Interacting Particle Systems: Almost sure uniqueness, pathwise duality, and the mean-field limit

Jan M. Swart

Lecture 3: The mean-field limit

A general set-up

Let $(\Lambda, \sim)$ be a countable graph.

For each $k \geq 1$, let $\mathcal{K}^k$ be the set of all words $i_1 \cdots i_k$, made from the alphabet Λ, such that:

$$\text{(i) } i_n \neq i_m \ \forall n \neq m \quad \text{(ii) } i_1 \sim i_2 \sim \cdots \sim i_k.$$

For each $g : S^k \to S$ and $i_1 \cdots i_k \in \mathcal{K}^k$, define $g_{i_1 \cdots i_k} : S^\Lambda \to S^\Lambda$ by

$$g_{i_1 \cdots i_k}(x)(j) := \begin{cases} g\big(x(i_1), \ldots, x(i_k)\big) & \text{if } j = i_1, \\ x(j) & \text{otherwise.} \end{cases}$$

For $g : S^0 \to S$ and $i_1 \in \Lambda$, define $g_{i_1} : S^\Lambda \to S$ by

$$g_{i_1}(x)(j) := \begin{cases} g(\emptyset) & \text{if } j = i_1, \\ x(j) & \text{otherwise.} \end{cases}$$

Examples

Let $S := \{0,1\}$ and define $\mathtt{dth} : S^0 \to S$, $\mathtt{bra} : S^2 \to S$, and $\mathtt{cob} : S^3 \to S$ by

$$\begin{aligned} \mathtt{dth}(\emptyset) &:= 0, \\ \mathtt{bra}\big(x(1),x(2)\big) &:= x(1) \vee x(2), \\ \mathtt{cob}\big(x(1),x(2),x(3)\big) &:= x(1) \vee \big(x(2) \wedge x(3)\big). \end{aligned}$$

Then $\mathtt{dth}_{i_1}$, $\mathtt{bra}_{i_1 i_2}$, and $\mathtt{cob}_{i_1 i_2 i_3}$ are the local maps defined in the first lecture.

A general set-up

(i) Polish space S *local state space.*

(ii) $(\Omega, \mathcal{B}, \mathbf{r})$ Polish space with Borel σ-field and finite measure: *source of external randomness.*

(iii) $\kappa : \Omega \to \mathbb{N}$ measurable function.

(iv) For each $\omega \in \Omega$, a measurable function $\gamma[\omega] : S^{\kappa(\omega)} \to S$.

We are interested in the interacting particle system that evolves as follows:

- We activate each site i with Poisson rate $|\mathbf{r}| := \mathbf{r}(\omega)$.
- We choose ω according to the law $|\mathbf{r}|^{-1}\,\mathbf{r}$.
- We uniformly choose $i = i_1 \sim \cdots \sim i_{\kappa(\omega)\vee 1}$, all different, if this is possible.
- We apply the map $\gamma_{i_1 \cdots i_{\kappa(\omega)\vee 1}}[\omega]$.

Examples

Let $S := \{0,1\}$ and $\Omega = \{1,2\}$. Then setting

$$\begin{aligned} \kappa(1) &:= 0, & \gamma[1] &:= \texttt{dth}, & \mathbf{r}(\{1\}) &:= 1, \\ \kappa(2) &:= 2, & \gamma[2] &:= \texttt{bra}, & \mathbf{r}(\{2\}) &:= \lambda, \end{aligned}$$

yields the contact process with infection rate λ. Similarly, setting

$$\begin{aligned} \kappa(1) &:= 0, & \gamma[1] &:= \texttt{dth}, & \mathbf{r}(\{1\}) &:= 1, \\ \kappa(2) &:= 3, & \gamma[2] &:= \texttt{cob}, & \mathbf{r}(\{2\}) &:= \lambda, \end{aligned}$$

yields the cooperative contact process with cooperative branching rate λ.

A general set-up

Our summability condition now reduces to

$$\int_\Omega \mathbf{r}(\mathrm{d}\omega)\,\kappa(\omega) < \infty \qquad \text{(SUM)}.$$

Let:

$(\mathbb{X}_{s,u})_{s\leq u}$	stochastic flow constructed from a Poisson set π,
X_0	S^Λ-valued random variable, independent of π.

Assuming (SUM), the process $(X_t)_{t\geq 0}$ defined as

$$X_t := \mathbb{X}_{s,s+t}(X_0) \quad (t \geq 0)$$

is the interacting particle system with generator G.

