Scaling limits for random trees and graphs

YEP VII "Probability, random trees and algorithms", 8th-12th March 2010
Christina Goldschmidt, University of Warwick

Preliminary version, 7th March 2010

INTRODUCTION

A taste of what's to come

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

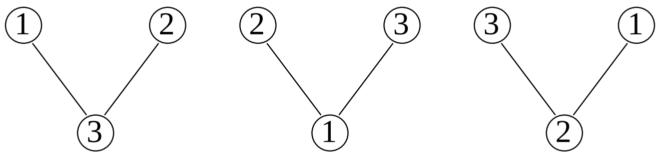

Unordered trees

Note that unordered means that these trees are all the same:

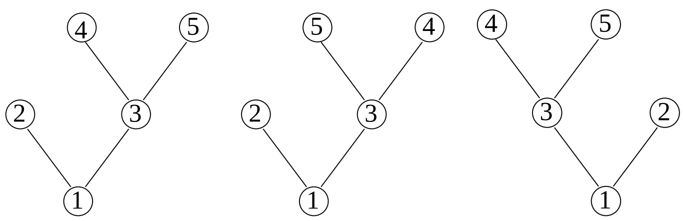

but this one is different:

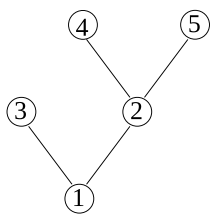

Uniform random trees

Cayley's formula tells us that $|\mathbb{T}_{[n]}| = n^{n-2}$.

Write T_n for a tree chosen uniformly from $\mathbb{T}_{[n]}$. When not otherwise qualified, this is what we mean by a uniform random tree.

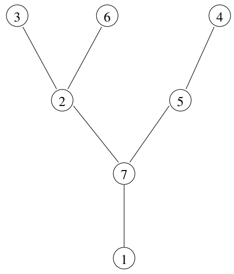

What happens as n grows?

An algorithm due to Aldous

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

$$V_i = \begin{cases} j & \text{with probability } 1/n,\ 1 \leq j \leq i-2, \\ i-1 & \text{with probability } 1-(i-2)/n. \end{cases}$$

4. Take a uniform random permutation of the labels.

[See Nicolas Broutin's lecture.]

An algorithm due to Aldous

Consider $n = 10$.

An algorithm due to Aldous

$V_2 = 1$ with probability 1

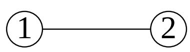

An algorithm due to Aldous

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

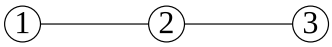

An algorithm due to Aldous

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

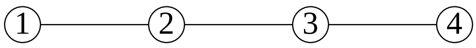

An algorithm due to Aldous

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

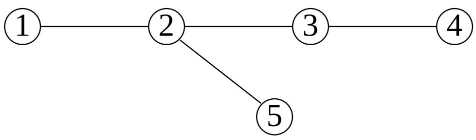

An algorithm due to Aldous

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

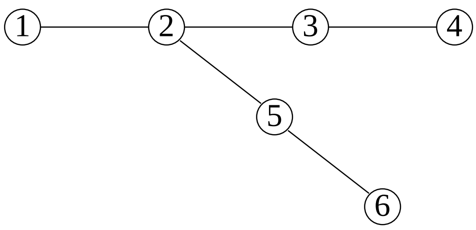

An algorithm due to Aldous

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

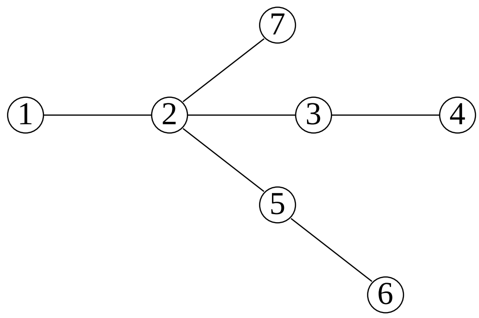

An algorithm due to Aldous

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

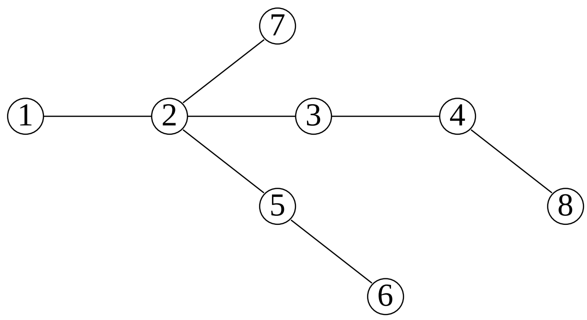

An algorithm due to Aldous

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

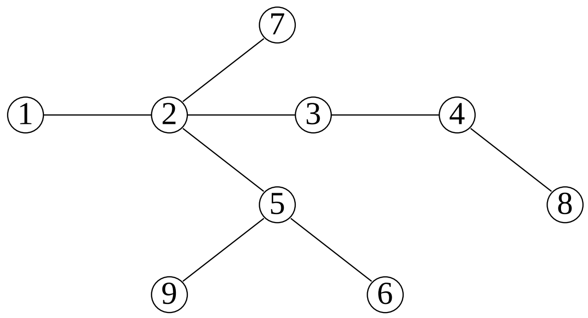

An algorithm due to Aldous

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

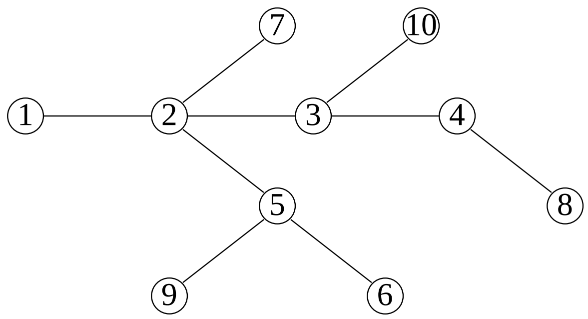

An algorithm due to Aldous

Permute.

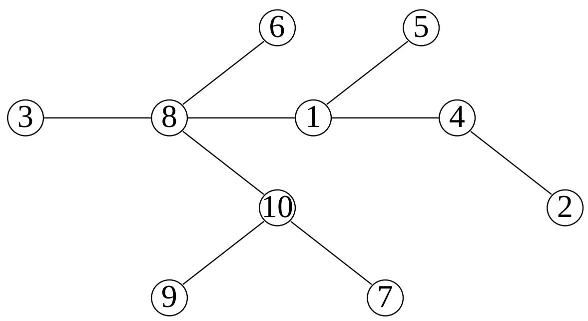

Typical distances

Consider the tree before we permute. Let $J_n = \inf\{i \geq 1 : V_{i+1} \neq i\}$. We can use J_n to give us an idea of typical distances in the tree.

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

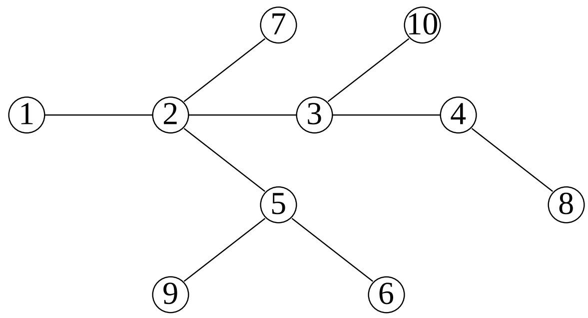

Typical distances

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

Imagine now that edges in the tree have length 1. This result 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.

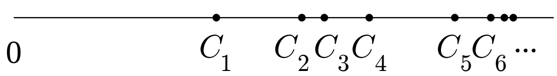

Consider the line-segments $[0, C_1)$, $[C_1, C_2), \ldots$.

