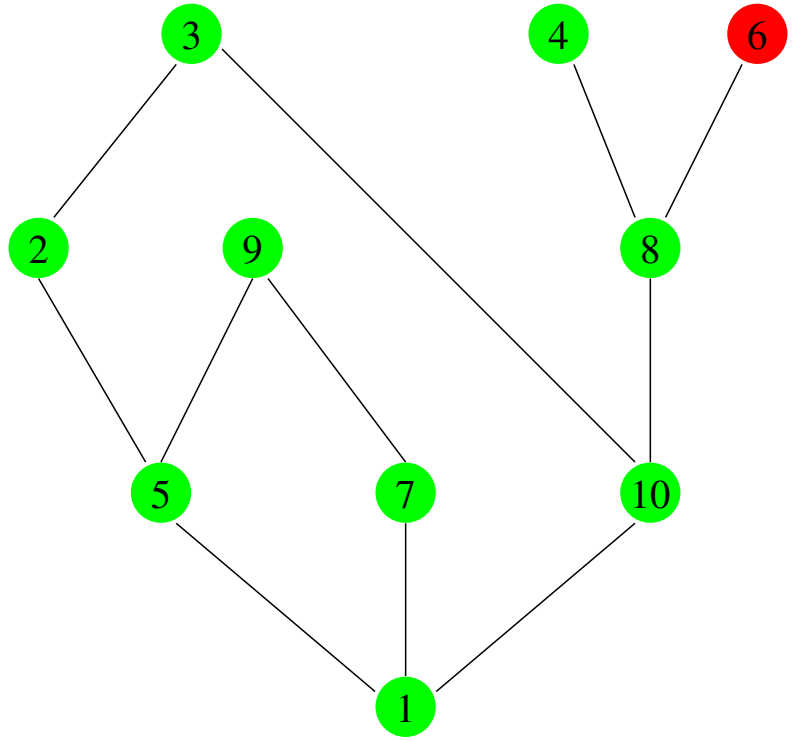

Current: 6 Alive: none Dead: 1, 5, 2, 3, 9, 7, 10, 8, 4

$X(9) = 0$.

Depth-first exploration: an example

We explored the graph as if the dashed edges weren't there:

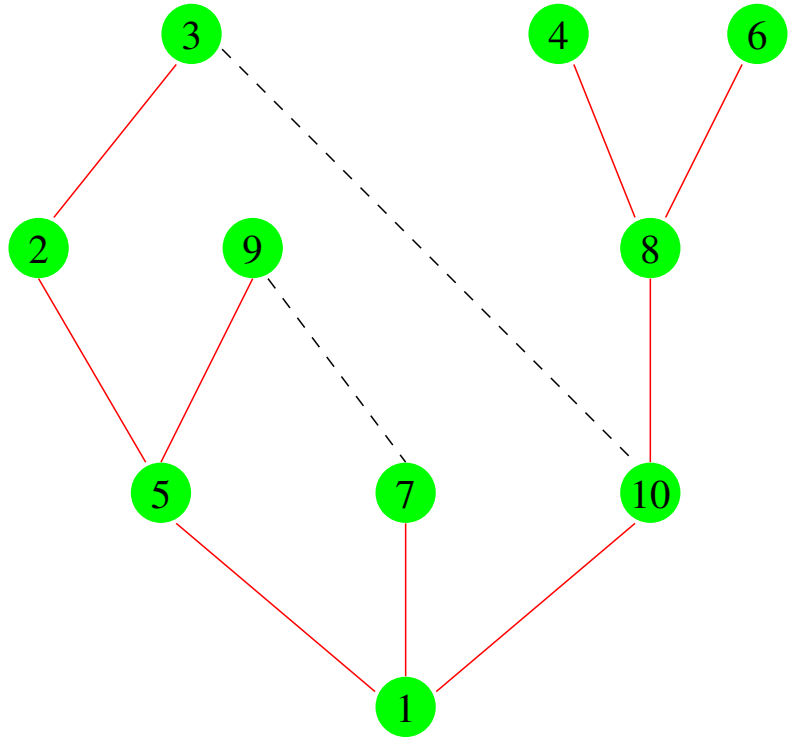

Depth-first walk

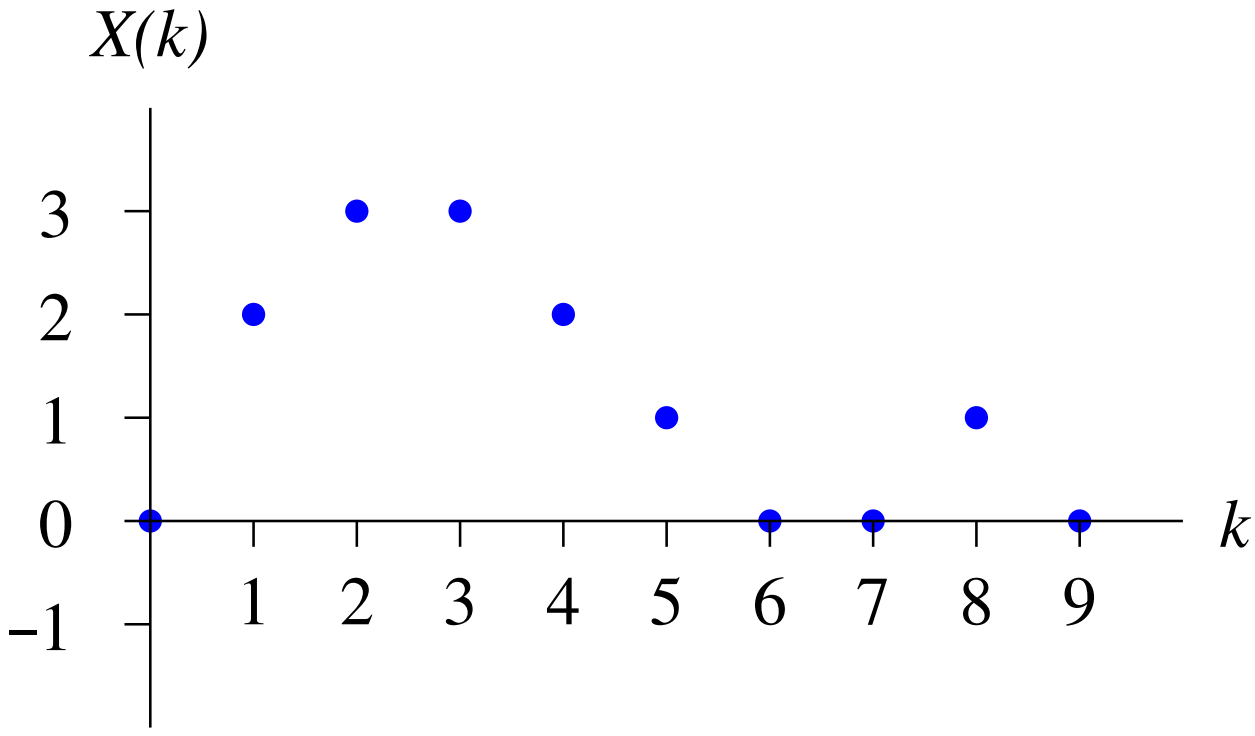

Depth-first walk

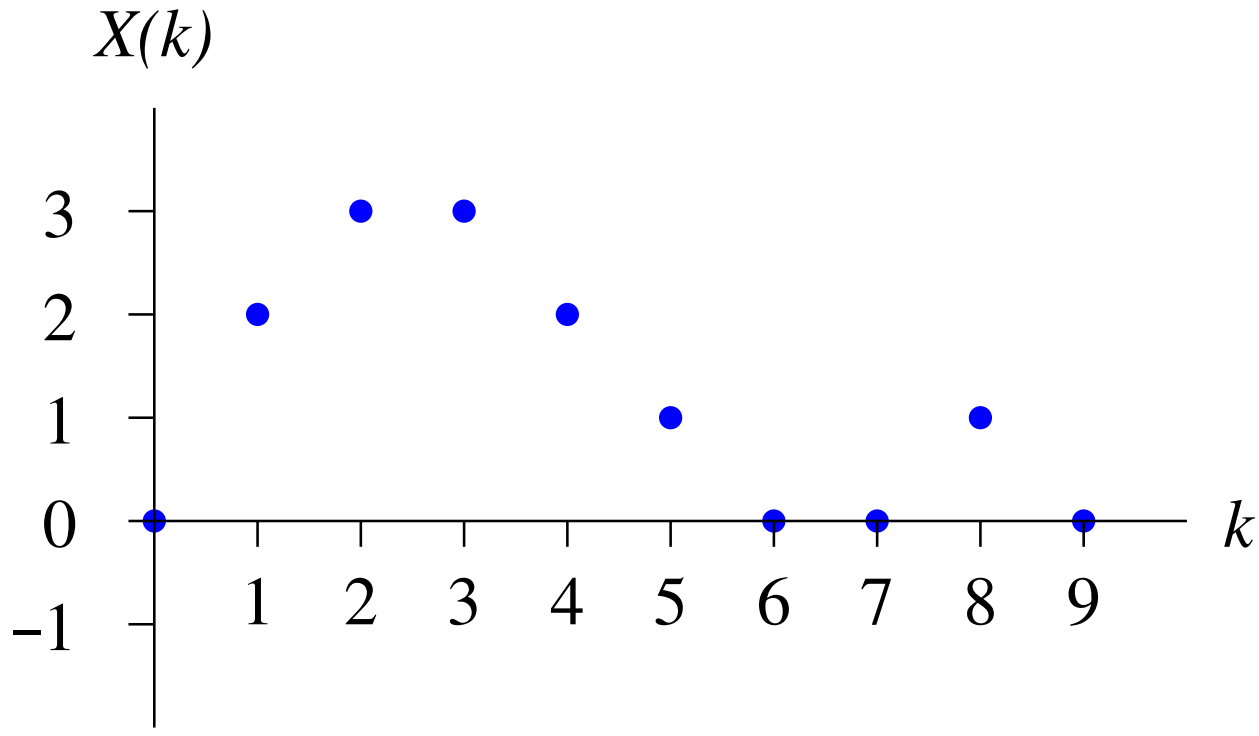

If there are several components, $T(k) = \inf\{i \geq 0 : X(i) = -k\}$ marks the beginning of the $(k+1)$th component. So the component sizes are $\{T(k+1) - T(k), k \geq 0\}$. This sequence can clearly be reconstructed from the path of $(X(i), i \geq 0)$.

Convergence of the depth-first walk

Let X_λ^n be the depth-first walk associated with $G\left(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}}\right)$.

Theorem. (Aldous (1997)) As $n \to \infty$,

$$(n^{-1/3} X_\lambda^n(\lfloor n^{2/3} t \rfloor), t \geq 0) \xrightarrow{d} (W_\lambda(t), t \geq 0),$$

uniformly on compact time-intervals.

Sketch of proof

Fortunately, X^n is a (time-inhomogeneous) Markov process. We need to understand its step distribution.

At time i,

- v_i is the current vertex;
- i vertices are dead;
- $X^n(i)$ vertices are alive;
- we want to know $k(v_i)$, the number of children of v_i.

We have not yet looked at the possible edges from v_i to any of the other $n - i - X^n(i)$ unexplored vertices in the graph. Each of these is present with probability $\frac{1}{n} + \frac{\lambda}{n^{4/3}}$ independently. So, given $X^n(i)$,

$$k(v_i) \sim \text{Bin}\left(n - i - X^n(i), \frac{1}{n} + \frac{\lambda}{n^{4/3}}\right).$$

As long as $X^n(i) = o(n)$ and $i = O(n^{2/3})$,

$$(n - i - X^n(i)) \left(\frac{1}{n} + \frac{\lambda}{n^{4/3}} \right) \approx 1 + \frac{\lambda}{n^{1/3}} - \frac{i}{n} + o(n^{-1/3}),$$

and so we approximately have

$$X^n(i+1) - X^n(i) \sim \text{Poisson} \left(1 + \frac{\lambda}{n^{1/3}} - \frac{i}{n} \right) - 1.$$

So X^n is close to being a random walk with a deterministic (but time-dependent) drift. Let

$$M^n(i) = X^n(i) - \sum_{j=0}^{i-1} \left(\frac{\lambda}{n^{1/3}} - \frac{i}{n} \right) \approx X^n(i) - \frac{\lambda i}{n^{1/3}} + \frac{i^2}{2n}.$$

$$X^n(i+1) - X^n(i) \sim \text{Poisson}\left(1 + \frac{\lambda}{n^{1/3}} - \frac{i}{n}\right) - 1$$

and so if

$$M^n(i) \approx X^n(i) - \frac{\lambda i}{n^{1/3}} + \frac{i^2}{2n}$$

then $(M^n(i), i \geq 0)$ is approximately a martingale.

Plug in $i = \lfloor tn^{2/3} \rfloor$:

$$n^{-1/3} M^n(\lfloor tn^{2/3} \rfloor) \approx n^{-1/3} X^n(\lfloor tn^{2/3} \rfloor) - \lambda t + \frac{t^2}{2}.$$

Since the Poisson distribution here has variance ≈ 1 for all i, we can apply the martingale functional CLT (a more general version of Donsker's theorem) to obtain

$$\left(n^{-1/3} X^n(\lfloor tn^{2/3} \rfloor) - \lambda t + \frac{t^2}{2}, t \geq 0\right) \xrightarrow{d} (W(t), t \geq 0).$$

Question

So we now understand the limiting sizes and surpluses of components of the critical random graph.

Question

So we now understand the limiting sizes and surpluses of components of the critical random graph.

But what do the limiting components look like?

Question

So we now understand the limiting sizes and surpluses of components of the critical random graph.

But what do the limiting components look like?

They are no longer (in general) trees. Again, the vertex-labels are irrelevant: we are really interested in what shapes and distances look like in the limit. So we will give a metric space answer, and convergence will be in the Gromov-Hausdorff distance.

Our approach

Consider the components one by one.

Simple but important fact: a component of $G(n,p)$ conditioned to have m vertices and s surplus edges is a uniform connected graph on those m vertices with $m+s-1$ edges.

Our approach

Consider the components one by one.

Simple but important fact: a component of $G(n,p)$ conditioned to have m vertices and s surplus edges is a uniform connected graph on those m vertices with $m+s-1$ edges.

General idea: to pick out a (well-chosen) spanning tree, and then understand where to put in the surplus edges.

Depth-first tree

In the depth-first exploration, we effectively explored this spanning tree; the dashed edges made no difference.

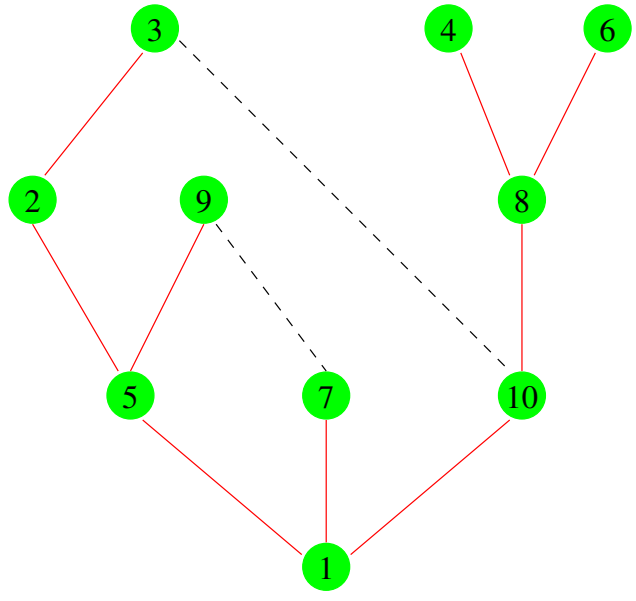

Call it the depth-first tree associated with the graph G, and write $T(G)$.

The tree case

There is one case which we already understand: when the surplus of a component is 0. Then the component is a uniform random tree (and is necessarily the same as its depth-first tree). In this case, it is clear that the scaling limit is the Brownian CRT.

The tree case

There is one case which we already understand: when the surplus of a component is 0. Then the component is a uniform random tree (and is necessarily the same as its depth-first tree). In this case, it is clear that the scaling limit is the Brownian CRT.

In the tree case, we should rescale distances by $1/\sqrt{m}$, where m is the number of vertices in the component. This is the correct distance rescaling for all of the big components in the random graph.

The tree case

There is one case which we already understand: when the surplus of a component is 0. Then the component is a uniform random tree (and is necessarily the same as its depth-first tree). In this case, it is clear that the scaling limit is the Brownian CRT.

In the tree case, we should rescale distances by $1/\sqrt{m}$, where m is the number of vertices in the component. This is the correct distance rescaling for all of the big components in the random graph. Since the big components have sizes of order $n^{2/3}$, we should rescale distances by $n^{-1/3}$.

The limit of the random graph: spanning tree

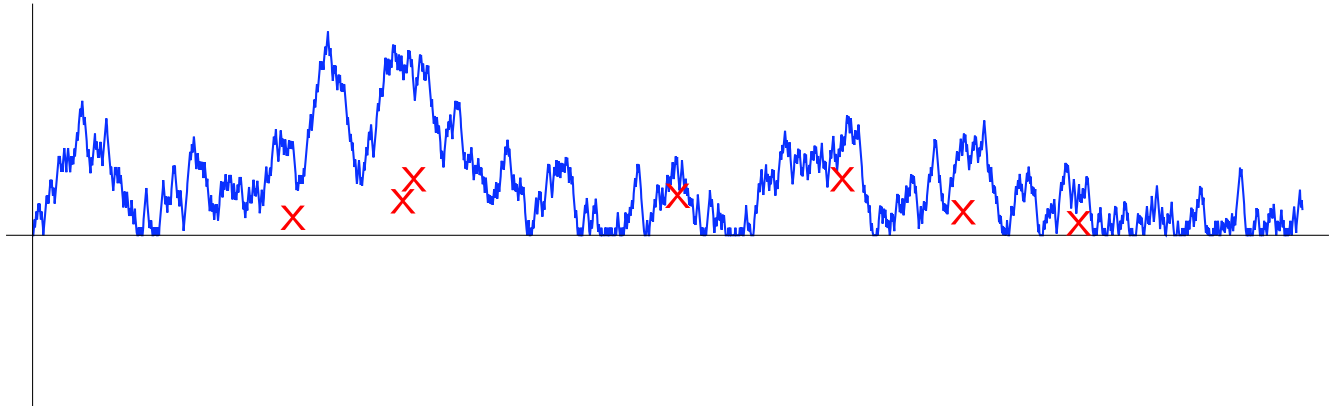

It turns out that each excursion of the process $(B^{\lambda}(t), t \geq 0)$ encodes a continuum random tree, which is a "spanning tree" for a limit component.

The limit of the random graph: spanning tree

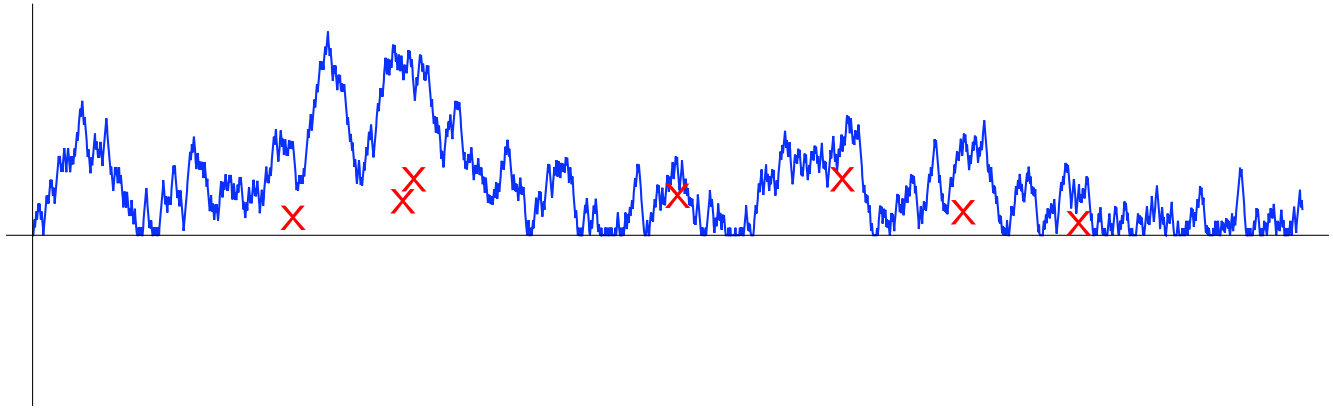

It turns out that each excursion of the process $(B^{\lambda}(t), t \geq 0)$ encodes a continuum random tree, which is a "spanning tree" for a limit component. These are *not* Brownian CRT's, but CRT's whose distribution has been "tilted":

The limit of the random graph: spanning tree

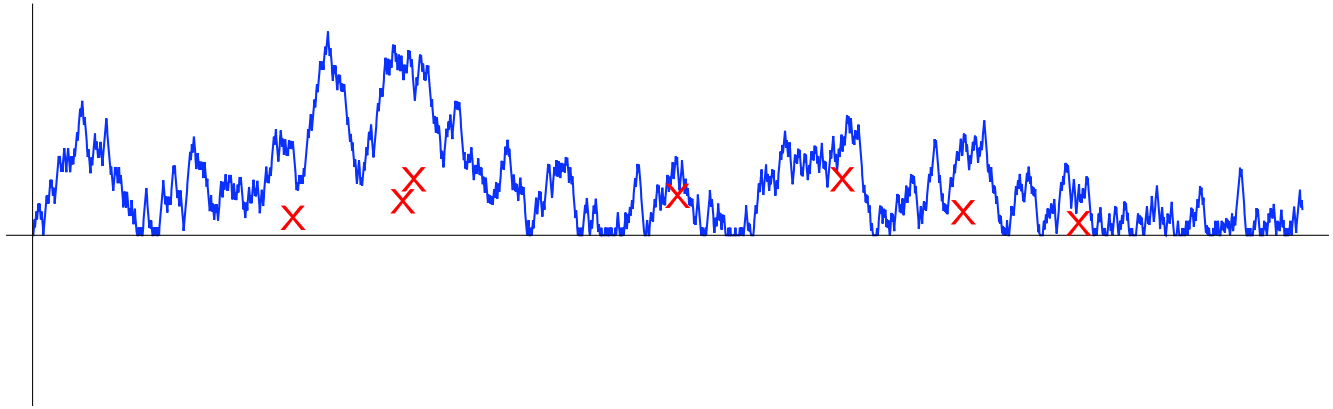

It turns out that each excursion of the process $(B^{\lambda}(t), t \geq 0)$ encodes a continuum random tree, which is a "spanning tree" for a limit component. These are *not* Brownian CRT's, but CRT's whose distribution has been "tilted": the corresponding excursion $\tilde{e}$ (normalized to have length 1) has a distribution specified by

$$\mathbb{E}\left[f\left(\tilde{e}\right)\right] = \frac{\mathbb{E}\left[f\left(e\right)\exp\left(\int_0^1 e(u)du\right)\right]}{\mathbb{E}\left[\exp\left(\int_0^1 e(u)du\right)\right]},$$

where f is any suitable test-function and e is a standard Brownian excursion.

