A binary branching model with Moran-type interactions

Emma Horton
University of Melbourne

Branching Processes and Applications, Angers
22-26 May 2023

Joint work with Alex Cox (University of Bath) and Denis Villemonais (Université de Lorraine)

Contents

Contents

A binary branching process

Consider a collection of particles $\{x_i(t) : i = 1, \ldots, N_t\}$ taking values in $E \cup \partial$ that evolve as follows.

- Particles move in E according to a continuous time Markov process $((\xi_t)_{t \geq 0}, P_x)$.
- When at $y \in E$, at rate $b(y)$ a particle is replaced by two new particles at positions $y_1, y_2 \in E$, i.e. branching occurs.
- When at $y \in E$, at rate $k(y)$ particles are sent to ∂, i.e. soft killing occurs.
- At time T_∂ a particle is sent to ∂, i.e. hard killing occurs.

A binary branching process

- The branching process is then defined via the atomic measures:

$$X_t := \sum_{i=1}^{N_t} \delta_{x_i(t)}, \qquad t \geq 0.$$

- Define its (linear) expectation semigroup,

$$\psi_t[g](x) := \mathbb{E}_{\delta_x}\left[\langle g, X_t\rangle\right] := \mathbb{E}_{\delta_x}\left[\sum_{i=1}^{N_t} g(x_i(t))\right], \quad t \geq 0, x \in E.$$

A Perron-Frobenius decomposition

Particularly interested in branching processes for which there exists

- $\lambda_* \in \mathbb{R}$,
- a positive function $\varphi \in L_\infty^+(E)$,
- a probability measure $\tilde{\varphi}$ on E

such that for all $g \in L_\infty^+(E)$, $x \in E$, we have

$$\psi_t[g](x) \sim \mathrm{e}^{\lambda_* t}\varphi(x)\langle\tilde{\varphi}, g\rangle, \quad t \to \infty.$$

A Perron-Frobenius decomposition

Particularly interested in branching processes for which there exists

- $\lambda_* \in \mathbb{R}$,
- a positive function $\varphi \in L_\infty^+(E)$,
- a probability measure $\tilde{\varphi}$ on E

such that for all $g \in L_\infty^+(E)$, $x \in E$, we have

$$\psi_t[g](x) \sim \mathrm{e}^{\lambda_* t}\varphi(x)\langle\tilde{\varphi}, g\rangle, \quad t \to \infty.$$

Aim: find efficient ways to estimate λ_*, φ and $\tilde{\varphi}$.

Example: Branching Brownian motion

- $E = [-L, L]$, $L > 0$
- ξ is Brownian motion
- $b = 1$, $k = 0$, killed on $\{-L, L\}$, local branching.

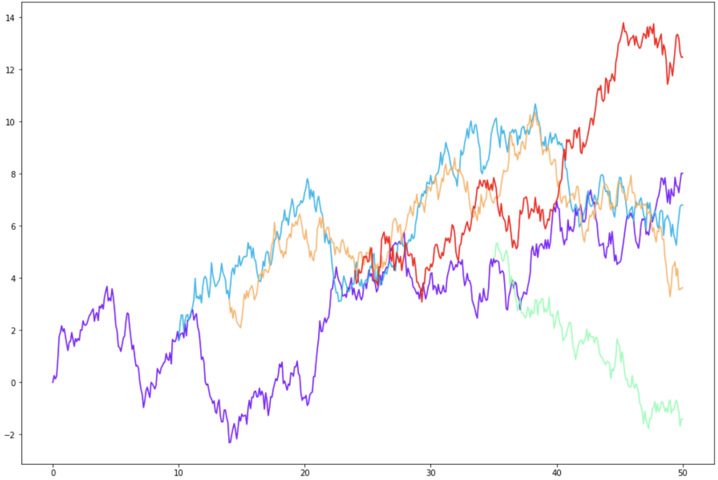

Example: Growth-fragmentation

- $E = [0, \infty)$.
- ξ an appropriate Markov process representing the mass of the particle.
- fragment at rate $b(x)$, mass of daughter particles are y and $x - y$ where $y \sim \kappa(x, \cdot)$.