The mean-field limit

Let $(\Lambda_N, \sim)$ be the complete graph with N vertices.
Let $(X_t^N)_{t\geq 0}$ be the particle system on Λ_N with generator G.
We are interested in the *empirical measure*

$$\mu_t^N := \frac{1}{N}\sum_{i=1}^{N} \delta_{X_t^N(i)} \qquad (t \geq 0).$$

We will prove that in the limit $N \to \infty$, the process $(\mu_t^N)_{t\geq 0}$ solves the *mean-field equation*

$$\frac{\partial}{\partial t}\mu_t = \int_\Omega \mathbf{r}(\mathrm{d}\omega)\big\{\mathbf{T}_{\gamma[\omega]}(\mu_t) - \mu_t\big\} \qquad (t \geq 0) \qquad \text{(MEAN)},$$

where for any measurable $g : S^k \to S$, we define
$\mathbf{T}_g : \mathcal{P}(S) \to \mathcal{P}(S)$ by

$$\mathbf{T}_g(\mu) := \text{ the law of } g(X_1, \ldots, X_k),$$

where $(X_i)_{i\geq 1}$ are i.i.d. with law μ.

The mean-field contact process

For the contact process, the mean-field equation takes the form

$$\frac{\partial}{\partial t}\mu_t = \lambda\big\{\mathsf{T}_{\mathtt{bra}}(\mu_t) - \mu_t\big\} + \big\{\mathsf{T}_{\mathtt{dth}}(\mu_t) - \mu_t\big\}.$$

Rewriting this in terms of $p_t := \mu_t(\{1\})$ yields

$$\frac{\partial}{\partial t}p_t = \lambda p_t(1 - p_t) - p_t =: F_\lambda(p_t) \qquad (t \geq 0).$$

The mean-field contact process

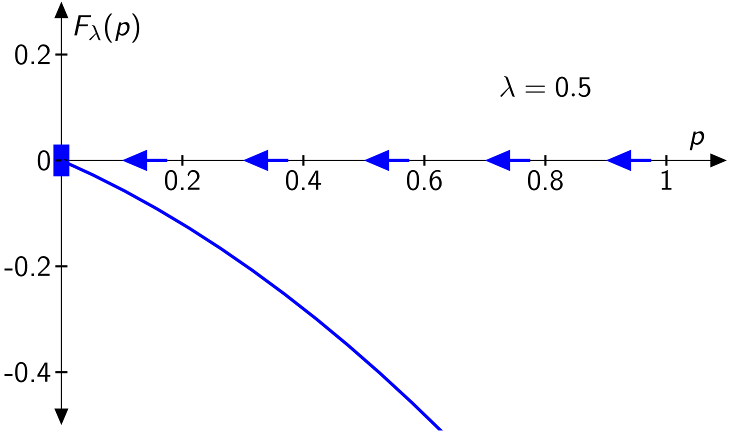

For $\lambda \leq 1$, the equation $\frac{\partial}{\partial t} p_t = F_\lambda(p_t)$ has a single, stable fixed point $p = 0$.

The mean-field contact process

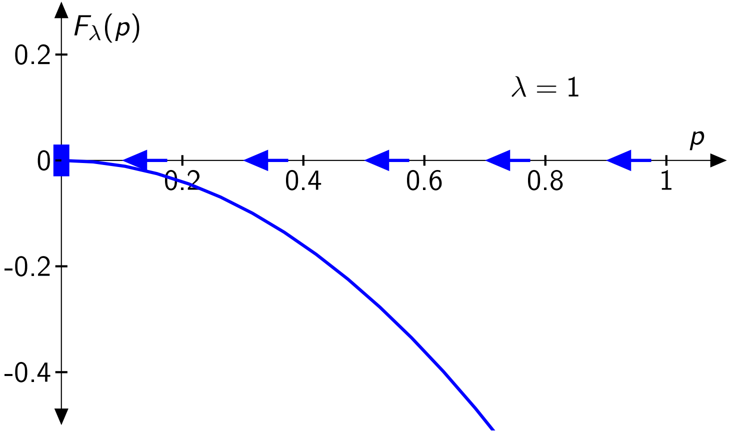

For $\lambda \leq 1$, the equation $\frac{\partial}{\partial t} p_t = F_\lambda(p_t)$ has a single, stable fixed point $p = 0$.

The mean-field contact process

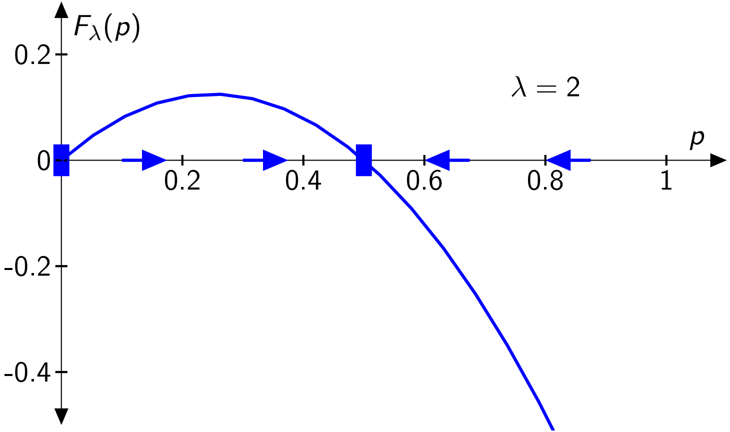

For $\lambda > 1$, the fixed point $p = 0$ becomes unstable and a new stable fixed point $p(\lambda) = 1 - 1/\lambda$ appears.