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

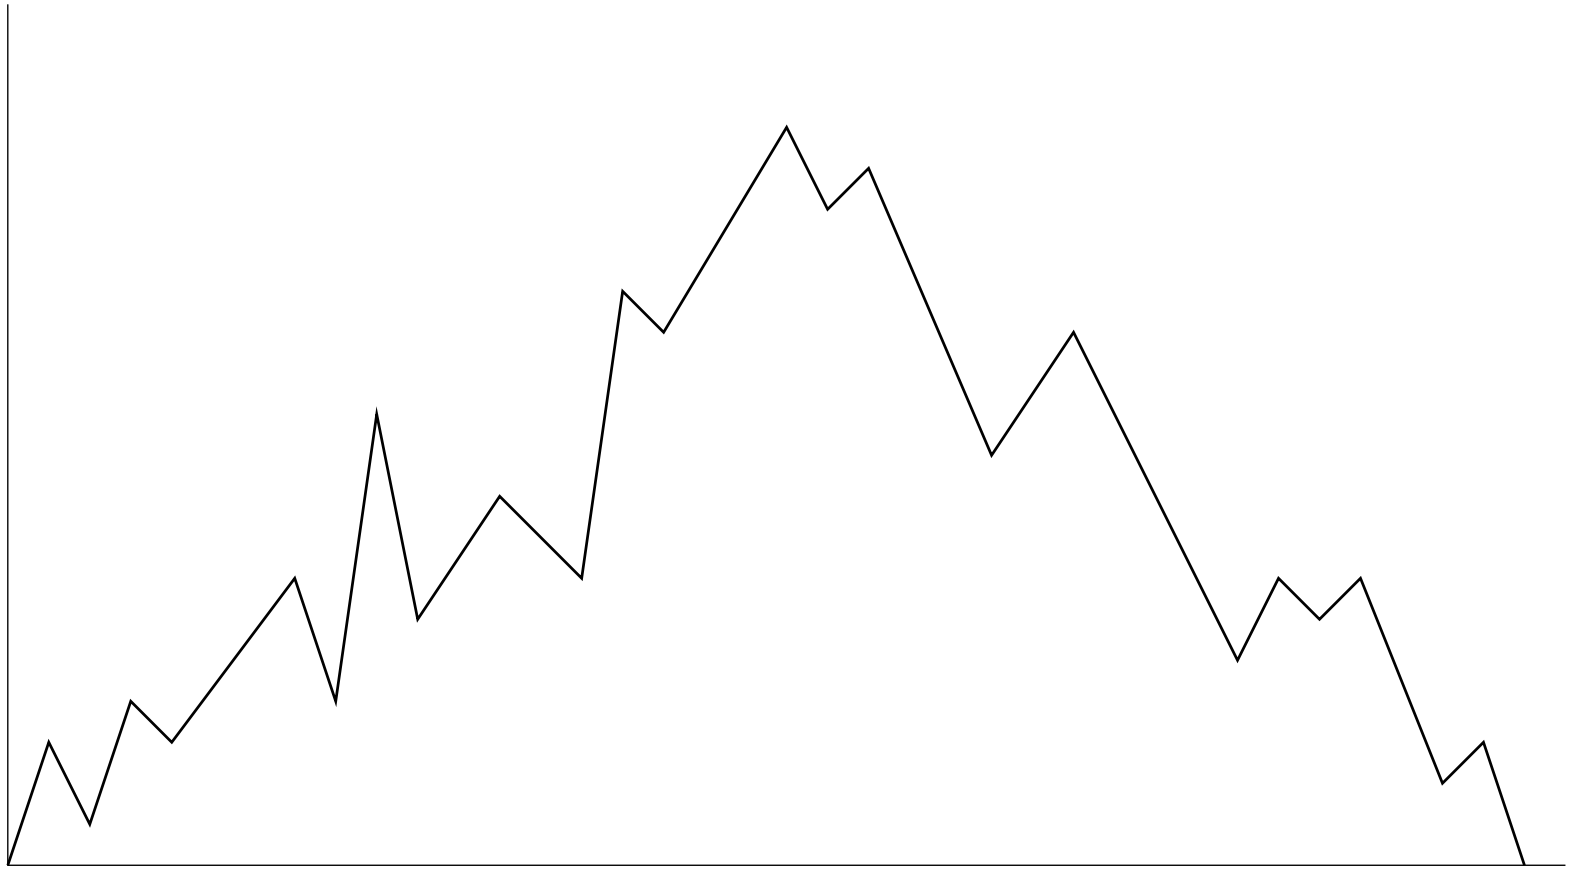

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

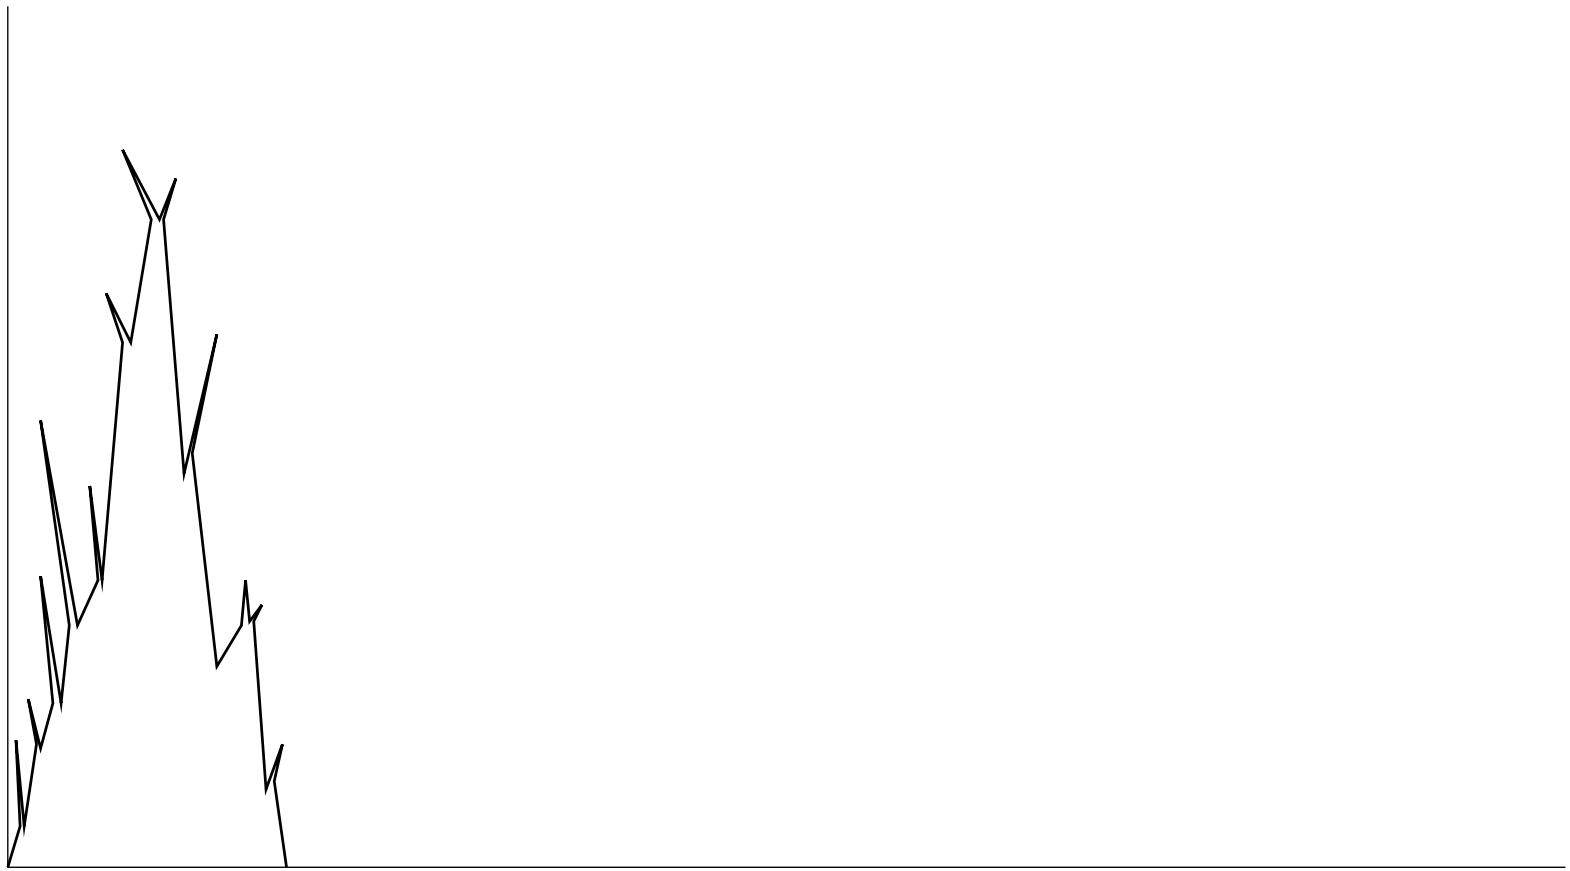

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

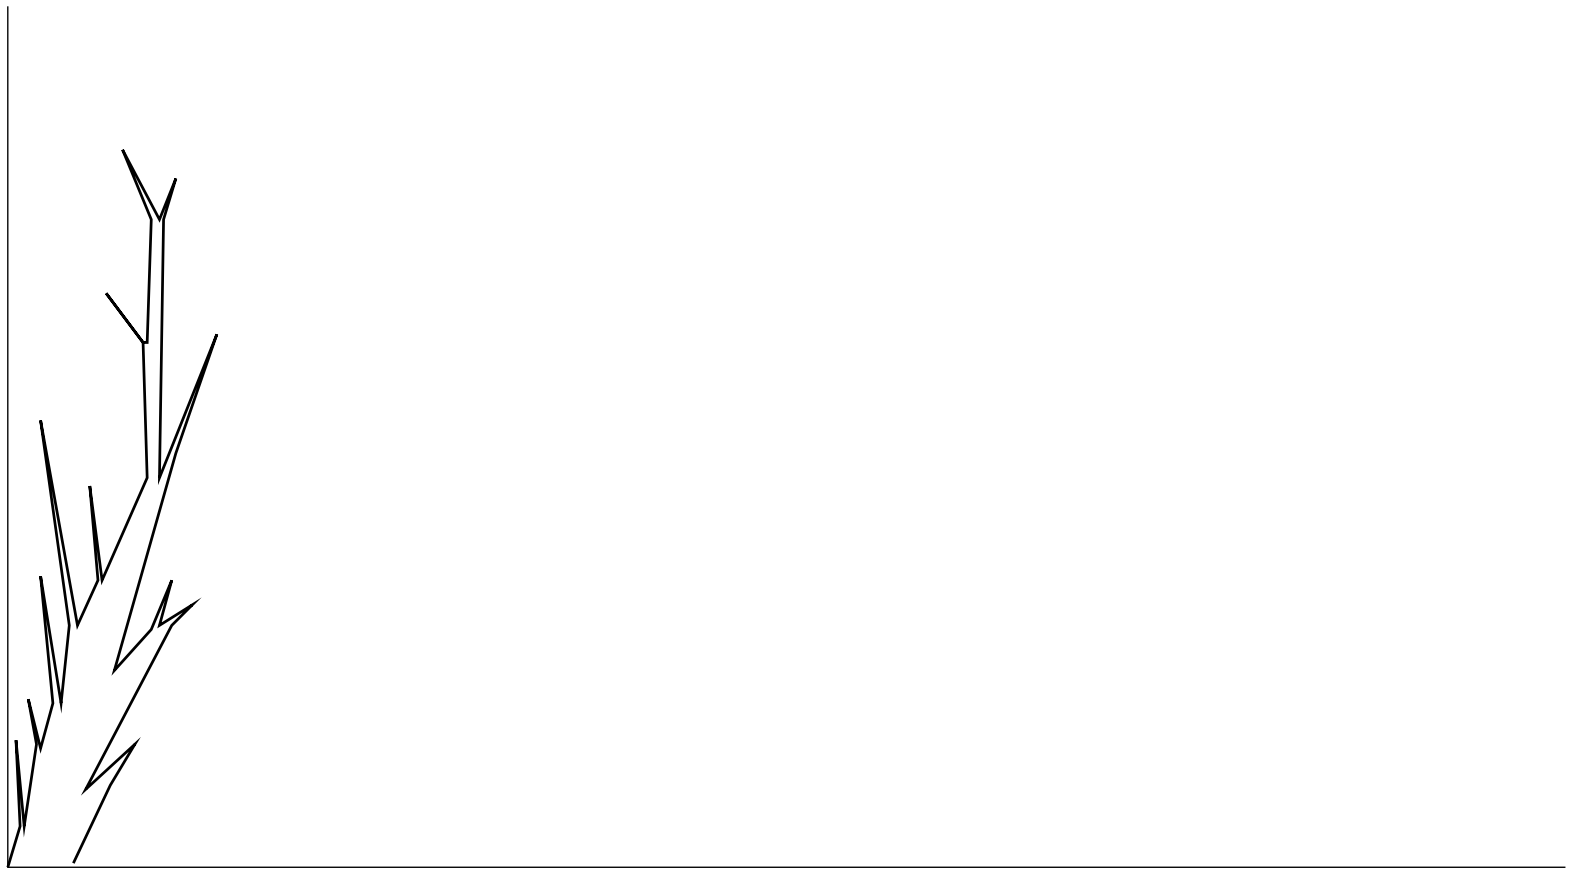

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

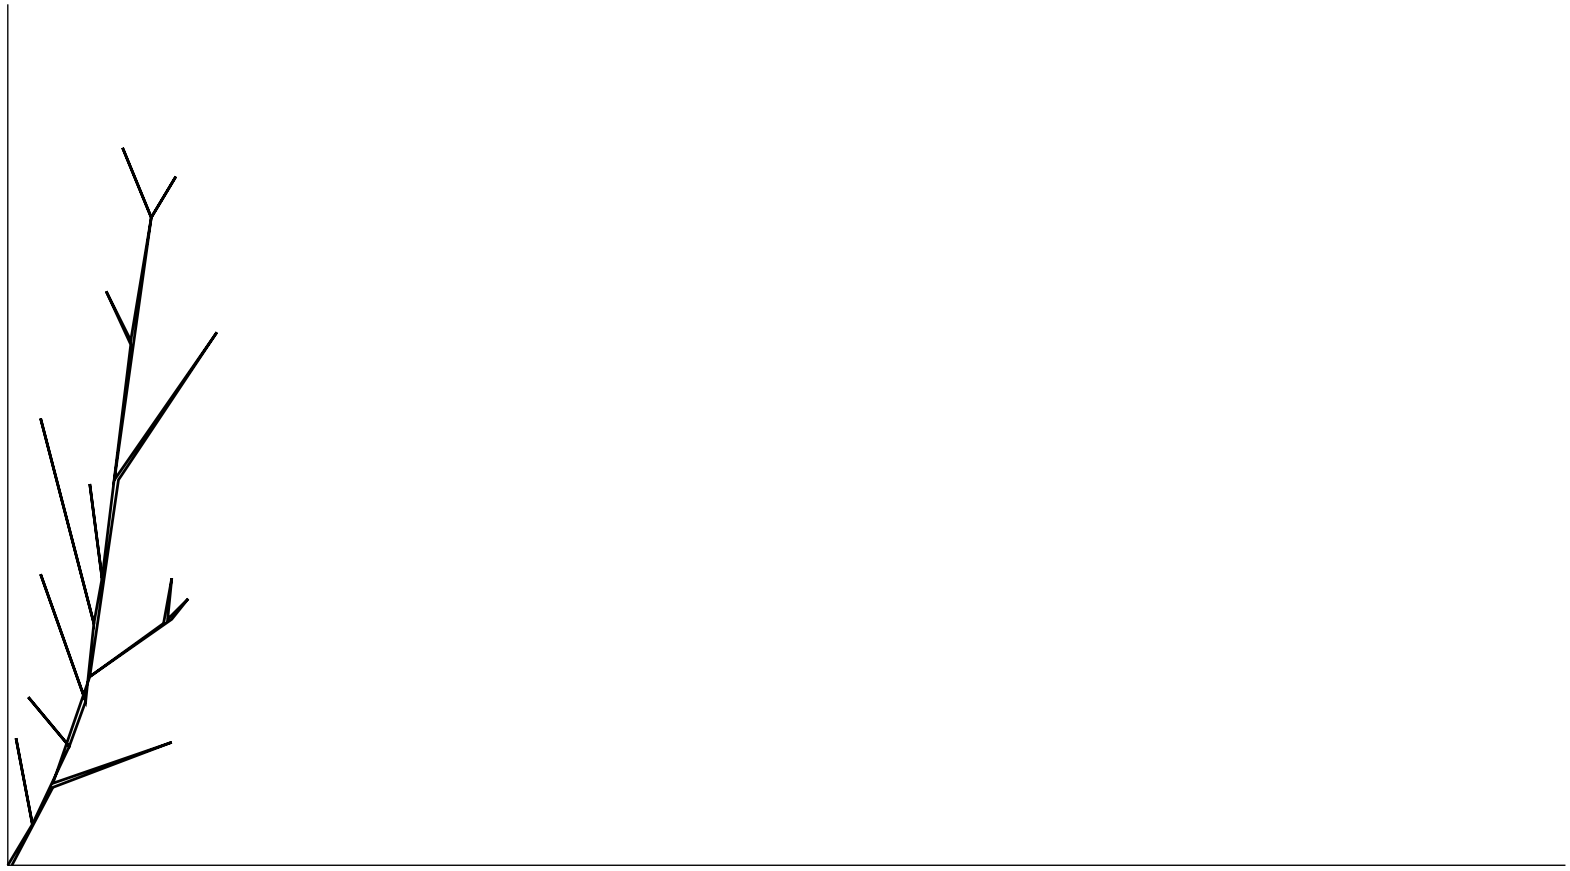

The limit of the random graph: spanning tree

Write $\tilde{\mathcal{T}}$ for the tree encoded by $2\tilde{e}$:

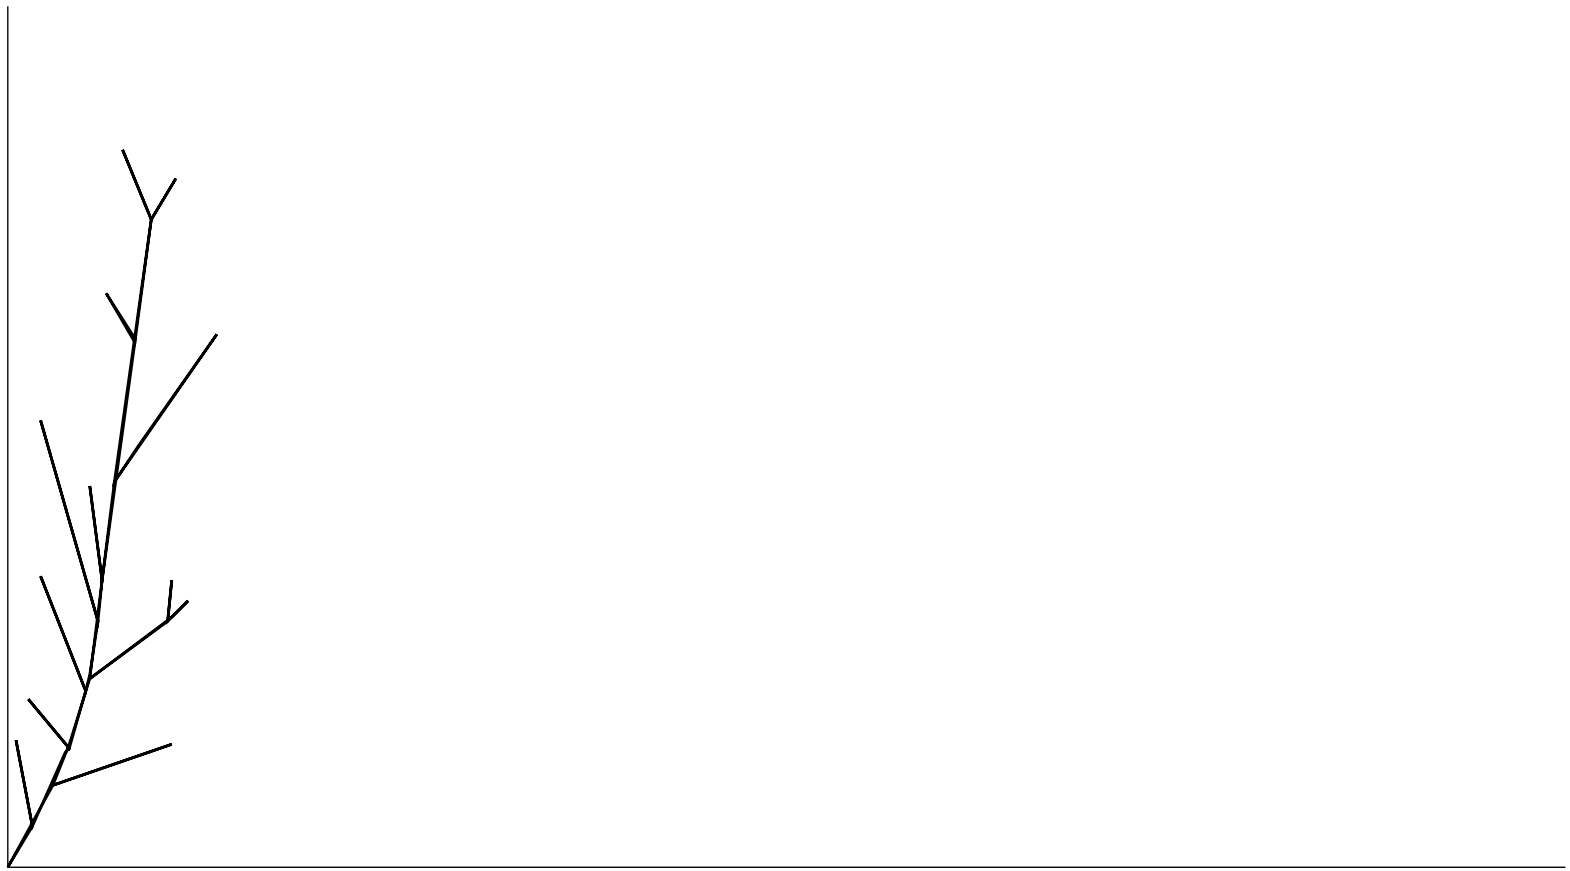

The limit of the random graph: vertex identifications

In the limit, surplus edges correspond to vertex-identifications (since edge-lengths have shrunk to 0). The points of the Poisson process tell us where these vertex-identifications should occur.