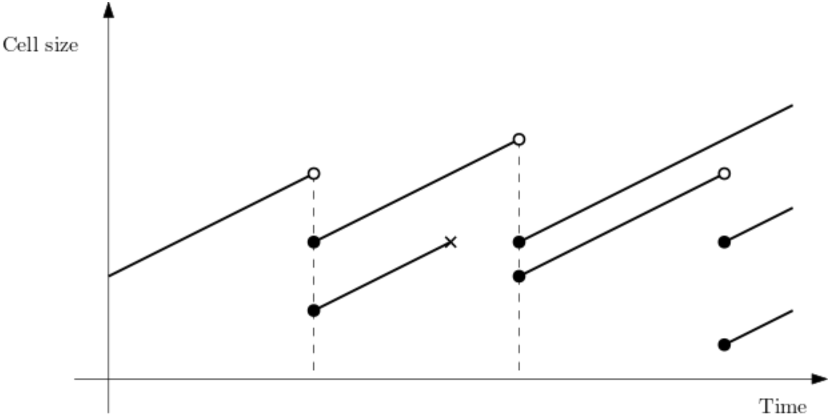

Example: Neutron transport

- $E = D \times V$
- ξ is a piecewise deterministic Markov process
- $b(r,v) = \sigma_{\mathtt{f}}(r,v)$, $k(r,v) = \sigma_{\mathtt{a}}(r,v)$, hard killing on $\{(r,v) : r \in \partial D, v \cdot \mathbf{n}_r > 0\}$, particles are produced at the same position but may have different velocities.

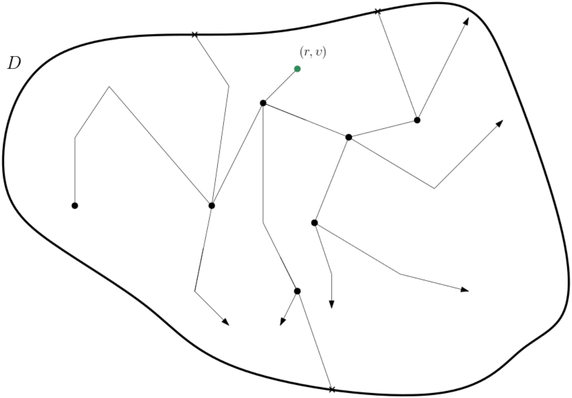

Monte Carlo methods: branching process

Recall the Perron Frobenius asymptotic,

$$\psi_t[g] \sim \mathrm{e}^{\lambda_* t}\langle\tilde{\varphi}, g\rangle\varphi, \quad t \to \infty.$$

Manipulation of this allows us to estimate the eigen-elements, e.g.

$$\lambda_* = \lim_{t\to\infty} \frac{1}{t} \log \psi_t[\mathbf{1}](x) = \lim_{t\to\infty} \frac{1}{t} \log \mathbb{E}_{\delta_x}[N_t].$$

The many-to-one states that there exists a process $(Y_t)_{t\geq 0}$ on E such that

$$\mathbb{E}_{\delta_x}[\langle g, X_t\rangle] = \mathbf{E}_x\left[\mathrm{e}^{\int_0^t b(Y_s)-k(Y_s)\mathrm{d}s} g(Y_t)\mathbf{1}_{t<\tau_\partial}\right], \quad t \geq 0, x \in E.$$

Thus, we can replace the branching process by the single weighted trajectory, e.g.

$$\lambda_* = \lim_{t\to\infty} \frac{1}{t} \log \psi_t[\mathbf{1}](x) = \lim_{t\to\infty} \frac{1}{t} \log \mathbf{E}_x\left[\mathrm{e}^{\int_0^t b(Y_s)-k(Y_s)\mathrm{d}s}\mathbf{1}_{t<\tau_\partial}\right].$$

Monte Carlo methods: many-to-one

The many-to-one states that there exists a process $(Y_t)_{t\geq 0}$ on E such that

$$\mathbb{E}_{\delta_x}[\langle g, X_t\rangle] = \mathbf{E}_x\left[\mathrm{e}^{\int_0^t b(Y_s)-k(Y_s)\mathrm{d}s} g(Y_t)\mathbf{1}_{t<\tau_\partial}\right], \quad t \geq 0, x \in E.$$

Thus, we can replace the branching process by the single weighted trajectory, e.g.

$$\lambda_* = \lim_{t\to\infty} \frac{1}{t} \log \psi_t[\mathbf{1}](x) = \lim_{t\to\infty} \frac{1}{t} \log \mathbf{E}_x\left[\mathrm{e}^{\int_0^t b(Y_s)-k(Y_s)\mathrm{d}s}\mathbf{1}_{t<\tau_\partial}\right].$$

Let $(Y, \mathbf{P}^{\dagger})$ be a sub-Markov process.

Simulate $N \geq 1$ independent copies of $(Y, \mathbf{P}^{\dagger})$ until one of the particles is absorbed.