The mean-field contact process

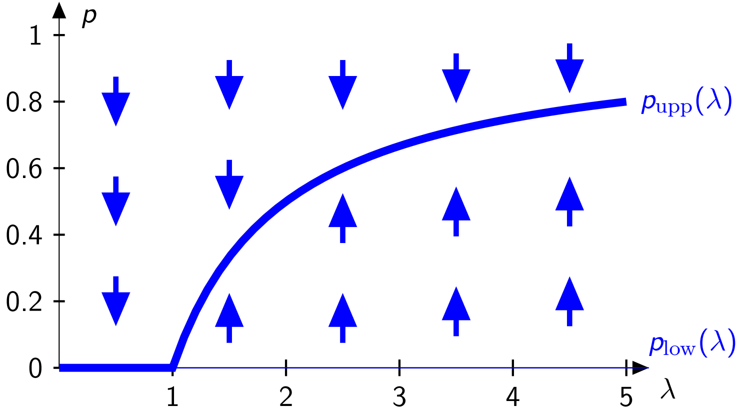

Fixed points and their domains of attraction as a function of λ.

The mean-field cooperative contact process

For the cooperative contact process, the mean-field equation takes the form

$$\frac{\partial}{\partial t}\mu_t = \lambda\big\{\mathbf{T}_{\mathtt{cob}}(\mu_t) - \mu_t\big\} + \big\{\mathbf{T}_{\mathtt{dth}}(\mu_t) - \mu_t\big\}.$$

Rewriting this in terms of $p_t := \mu_t(\{1\})$ yields

$$\frac{\partial}{\partial t}p_t = \lambda p_t^2(1 - p_t) - p_t =: F_\lambda(p_t) \qquad (t \geq 0).$$

The mean-field cooperative contact process

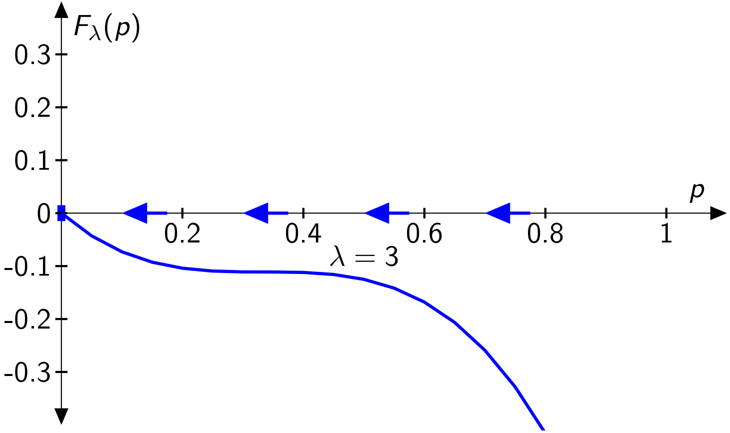

For $\lambda < 4$, the equation $\frac{\partial}{\partial t} p_t = F_\lambda(p_t)$ has a single, stable fixed point $p = 0$.

The mean-field cooperative contact process

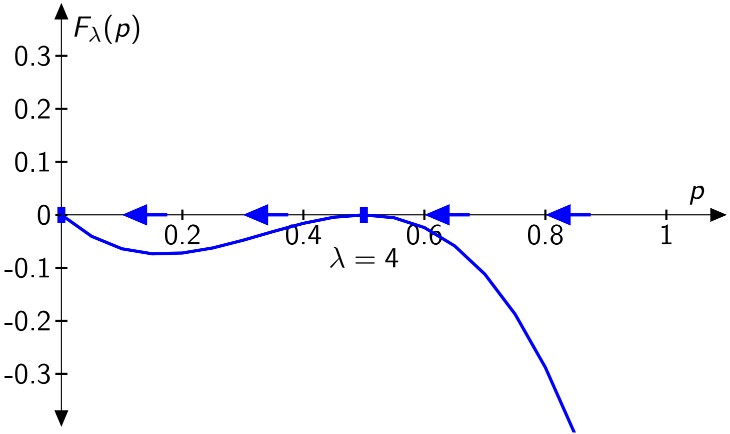

For $\lambda = 4$, a second fixed point appears at $p = 0.5$.