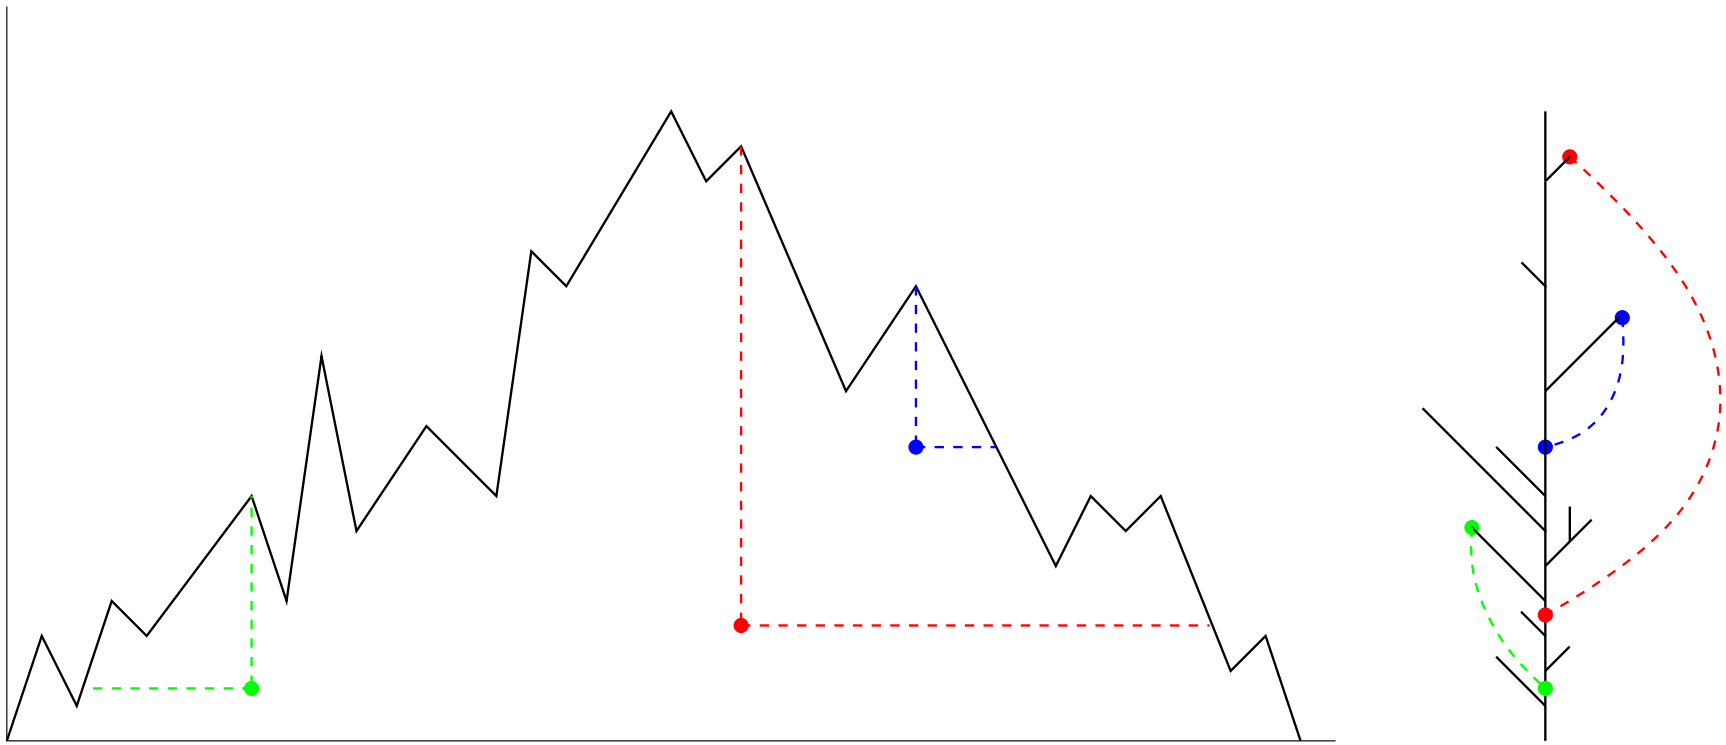

A point at (x, y) under $2\tilde{e}$ identifies the vertex v at height $2\tilde{e}(x)$ with the vertex at distance y along the path from the root to v in the tree $\tilde{\mathcal{T}}$.

A limiting component

What about components which don't have size exactly $n^{2/3}$ but, say, $xn^{2/3}$?

A limiting component

What about components which don't have size exactly $n^{2/3}$ but, say, $xn^{2/3}$?

Such a component, in the limit, is encoded by a tilted excursion of length x, $\tilde{e}^{(x)}$, which has a Brownian scaling relationship with the standard tilted excursion:

$$\mathbb{E}\left[f\left(\tilde{e}^{(x)}\right)\right] = \mathbb{E}\left[f\left(\sqrt{x}\tilde{e}(t/x), 0 \leq t \leq x\right)\right].$$

A limiting component

What about components which don't have size exactly $n^{2/3}$ but, say, $xn^{2/3}$?

Such a component, in the limit, is encoded by a tilted excursion of length x, $\tilde{e}^{(x)}$, which has a Brownian scaling relationship with the standard tilted excursion:

$$\mathbb{E}\left[f\left(\tilde{e}^{(x)}\right)\right] = \mathbb{E}\left[f\left(\sqrt{x}\tilde{e}(t/x), 0 \leq t \leq x\right)\right].$$

Note that it follows from properties of the tilted trees and of the Poisson process that we may equivalently describe the limit of a component on $\sim xn^{2/3}$ vertices as follows.

A limiting component

Sample a tilted excursion $\tilde{e}^{(x)}$ of length x and use it to create a CRT $\tilde{\mathcal{T}}$.

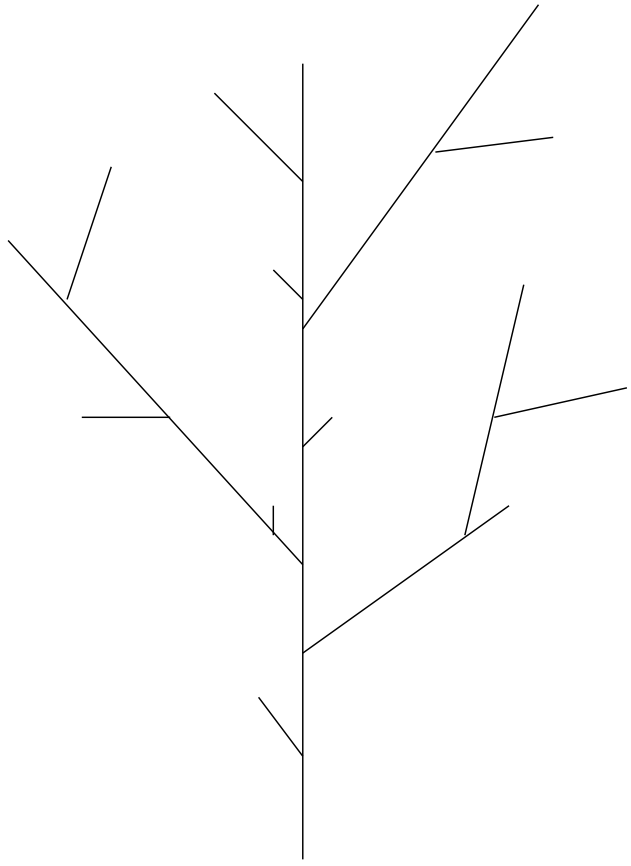

A limiting component

Sample a tilted excursion $\tilde{e}^{(x)}$ of length x and use it to create a CRT $\tilde{\mathcal{T}}$.

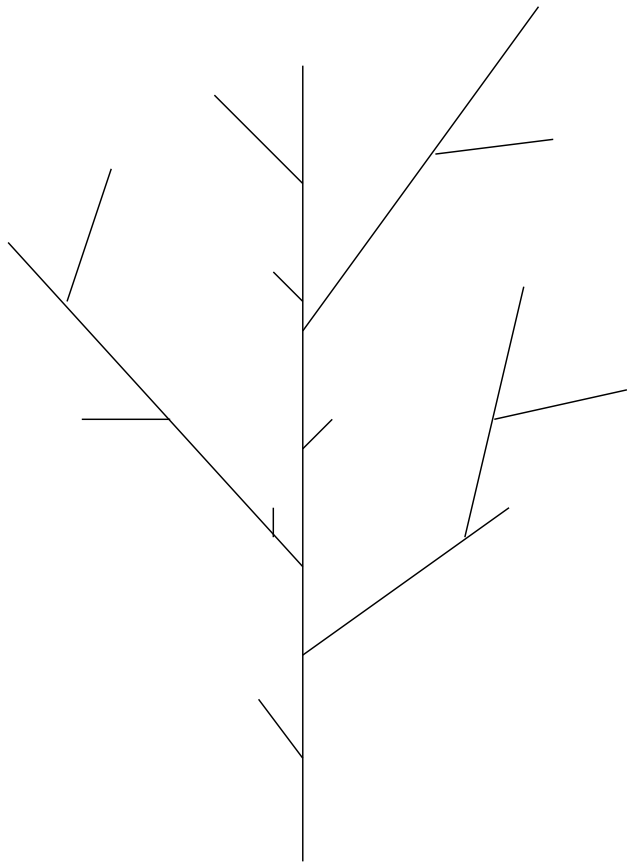

Conditional on $\tilde{e}^{(x)}$, sample a random variable P with Poisson $\left(\int_0^x \tilde{e}^{(x)}(u)du\right)$ distribution.

A limiting component

Conditional on $P = s$, pick s vertices of the tree $\tilde{\mathcal{T}}$ independently with density proportional to their height. (These will almost surely be leaves.)

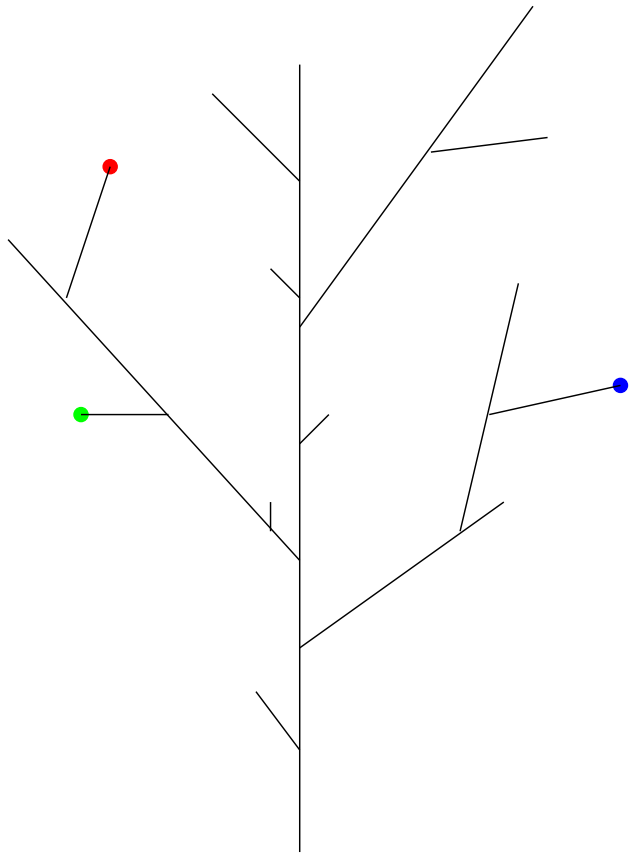

A limiting component

For each of the selected leaves, pick a uniform point on the path from the leaf to the root.

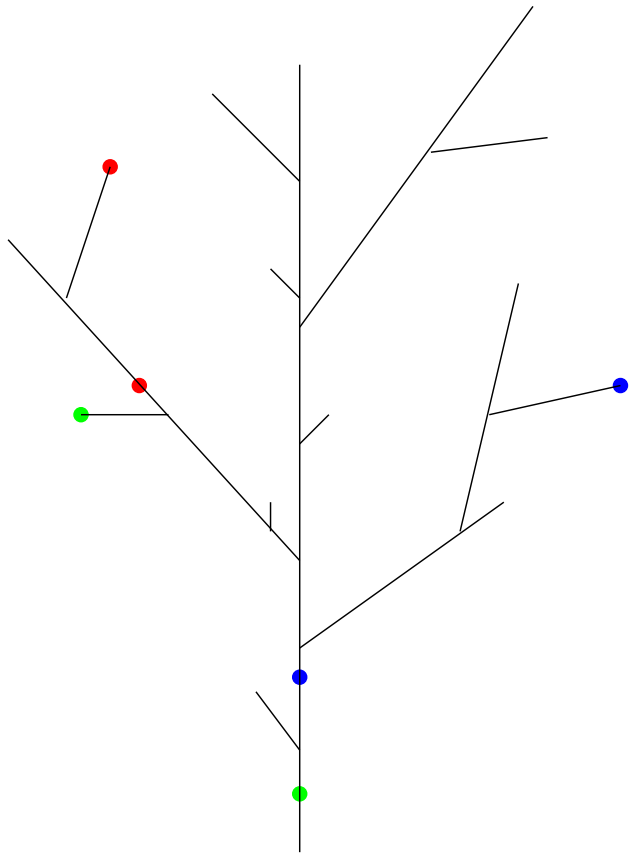

A limiting component

Identify each of the selected leaves with its chosen point.

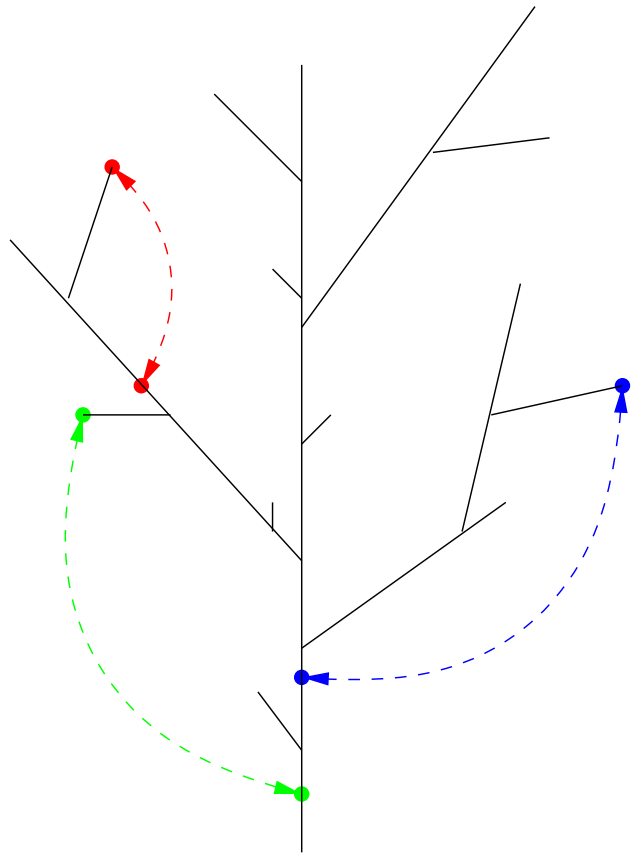

Convergence result

Let $\mathcal{C}_1^n, \mathcal{C}_2^n, \ldots$ be the sequence of components of $G\left(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}}\right)$ in decreasing order of size, considered as metric spaces with the graph distance.

Theorem. As $n \to \infty$,

$$n^{-1/3}(\mathcal{C}_1^n, \mathcal{C}_2^n, \ldots) \xrightarrow{d} (\mathcal{C}_1, \mathcal{C}_2, \ldots),$$

where $\mathcal{C}_1, \mathcal{C}_2, \ldots$ is the sequence of metric spaces corresponding to the excursions of Aldous' marked limit process B_λ in decreasing order of length.

Convergence is with respect to the distance

$$\mathrm{d}(\mathcal{A}, \mathcal{B}) := \left(\sum_{i=1}^{\infty} \mathrm{d_{GH}}(\mathcal{A}_i, \mathcal{B}_i)^4\right)^{1/4}.$$

Idea of proof

The key idea turns out to be study a component of $G(n, p)$ conditioned on its size (but *not* on its surplus).

Depth-first tree

Take an arbitrary component G of $G(n,p)$ of size m. Recall that $T(G)$ is the depth-first tree associated with G

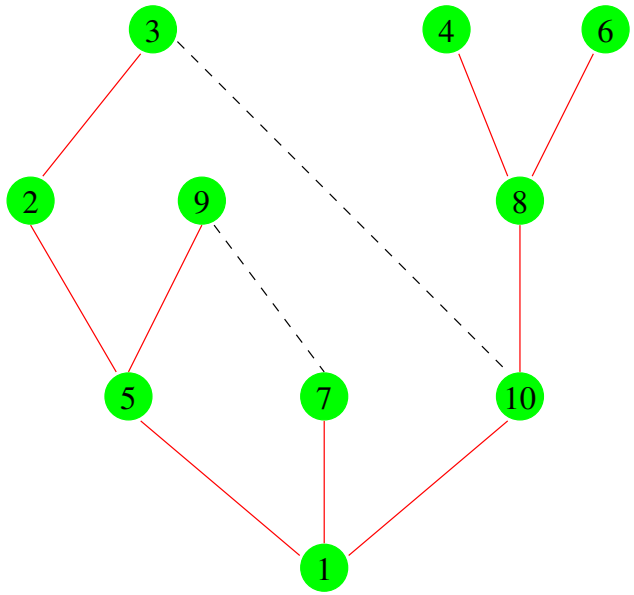

and that $(X(k), 0 \leq k \leq m)$ is the depth-first walk of $T(G)$.

Permitted edges

Look at things the other way round: for a given tree T, which connected graphs G have depth-first tree $T(G) = T$?

In other words, where can we put surplus edges so that they don't change T?

Call such edges permitted.

Depth-first walk and permitted edges

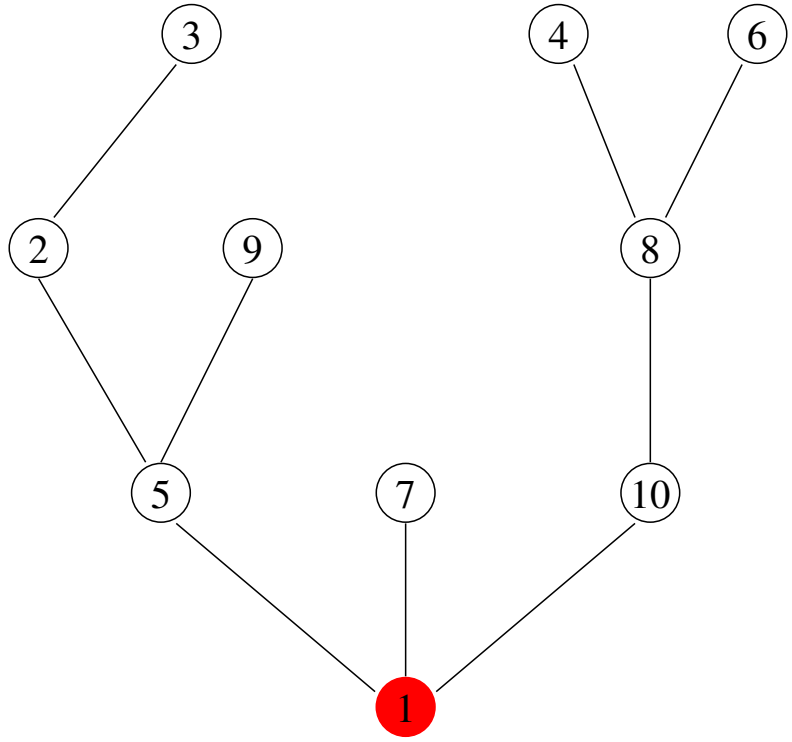

Step 0: $X(0) = 0$.

Depth-first walk and permitted edges

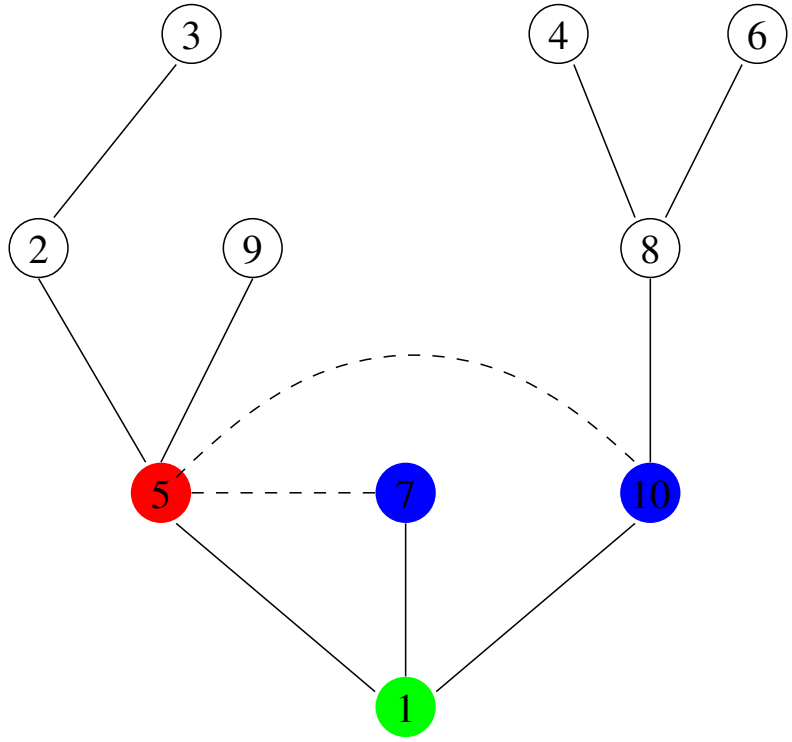

Step 1: $X(1) = 2$.