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

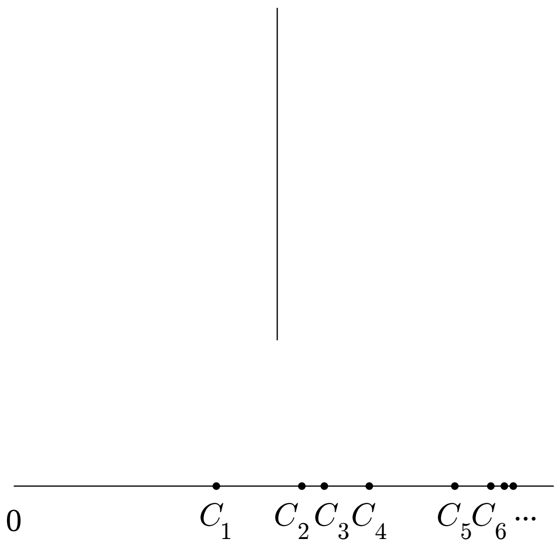

Stick-breaking procedure

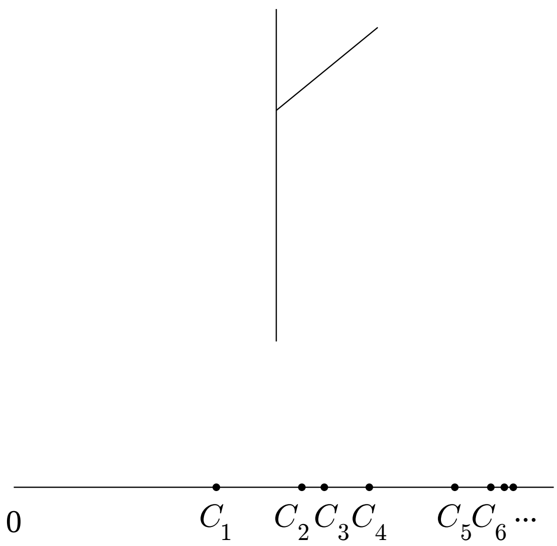

Stick-breaking procedure

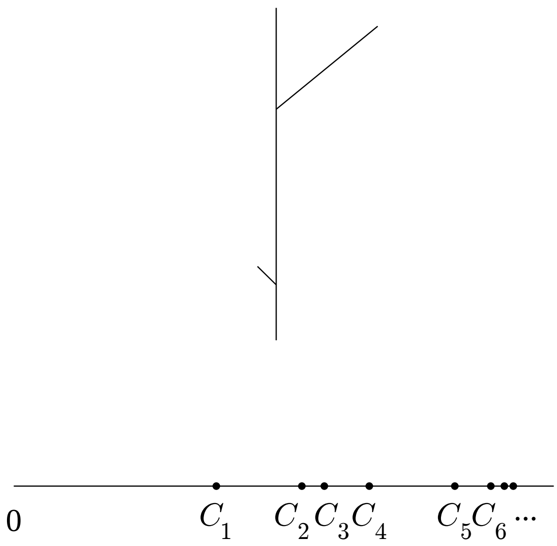

Stick-breaking procedure

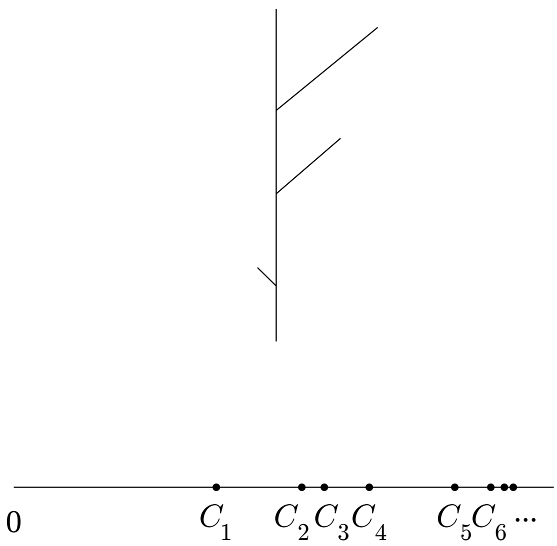

Stick-breaking procedure

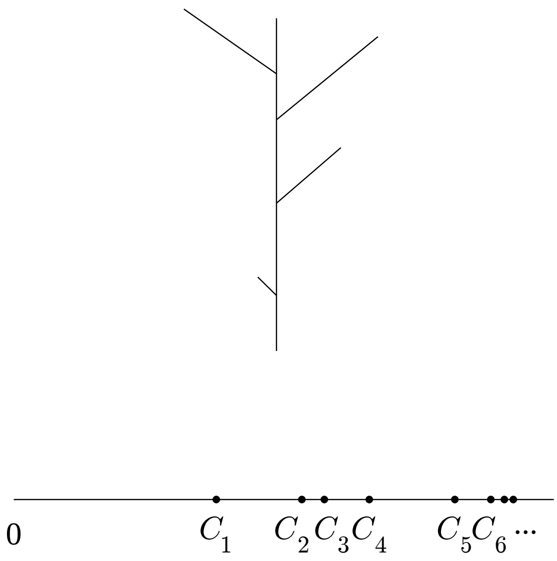

Stick-breaking procedure

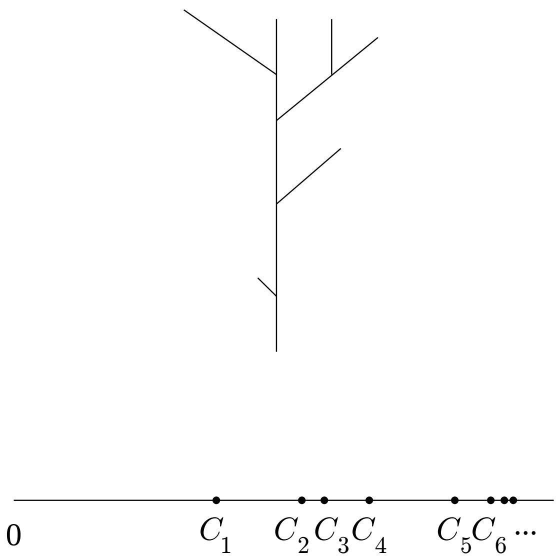

Stick-breaking procedure

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 closure of the union of all the branches.

This procedure gives (a rather informally expressed) definition of Aldous' Brownian continuum random tree (CRT).

The Brownian continuum random tree

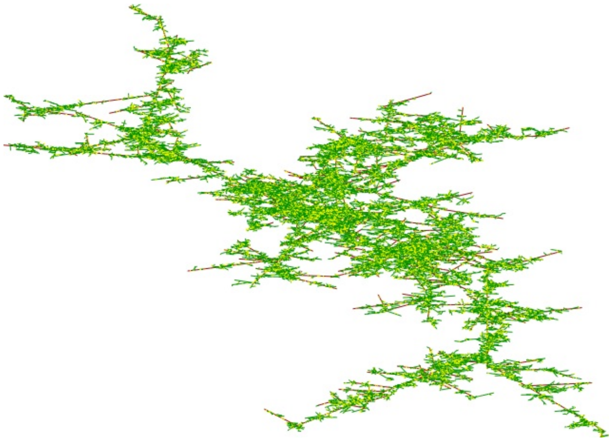

[Picture by Grégory Miermont]

DISCRETE TREES

Based in large part on **Random trees and applications** by Jean-François Le Gall.

Ordered trees

It turns out to be more natural to work with rooted, ordered trees (also called plane trees).

Ordered trees

We will use the Ulam-Harris labelling. Let $\mathbb{N} = \{1, 2, 3, \dots\}$ and

$$\mathcal{U} = \bigcup_{n=0}^{\infty} \mathbb{N}^n,$$

where $\mathbb{N}^0 = \{\emptyset\}$. An element $u \in \mathcal{U}$ is a sequence $u = (u^1, u^2, \dots, u^n)$ of natural numbers representing a point in an infinitary tree:

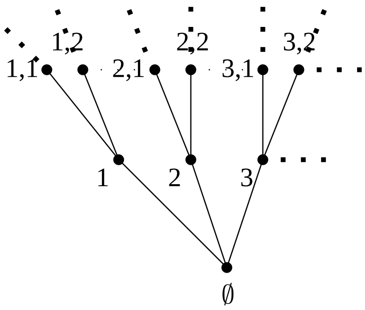

Thus the label of a vertex indicates its genealogy.

Ordered trees

Write $|u| = n$ for the generation of u.

u has parent $p(u) = (u^1, u^2, \dots, u^{n-1})$.

u has children $u1, u2, \dots$ where, in general, $uv = (u^1, u^2, \dots, u^n, v^1, v^2, \dots, v^m)$ is the concatenation of sequences $u = (u^1, u^2, \dots, u^n)$ and $v = (v^1, v^2, \dots, v^m)$.

We root the tree at $\emptyset$.

Ordered trees

A (finite) rooted, ordered tree $\mathbf{t}$ is a finite subset of $\mathcal{U}$ such that

- $\emptyset \in \mathbf{t}$
- for all $u \in \mathbf{t}$ such that $u \neq \emptyset$, $p(u) \in \mathbf{t}$
- for all $u \in \mathbf{t}$, there exists $k(u) \in \mathbb{Z}_+$ such that for $j \in \mathbb{N}$, $uj \in \mathbf{t}$ iff $1 \leq j \leq k(u)$.

$k(u)$ is the number of children of u in $\mathbf{t}$.

Write $\#(\mathbf{t})$ for the size (number of vertices) of $\mathbf{t}$ and note that

$$\#(\mathbf{t}) = 1 + \sum_{u \in \mathbf{t}} k(u).$$

Write $\mathbf{T}$ for the set of all rooted ordered trees.

Two ways of encoding a tree

Consider a rooted ordered tree $\mathbf{t} \in \mathbf{T}$.

It will be convenient to encode this tree in terms of discrete functions which are easier to manipulate.

We will do this is two different ways:

- the height function
- the depth-first walk.

Height function

Suppose that $\mathbf{t}$ has n vertices. Let them be $v_0, v_1, \dots, v_{n-1}$, listed in lexicographical order.

Then the height function is defined by

$$H(k) = |v_k|, \quad 0 \leq k \leq n-1.$$

Height function

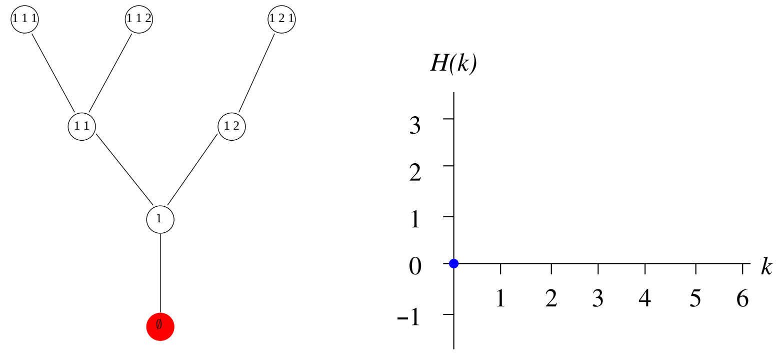

Height function

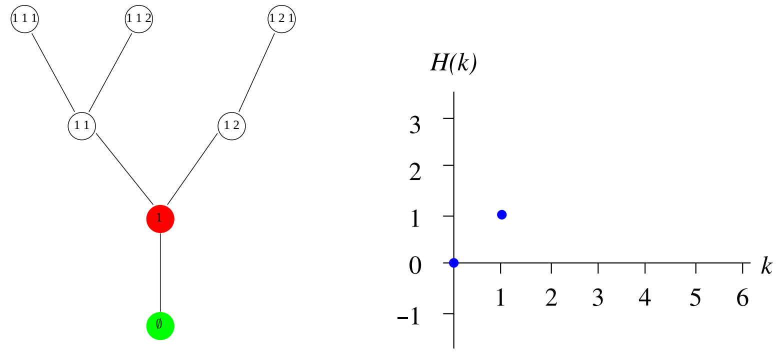

Height function

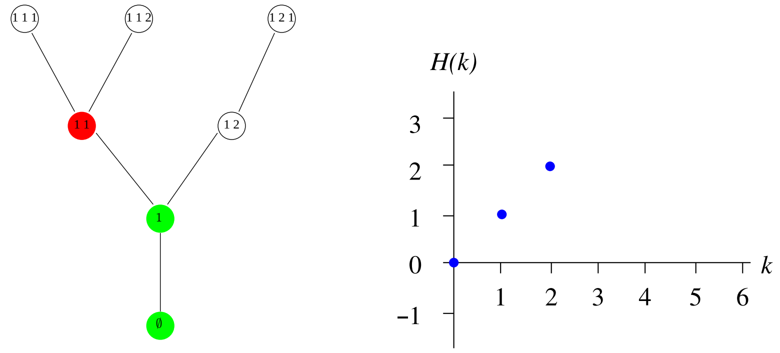

Height function

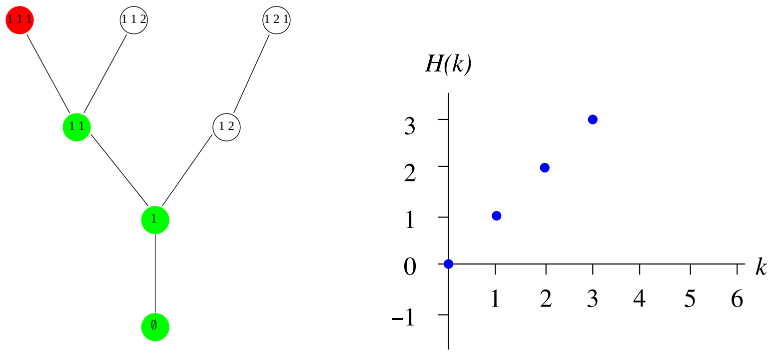

Height function

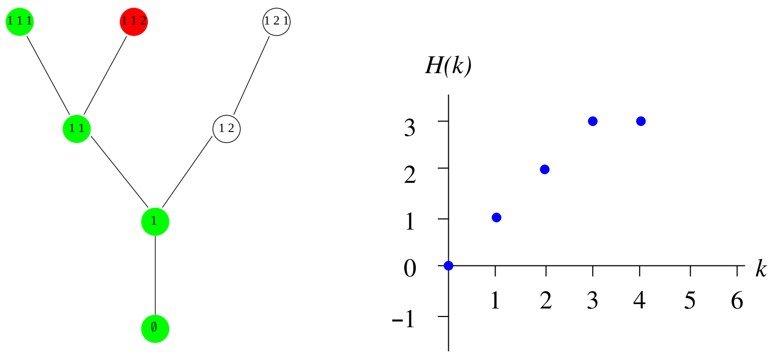

Height function

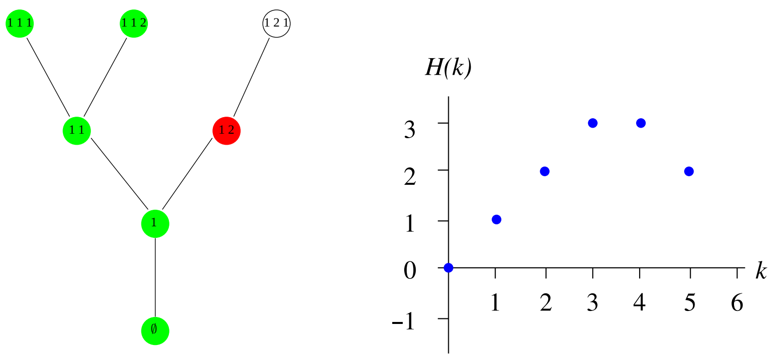

Height function

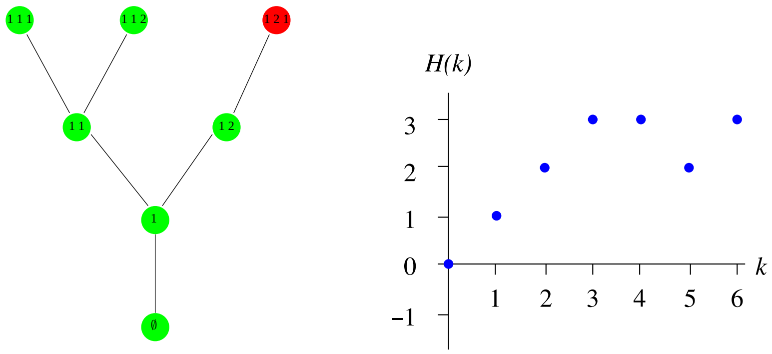

Height function

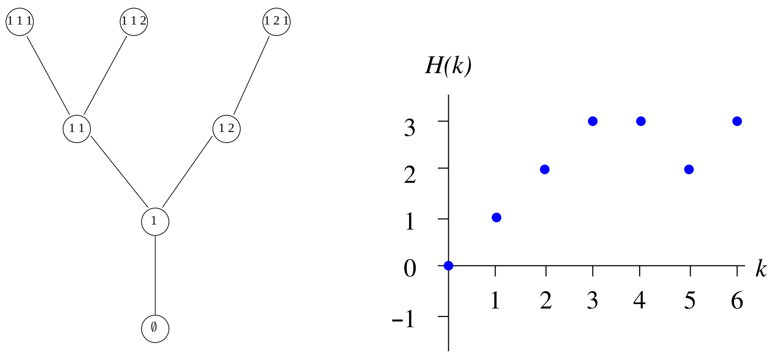

We can recover the tree from its height function (after a little thought!).