The mean-field cooperative contact process

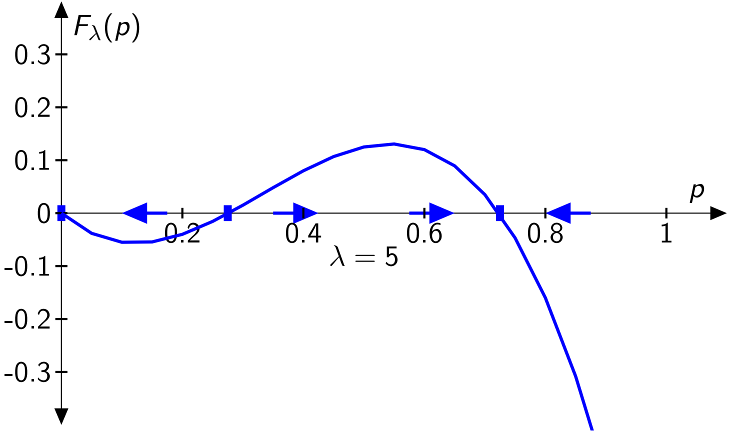

For $\lambda > 4$, there are two stable fixed points and one unstable fixed point, which separates the domains of attraction of the other two.

The mean-field cooperative contact process

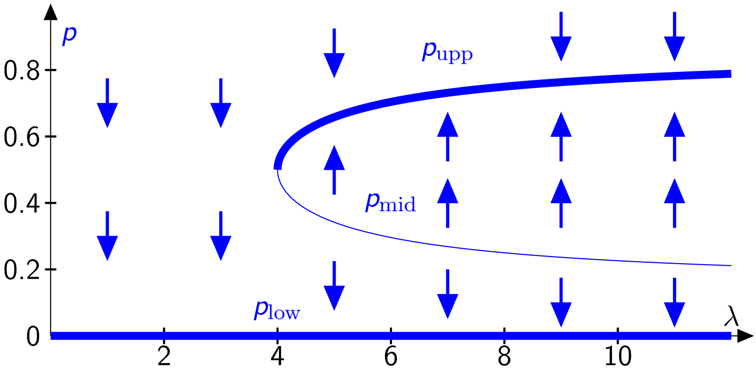

Fixed points of $\frac{\partial}{\partial t} p_t = F_\lambda(p_t)$ for different values of λ.

Using the notation $|\mathbf{r}| := \mathbf{r}(\omega)$ and

$$\mathbf{T}(\mu) := |\mathbf{r}|^{-1} \int_\Omega \mathbf{r}(\mathrm{d}\omega) \mathbf{T}_{\gamma[\omega]}(\mu),$$

we can rewrite the mean-field equation as

$$\tfrac{\partial}{\partial t}\mu_t = |\mathbf{r}|\big\{\mathbf{T}(\mu_t) - \mu_t\big\} \qquad (t \geq 0) \qquad \text{(MEAN)}.$$

Recall

$$\int_\Omega \mathbf{r}(\mathrm{d}\omega)\, \kappa(\omega) < \infty \qquad \text{(SUM)}.$$

Theorem [Mach, Sturm, S. '20] Under the condition (SUM), the mean-field equation (MEAN) has a unique solution for each initial state $\mu_0 \in \mathcal{P}(S)$.

The abstract setting

We define a (nonlinear) semigroup $(\mathbf{T}_t)_{t\geq 0}$ of operators acting on probability measures by

$$\mathbf{T}_t(\mu) := \mu_t \quad \text{where } (\mu_t)_{t\geq 0} \text{ solves (MEAN) with } \mu_0 = \mu.$$

This is a sort of continuous-time version of the discrete evolution $\mu \mapsto \mathbf{T}(\mu) \mapsto \mathbf{T}^2(\mu) \mapsto \cdots$.

Assuming that, for all $k \geq 0$ and $x \in S^k$,

$$\mathbf{r}\big(\{\omega : \kappa(\omega) = k,\ \gamma[\omega] \text{ is discontinuous at x}\}\big) = 0 \qquad \text{(CONT)},$$

one can show that the operators $\mathbf{T}^n$ and $\mathbf{T}_t$ are continuous w.r.t. weak convergence.

The abstract setting

Let d be any metric that generates the topology of weak convergence and let $\|\cdot\|$ denote the total variation norm.

Theorem [Mach, Sturm, S. '20] Assume (SUM) and at least one of the following conditions:

(i) $\mathbb{P}\big[d(\mu_0^N,\mu_0)\geq\varepsilon\big]\underset{N\to\infty}{\longrightarrow}0$ for all $\varepsilon>0$, and (CONT) holds.