Depth-first walk and permitted edges

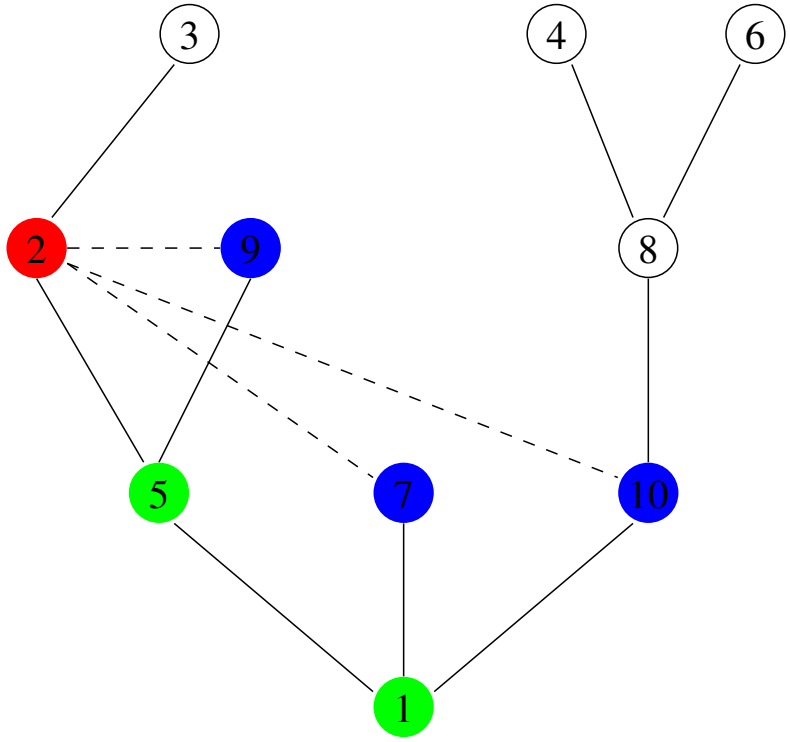

Step 2: $X(2) = 3$.

Depth-first walk and permitted edges

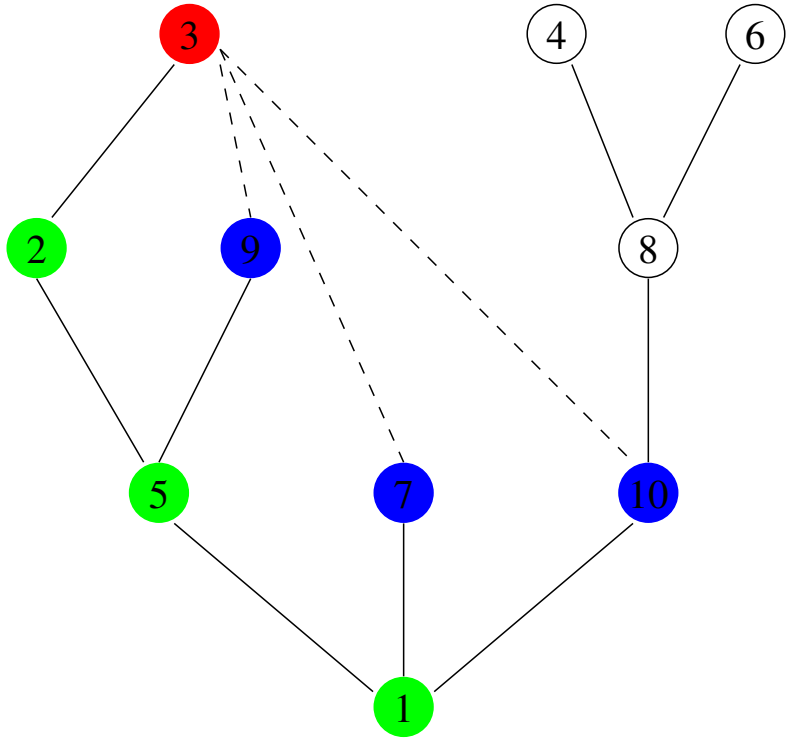

Step 3: $X(3) = 3$.

Depth-first walk and permitted edges

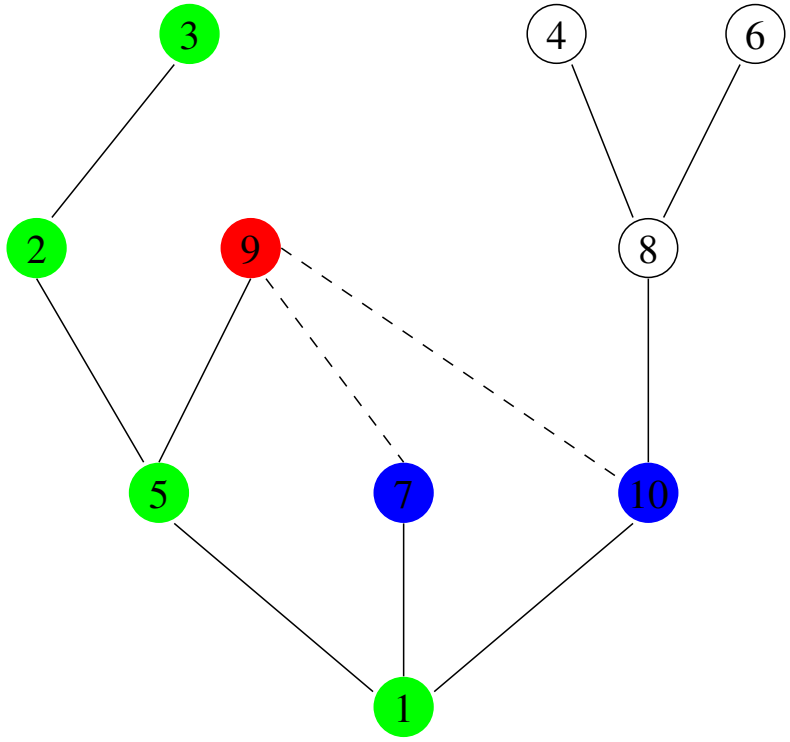

Step 4: $X(4) = 2$.

Depth-first walk and permitted edges

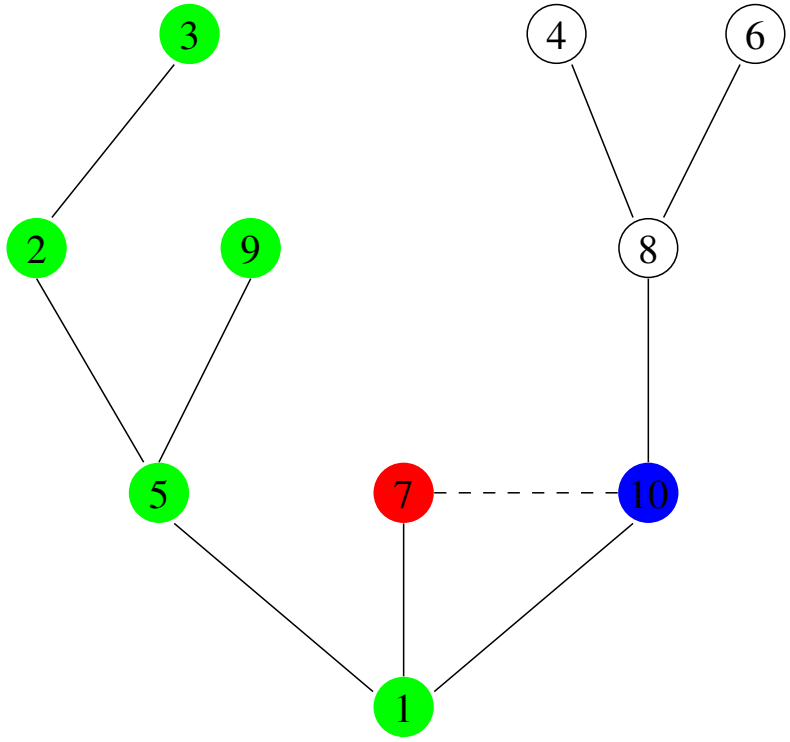

Step 5: $X(5) = 1$.

Depth-first walk and permitted edges

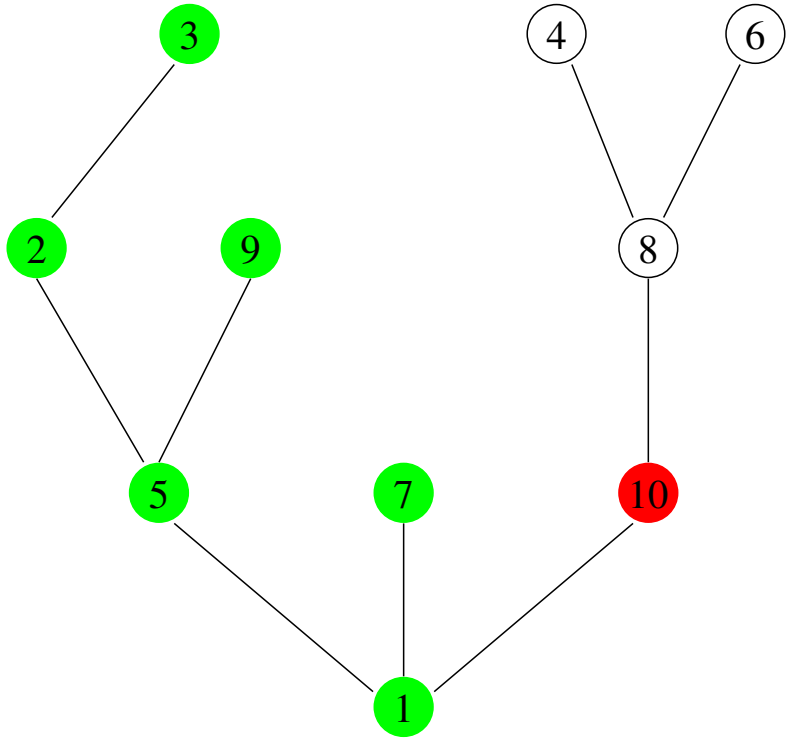

Step 6: $X(6) = 0$.

Depth-first walk and permitted edges

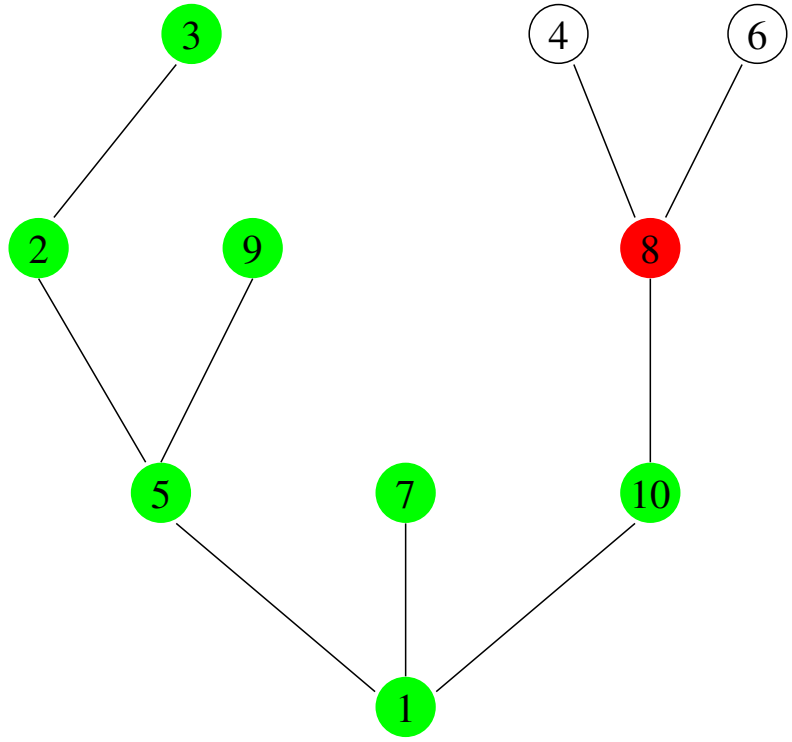

Step 7: $X(7) = 0$.

Depth-first walk and permitted edges

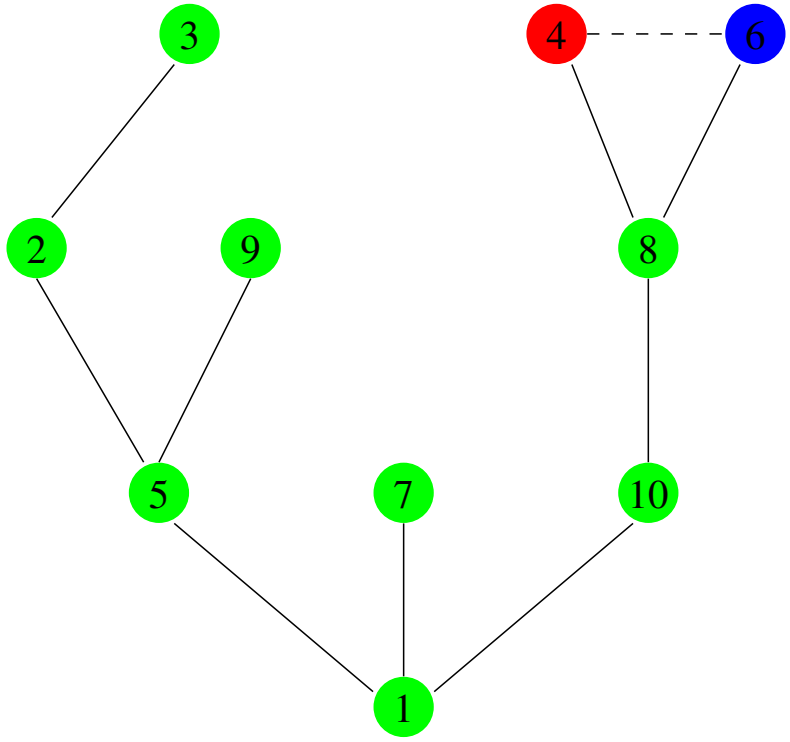

Step 8: $X(8) = 1$.

Depth-first walk and permitted edges

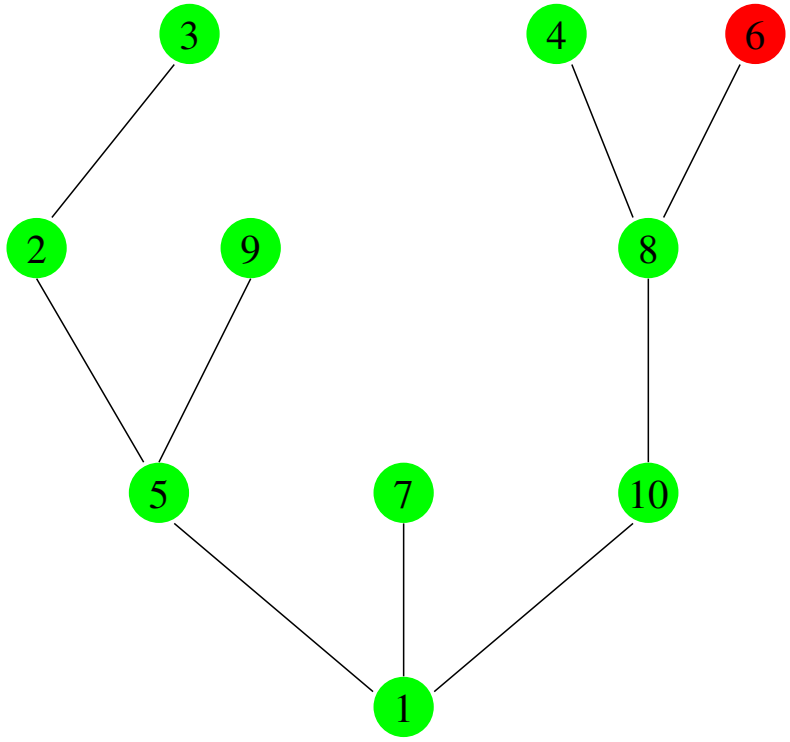

Step 10: $X(9) = 0$.

Area

At step $k \geq 0$ there are $X(k)$ permitted edges. So the total number is

$$a(T) = \sum_{k=0}^{m-1} X(k).$$

We call this the area of T.

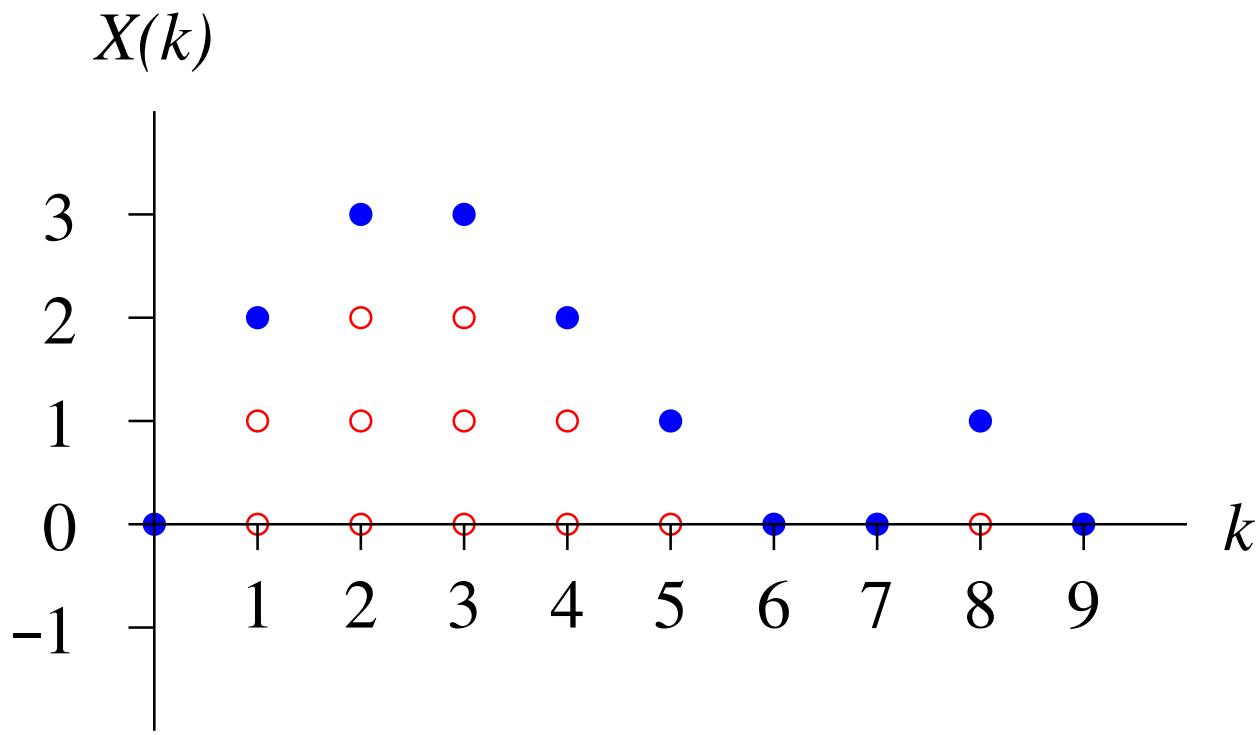

Classifying graphs by depth-first tree

Let $\mathbb{G}_T$ be the set of graphs G such that $T(G) = T$. It follows that $|\mathbb{G}_T| = 2^{a(T)}$, since each permitted edge may either be included or not.

Classifying graphs by depth-first tree

Let $\mathbb{G}_T$ be the set of graphs G such that $T(G) = T$. It follows that $|\mathbb{G}_T| = 2^{a(T)}$, since each permitted edge may either be included or not.

Recall that $\mathbb{T}_m$ is the set of trees with label-set $[m] = \{1, 2, \ldots, m\}$. Then

$$\{\mathbb{G}_T : T \in \mathbb{T}_m\}$$

is a partition of the set of connected graphs on $[m]$.

Recipe for creating a connected graph on $[m]$

Create a connected graph $\tilde{G}_m^p$ as follows.

Recipe for creating a connected graph on $[m]$

Create a connected graph $\tilde{G}_m^p$ as follows.

- Pick a random labelled tree $\tilde{T}_m^p$ such that

$$\mathbb{P}\left(\tilde{T}_m^p = T\right) \propto (1-p)^{-a(T)}, \quad T \in \mathbb{T}_{[m]}.$$

Recipe for creating a connected graph on $[m]$

Create a connected graph $\tilde{G}_m^p$ as follows.

- Pick a random labelled tree $\tilde{T}_m^p$ such that

$$\mathbb{P}\left(\tilde{T}_m^p = T\right) \propto (1-p)^{-a(T)}, \quad T \in \mathbb{T}_{[m]}.$$

- Add each of the $a(\tilde{T}_m^p)$ permitted edges to $\tilde{T}_m^p$ independently with probability p.

Recipe for creating a connected graph on $[m]$

Lemma. $\tilde{G}_m^p$ has the same distribution as G_m^p, a component of $G(n,p)$ conditioned to have vertex-set $[m]$.

Recipe for creating a connected graph on $[m]$

Lemma. $\tilde{G}_m^p$ has the same distribution as G_m^p, a component of $G(n,p)$ conditioned to have vertex-set $[m]$.

Proof. For a connected graph G on $[m]$ which has $T(G) = T$ and surplus s,

$$\mathbb{P}\left(\tilde{G}_m^p = G\right) \propto (1-p)^{-a(T)} p^s (1-p)^{a(T)-s} = (p/(1-p))^s.$$

Recipe for creating a connected graph on $[m]$

Lemma. $\tilde{G}_m^p$ has the same distribution as G_m^p, a component of $G(n,p)$ conditioned to have vertex-set $[m]$.

Proof. For a connected graph G on $[m]$ which has $T(G) = T$ and surplus s,

$$\mathbb{P}\left(\tilde{G}_m^p = G\right) \propto (1-p)^{-a(T)} p^s (1-p)^{a(T)-s} = (p/(1-p))^s.$$

Likewise, by the definition of $G(n,p)$,

$$\begin{aligned}\mathbb{P}\left(G_m^p = G\right) &\propto \mathbb{P}\left(G(m,p) = G\right) = p^{m+s-1}(1-p)^{\binom{m}{2}-m-s+1} \\ &\propto (p/(1-p))^s.\end{aligned}$$

Taking limits

When $m \sim xn^{2/3}$ and $p = \frac{1}{n} + \frac{\lambda}{n^{4/3}}$, we need to show that

- the tree $\tilde{T}_m^p$ converges to a CRT coded by a tilted excursion;
- the locations of the surplus edges converge to the locations in our limiting picture.

Taking limits

When $m \sim xn^{2/3}$ and $p = \frac{1}{n} + \frac{\lambda}{n^{4/3}}$, we need to show that

- the tree $\tilde{T}_m^p$ converges to a CRT coded by a tilted excursion;
- the locations of the surplus edges converge to the locations in our limiting picture.

For simplicity, take $p = m^{-3/2}$ (i.e. $x = 1$, $\lambda = 0$); the general case is similar.