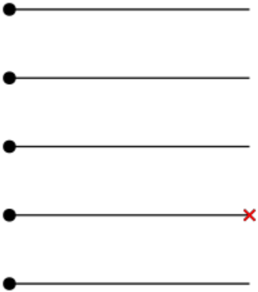

Figure: Fleming Viot particle system

Monte Carlo methods: (Fake) Fleming Viot

When this happens, duplicate one of the remaining $N-1$ particles and return to the previous step.

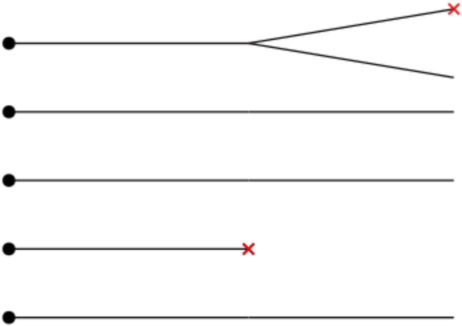

Figure: Fleming Viot particle system

Monte Carlo methods: (Fake) Fleming Viot

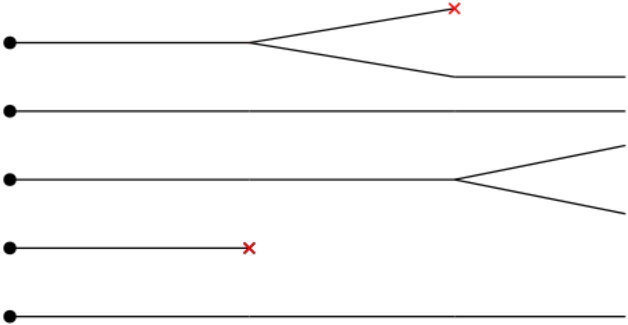

Figure: Fleming Viot particle system

Monte Carlo methods: (Fake) Fleming Viot

- Let A_t denote the number of resampling events up to time t.
- Then

$$\mathbf{E}_x^{\dagger}[g(Y_t)] = \mathbb{E}\left[\left(\frac{N-1}{N}\right)^{A_t}\sum_{i=1}^{N} f(Y_t^i)\right].$$

- Thus, if $(Y, \mathbf{P}^{\dagger})$ admits a QSD with associated rate λ_0, we can use the particle system to estimate these.

- With $\beta := \sup_{x \in E}(b(x) - k(x))$, consider

$$\mathrm{e}^{-\beta t}\psi_t[g](x) = \mathbf{E}_x\left[\mathrm{e}^{\int_0^t (b(Y_s)-k(Y_s)-\beta)\mathrm{d}s} g(Y_t)\mathbf{1}_{\{t<\tau_\partial\}}\right] =: \mathbf{E}_x^\dagger[g(Y_t)].$$

- In this case, we have

$$\mathbb{E}_{\delta_x}[\langle g, X_t\rangle] = \mathrm{e}^{\beta t}\mathbb{E}\left[\left(\frac{N-1}{N}\right)^{A_t}\sum_{i=1}^N f(X_t^i)\right]$$

and

$$\lambda_* = \beta + \lim_{t\to\infty}\frac{1}{t}\log\mathbb{E}\left[\left(\frac{N-1}{N}\right)^{A_t}\right].$$

Monte Carlo methods: (Fake) Fleming Viot

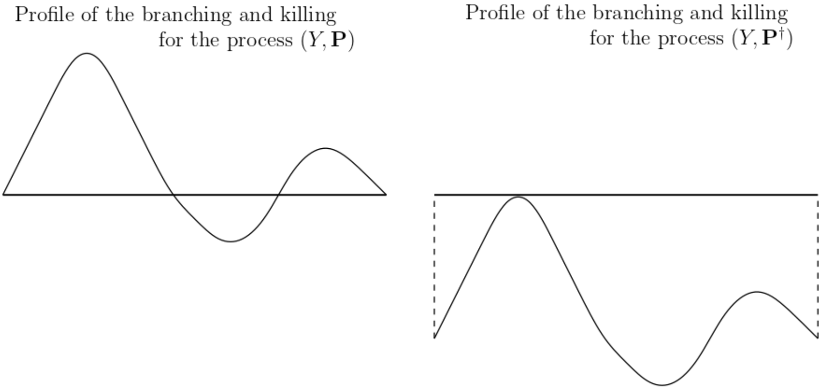

Contents

A binary branching model with Moran interactions (BBMMI)