(ii) $\big\|\mathbb{E}[(\mu_0^N)^{\otimes n}]-\mu_0^{\otimes n}\big\|\underset{N\to\infty}{\longrightarrow}0$ for all $n\geq 1$.

Then

$$\mathbb{P}\Big[\sup_{0\leq t\leq T}d\big(\mu_t^N,\mathbf{T}_t(\mu_0)\big)\geq\varepsilon\Big]\underset{N\to\infty}{\longrightarrow}0\qquad(\varepsilon>0,\ T<\infty).$$

Proof By a probabilistic representation of the semigroup $(\mathbf{T}_t)_{t\geq 0}$.

The backward in time process

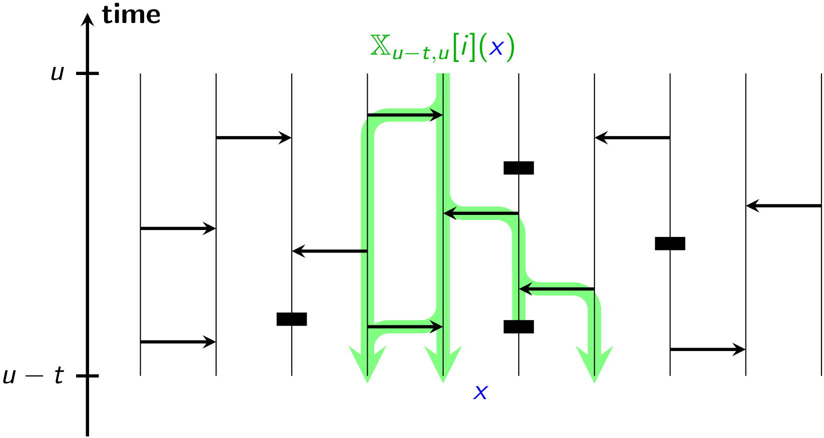

Recall the "backward in time" Markov process

$$\big(\mathcal{R}(\mathbb{X}_{u-t,u}[i]), \mathbb{X}_{u-t,u}[i]\big)_{t\geq 0}.$$

The backward in time process

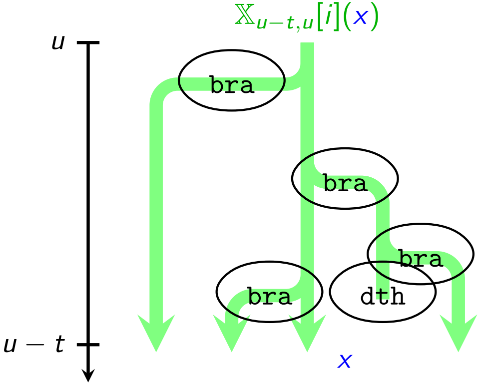

In the mean-field limit, the "backward in time" process becomes a random tree with maps attached to its nodes.

Aim Develop a stochastic representation for the nonlinear semigroup $(\mathbf{T}_t)_{t\geq 0}$, and also for the discrete-time evolution maps $\mathbf{T}^n$, in terms of a random tree with maps attached to its nodes.

Aldous & Banyopadyay (2005) (discrete time),
Mach, Sturm & S. (2020) (continuous time).

Fix $d \in \mathbb{N}_+ \cup \{\infty\}$ such that $\kappa(\omega) \leq d$ for all $\omega \in \Omega$. Let $\mathbb{T} = \mathbb{T}^d$ denote the space of all words $\mathbf{i} = i_1 \cdots i_n$ made from the alphabet $\{1, \ldots, d\}$ (if $d < \infty$) resp. $\mathbb{N}_+$ (if $d = \infty$).

A recursive tree representation

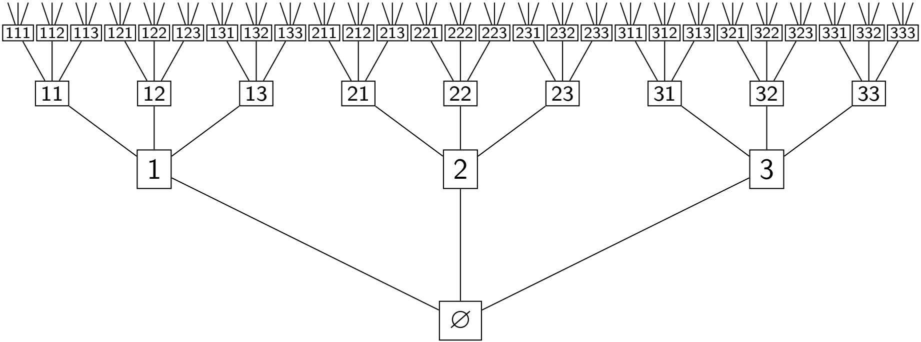

We view $\mathbb{T} = \mathbb{T}^d$ as a tree with root $\varnothing$, the word of length zero.