Depth-first walk

Recall that $k(v)$ is the number of children of v, and that $v_0, v_1, \ldots, v_{n-1}$ is a list of the vertices of $\mathbf{t}$ in lexicographical order.

Define

$$X(0) = 0,$$
$$X(i) = \sum_{j=0}^{i-1}(k(v_j) - 1), \text{ for } 1 \leq i \leq n.$$

In other words,

$$X(i+1) = X(i) + k(v_i) - 1, \quad 0 \leq i \leq n-1.$$

Depth-first walk

Step 0

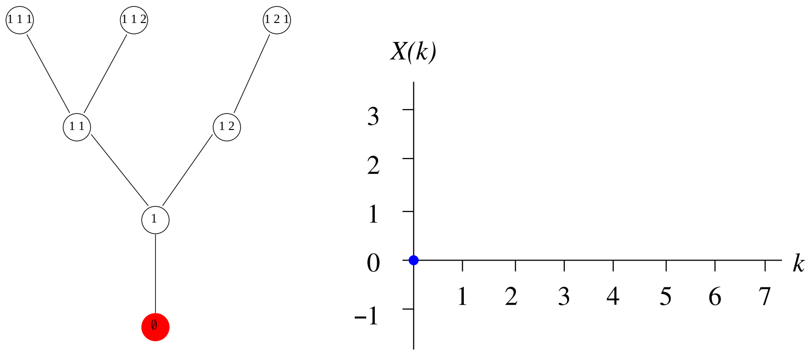

Current: ∅ Alive: none Dead: none

Depth-first walk

Step 1

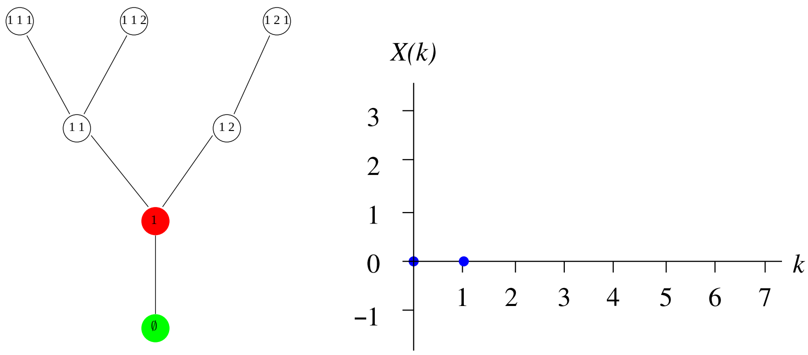

Current: 1 Alive: none Dead: ∅

Depth-first walk

Step 2

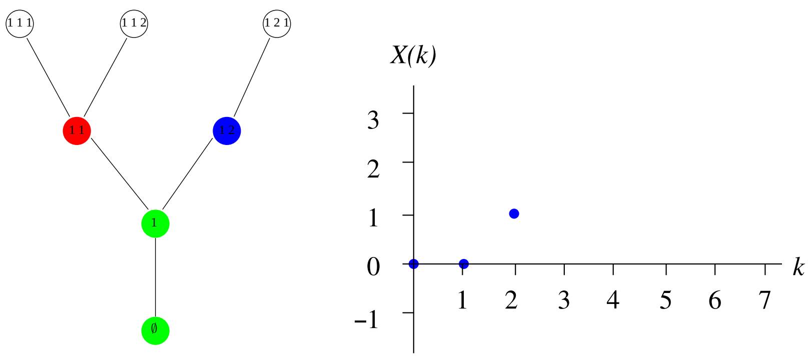

Current: 1 1 Alive: 1 2 Dead: ∅, 1

Depth-first walk

Step 3

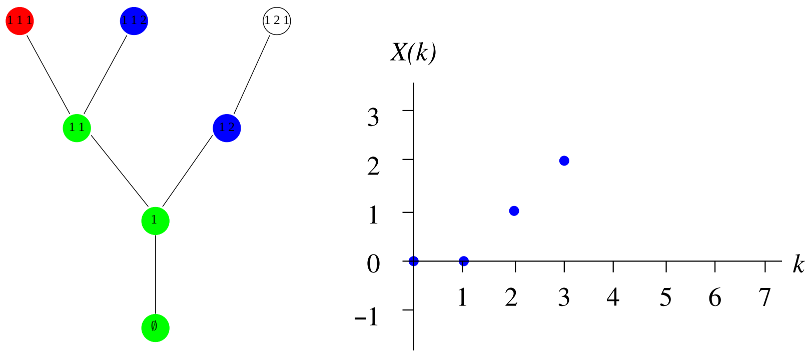

Current: 1 1 1 Alive: 1 1 2, 1 2 Dead: ∅, 1, 1 1

Depth-first walk

Step 4

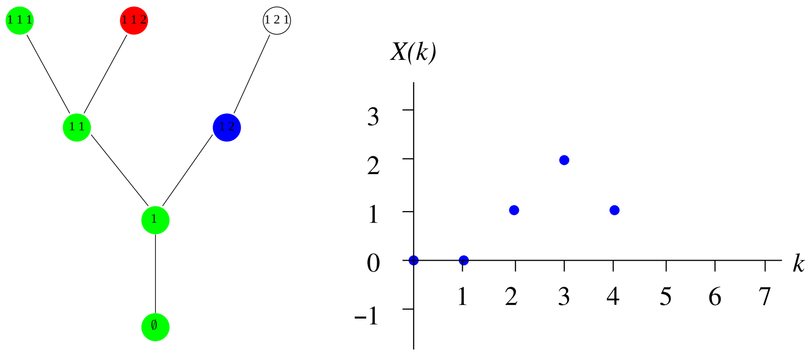

Current: 1 1 2 Alive: 1 2 Dead: ∅, 1, 1 1, 1 1 1

Depth-first walk

Step 5

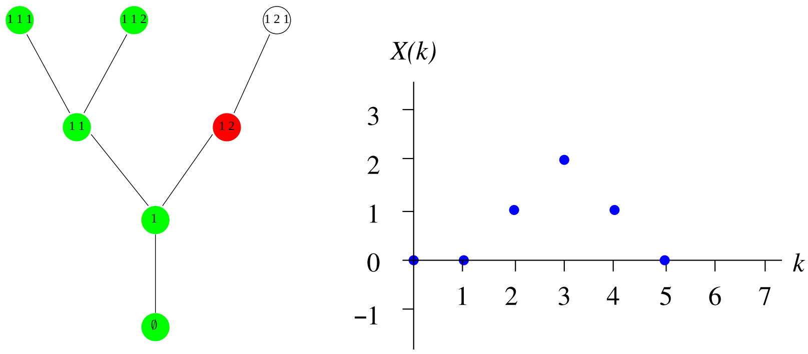

Current: 1 2 Alive: none Dead: ∅, 1, 1 1, 1 1 1, 1 1 2

Depth-first walk

Step 6

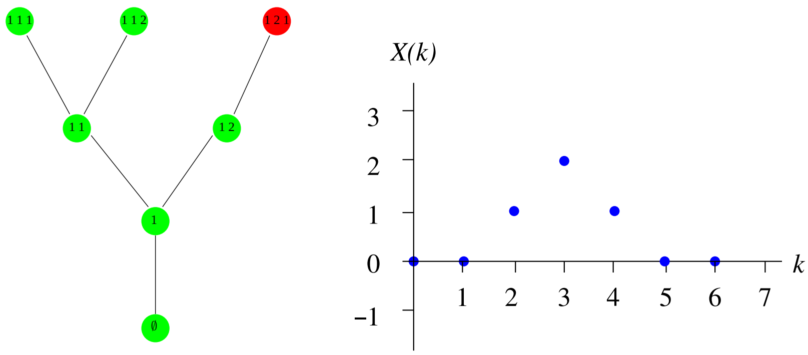

Current: 1 2 1 Alive: none Dead: ∅, 1, 1 1, 1 1 1, 1 1 2, 1 2

Depth-first walk

Step 7

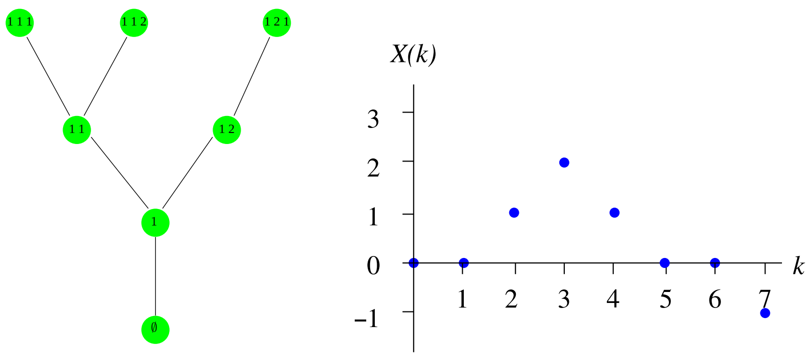

Dead: ∅, 1, 1 1, 1 1 1, 1 1 2, 1 2, 1 2 1

Depth-first walk

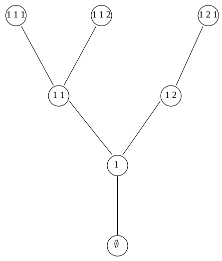

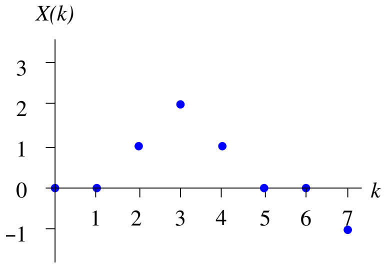

It is less easy to see that the depth-first walk also encodes the tree.

Proposition 2. *For* $0 \leq i \leq n-1$,

$$H(i) = \#\left\{0 \leq j \leq i-1 : X(j) = \min_{j \leq k \leq i} X(k)\right\}.$$ ▶

Galton-Watson process

A Galton-Watson branching process $(Z_n)_{n \geq 0}$ 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 $\mu(k)$ gives the probability of k children, $k \geq 0$.
- Each child reproduces as an independent copy of the original individual.

Z_n gives the number of individuals in generation n (in particular, $Z_0 = 1$).

Galton-Watson trees

A Galton-Watson tree is the family tree arising from a Galton-Watson branching process. We will think of this as a rooted ordered tree.

We will consider the case where the offspring distribution μ is critical or subcritical i.e.

$$\sum_{k=1}^{\infty} k\mu(k) \leq 1.$$

This ensures that the resulting tree, T, is finite.

Since the tree is random, we will refer to the height *process* rather than function.