Taking limits for the tree

Write $\tilde{X}^m$ for the depth-first walk associated with $\tilde{T}_m^p$. Then

$$a\left(\tilde{T}_m^p\right) = \int_0^m \tilde{X}^m(s)ds.$$

Taking limits for the tree

Write $\tilde{X}^m$ for the depth-first walk associated with $\tilde{T}_m^p$. Then

$$a\left(\tilde{T}_m^p\right) = \int_0^m \tilde{X}^m(s)ds.$$

If T_m is a uniform random tree on $[m]$ and X^m is its depth-first walk, then

$$(m^{-1/2}X^m(mt), 0 \leq t \leq 1) \xrightarrow{d} (e(t), 0 \leq t \leq 1).$$

Taking limits for the tree

Use the change of measure to get from $\tilde{X}^m$ to X^m: for any bounded continuous function f,

$$\mathbb{E}\left[f\left(m^{-1/2}\tilde{X}^m(mt), 0 \le t \le 1\right)\right]$$

$$= \frac{\mathbb{E}\left[f\left(m^{-1/2}X^m(mt), 0 \le t \le 1\right)(1-p)^{-\int_0^m X^m(u)du}\right]}{\mathbb{E}\left[(1-p)^{-\int_0^m X^m(u)du}\right]}$$

Taking limits for the tree

Use the change of measure to get from $\tilde{X}^m$ to X^m: for any bounded continuous function f,

$$\mathbb{E}\left[f\left(m^{-1/2}\tilde{X}^m(mt), 0 \le t \le 1\right)\right]$$
$$= \frac{\mathbb{E}\left[f\left(m^{-1/2}X^m(mt), 0 \le t \le 1\right)(1-p)^{-\int_0^m X^m(u)du}\right]}{\mathbb{E}\left[(1-p)^{-\int_0^m X^m(u)du}\right]}$$

Note that by a change of variable in the integral,

$$(1-p)^{-\int_0^m X^m(u)du} = (1-p)^{-m^{3/2}\int_0^1 m^{-1/2}X^m(ms)ds}.$$

Taking limits for the tree

Use the change of measure to get from $\tilde{X}^m$ to X^m: for any bounded continuous function f,

$$\mathbb{E}\left[f\left(m^{-1/2}\tilde{X}^m(mt), 0 \le t \le 1\right)\right]$$
$$= \frac{\mathbb{E}\left[f\left(m^{-1/2}X^m(mt), 0 \le t \le 1\right)(1-p)^{-\int_0^m X^m(u)du}\right]}{\mathbb{E}\left[(1-p)^{-\int_0^m X^m(u)du}\right]}$$

Note that by a change of variable in the integral,

$$(1-p)^{-\int_0^m X^m(u)du} = (1-p)^{-m^{3/2}\int_0^1 m^{-1/2}X^m(ms)ds}.$$

Since $p = m^{-3/2}$,

$$(1-p)^{-m^{3/2}\int_0^1 m^{-1/2}X^m(ms)ds} \xrightarrow{d} \exp\left(\int_0^1 e(u)du\right).$$

Taking limits for the tree

Taking care with the limits, we obtain

$$\mathbb{E}\left[f\left(m^{-1/2}\tilde{X}^m(mt), 0 \le t \le 1\right)\right] \to \frac{\mathbb{E}\left[f(e)\exp\left(\int_0^1 e(u)du\right)\right]}{\mathbb{E}\left[\exp\left(\int_0^1 e(u)du\right)\right]} = \mathbb{E}\left[f(\tilde{e})\right].$$

Taking limits for the tree

Taking care with the limits, we obtain

$$\mathbb{E}\left[f\left(m^{-1/2}\tilde{X}^m(mt), 0 \le t \le 1\right)\right] \to \frac{\mathbb{E}\left[f(e)\exp\left(\int_0^1 e(u)du\right)\right]}{\mathbb{E}\left[\exp\left(\int_0^1 e(u)du\right)\right]} = \mathbb{E}\left[f(\tilde{e})\right].$$

This (after a bit more work) entails that

$$\frac{1}{\sqrt{m}}\tilde{T}_m^p \xrightarrow{d} \tilde{\mathcal{T}}.$$

Taking limits for the surplus edges

The permitted edges are in bijective correspondence with the integer points under the graph of the depth-first walk.

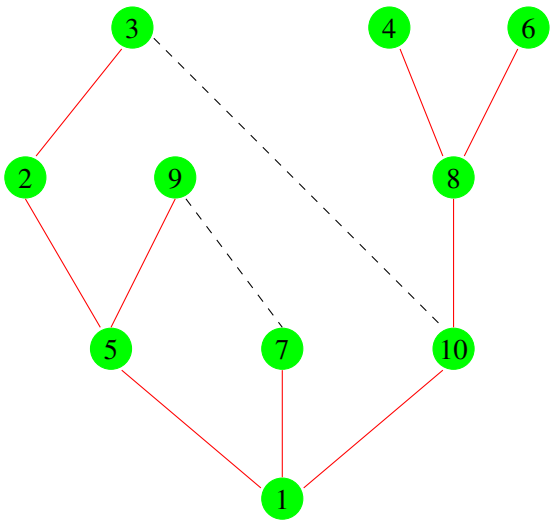

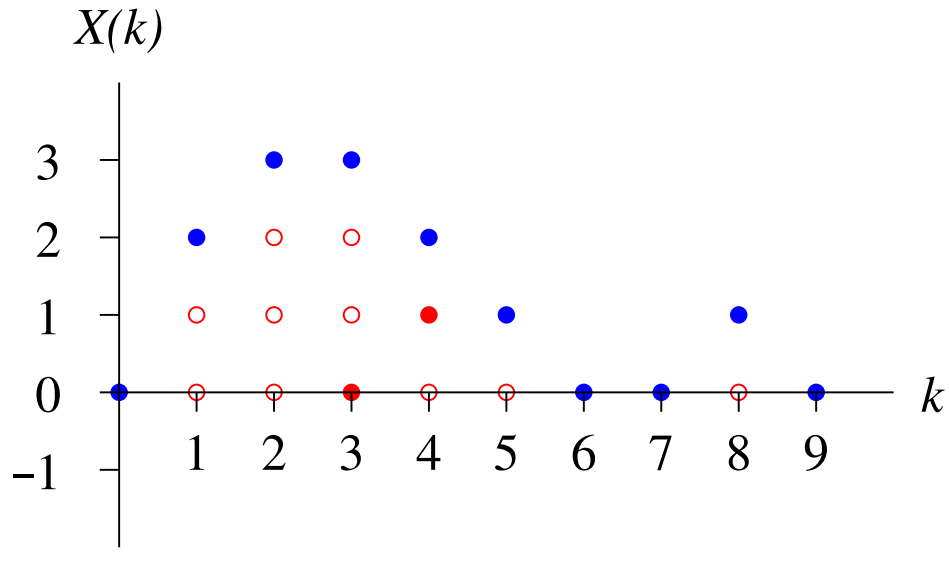

Taking limits for the surplus edges

The permitted edges are in bijective correspondence with the integer points under the graph of the depth-first walk. Since each permitted edge is included independently with probability p, the surplus edges form a Binomial point process.

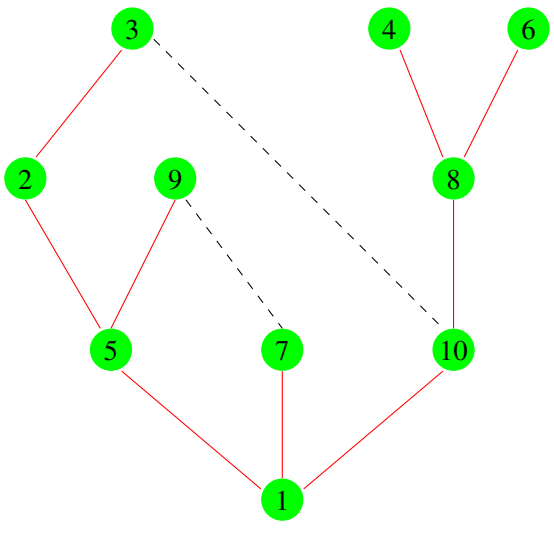

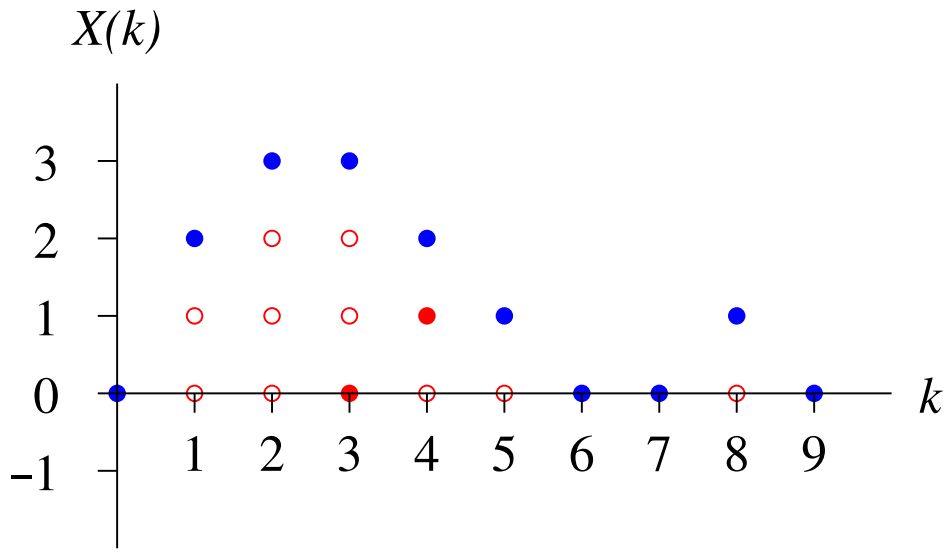

Taking limits for the surplus edges

Surplus edges almost go to ancestors... In fact, they always go to younger children of ancestors of the current vertex.

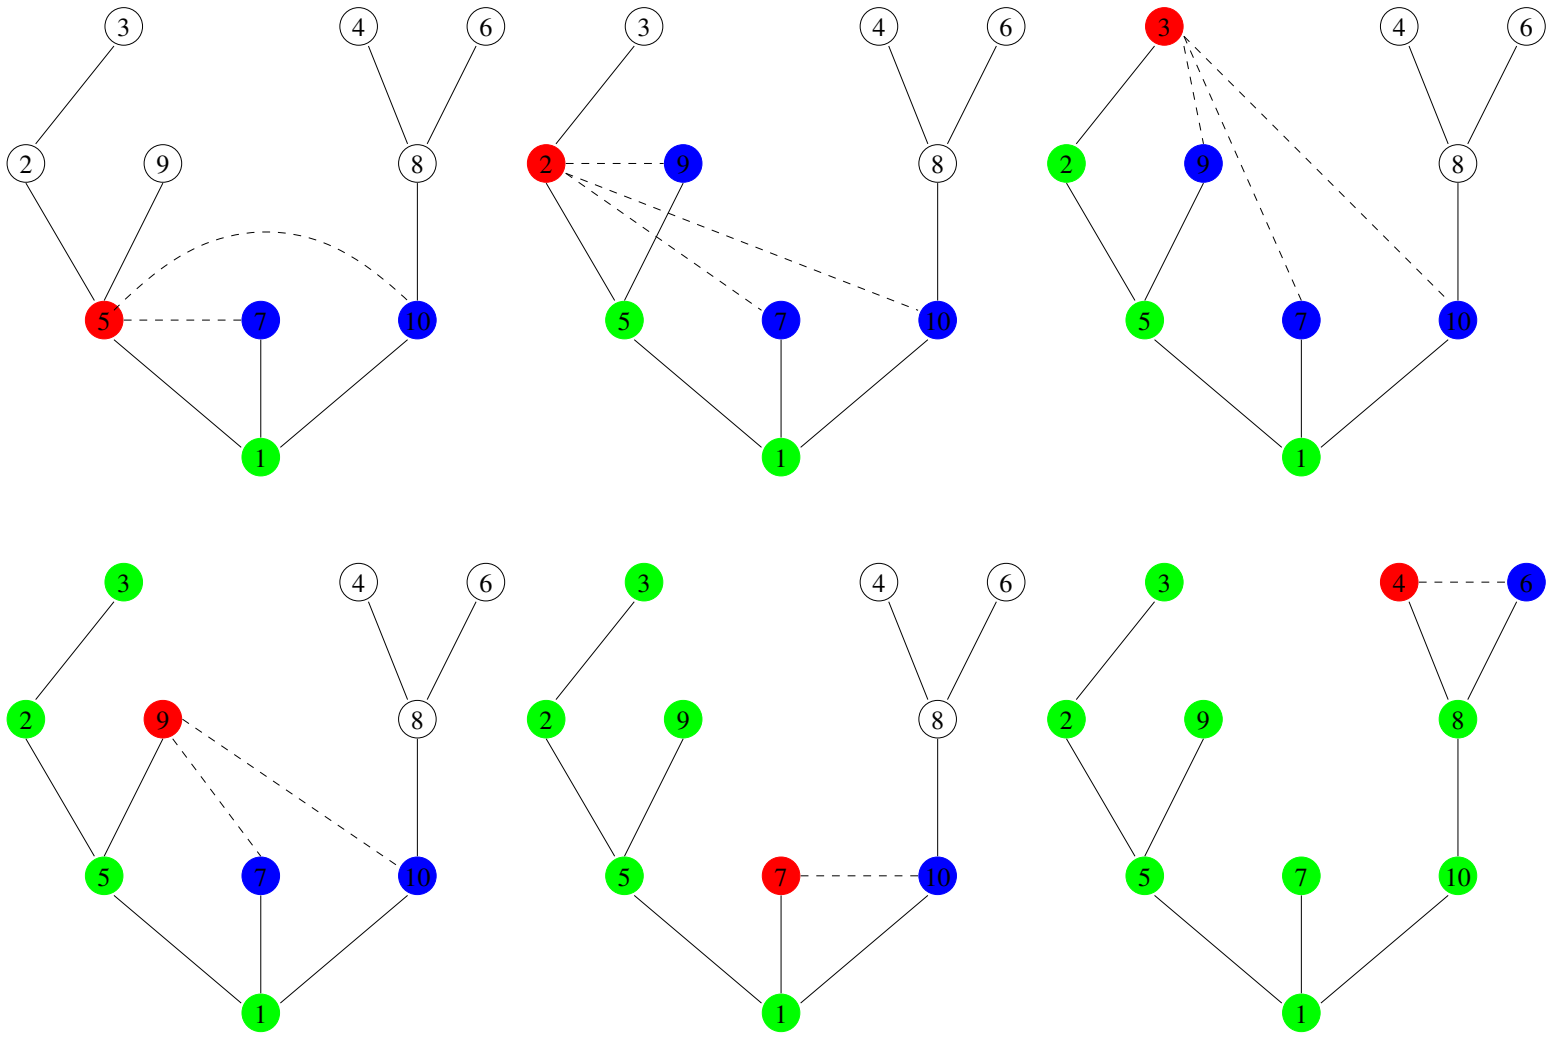

Taking limits for the surplus edges

When we rescale, the distance between a vertex and one of its children vanishes and so, in the limit, surplus "edges" do go to ancestors of the current vertex (i.e. vertices on the path down to the root).

Taking limits for the surplus edges

When we rescale, the distance between a vertex and one of its children vanishes and so, in the limit, surplus "edges" do go to ancestors of the current vertex (i.e. vertices on the path down to the root).

The Binomial point process of surplus edges, when rescaled, straightforwardly converges to the required Poisson point process.

Taking limits for the surplus edges

When we rescale, the distance between a vertex and one of its children vanishes and so, in the limit, surplus "edges" do go to ancestors of the current vertex (i.e. vertices on the path down to the root).

The Binomial point process of surplus edges, when rescaled, straightforwardly converges to the required Poisson point process. (This gives another proof of Aldous' result on the limiting number of surplus edges.)

Taking care over the details, this completes the proof.

Cycle structure of a limit component

A limiting component can have quite a complicated cycle structure:

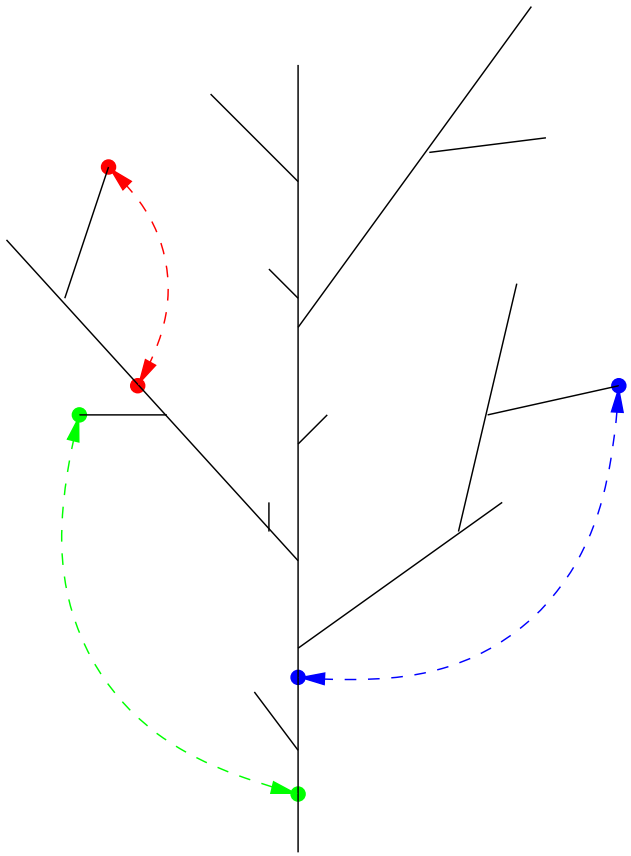

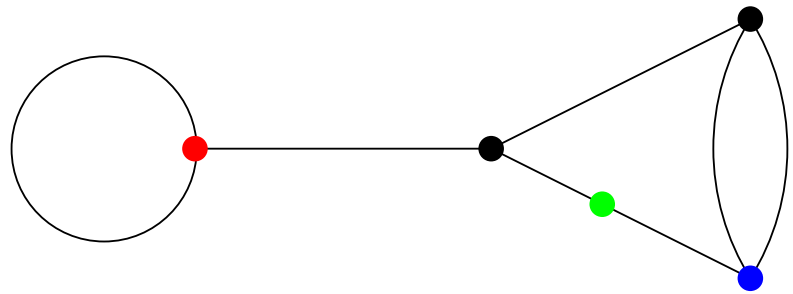

Cycle structure of a limit component

A limiting component can have quite a complicated cycle structure:

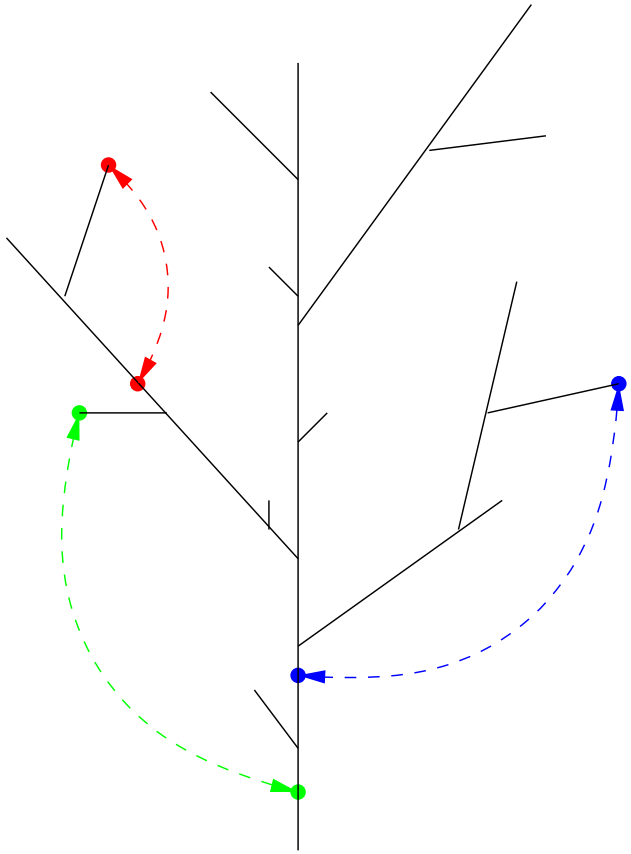

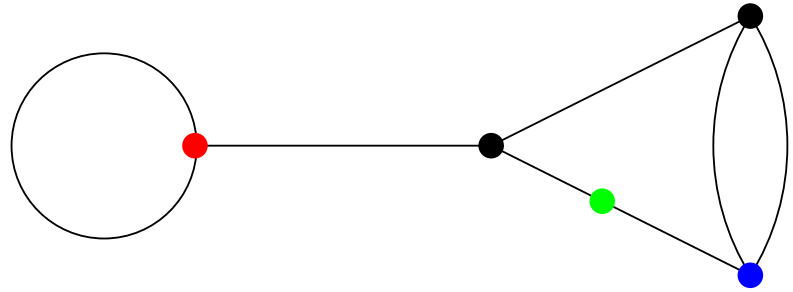

What more can we say about it?

Cycle structure of a graph

Fix a connected graph G.

Cycle structure of a graph

Fix a connected graph G.

The core $C(G)$ consists of the edges in cycles and those joining the cycles.

Cycle structure of a graph

Fix a connected graph G.

The core $C(G)$ consists of the edges in cycles and those joining the cycles.

If G is a tree, $C(G)$ is empty.