A recursive tree representation

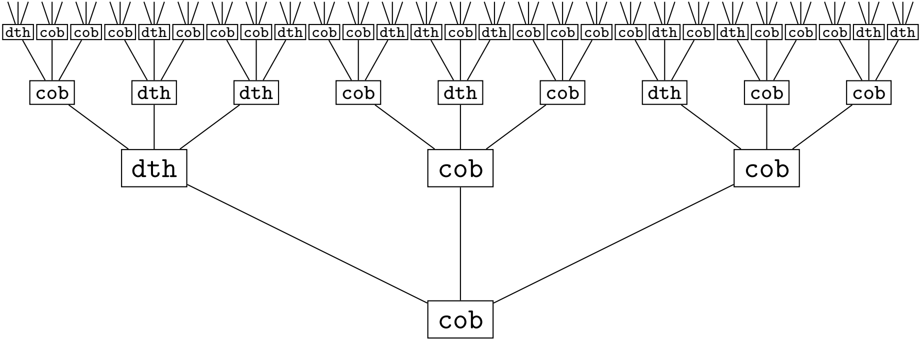

We attach i.i.d. $(\omega_{\mathbf{i}})_{\mathbf{i}\in\mathbb{T}}$ with law $|\mathbf{r}|^{-1}\mathbf{r}$ to each node, which translate into maps $(\gamma[\omega_{\mathbf{i}}])_{\mathbf{i}\in\mathbb{T}}$.

A recursive tree representation

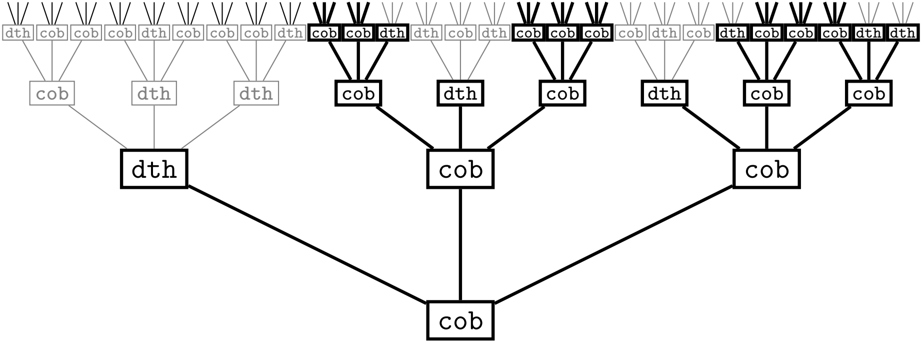

Let $\mathbb{S}$ be the random subtree of $\mathbb{T}$ defined as

$$\mathbb{S} := \{ i_1 \cdots i_n \in \mathbb{T} : i_m \leq \kappa(\omega_{i_1 \cdots i_{m-1}})\ \forall 1 \leq m \leq n \}.$$

A recursive tree representation

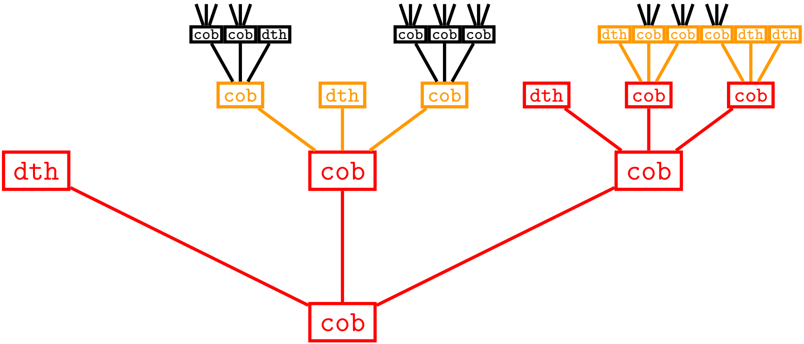

For any rooted subtree $\mathbb{U} \subset \mathbb{S}$, let

$$\nabla\mathbb{U} := \big\{ i_1 \cdots i_n \in \mathbb{S} : i_1 \cdots i_{n-1} \in \mathbb{U},\ \ i_1 \cdots i_n \not\in \mathbb{U} \big\}$$

denote the boundary of $\mathbb{U}$ relative to $\mathbb{S}$.