Uniform random trees revisited

Proposition 3. *Let P be a (rooted, ordered) Galton-Watson tree, with Poisson(1) offspring distribution and total progeny N. Assign the vertices labels uniformly at random from $\{1, 2, \ldots, N\}$ and then forget the ordering and the root. Let $\tilde{P}$ be the labelled tree obtained. Then, conditional on $N = n$, $\tilde{P}$ has the same distribution as T_n, a uniform random tree on n vertices.* ▶

Other combinatorial trees (in disguise)

Let T be a Galton-Watson tree with offspring distribution μ and total progeny N.

- If $\mu(k) = 2^{-k-1}$, $k \geq 0$ (i.e. Geometric(1/2) offspring distribution) then conditional on $N = n$, the tree is uniform on the set of ordered trees with n vertices.
- If $\mu(k) = \frac{1}{2}(\delta_0(k) + \delta_2(k))$, $k \geq 0$ then conditional on $N = n$, for n odd, the tree is uniform on the set of (complete) binary trees.

The depth-first walk of a Galton-Watson process is a stopped random walk

Recall that μ is a distribution on $\mathbb{Z}_+$ such that $\sum_{k=1}^{\infty} k\mu(k) \leq 1$.

Proposition 4. *Let $(R(k), k \geq 0)$ be a random walk with initial value 0 and step distribution $\nu(k) = \mu(k+1), k \geq -1$. Set*

$$M = \inf\{k \geq 0 : R(k) = -1\}.$$

Now suppose that T is a Galton-Watson tree with offspring distribution μ and total progeny N. Then

$$(X(k), 0 \leq k \leq N) \stackrel{d}{=} (R(k), 0 \leq k \leq M).$$

[Careful proof: see Le Gall.]

Galton-Watson forest

It turns out to be technically easier to deal with a sequence of i.i.d. Galton-Watson trees rather than a single tree. We can concatenate their height processes in order to encode the whole Galton-Watson forest.

For the depth-first walks, we retain the relation $X(i+1) = X(i) + c(v_i) - 1$, so that the first tree ends when the walk first hits -1, the second tree ends when we first hit -2 and so on.

It can be checked that we still have

$$H(i) = \#\left\{0 \leq j \leq i-1 : X(j) = \min_{j \leq k \leq i} X(k)\right\}, i \geq 0.$$

Convergence of the depth-first walk

Now specialise to the case where μ is critical and has finite offspring variance $\sigma^2 > 0$.

Then $(X(k), k \geq 0)$ is a random walk with no drift and finite-variance step sizes.

Proposition 5 (Donsker's theorem)**.** *As $n \to \infty$,*

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

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

Convergence of the height process

Theorem 6. *As $n \to \infty$,*

$$\left(\frac{1}{\sqrt{n}} H(\lfloor nt \rfloor), t \geq 0 \right) \xrightarrow{d} \frac{2}{\sigma} \left(|B(t)|, t \geq 0 \right),$$

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

Convergence of the height process: finite-dimensional distributions

Lemma 7. *For any $m \geq 1$ and $0 \leq t_1 \leq t_2 \leq \ldots \leq t_m < \infty$,*

$$\frac{1}{\sqrt{n}} (H(\lfloor nt_1 \rfloor), H(\lfloor nt_2 \rfloor), \ldots, H(\lfloor nt_m \rfloor)) \xrightarrow{d} \frac{2}{\sigma} (|B_{t_1}|, |B_{t_2}|, \ldots, |B_{t_m}|)$$

as $n \to \infty$. ▶

In order to get the functional convergence stated in the theorem, it remains to demonstrate that we have tightness. [Proof: see Le Gall.]

Galton-Watson trees conditioned on their total progeny

Each excursion above 0 of the height process of the Galton-Watson forest corresponds to a tree, and the length of the excursion corresponds to the total progeny of that tree. If we condition the total progeny of the tree to be n, and let $n \to \infty$, intuitively we should obtain something like an excursion of the limit process.

We need to make rigorous sense of what we mean by "an excursion of the limit process" before we can proceed.

A BRIEF INTRODUCTION TO EXCURSION THEORY

Partly based on **A guided tour through excursions** by Chris Rogers.

A tool: Itô's formula

Recall that for $f \in \mathcal{C}^2(\mathbb{R}, \mathbb{R})$,

$$f(B_t) = f(B_0) + \int_0^t f'(B_s) dB_s + \frac{1}{2} \int_0^t f''(B_s) ds.$$

A motivating example

Consider a simple symmetric random walk $(X(n), n \geq 0)$. Let $T_0 = 0$ and, for $n \geq 1$,

$$T_n = \inf\{m > T_{n-1} : X(m) = 0\}.$$

For $n \geq 1$, let

$$\xi^n(k) = \begin{cases} X(T_{n-1} + k) & \text{for } 0 \leq k \leq T_n - T_{n-1} \\ 0 & \text{for } k > T_n - T_{n-1}. \end{cases}$$

Then ξ^n is the nth excursion of X away from 0.

By the Strong Markov property, $\xi^1, \xi^2, \ldots$ are i.i.d.

In other words, the path of the random walk can be cut up into i.i.d. excursions away from 0.

Brownian excursions

Since the path of a Brownian motion $(B_t, t \geq 0)$ is continuous, the set $\{t : B_t \neq 0\}$ is open and so we can express it as a disjoint countable union of maximal open intervals $\cup_{i=1}^{\infty}(g_i, d_i)$ during which B makes an excursion away from 0.

Let $\mathcal{Z} = \{t : B_t = 0\}$. It turns out to be essential to have a measure of how much time B spends at 0. The obvious one doesn't work:

Proposition 8. $\mathrm{Leb}(\mathcal{Z}) = 0$ *a.s.* ▶

The zero set of a Brownian motion

$\mathcal{Z} = \{t : B_t = 0\}$.

- $0 \in \mathcal{Z}$ and $\inf\{t > 0 : B_t = 0\} = 0$, so there is a sequence of points in $\mathcal{Z}$ whose limit is 0.
- By the Strong Markov property, *any* point in $\mathcal{Z}$ is a limit of other points in $\mathcal{Z}$.
- $\mathcal{Z}$ is closed by the continuity of the Brownian path.
- $\mathcal{Z}$ is almost surely uncountable.

Think of $\mathcal{Z}$ as being similar to the Cantor set (only random).

Local time

We want a process $(L_t)_{t \geq 0}$ which increases on $\mathcal{Z}$ and is constant off it.

Definition 9 (Brownian local time)**.**

$$L_t = \lim_{\epsilon \downarrow 0} \frac{1}{2\epsilon} \int_0^t \mathbb{1}_{\{|B_s| \leq \epsilon\}} ds.$$

Note that $(L_t, t \geq 0)$ is clearly increasing.

Why is this the right definition?

Local time

Consider again the simple symmetric random walk on $\mathbb{Z}$, started from 0. For $m \in \mathbb{Z}$, let $\operatorname{sgn}(m) = \begin{cases} 1 & \text{if } m > 0 \\ 0 & \text{if } m = 0 \\ -1 & \text{if } m < 0 \end{cases}$. Then for $n \geq 1$,

$$|X(n)| = \sum_{k=0}^{n-1} \operatorname{sgn}(X(k))(X(k+1) - X(k)) + \sum_{k=0}^{n-1} \mathbb{1}_{\{X(k)=0\}}$$

and so

$$\sum_{k=0}^{n-1} \mathbb{1}_{\{X(k)=0\}} = |X(n)| - \sum_{k=0}^{n-1} \operatorname{sgn}(X(k))(X(k+1) - X(k)).$$

This can be easily understood as an actual measure of how much time the random walk spends at the origin. Now imagine rescaling and using Donsker's theorem. There should be a limiting version of this equation for Brownian motion.

Tanaka's formula

Theorem 10.

$$L_t = |B_t| - \int_0^t \operatorname{sgn}(B_s) dB_s.$$

▶

Note that this entails that $(L_t, t \geq 0)$ is continuous.

Local time measures the time spent at 0

Theorem 11. $(L_t, t \geq 0)$ *almost surely increases only on the set* $\mathcal{Z}$. ▶

So $(L_t, t \geq 0)$ is constant during excursions away from 0.

Excursions

Recall that we can write $\{t : B_t \neq 0\} = \bigcup_{i=1}^{\infty} (g_i, d_i)$. For each i, the excursion is $\xi^i = (B_{(g_i+t)\vee d_i}, t \geq 0)$, which takes values in

$$\begin{aligned} \mathcal{E} = \{ f \in \mathcal{C}([0, \infty), \mathbb{R}) : f(0) = 0, f(t) \neq 0 \text{ for } t \in (0, \zeta), \\ f(t) = 0 \text{ for } t \geq \zeta, \text{ some } 0 < \zeta < \infty \}. \end{aligned}$$

Because the Brownian path oscillates so wildly, there is no first excursion. Although we cannot give a sense to the idea of the first, second etc excursion, they are ordered: (g_i, d_i) comes before (g_j, d_j) if $g_i < g_j$.

The ordering cannot be captured by $\mathbb{N}$, but it turns out that it can be captured by the local time: we can think of the excursion straddling (g_i, d_i) as the excursion at local time ℓ for some ℓ, which occurs before the excursion straddling (g_j, d_j), the excursion at local time $\ell' > \ell$.

A point process of excursions

Let $\tau_\ell = \inf\{t \geq 0 : L_t > \ell\}$. $(\tau_\ell, \ell \geq 0)$ is clearly right-continuous and increasing since $(L_t, t \geq 0)$ is continuous and increasing.

Write δ for the zero excursion (i.e. $\delta(t) = 0$ for all $t \geq 0$) and $\mathcal{E}_\delta = \mathcal{E} \cup \{\delta\}$.

Let $(\Xi_\ell, \ell \geq 0)$ be a $\mathcal{E}_\delta$-valued point process defined as follows:

- if $\tau_\ell - \tau_{\ell-} > 0$ then $\Xi_\ell(t) = B_{(\tau_{\ell-}+t)\vee\tau_\ell}$
- if $\tau_\ell - \tau_{\ell-} = 0$ then $\Xi_\ell = \delta$.

In other words, $\Xi_\ell = \xi$ iff B makes an excursion ξ at local time ℓ.

There are only countably many values of ℓ such that $\Xi_\ell \neq \delta$, but there are infinitely many of them in (a, b) for $0 \leq a < b$.

A Poisson point process of excursions

Theorem 12 (Itô (1970))**.** *Ξ is a Poisson point process with intensity measure* Leb $\times$ n *where* n *is a σ-finite measure on $\mathcal{E}$ called the excursion measure.* ▶

Proposition 13. $n(\{f \in \mathcal{E} : \sup_t f(t) > a\}) = \frac{1}{2a}$. ▶

Proposition 14. $\mathrm{n}(\{f \in \mathcal{E} : \zeta \geq x\}) = \sqrt{\frac{2}{\pi x}}$.

[See Kallenberg **Foundations of modern probability** for a nice proof.]

Scaling property

Recall that Brownian motion has a scaling property:

$$(\lambda^{-1/2} B_{\lambda t}, t \geq 0) \stackrel{d}{=} (B_t, t \geq 0).$$

It turns out that this carries over to its excursions.

Scaling property

Let $\mathcal{E}_x = \{f \in \mathcal{E} : \zeta = x\}$. For $f \in \mathcal{E}$ with duration ζ, put

$$\nu_x(f) = ((x/\zeta)^{1/2} f(\zeta t/x), t \geq 0)$$

Then $\nu_x(f) \in \mathcal{E}_x$.

Proposition 15. *For any $A \subseteq \mathcal{E}_x$,*

$$\mathrm{n}(\nu_x^{-1}(A) | \zeta \geq c) := \frac{\mathrm{n}(\nu_x^{-1}(A) \cap \{\zeta \geq c\})}{\mathrm{n}(\zeta \geq c)}$$

does not depend on $c > 0$. ▶

Scaling property

We can interpret the proposition as saying that the shape of the excursion and its length are "independent".