Cycle structure: an example

Graph G

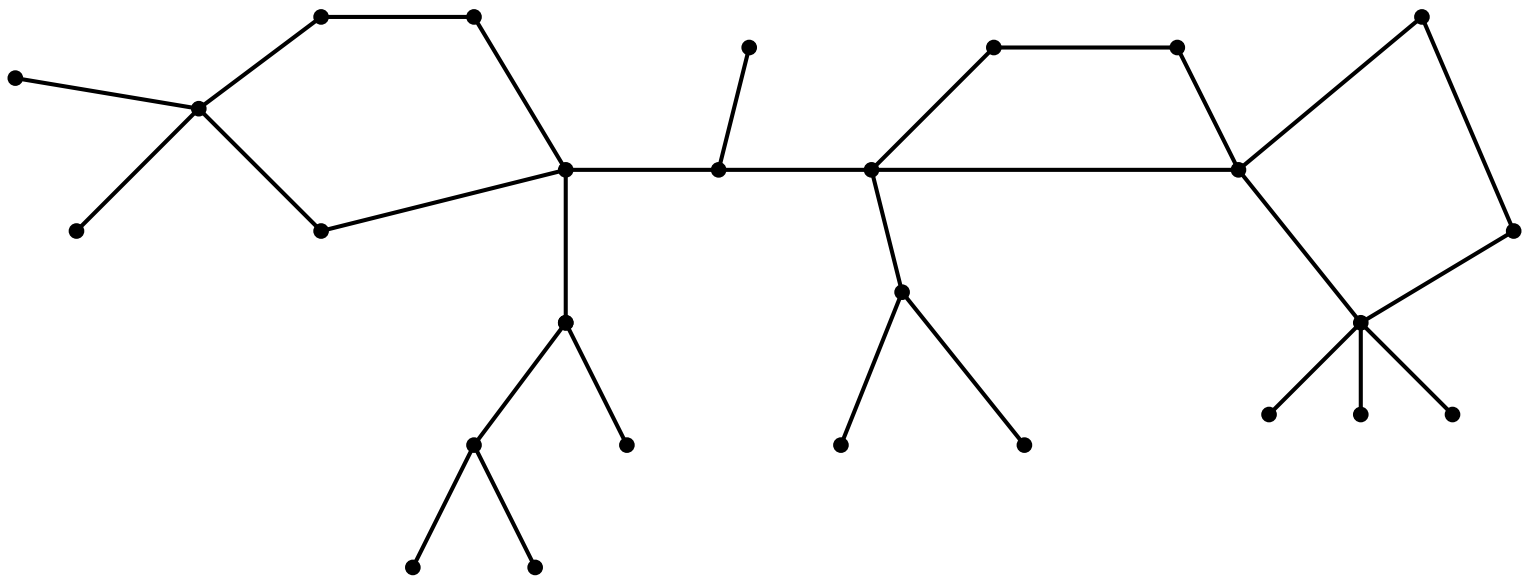

Cycle structure: an example

Core $C(G)$

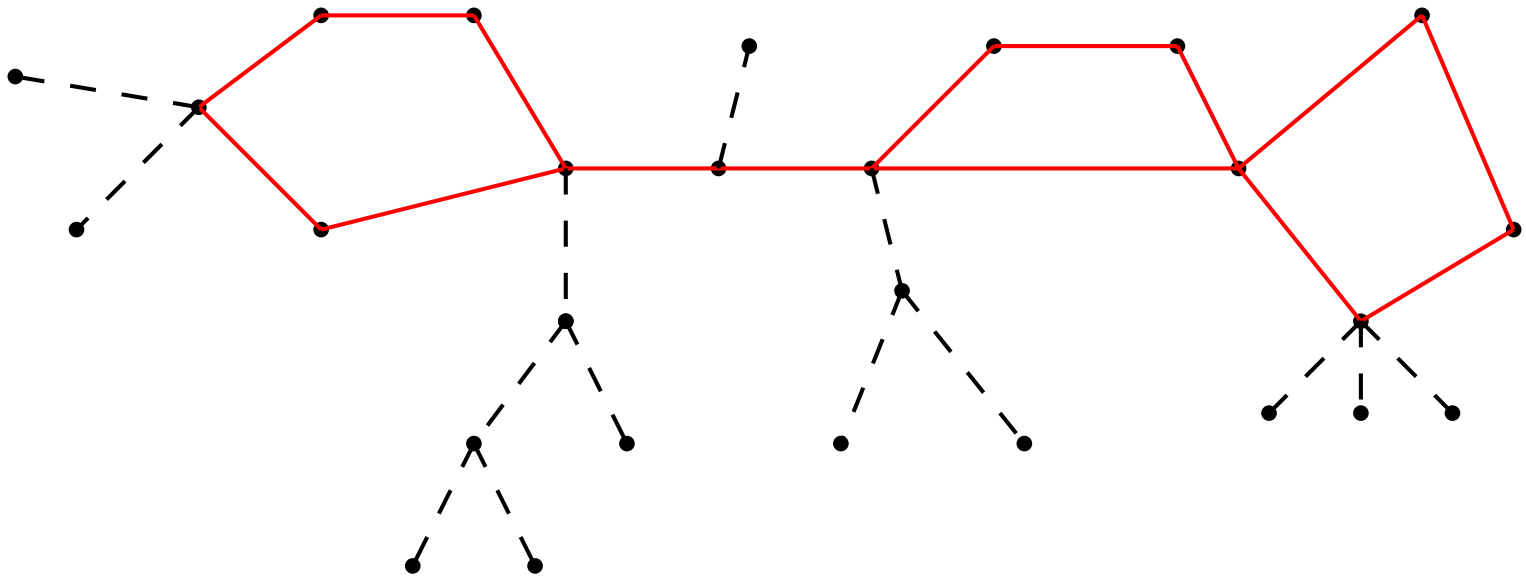

Cycle structure of a graph

The kernel $K(G)$ is the multigraph which gives the "shape of the core":

Cycle structure of a graph

The kernel $K(G)$ is the multigraph which gives the "shape of the core": take the vertices of the core of degree 3 or more; contract the paths between them to a single edge.

Cycle structure of a graph

The kernel $K(G)$ is the multigraph which gives the "shape of the core": take the vertices of the core of degree 3 or more; contract the paths between them to a single edge.

By convention, the kernel of a tree or unicyclic component is empty.

Cycle structure: an example

Vertices of degree at least 3 in the core

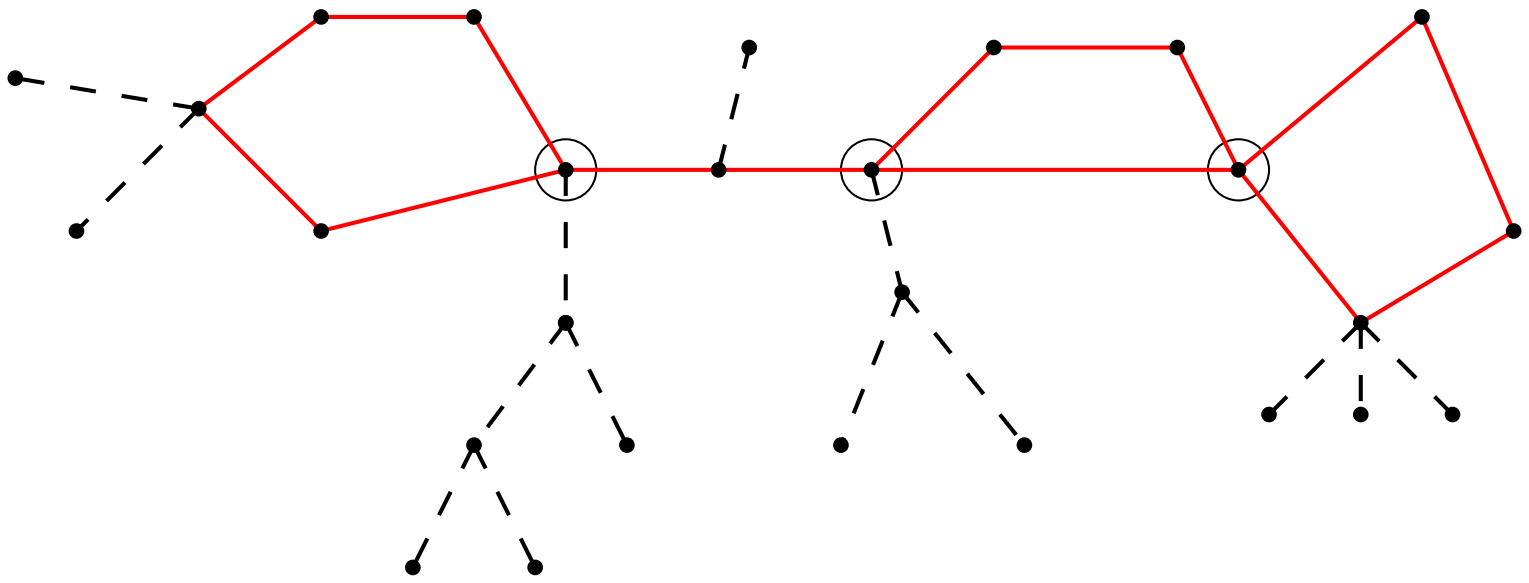

Cycle structure: an example

Contract paths between them

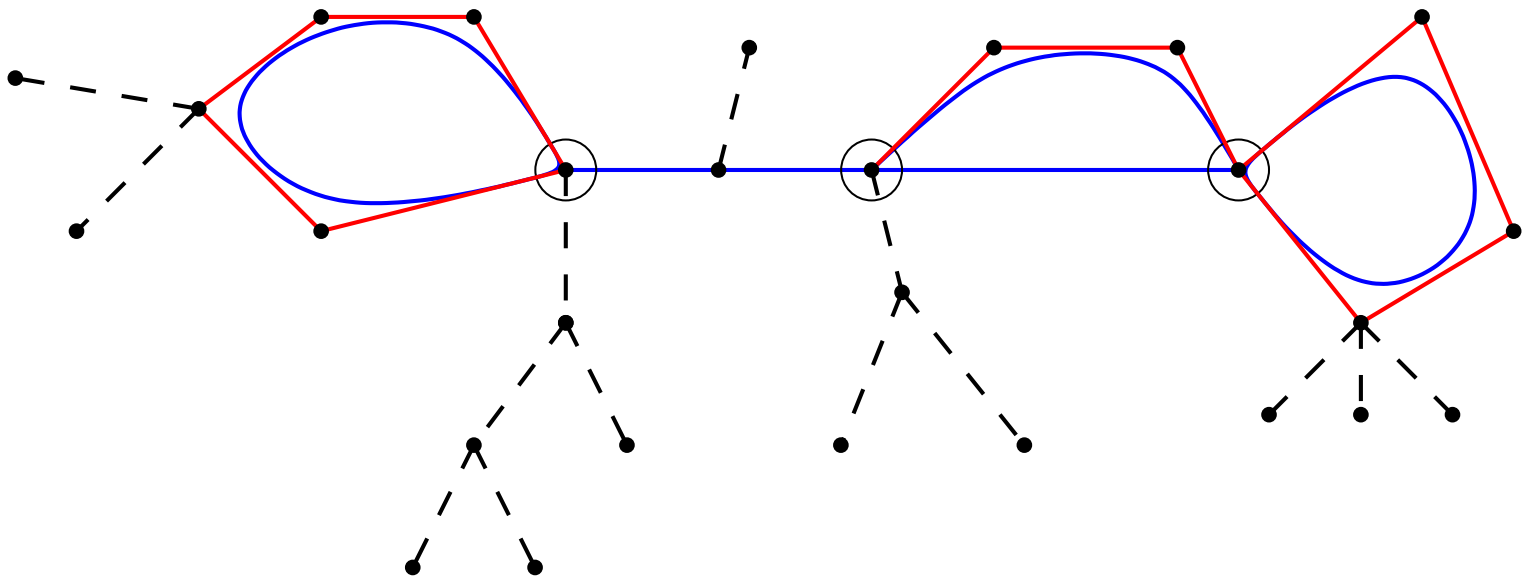

Cycle structure: an example

Kernel $K(G)$

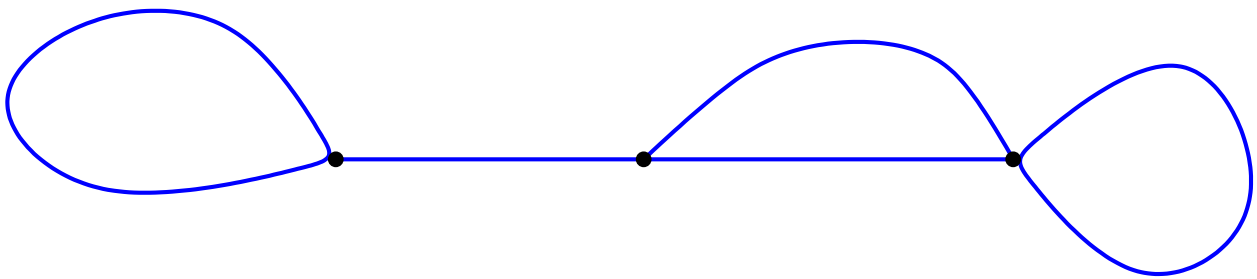

Cycle structure: an example

Kernel $K(G)$

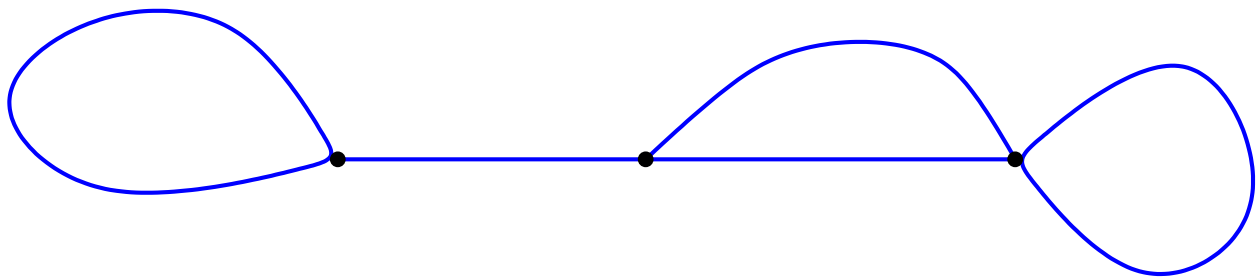

Note that the kernel has the same surplus as the original graph.

Cycle structure of a real tree with vertex identifications

It still makes sense to talk about the degree of a point in a real tree with vertex identifications.

Cycle structure of a real tree with vertex identifications

It still makes sense to talk about the degree of a point in a real tree with vertex identifications.

It's not hard to see that the core and kernel also make sense in the real tree context as a path metric space and a discrete multigraph respectively.

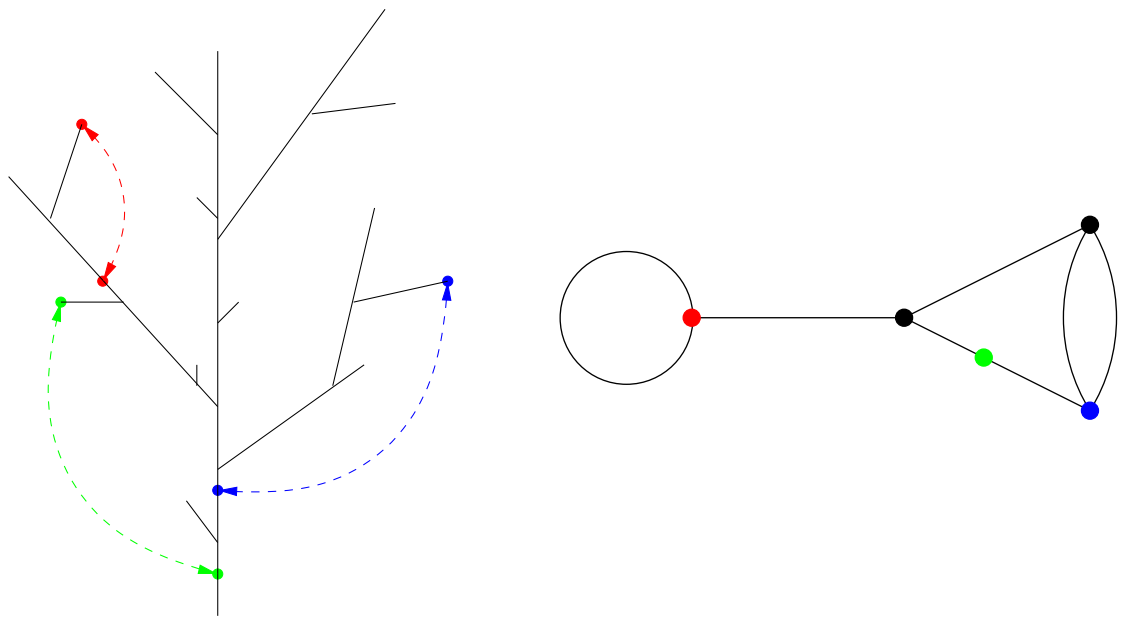

Convergence

The kernels of the components of $G\left(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}}\right)$ converge in distribution.

Alternative construction of a limit component

Condition on the size (= total mass) and surplus, s.

Alternative construction of a limit component

Condition on the size (= total mass) and surplus, s.

Once we condition on the size of such a component, it only appears as a scaling factor so, for ease of presentation, I will take it to be 1.

Alternative construction of a limit component

Condition on the size (= total mass) and surplus, s.

Once we condition on the size of such a component, it only appears as a scaling factor so, for ease of presentation, I will take it to be 1.

The cases $s = 0$ and $s = 1$ are a little different, so I'll skip them in order to concentrate on the generic case (as represented by $s = 2$!).

Alternative construction of a limit component

Sample a kernel according to the correct distribution, conditioned to have surplus s.

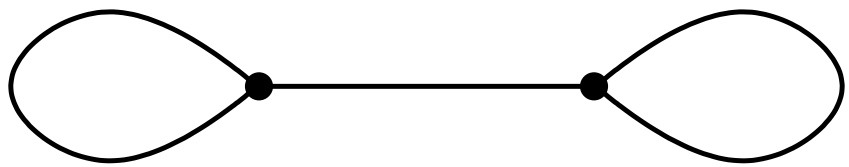

Alternative construction of a limit component

Sample independent rooted Brownian CRT's $\mathcal{T}_1, \mathcal{T}_2, \ldots, \mathcal{T}_{3(s-1)}$.

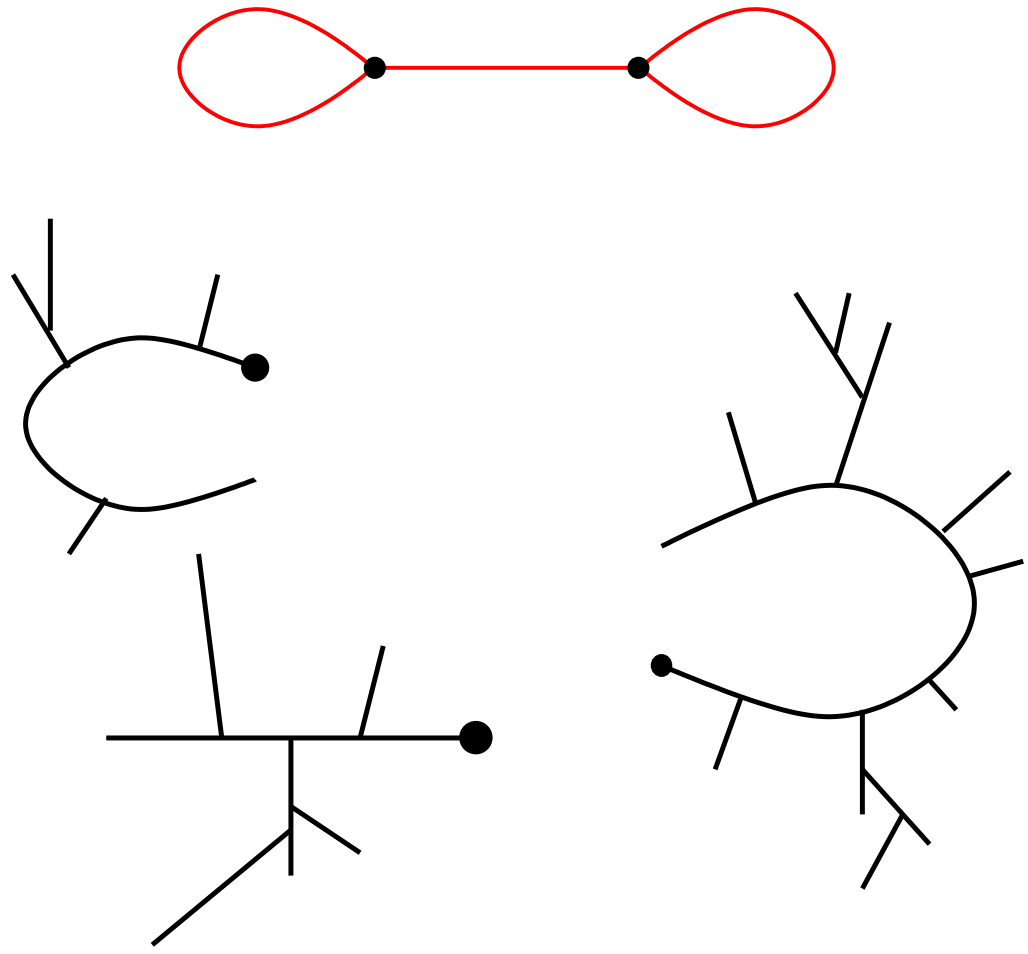

Alternative construction of a limit component

Sample a uniform point in each.

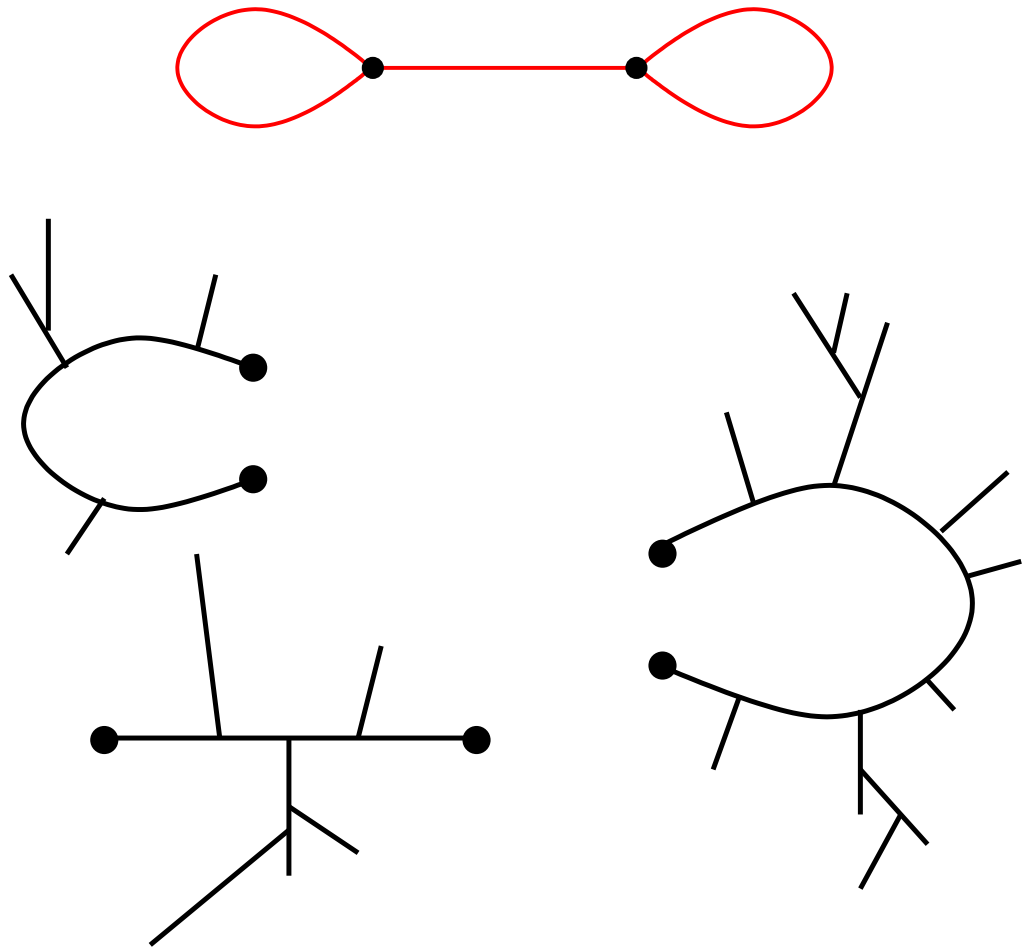

Alternative construction of a limit component

Randomly rescale so that the mass of $\mathcal{T}_i$ becomes X_i, where $(X_1, X_2, \ldots, X_{3(s-1)}) \sim \text{Dirichlet}(\frac{1}{2}, \ldots, \frac{1}{2})$.