A recursive tree representation

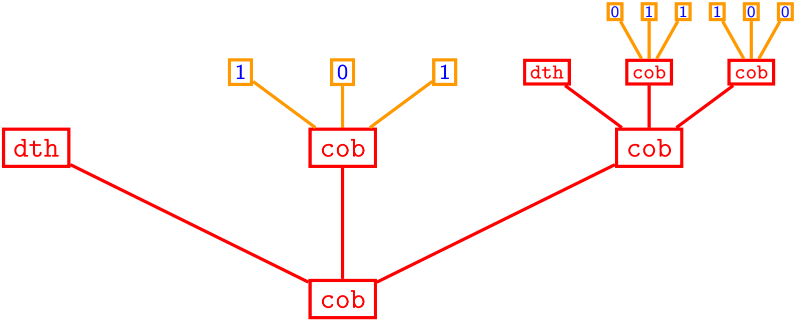

Given $(X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{U}}$, we inductively define $(X_{\mathbf{i}})_{\mathbf{i}\in\mathbb{U}}$ by

$$X_{\mathbf{i}} = \gamma[\omega_{\mathbf{i}}]\big(X_{\mathbf{i}1}, \ldots, X_{\mathbf{i}\kappa(\omega)}\big) \qquad (\mathbf{i} \in \mathbb{U}).$$

A recursive tree representation

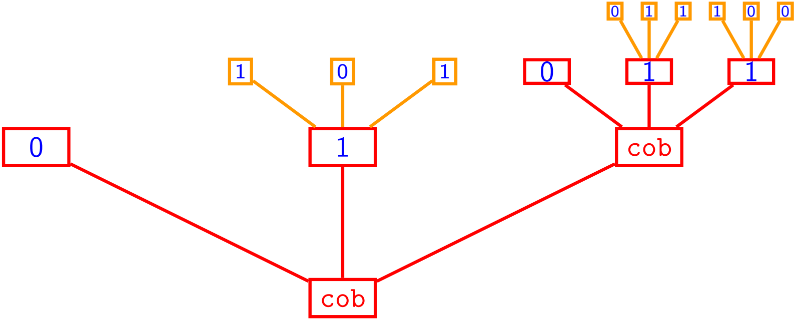

Given $(X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{U}}$, we inductively define $(X_{\mathbf{i}})_{\mathbf{i}\in\mathbb{U}}$ by

$$X_{\mathbf{i}} = \gamma[\omega_{\mathbf{i}}]\big(X_{\mathbf{i}1},\ldots,X_{\mathbf{i}\kappa(\omega)}\big) \qquad (\mathbf{i}\in\mathbb{U}).$$

A recursive tree representation

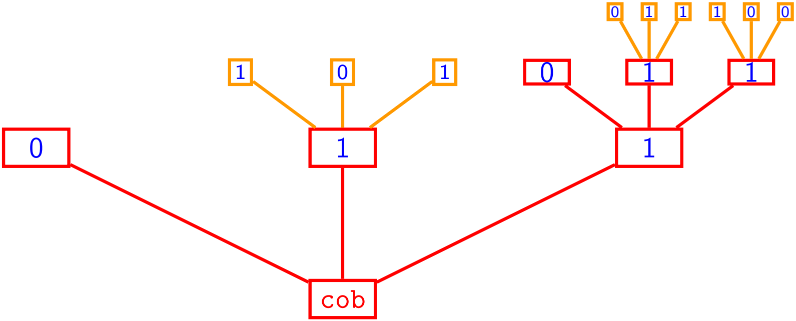

Given $(X_\mathbf{i})_{\mathbf{i}\in\nabla\mathbb{U}}$, we inductively define $(X_\mathbf{i})_{\mathbf{i}\in\mathbb{U}}$ by

$$X_\mathbf{i} = \gamma[\omega_\mathbf{i}]\big(X_{\mathbf{i}1},\ldots,X_{\mathbf{i}\kappa(\omega)}\big) \qquad (\mathbf{i}\in\mathbb{U}).$$

A recursive tree representation

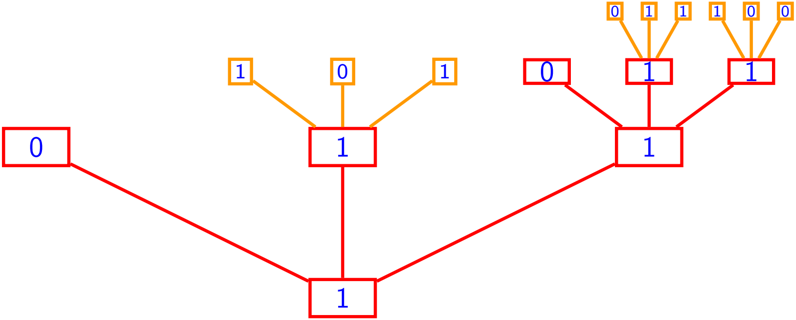

Given $(X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{U}}$, we inductively define $(X_{\mathbf{i}})_{\mathbf{i}\in\mathbb{U}}$ by

$$X_{\mathbf{i}} = \gamma[\omega_{\mathbf{i}}]\big(X_{\mathbf{i}1},\ldots,X_{\mathbf{i}\kappa(\omega)}\big) \qquad (\mathbf{i}\in\mathbb{U}).$$

A recursive tree representation

Define $\Gamma_{\mathbb{U}} : S^{\nabla\mathbb{U}} \to S$ by $\Gamma_{\mathbb{U}}\big((X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{U}}\big) := X_{\varnothing}$.