A little more work shows that we can make sense of $\mathrm{n}^{(x)}(A) := \mathrm{n}(|f| \in A | \zeta = x)$ as a probability measure on $\mathcal{E}_x^+ = \{f \in \mathcal{E}_x : f \geq 0\}$, the law of a process called a Brownian excursion of length x,

$$(e^{(x)}(t), 0 \leq t \leq x).$$

Excursions of different lengths are related via

$$(\sqrt{\lambda} e^{(x)}(t/\lambda), 0 \leq t \leq \lambda x) \stackrel{d}{=} (e^{(\lambda x)}(t), 0 \leq t \leq \lambda x).$$

We refer to $(e^{(1)}(t), 0 \leq t \leq 1)$ as a standard Brownian excursion (and usually omit the superscript in this case).

Standard Brownian excursion, $(e(t), 0 \leq t \leq 1)$

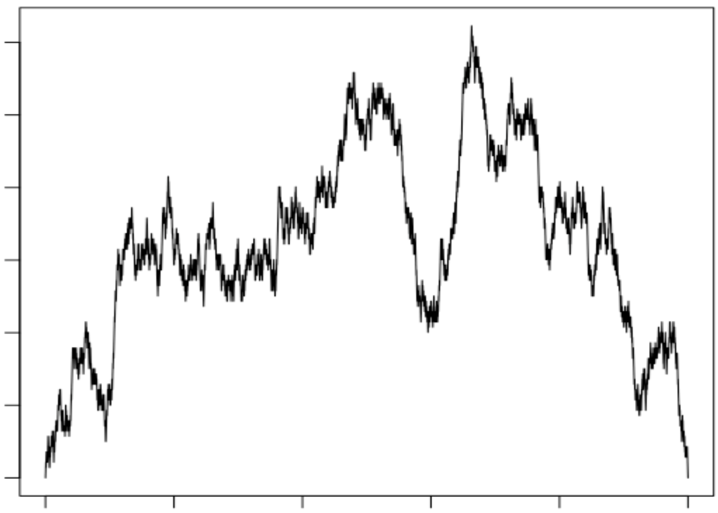

Two-stage description of the excursion process

This also means that we think of the Poisson process of excursions in two steps. For simplicity, we describe the Poisson process which gives $(|B_t|, t \geq 0)$ rather than $(B_t, t \geq 0)$.

- Take a Poisson point process Θ on $[0, \infty) \times [0, \infty)$ of intensity $\mathrm{Leb} \times \mathrm{m}$, where $\mathrm{m}(dx) = \mathrm{n}(\zeta \in dx) = (2\pi)^{-1/2} x^{-3/2} dx$.
- For a point at (ℓ, ζ), sample a standard Brownian excursion e_ℓ. Then $(\sqrt{\zeta} e_\ell(t/\zeta), t \geq 0)$ gives the excursion straddling local time ℓ.

Some loose ends: Galton-Watson trees

Recall that we showed that a Galton-Watson forest can be coded by its depth-first walk and height process. We showed that as $n \to \infty$,

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

and

$$\left(\frac{1}{\sqrt{n}} H(\lfloor nt \rfloor), t \geq 0\right) \xrightarrow{d} \frac{2}{\sigma}\left(|B(t)|, t \geq 0\right)$$

Galton-Watson trees conditioned on their total progeny

Recall that the depth-first walk X of a critical Galton-Watson tree with offspring variance $\sigma^2 > 0$ is a random walk with step mean 0 and variance σ^2. The total progeny N is equal to $\inf\{k \geq 0 : X(k) = -1\}$. Write $(X^n(k), 0 \leq k \leq n)$ for the depth-first walk conditioned on $N = n$. Then there is a conditional version of Donsker's theorem:

Lemma 16. *As $n \to \infty$,*

$$(n^{-1/2} X^n(\lfloor nt \rfloor), 0 \leq t \leq 1) \xrightarrow{d} \sigma(e(t), 0 \leq t \leq 1).$$

[See W.D. Kaigh **An invariance principle for random walk conditioned by a late return to zero** *Annals of Probability* **4** (1976) pp.115-121.]

Convergence of the coding processes

Let $(X^n(i), 0 \leq i \leq n)$ and $(H^n(i), 0 \leq i \leq n)$ be the depth-first walk and height process respectively of a critical Galton-Watson tree with offspring variance $\sigma^2 > 0$, conditioned to have total progeny n.

Theorem 17. *As $n \to \infty$,*

$$(n^{-1/2} X^n(\lfloor n \cdot \rfloor), n^{-1/2} H^n(\lfloor n \cdot \rfloor)) \xrightarrow{d} \left(\sigma e, \frac{2}{\sigma} e\right),$$

where $e = (e(t), 0 \leq t \leq 1)$ is a standard Brownian excursion.

[Proof: see Le Gall.] This result suggests the existence of some sort of limiting tree, which is "coded" by the Brownian excursion.

REAL TREES

Real trees

Definition 18. *A compact metric space $(\mathcal{T}, d)$ is a real tree if for all $x, y \in \mathcal{T}$,*

- *There exists a unique shortest path $[[x, y]]$ from x to y (of length $d(x, y)$). (There is a unique isometric map $f_{x,y}$ from $[0, d(x, y)]$ into $\mathcal{T}$ such that $f(0) = x$ and $f(d(x, y)) = y$. We write $f_{x,y}([0, d(x, y)]) = [[x, y]]$.)*

- *The only non-self-intersecting path from x to y is $[[x, y]]$. (If g is a continuous injective map from $[0, 1]$ into $\mathcal{T}$, such that $g(0) = x$ and $g(1) = y$, then $g([0, 1]) = [[x, y]]$.)*

An element $v \in \mathcal{T}$ is called a vertex. A rooted real tree has a distinguished vertex ρ called the root. The height of a vertex v is its distance $d(\rho, v)$ from the root. A leaf is a vertex v such that $v \notin [[\rho, w]]$ for any $w \neq v$.

Coding real trees

Suppose that $h : [0, \infty) \to [0, \infty)$ is a continuous function of compact support such that $h(0) = 0$. h will play the role of the height process for a real tree.

Coding real trees

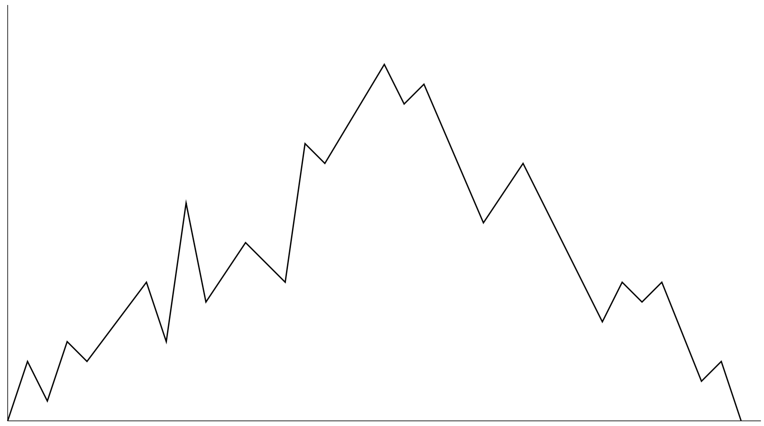

Coding real trees

Use h to define a distance:

$$d_h(x, y) = h(x) + h(y) - 2 \inf_{x \wedge y \leq z \leq x \vee y} h(z).$$

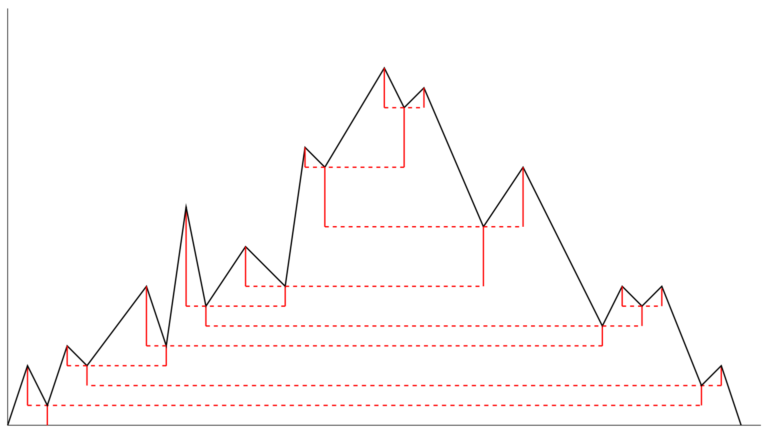

Coding real trees

Let $y \sim y'$ if $d_h(y, y') = 0$ and take the quotient $\mathcal{T}_h = [0, \infty)/\sim$.

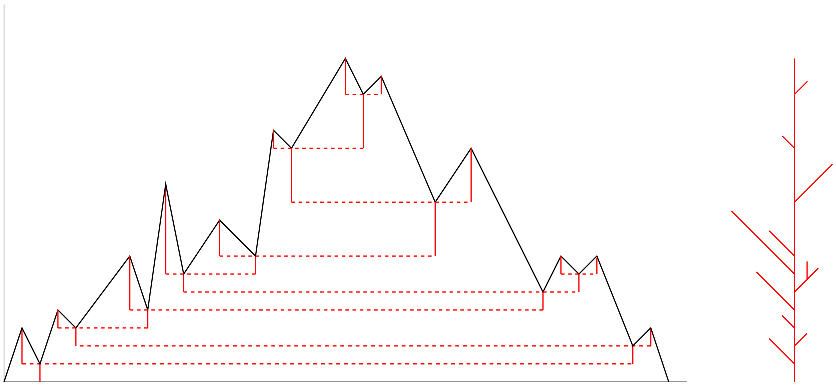

Coding real trees

Theorem 19. *$(\mathcal{T}_h, d_h)$ is a real tree.*

[Proof: see Le Gall.]

We will always take the equivalence class of 0 to be the root, ρ.

Definition 20. *The Brownian continuum random tree is $\mathcal{T}_{2e}$, where e is a standard Brownian excursion.*

The Brownian continuum random tree $\mathcal{T}_{2e}$

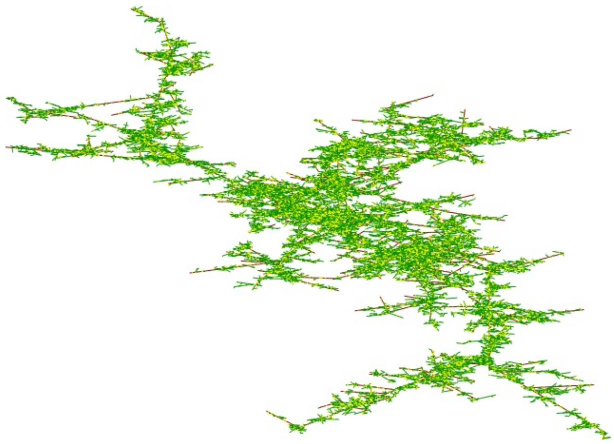

[Picture by Grégory Miermont]

Measuring the distance between metric spaces

The Hausdorff distance between two compact subsets K and K' of a metric space (M, δ) is

$$d_H(K, K') = \inf\{\epsilon > 0 : K \subseteq F_\epsilon(K'), K' \subseteq F_\epsilon(K)\},$$

where $F_\epsilon(K) := \{x \in M : \delta(x, K) \leq \epsilon\}$ is the ϵ-fattening of K.

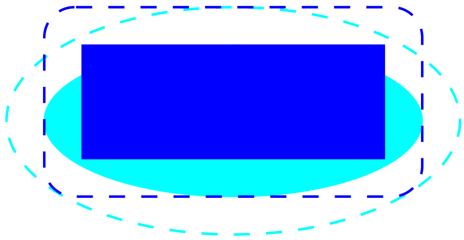

Measuring the distance between metric spaces

To measure the distance between two compact metric spaces (X, d) and (X', d'), the idea is to embed them (isometrically) into a single larger metric space and then compare them using the Hausdorff distance.

So define the Gromov-Hausdorff distance

$$d_{GH}(X, X') = \inf\{d_H(\phi(X), \phi'(X'))\},$$

where the infimum is taken over all choices of metric space (M, δ) and all isometric embeddings $\phi : X \to M$, $\phi' : X' \to M$.