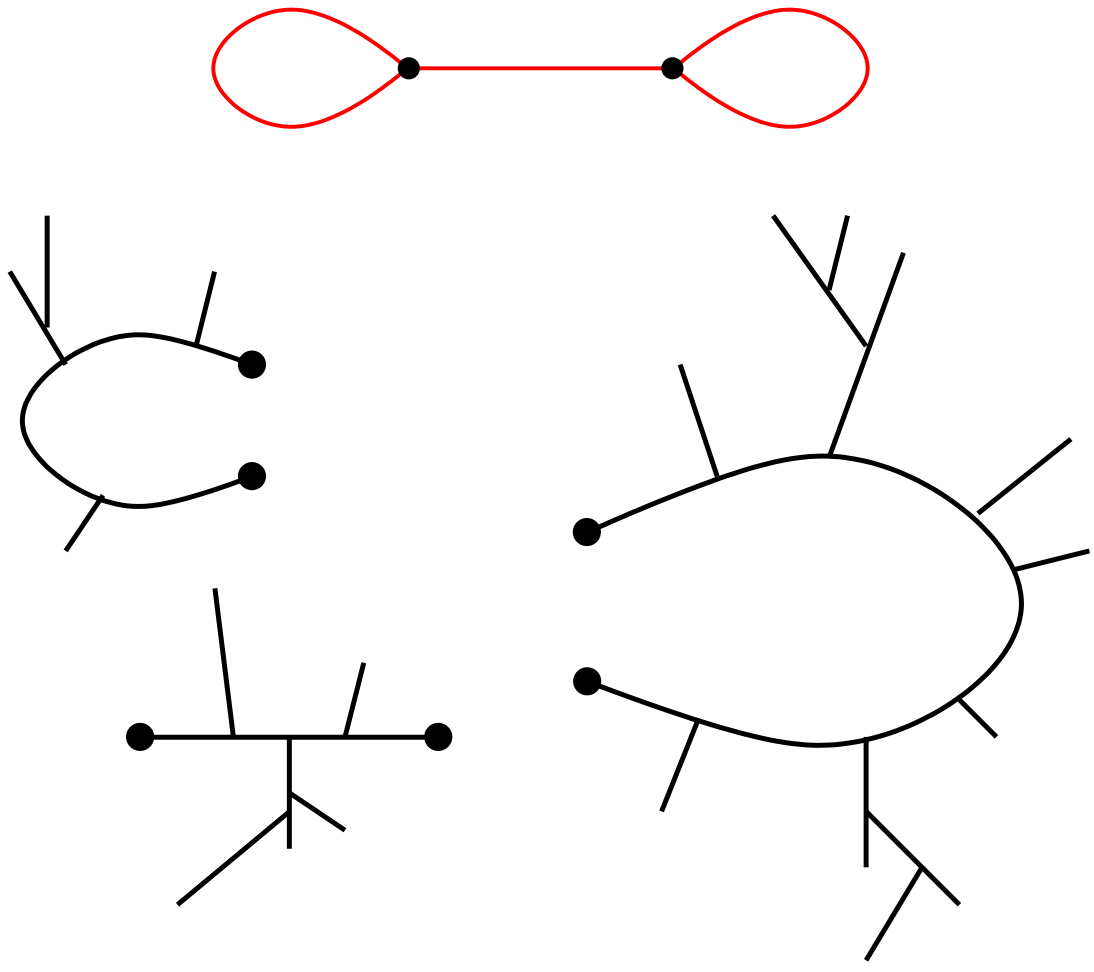

Alternative construction of a limit component

Glue the trees to the kernel.

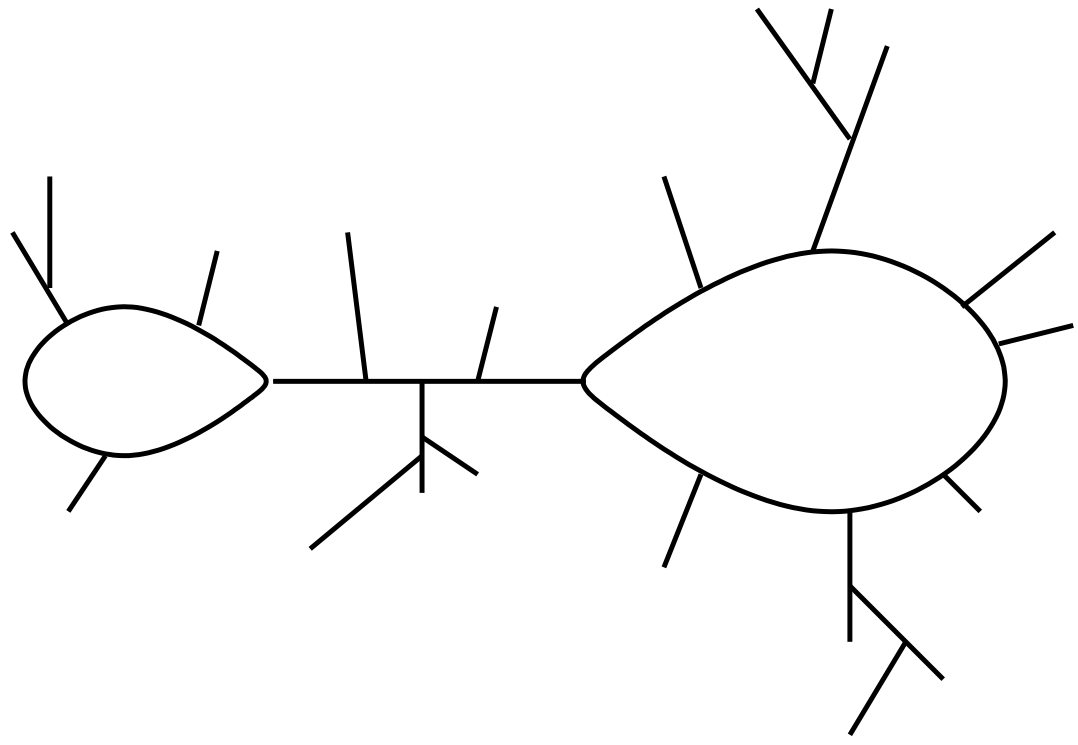

Alternative construction of a limit component

Glue the trees to the kernel.

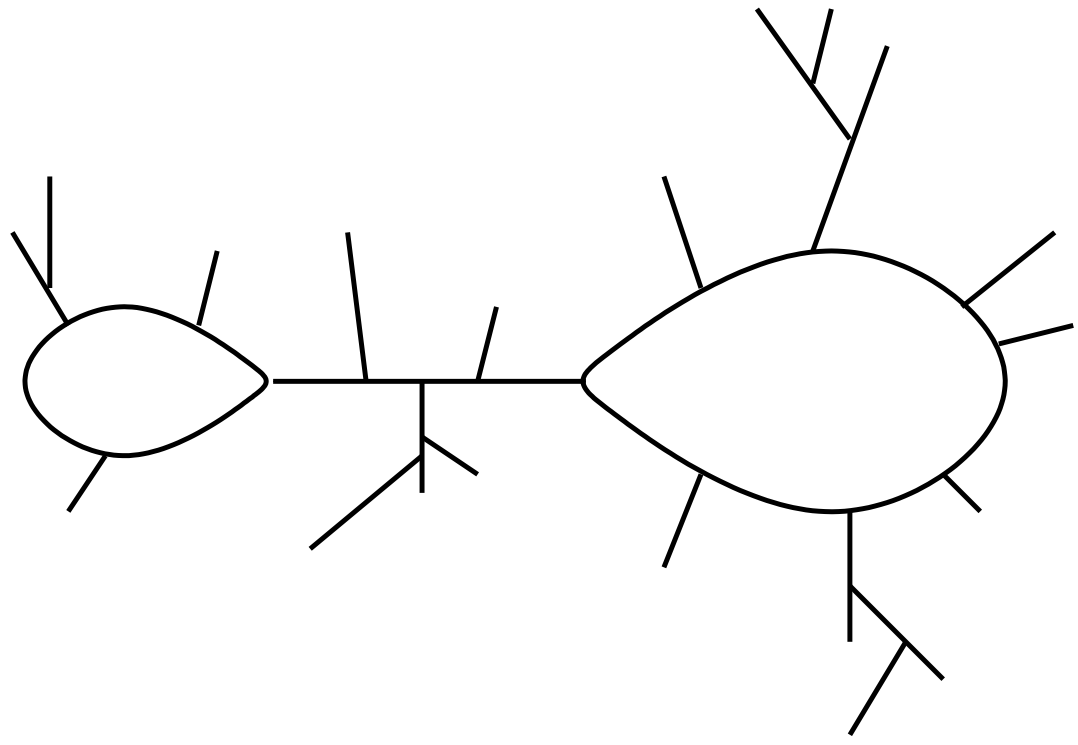

This has the same distribution as our real tree with vertex-identifications, conditioned to have mass 1 and surplus s.

The core

We get core paths of relative lengths Dirichlet$(1, 1, \ldots, 1)$.

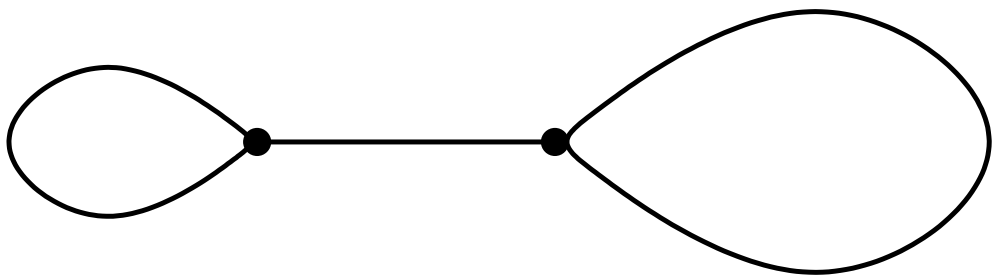

Moreover, the total length of the core is $\sqrt{\Gamma}$ where $\Gamma \sim \text{Gamma}((3s-2)/2, 1/2)$ (independently of the relative core lengths).

Stick-breaking construction

Starting from the core, it turns out we can give a stick-breaking construction for the rest of the limit component.

Stick-breaking construction

Starting from the core, it turns out we can give a stick-breaking construction for the rest of the limit component.

Take an inhomogeneous Poisson process of rate t at time t, conditioned to have its first point at $\sqrt{\Gamma}$. Write $C_0 = \sqrt{\Gamma}$; subsequent points occur at times $C_1, C_2, \ldots$.

Stick-breaking construction

Starting from the core, it turns out we can give a stick-breaking construction for the rest of the limit component.

Take an inhomogeneous Poisson process of rate t at time t, conditioned to have its first point at $\sqrt{\Gamma}$. Write $C_0 = \sqrt{\Gamma}$; subsequent points occur at times $C_1, C_2, \ldots$.

Consider the line-segments $[C_0, C_1)$, $[C_1, C_2), \ldots$ and proceed inductively.

Stick-breaking construction

Starting from the core, it turns out we can give a stick-breaking construction for the rest of the limit component.

Take an inhomogeneous Poisson process of rate t at time t, conditioned to have its first point at $\sqrt{\Gamma}$. Write $C_0 = \sqrt{\Gamma}$; subsequent points occur at times $C_1, C_2, \ldots$.

Consider the line-segments $[C_0, C_1)$, $[C_1, C_2), \ldots$ and proceed inductively.

For $i \geq 0$, attach $[C_i, C_{i+1})$ at a random point chosen uniformly over the existing structure.

Stick-breaking construction

Starting from the core, it turns out we can give a stick-breaking construction for the rest of the limit component.

Take an inhomogeneous Poisson process of rate t at time t, conditioned to have its first point at $\sqrt{\Gamma}$. Write $C_0 = \sqrt{\Gamma}$; subsequent points occur at times $C_1, C_2, \ldots$.

Consider the line-segments $[C_0, C_1)$, $[C_1, C_2), \ldots$ and proceed inductively.

For $i \geq 0$, attach $[C_i, C_{i+1})$ at a random point chosen uniformly over the existing structure.

Take the closure of the metric space obtained.

Stick-breaking construction

Random core.

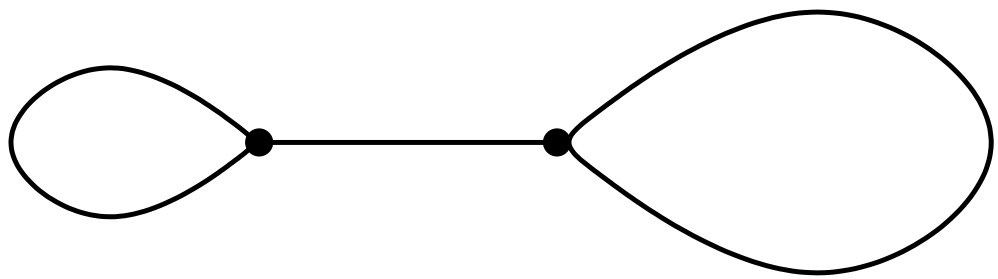

Stick-breaking construction

Run Poisson process...

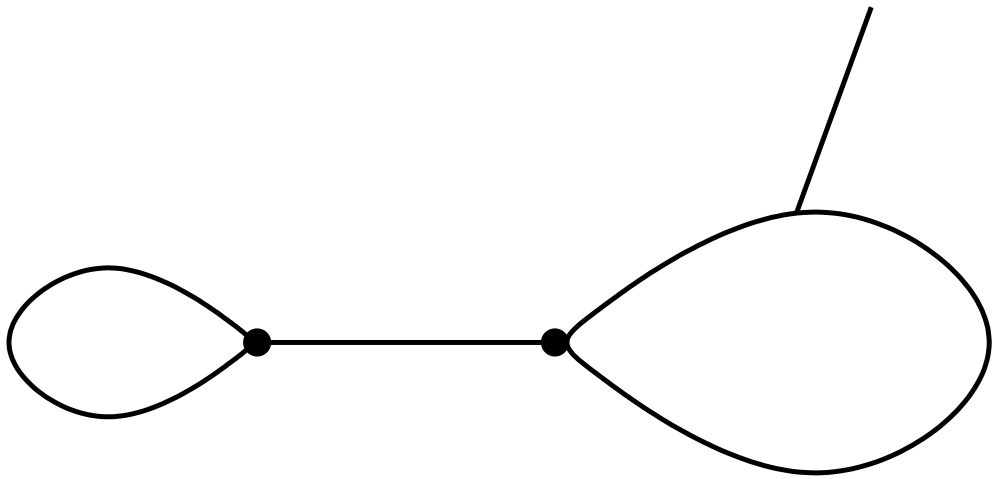

Stick-breaking construction

Run Poisson process...

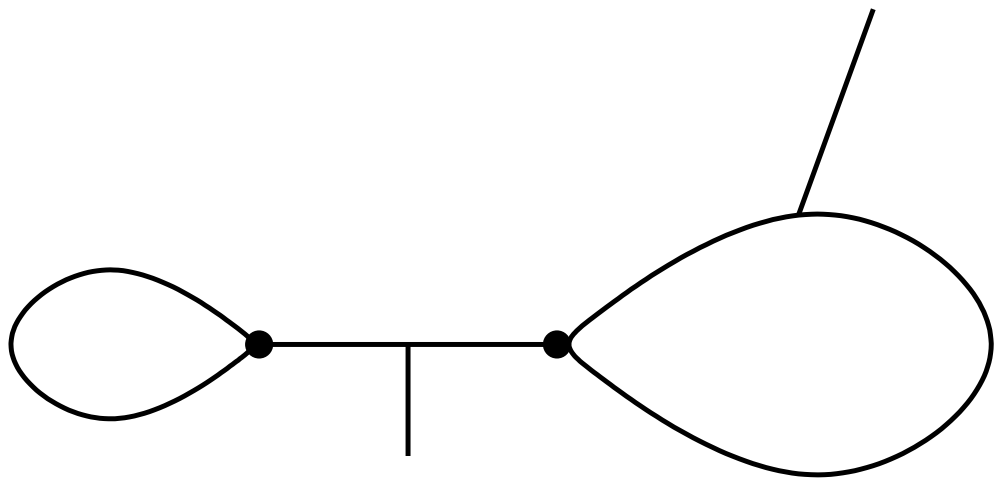

Stick-breaking construction

Run Poisson process...

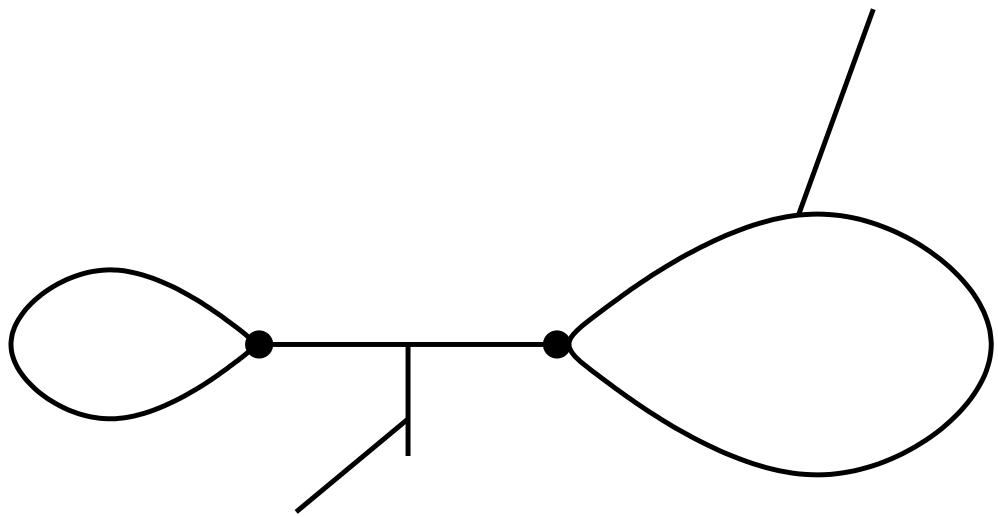

Stick-breaking construction

Run Poisson process...

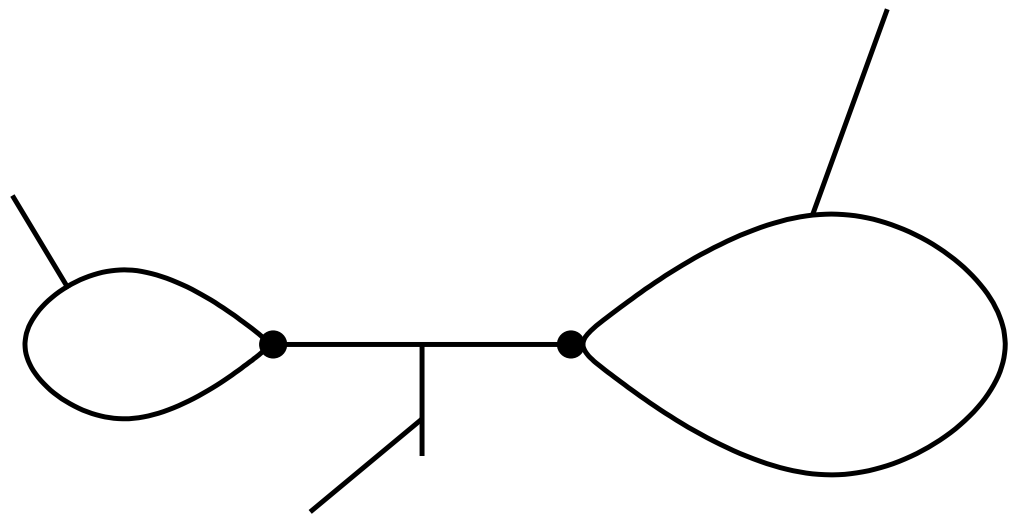

Stick-breaking construction

Run Poisson process...

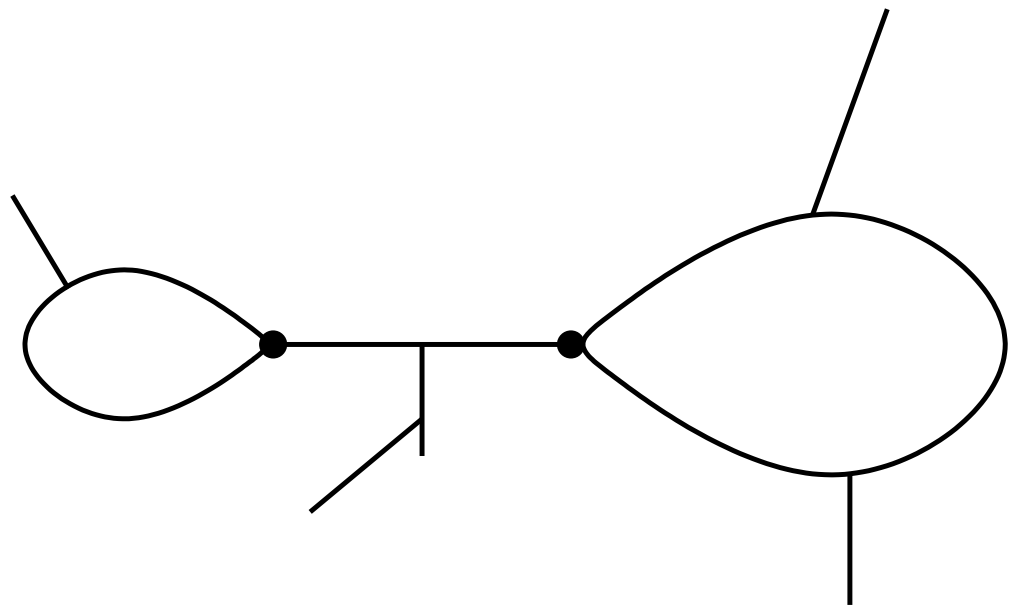

Stick-breaking construction

Run Poisson process...