$\Gamma_{\mathbb{U}}$ is the concatenation of the maps $\big(\gamma[\omega_{\mathbf{i}}]\big)_{\mathbf{i}\in\mathbb{U}}$ according to the tree structure of $\mathbb{U}$.

Let $|i_1\cdots i_n| := n$ denote the length of a word $\mathbf{i}$ and set

$$\mathbb{S}_{(n)} := \{\mathbf{i}\in\mathbb{S} : |\mathbf{i}| < n\} \quad \text{and} \quad \nabla\mathbb{S}_{(n)} = \{\mathbf{i}\in\mathbb{S} : |\mathbf{i}| = n\}.$$

Aldous and Bandyopadyay (2005) proved that

$$\mathbf{T}^n(\mu) := \text{ the law of } \Gamma_{\mathbb{S}_{(n)}}\big((X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{S}_{(n)}}\big),$$

with $(X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{S}_{(n)}}$ i.i.d. with law μ and independent of $(\omega_{\mathbf{i}})_{\mathbf{i}\in\mathbb{S}_{(n)}}$.

A recursive tree representation

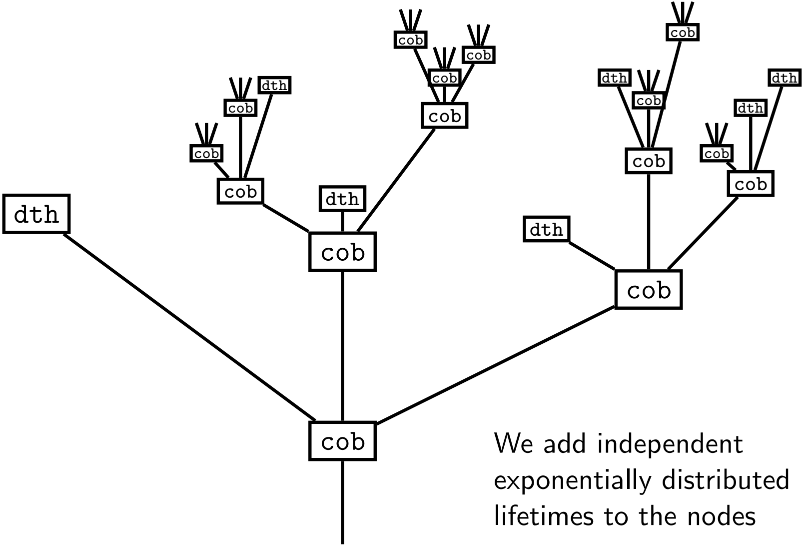

We add independent exponentially distributed lifetimes to the nodes

A recursive tree representation

Let $(\sigma_{\mathbf{i}})_{\mathbf{i}\in\mathbb{T}}$ be i.i.d. exponentially distributed with mean $|\mathbf{r}|^{-1}$, independent of $(\omega_{\mathbf{i}})_{\mathbf{i}\in\mathbb{T}}$, and set

$$\tau^*_{\mathbf{i}} := \sum_{m=1}^{n-1} \sigma_{i_1\cdots i_m} \quad \text{and} \quad \tau^\dagger_{\mathbf{i}} := \tau^*_{\mathbf{i}} + \sigma_{\mathbf{i}} \qquad (\mathbf{i} = i_1\cdots i_n),$$

$$\mathbb{S}_t := \big\{\mathbf{i}\in\mathbb{S} : \tau^\dagger_{\mathbf{i}} \leq t\big\} \quad \text{and} \quad \nabla\mathbb{S}_t = \big\{\mathbf{i}\in\mathbb{S} : \tau^*_{\mathbf{i}} \leq t < \tau^\dagger_{\mathbf{i}}\big\}.$$

Let $\mathcal{F}_t$ be the filtration

$$\mathcal{F}_t := \sigma\big(\nabla\mathbb{S}_t, (\omega_{\mathbf{i}}, \sigma_{\mathbf{i}})_{\mathbf{i}\in\mathbb{S}_t}\big) \qquad (t\geq 0).$$

Theorem [Mach, Sturm, S. '20]

$$\mathbf{T}_t(\mu) := \text{ the law of } \Gamma_{\mathbb{S}_t}\big((X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{S}_t}\big),$$

where $(X_{\mathbf{i}})_{\mathbf{i}\in\nabla\mathbb{S}_t}$ are i.i.d. with law μ and independent of $\mathcal{F}_t$.

A recursive tree representation