Measuring the distance between metric spaces

If the metric spaces are rooted, at ρ and ρ' respectively, we take

$$d_{GH}(X, X') = \inf\{d_H(\phi(X), \phi'(X')) \vee \delta(\phi(\rho), \phi'(\rho')\}$$

Fortunately, we do not have to seek an optimal embedding!

For compact metric spaces (X, d) and (X', d'), a correspondence between X and X' is a subset $\mathcal{R}$ of $X \times X'$ such that for each $x \in X$, there exists at least one $x' \in X'$ such that $(x, x') \in \mathcal{R}$ and vice versa.

The distortion of a correspondence $\mathcal{R}$ is defined by

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

Measuring the distance between metric spaces

Proposition 21. *If X and X' are compact metric spaces rooted at ρ and ρ' respectively then*

$$d_{GH}(X, X') = \frac{1}{2} \inf \mathrm{dis}(\mathcal{R}),$$

where the infimum is taken over all correspondences $\mathcal{R}$ between X and X' such that $(\rho, \rho') \in \mathcal{R}$.

[Proof: see D. Burago, Y. Burago and S. Ivanov **A course in metric geometry**.]

Convergence to the CRT

Let T_n be our Galton-Watson tree conditioned to have size n.

Theorem 22. *As $n \to \infty$,*

$$\frac{\sigma}{\sqrt{n}} T_n \xrightarrow{d} \mathcal{T}_{2e},$$

where convergence is in the Gromov-Hausdorff sense. ▶

[Approach due to Grégory Miermont.]

The mass measure of the CRT

Consider now a uniform random tree T_n. Put mass $1/n$ at each vertex. Call the resulting probability measure μ_n. It should be intuitively clear that

$$\left(\frac{1}{\sqrt{n}} T_n, \mu_n \right) \xrightarrow{d} (\mathcal{T}_{2e}, \mu),$$

where the probability measure μ is the image of Lebesgue measure on $[0,1]$ on the tree $\mathcal{T}_{2e}$.

Lemma 23. *Let $\mathcal{L}$ be the set of leaves of $\mathcal{T}_{2e}$. Then*

$$\mu(\mathcal{L}) = 1.$$

[Intuition: non-leaf vertices of T_n are typically at distance $o(\sqrt{n})$ from a leaf. Proof: see Aldous (1991).]

RANDOM GRAPHS

The Erdős-Rényi random graph

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

Example: $n = 10$, $p = 0.4$ (vertex labels omitted).

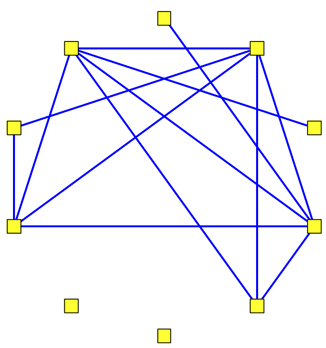

Connected components

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

Below, there are three of them.

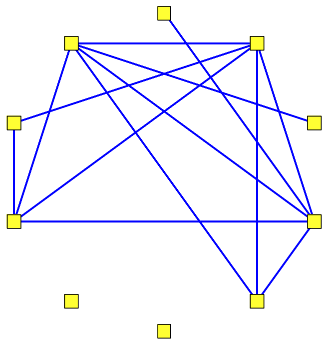

The phase transition

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

$n = 200$, $c = 0.4$

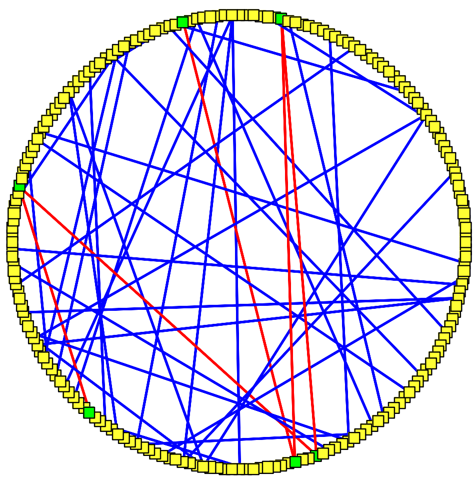

The phase transition

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

$n = 200$, $c = 0.8$

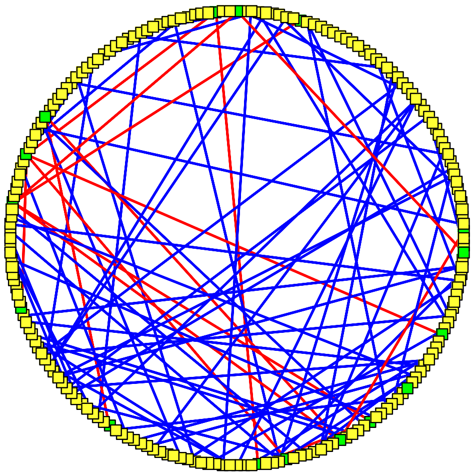

The phase transition

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

$n = 200$, $c = 1.2$

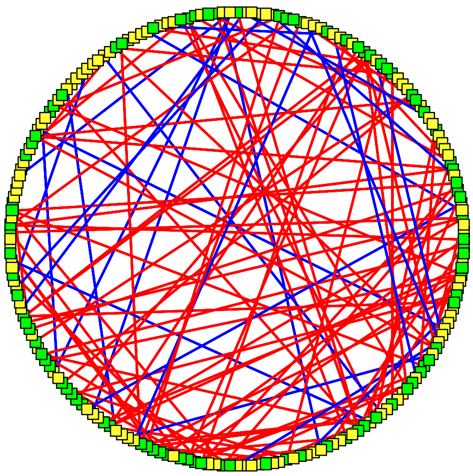

The phase transition (Erdős and Rényi (1960))

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

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 $n \to \infty$.]

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 $\Theta(n^{2/3})$.

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:

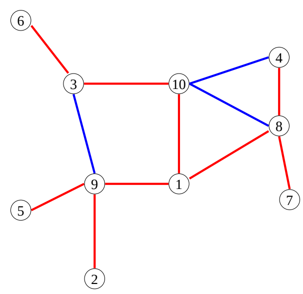

Convergence of the sizes and surpluses

Fix λ and let $C_1^n, C_2^n, \ldots$ be the sequence of component sizes in decreasing order, and let $S_1^n, S_2^n, \ldots$ be their surpluses.

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

Theorem 24 (Aldous (1997))**.** *As $n \to \infty$,*

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

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.

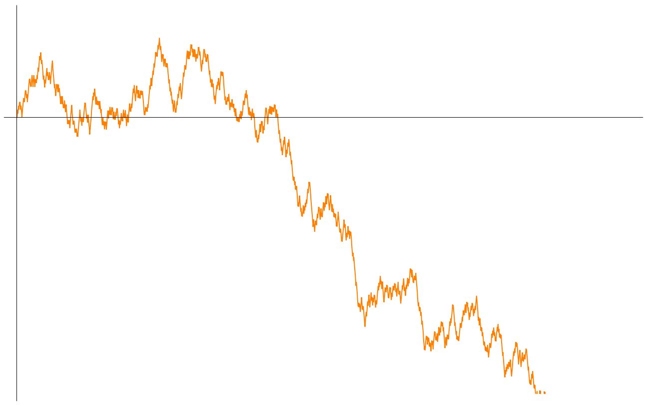

Limiting sizes and surpluses

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.

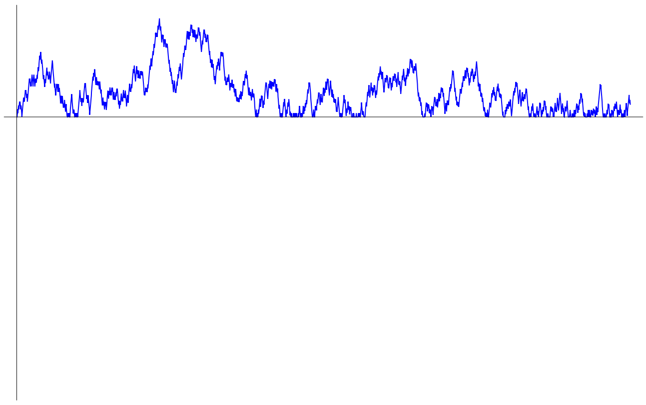

Limiting sizes and surpluses

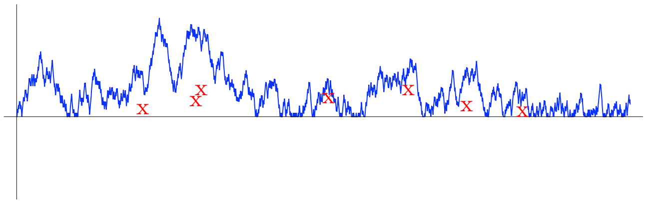

Decorate the picture with the points of a rate one Poisson process which fall above the x-axis and below the graph.

$\mathbf{C}$ is the sequence of excursion-lengths of this process, in decreasing order.

$\mathbf{S}$ is the sequence of numbers of points falling in the corresponding excursions.

Convergence of the sizes and surpluses

Theorem 25 (Aldous (1997)). *As $n \to \infty$,*

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

where $\mathbf{C}$ *is the sequence of excursion-lengths of* B^λ *in decreasing order, and* $\mathbf{S}$ *is the sequence of numbers of Poisson points falling in the corresponding excursions.*

Here, convergence in the first co-ordinate takes place in

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

Proof technique: depth-first exploration

As for the discrete trees at the beginning of the course, a key tool is a depth-first exploration. We previously defined the depth-first walk by $X(0) = 0$ and, for $1 \leq k \leq n$,

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

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

There are two problems with this definition: the components of a random graph are are labelled but not ordered, and they are not trees.

Depth-first exploration

These problems are resolved by stepping through the graph vertex by vertex, using the natural ordering of the labels, and ignoring non-tree edges. Exactly how we do this is best explained on an example.

It's useful to say that vertices can have four states: current, alive, dead or unexplored. For the first component, $X(k)$ will turn out to be the number of alive vertices at step k. Thereafter, it will be the number of vertices alive minus the number of components already fully explored.

Depth-first exploration: an example

Step 0

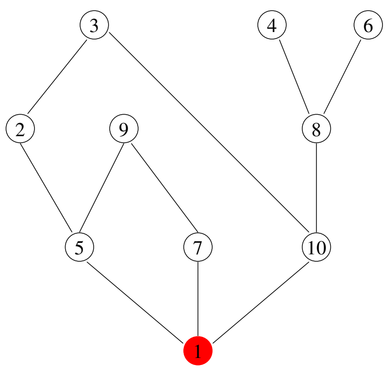

Current: 1   Alive: none   Dead: none $X(0) = 0$.

Depth-first exploration: an example

Step 1

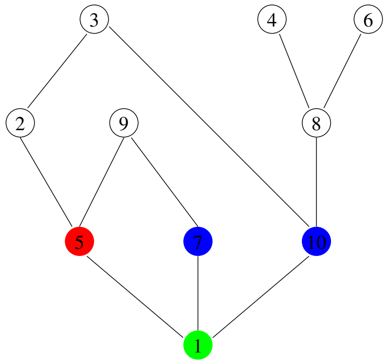

Current: 5   Alive: 7, 10   Dead: 1 $X(1) = 2$.

Depth-first exploration: an example

Step 2

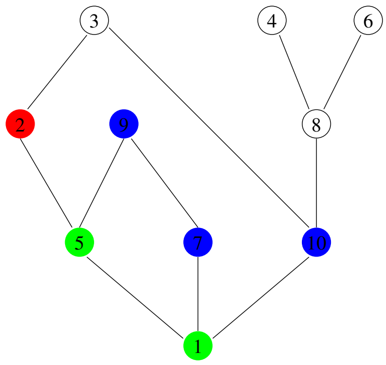

Current: 2   Alive: 9, 7, 10   Dead: 1, 5 $X(2) = 3$.