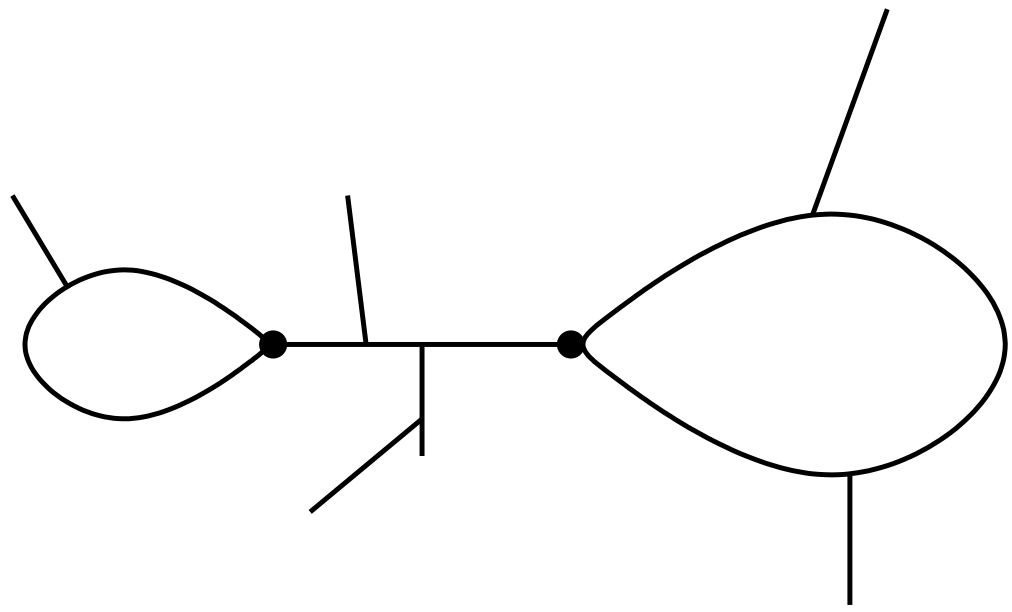

Stick-breaking construction

Run Poisson process...

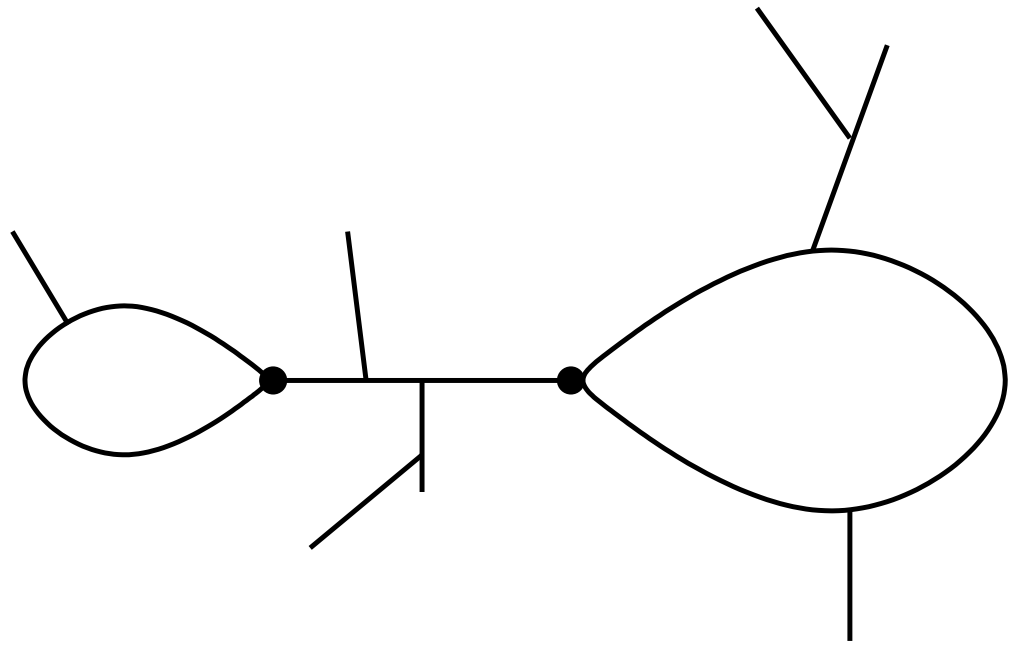

Stick-breaking construction

Run Poisson process...

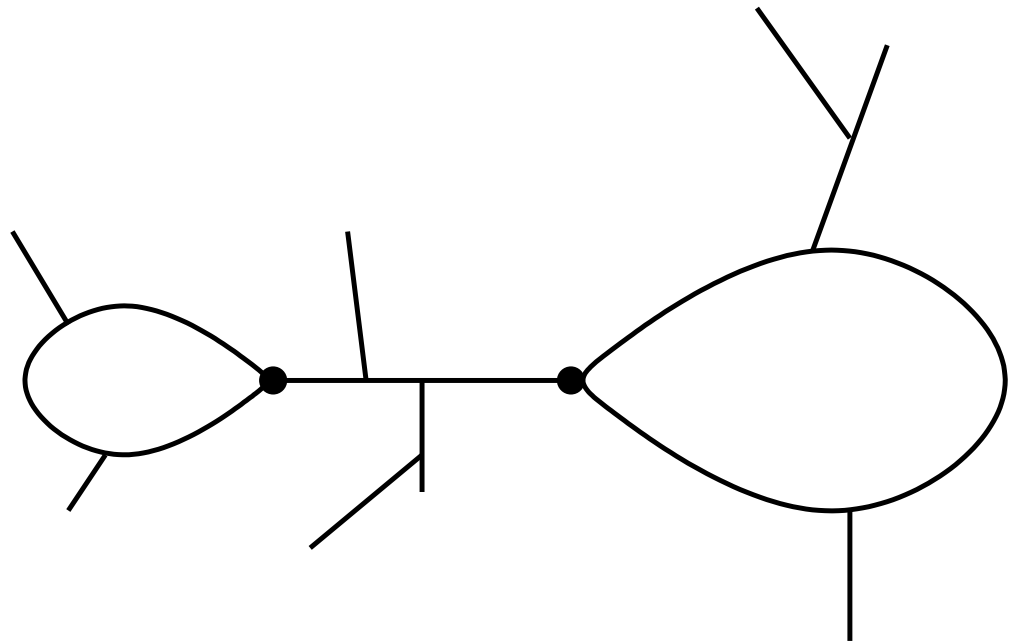

Stick-breaking construction

Final component

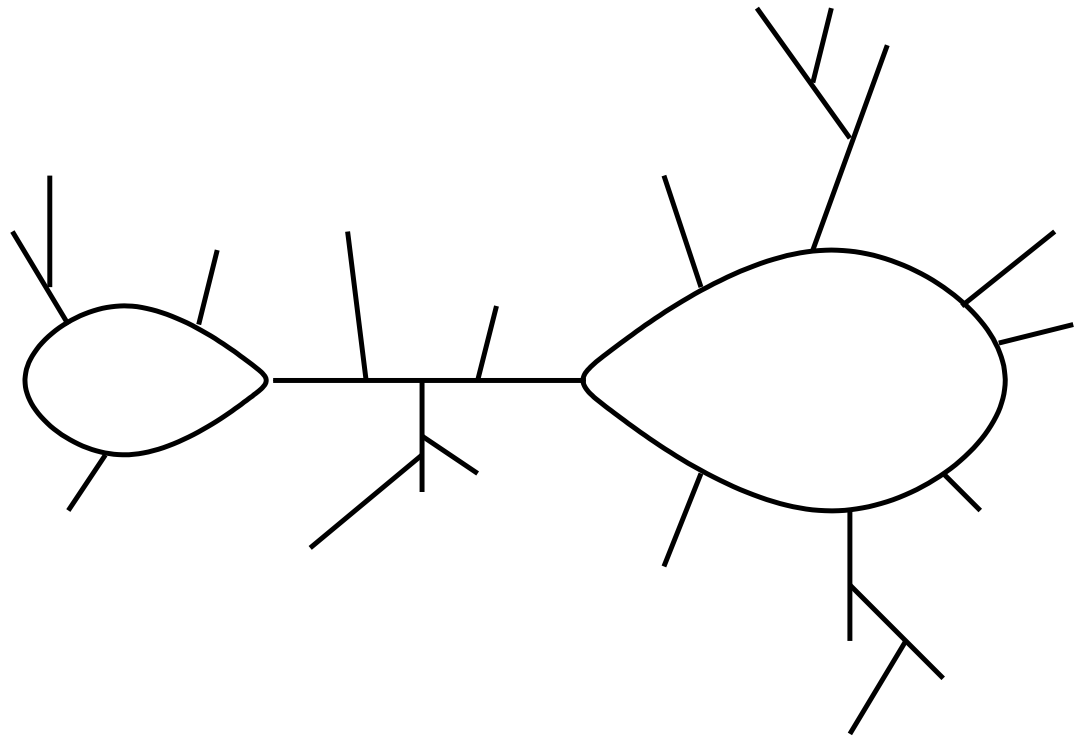

Universality

The Brownian CRT is the scaling limit for any conditioned Galton-Watson branching process family tree with offspring mean 1 and finite variance (and for several other tree models besides).

Universality

The Brownian CRT is the scaling limit for any conditioned Galton-Watson branching process family tree with offspring mean 1 and finite variance (and for several other tree models besides).

We expect that our limit for the critical random graph should be universal for a whole range of random graph models, including those with fixed degree sequences satisfying certain conditions.

The minimum spanning tree of the complete graph

Consider the complete graph on n vertices with independent edge-weights which are uniformly distributed on $[0, 1]$.

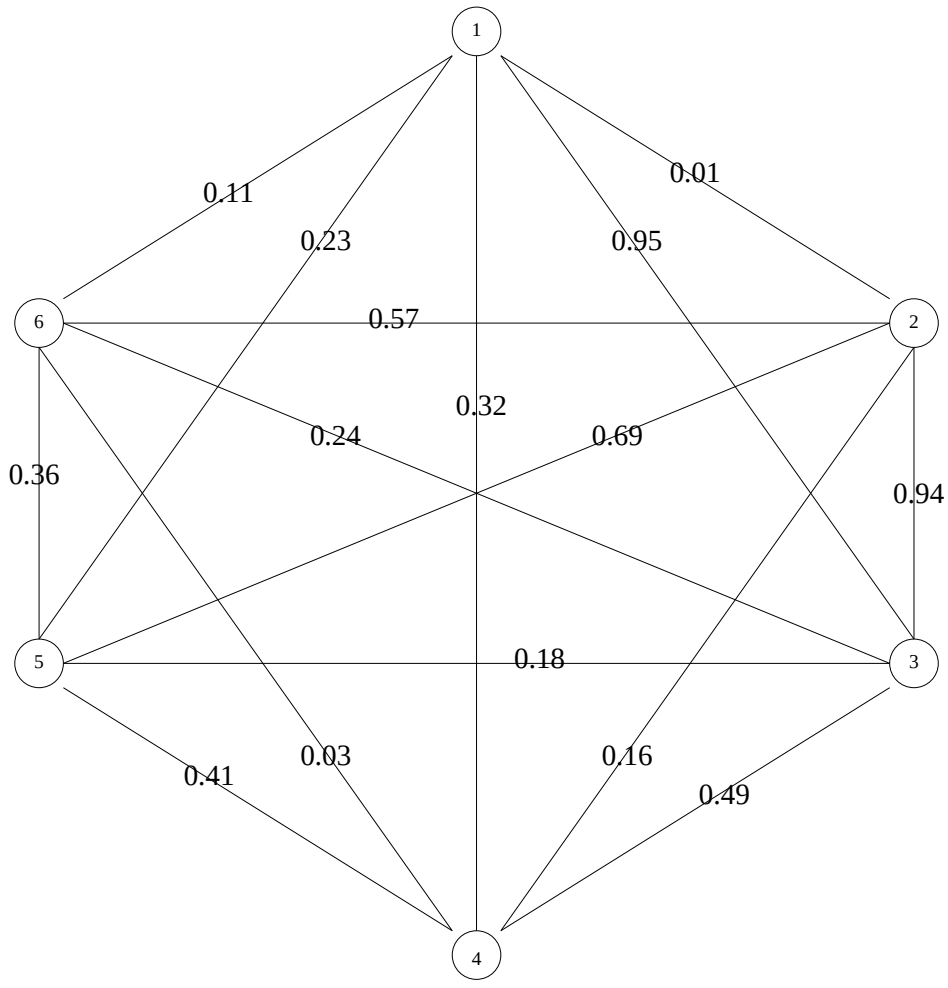

The minimum spanning tree of the complete graph

Find the minimum spanning tree (MST).

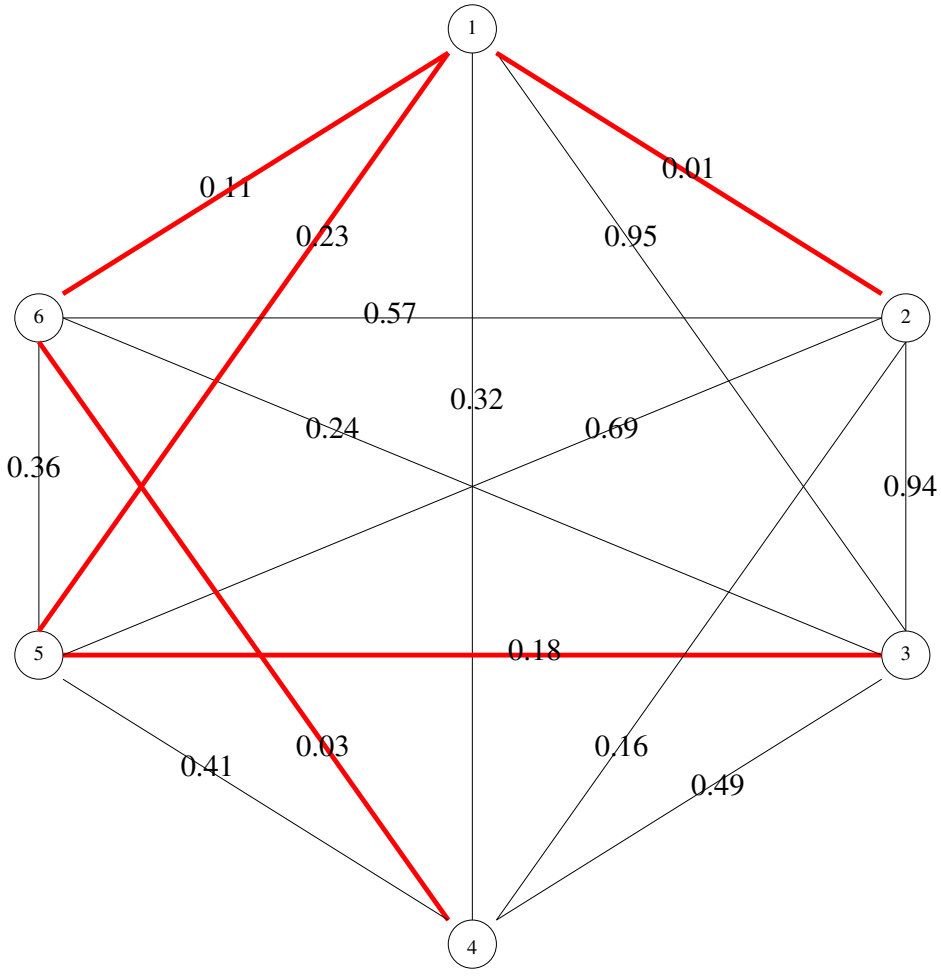

Question

Does the MST of the complete graph on n vertices possess a scaling limit?

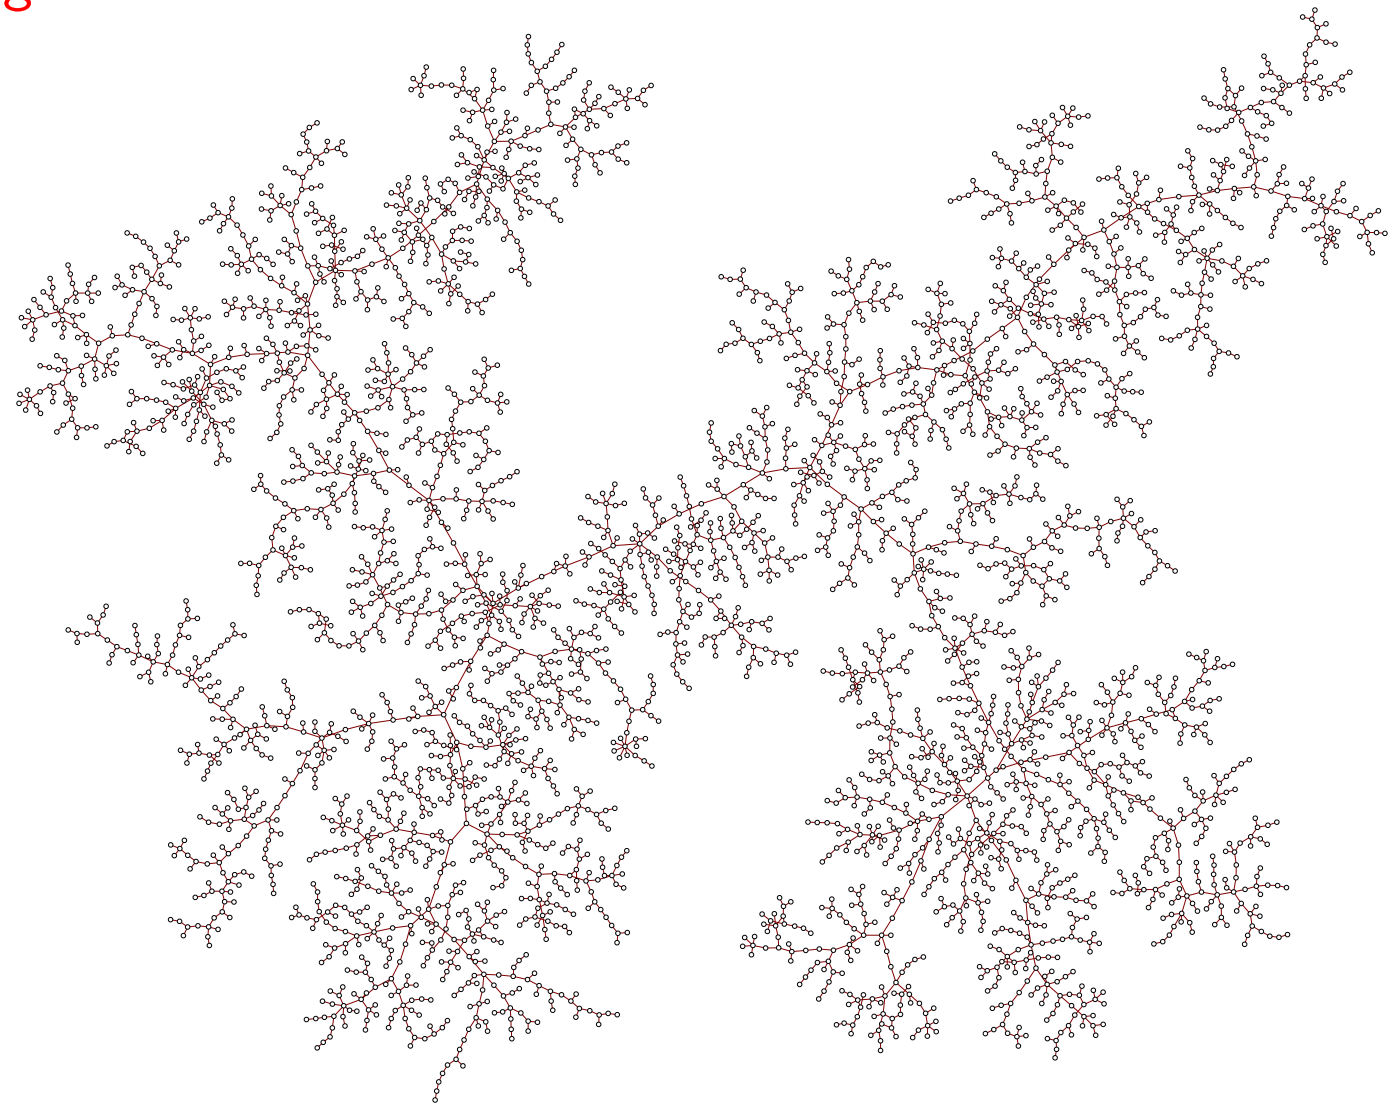

[Picture by Louigi Addario-Berry]

The scaling limit of the MST

Let M_n be the MST of the complete graph on n vertices, considered as a metric space using the graph distance.

Theorem. (Addario-Berry, Broutin, G. & Miermont) There exists a random compact real tree $\mathcal{M}$ such that

$$\frac{1}{n^{1/3}} M_n \xrightarrow{d} \mathcal{M}$$

as $n \to \infty$, in the sense of $\mathrm{d_{GH}}$. Moreover, although $\mathcal{M}$ is almost surely binary, it is not the Brownian CRT.

The scaling limit of the MST

Let M_n be the MST of the complete graph on n vertices, considered as a metric space using the graph distance.

Theorem. (Addario-Berry, Broutin, G. & Miermont) There exists a random compact real tree $\mathcal{M}$ such that

$$\frac{1}{n^{1/3}} M_n \xrightarrow{d} \mathcal{M}$$

as $n \to \infty$, in the sense of $\mathrm{d_{GH}}$. Moreover, although $\mathcal{M}$ is almost surely binary, it is not the Brownian CRT.

The key to understanding this result is a connection between the Erdős-Rényi random graph and Kruskal's algorithm for constructing the MST.

Bibliography: random graphs

Brownian excursions, critical random graphs and the multiplicative coalescent
D. Aldous, *Annals of Probability* **25** (1997) pp.812-854.

The continuum limit of critical random graphs
L. Addario-Berry, N. Broutin and C. Goldschmidt, *Probability Theory and Related Fields* **152** (2012), pp.367-406.

Critical random graphs: limiting constructions and distributional properties
L. Addario-Berry, N. Broutin and C. Goldschmidt, *Electronic Journal of Probability* **15** (2010), pp.741-775.

The scaling limit of the minimum spanning tree of the complete graph
L. Addario-Berry, N. Broutin, C. Goldschmidt and G. Miermont in preparation.

Thank you for your attention!