Depth-first exploration: an example

Step 3

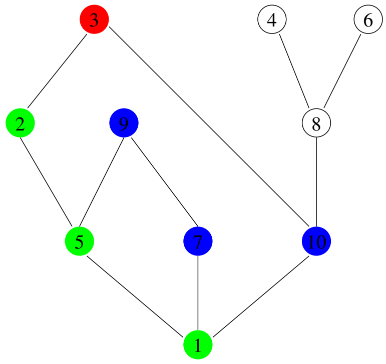

Current: 3 Alive: 9, 7, 10 Dead: 1, 5, 2 $X(3) = 3$.

Depth-first exploration: an example

Step 4

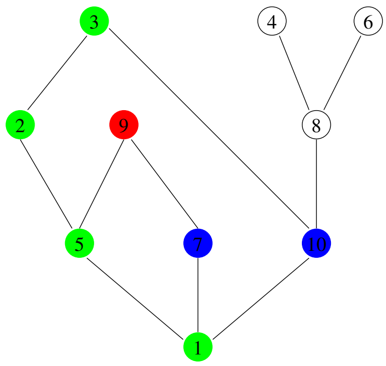

Current: 9 Alive: 7, 10 Dead: 1, 5, 2, 3 $X(4) = 2$.

Depth-first exploration: an example

Step 5

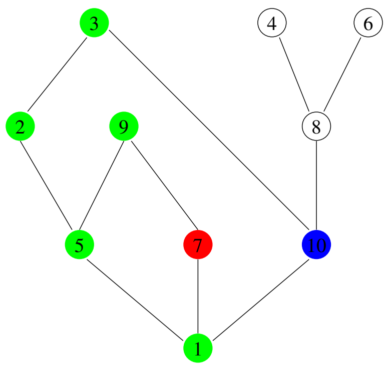

Current: 7   Alive: 10   Dead: 1, 5, 2, 3, 9 $X(5) = 1.$

Depth-first exploration: an example

Step 6

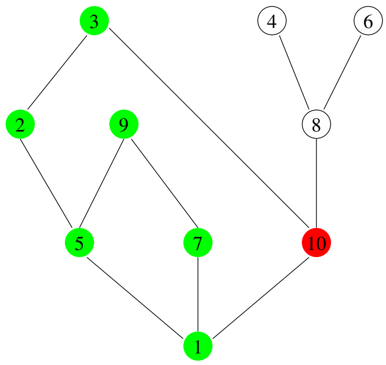

Current: 10   Alive: none   Dead: 1, 5, 2, 3, 9, 7 $X(6) = 0.$

Depth-first exploration: an example

Step 7

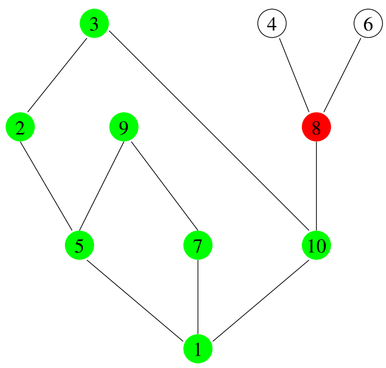

Current: 8   Alive: none   Dead: 1, 5, 2, 3, 9, 7, 10 $X(7) = 0.$

Depth-first exploration: an example

Step 8

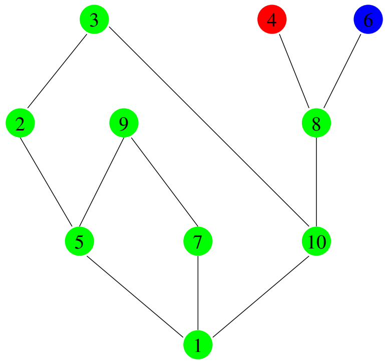

Current: 4 Alive: 6 Dead: 1, 5, 2, 3, 9, 7, 10, 8 $X(8) = 1$.

Depth-first exploration: an example

Step 9

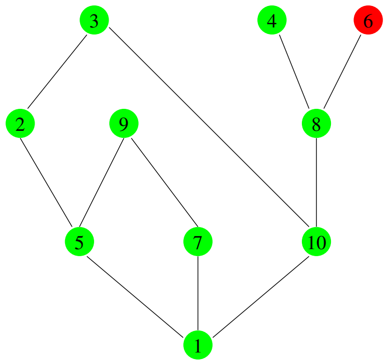

Current: 6 Alive: none Dead: 1, 5, 2, 3, 9, 7, 10, 8, 4 $X(9) = 0$.

Depth-first exploration: an example

We explored the graph on the left as if it were the tree on the right:

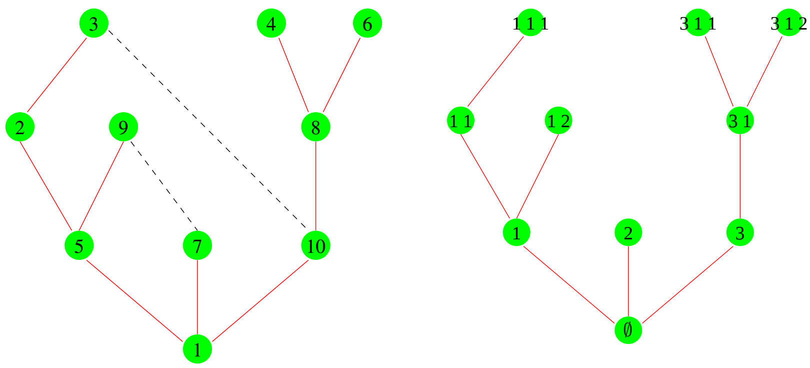

Depth-first walk

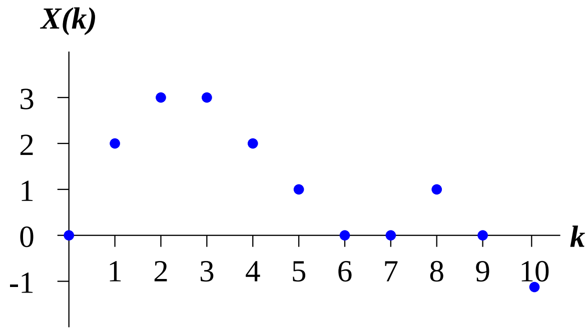

As for a forest, if there are several components, $T(k) = \inf\{i \geq 0 : X(i) = -k\}$ marks the beginning of the $(k+1)$th component. So the component sizes are $\{T(k+1) - T(k), k \geq 0\}$. This sequence can clearly be reconstructed from the path of $(X(i), i \geq 0)$.

Convergence of the depth-first walk

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

Theorem 26. *As $n \to \infty$,*

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

▶

The convergence here is uniform on compact time-intervals.

To finish the proof...

A little care needs to be taken to check that the lengths of excursions above past-minima of X_n^λ converge to lengths of excursions above past-minima of W^λ, and that we don't miss any excursions of length $\Omega(n^{2/3})$. [Proof: see Aldous (1997).]

We will deal with the surplus edges a little later.

Question

What do the limiting components look like?

The vertex-labels are irrelevant: we are really interested in what 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

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 general approach is to pick out a spanning tree, and then 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.

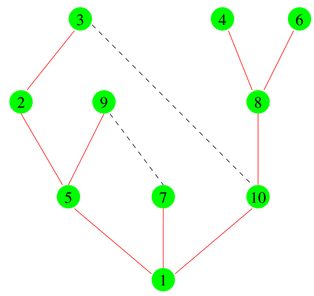

Call it the depth-first tree associated with the graph G, and write $T(G)$.

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.

Overview: the limit of the random graph

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 $n^{-1/3}$.

Each excursion of the process $(B^\lambda(t), t \geq 0)$ of length x corresponds to the limit of a component on $\sim xn^{2/3}$ vertices. Such an excursion codes a continuum random tree, which is a "spanning tree" for that limit component. These CRT's are *not* rescaled Brownian CRT's, but CRT's whose distribution has been "tilted" in a way which we will make precise in a moment.

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

Excursions of the limit process

Consider the process $(B^\lambda(t), t \geq 0)$.

Lemma 27. *An excursion $\tilde{e}^{(x)}$ of $(B^\lambda(t), t \geq 0)$, conditioned to have length x, has a distribution specified by*

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

where f is any suitable test-function and $e^{(x)}$ is a Brownian excursion of length x. ▶

Note that this holds independently of λ. We refer to $\tilde{e}^{(x)}$ as a tilted excursion and to the tree $\tilde{\mathcal{T}}$ that it encodes as a tilted tree.

Vertex identifications

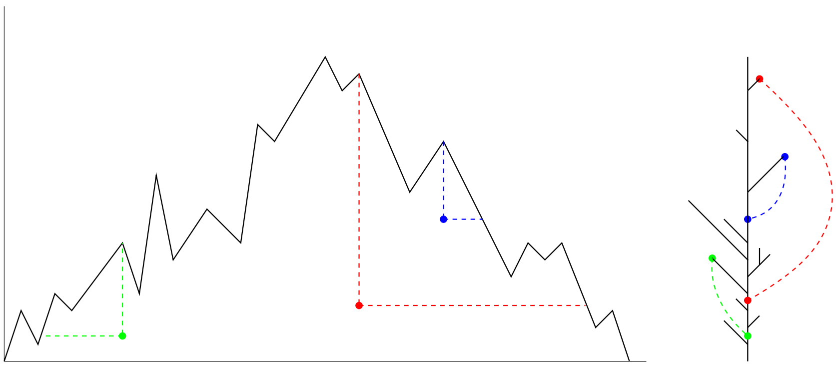

A point at (x, y) identifies the vertex v at height $h(x)$ with the vertex at distance y along the path from the root to v.

A limiting component

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 $\tilde{\mathcal{T}}$.

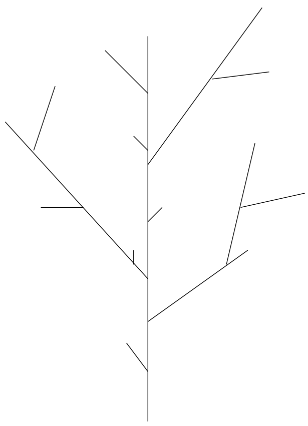

Conditional on $\tilde{e}^{(x)}$, sample a random variable P with Poisson $\left(\int_0^x \tilde{e}^{(x)}(u)du\right)$ distribution.

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.)

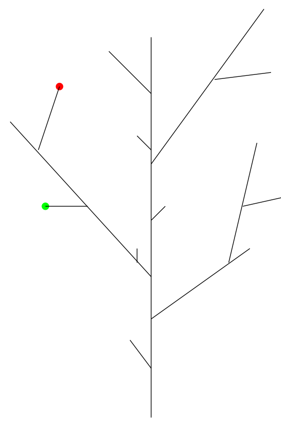

A limiting component

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

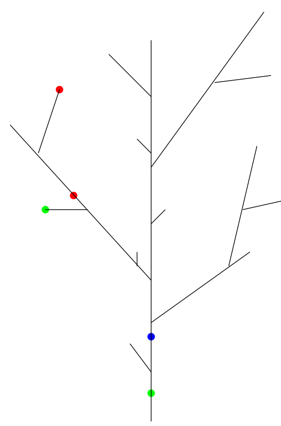

A limiting component

Identify each of the selected leaves with its chosen point.

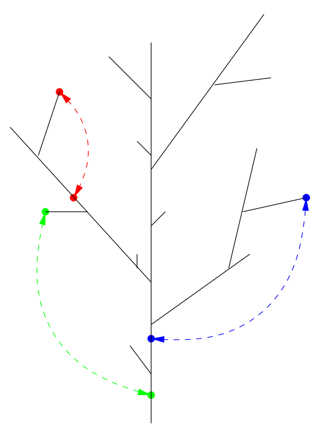

Convergence result

Let $\mathcal{C}_1^n, \mathcal{C}_2^n, \ldots$ be the sequence of components of $G(n,p)$ in decreasing order of size, considered as metric spaces with the graph distance.

Theorem 28. *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 the marked limit process B^λ in decreasing order of length.

Here, convergence is with respect to the metric
$$d(\mathcal{A}, \mathcal{B}) := \left(\sum_{i=1}^{\infty} d_{GH}(\mathcal{A}_i, \mathcal{B}_i)^4 \right)^{1/4}.$$

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 $G(n,p)$. Recall that $T(G)$ is the depth-first tree associated with G

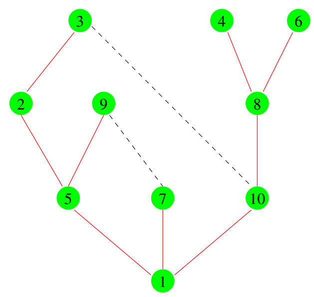

and that $(X(k), 0 \le k \le n)$ is the depth-first walk of $T(G)$.

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

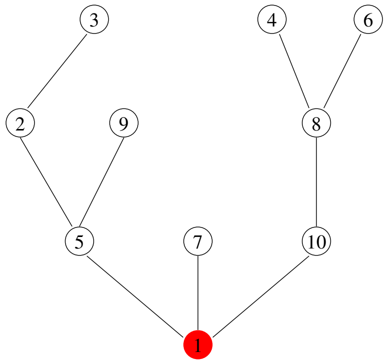

Step 0: $X(0) = 0$.

Depth-first walk and permitted edges

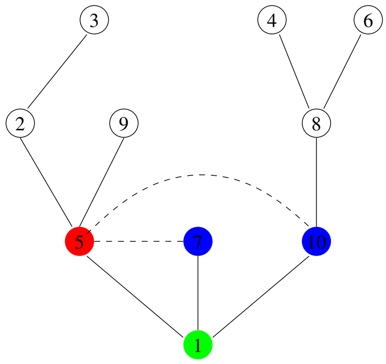

Step 1: $X(1) = 2$.

Depth-first walk and permitted edges

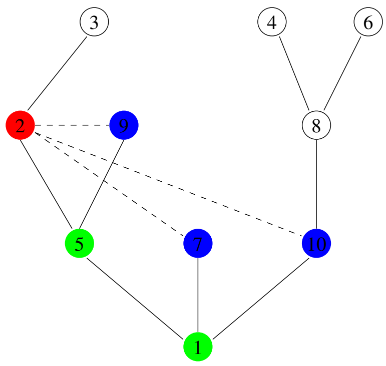

Step 2: $X(2) = 3$.

Depth-first walk and permitted edges

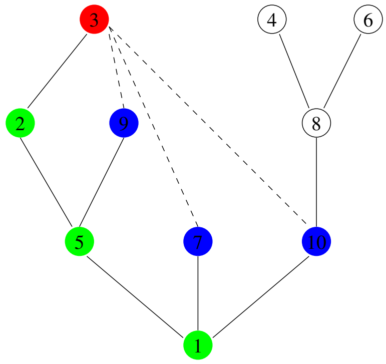

Step 3: $X(3) = 3$.

Depth-first walk and permitted edges

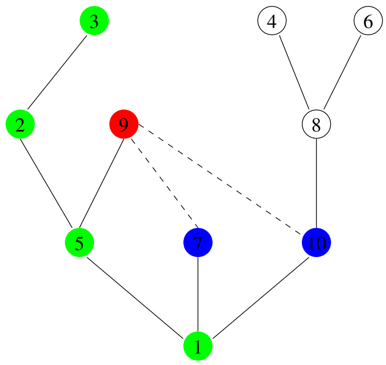

Step 4: $X(4) = 2$.

Depth-first walk and permitted edges

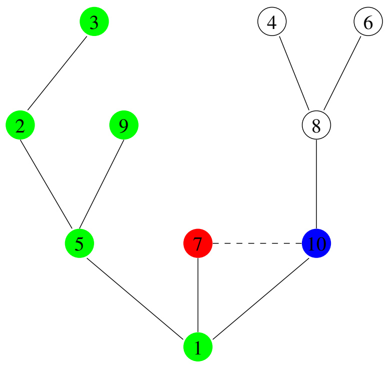

Step 5: $X(5) = 1$.

Depth-first walk and permitted edges

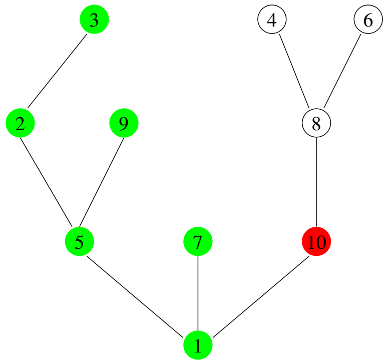

Step 6: $X(6) = 0$.

Depth-first walk and permitted edges

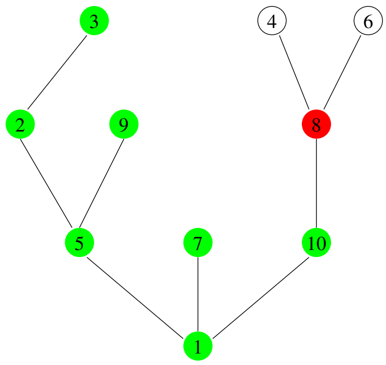

Step 7: $X(7) = 0$.

Depth-first walk and permitted edges

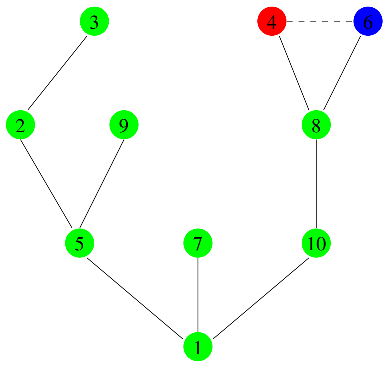

Step 8: $X(8) = 1$.

Depth-first walk and permitted edges

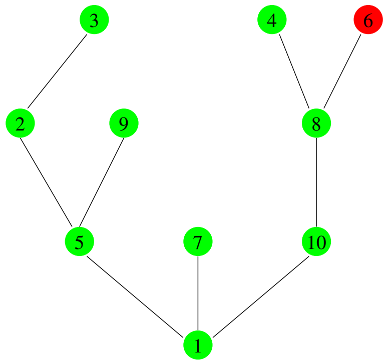

Step 10: $X(9) = 0$.

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.

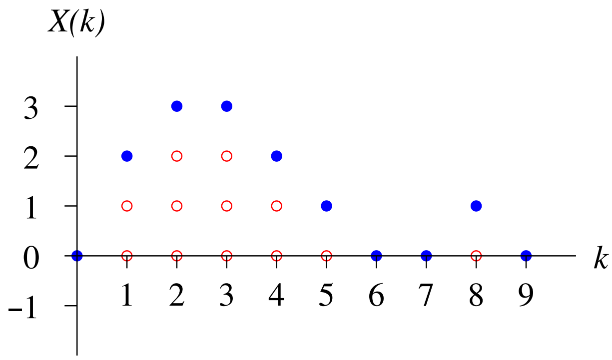

Classifying graphs by depth-first tree

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.

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 $\tilde{G}^p_m$ as follows.

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

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

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

Lemma 29. *$\tilde{G}^p_m$ has the same distribution as G^p_m, a component of $G(n,p)$ conditioned to have vertex-set $[m]$.* ▶

Taking limits

So we need to prove that

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

We will deal with the tree first. For simplicity, we will take $p = m^{-3/2}$; the general case is similar.

Convergence of the tree

Theorem 30. *Suppose $p = m^{-3/2}$. Then*

$$\frac{1}{\sqrt{m}}\tilde{T}^p_m \xrightarrow{d} \tilde{\mathcal{T}}$$

as $m \to \infty$. ▶

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.

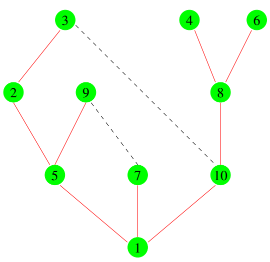

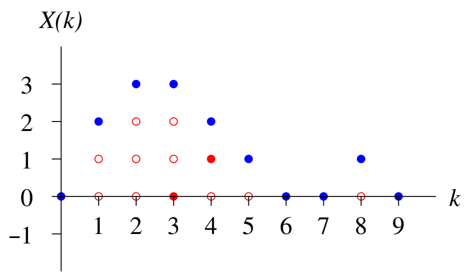

Surplus edges

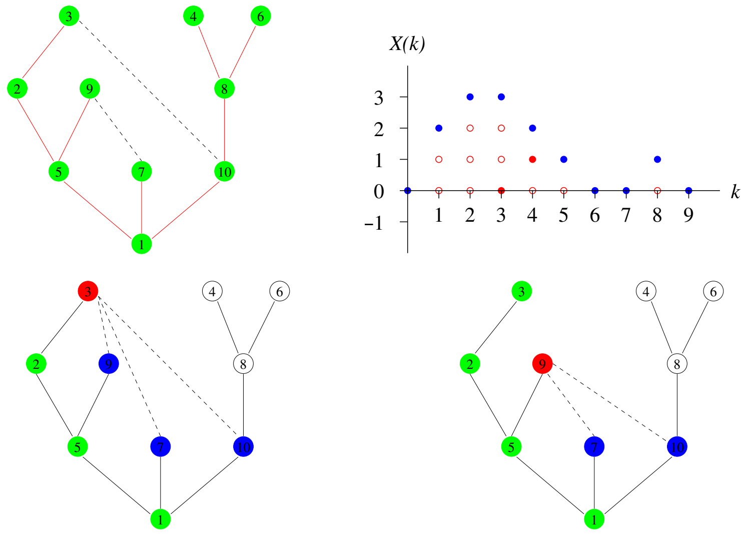

A point at (k, j) means "put an edge between the current vertex at step k and the vertex at distance j from the bottom of the list of alive vertices".

Surplus edges

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

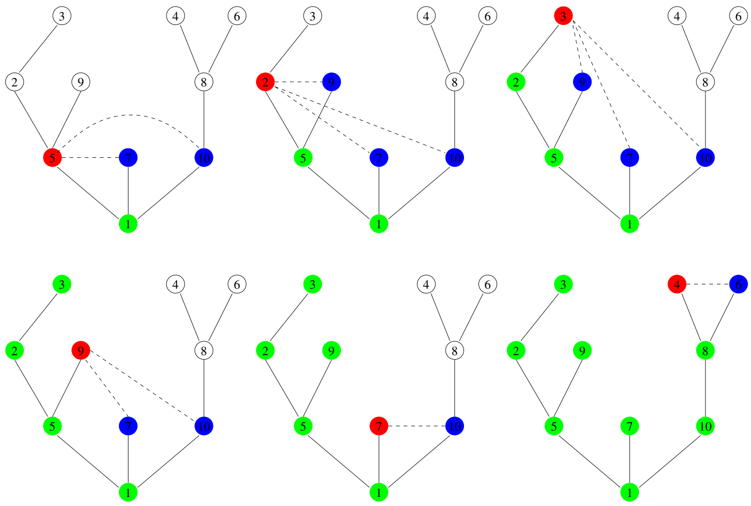

Surplus edges

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.

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.)

The difference between the depth-first walk and the height process is also small, and so the locations of the surplus "edges" are essentially as described in our limit process.

Further reading: limits of 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).

Further reading: excursion theory

A guided tour through excursions L.G.C. Rogers, *Bulletin of the London Mathematical Society* **21** (1989) pp.305–341.

Continuous martingales and Brownian motion D. Revuz and M. Yor, 3rd edition (1999), Springer.

Diffusions, Markov processes and martingales L.G.C. Rogers and D. Williams, 2nd edition (2000), Cambridge University Press.

Further reading: random graphs

Random graphs S. Janson, T. Łuczak and A. Ruciński (2000), Wiley.

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, `arXiv:0903.4730 [math.PR]`.

Critical random graphs: limiting constructions and distributional properties L. Addario-Berry, N. Broutin and C. Goldschmidt, `arXiv:0908.3629 [math.PR]`.