- $\mathcal{P}_f(\mathbb{N})$ is the collection of finite subsets of $\mathbb{N}$.
- Fix $N_0 \geq 2$. We consider a particle system $(S_t, (X_t^i)_{i \in S_t})_{t \geq 0}$, where
  - $S_t \in \mathcal{P}_f(\mathbb{N})$ is the set enumerating the particles in the system at time $t \geq 0$,
  - X_t^i denotes the position of the i-th particle in the system at time $t \geq 0$.
- Let $b_i : (E \cup \partial)^{\mathcal{P}_f(\mathbb{N})} \to [0, \infty)$ and $\kappa_i : (E \cup \partial)^{\mathcal{P}_f(\mathbb{N})} \to [0, \infty)$, be bounded measurable functions, $i \in \mathbb{N}$.
- Let $p_i : (E \cup \partial)^{\mathcal{P}_f(\mathbb{N})} \to [0, 1]$ and $q_i : (E \cup \partial)^{\mathcal{P}_f(\mathbb{N})} \to [0, 1]$ be measurable functions for $i \in \mathbb{N}$.

1. The particle $X_0^i, i \in S_0$, evolves as an independent copy of Y.

2. Set

$$\tau_b^i = \inf\{t > 0 : \int_0^t b_i(X_s^i, i \in S_0)\mathrm{d}s > \mathbf{e}_i^b\},$$
$$\tau_\kappa^i = \inf\{t > 0 : \int_0^t \kappa_i(X_s^i, i \in S_0)\mathrm{d}s > \mathbf{e}_i^\kappa\},$$
$$\tau_\partial^i = \inf\{t > 0 : X_t^i \in \partial\}.$$

3. Further set $\tau_1 = \inf_{i \in S_0}(\tau_b^i \wedge \tau_\kappa^i \wedge \tau_\partial^i)$ and let i_0 denote the index of the (unique) particle such that $\tau_1 = \tau_b^{i_0} \wedge \tau_\kappa^{i_0} \wedge \tau_\partial^{i_0}$.

BBMMI: Evolution of the system

1. The particle $X_0^i, i \in S_0$, evolves as an independent copy of Y.

2. Set

$$\tau_b^i = \inf\{t > 0 : \int_0^t b_i(X_s^i, i \in S_0)\mathrm{d}s > \mathbf{e}_i^b\},$$
$$\tau_\kappa^i = \inf\{t > 0 : \int_0^t \kappa_i(X_s^i, i \in S_0)\mathrm{d}s > \mathbf{e}_i^\kappa\},$$
$$\tau_\partial^i = \inf\{t > 0 : X_t^i \in \partial\}.$$

3. Further set $\tau_1 = \inf_{i \in S_0} (\tau_b^i \wedge \tau_\kappa^i \wedge \tau_\partial^i)$ and let i_0 denote the index of the (unique) particle such that $\tau_1 = \tau_b^{i_0} \wedge \tau_\kappa^{i_0} \wedge \tau_\partial^{i_0}$.

1. The particle $X_0^i, i \in S_0$, evolves as an independent copy of Y.
2. Set
$$\tau_b^i = \inf\{t > 0 : \int_0^t b_i(X_s^i, i \in S_0)\mathrm{d}s > \mathbf{e}_i^b\},$$
$$\tau_\kappa^i = \inf\{t > 0 : \int_0^t \kappa_i(X_s^i, i \in S_0)\mathrm{d}s > \mathbf{e}_i^\kappa\},$$
$$\tau_\partial^i = \inf\{t > 0 : X_t^i \in \partial\}.$$
3. Further set $\tau_1 = \inf_{i \in S_0}(\tau_b^i \wedge \tau_\kappa^i \wedge \tau_\partial^i)$ and let i_0 denote the index of the (unique) particle such that $\tau_1 = \tau_b^{i_0} \wedge \tau_\kappa^{i_0} \wedge \tau_\partial^{i_0}$.

BBMMI: Evolution of the system

4. If $\tau_1 \in \{\tau_\partial^{i_0}, \tau_\kappa^{i_0}\}$, we say a **killing** event occurred.
   $\rightsquigarrow$ Remove particle i_0 from the system
   $\rightsquigarrow$ With probability $p^{i_0}(X^i_{\tau_1}, i \in S_0)$ a **resampling** event occurs: another particle is chosen uniformly at random and is duplicated.

5. If $\tau_1 = \tau_b^{i_0}$, we say a **branching** event occurred.
   $\rightsquigarrow$ Duplicate particle i_0.
   $\rightsquigarrow$ With probability $q_{i_0}(X^i_{\tau_1}, i \in S_0)$ a **selection** event occurs: one of the particles is selected uniformly at random and removed from the system.

Repeat to generate a sequence $\tau_0 := 0 < \tau_1 < \tau_2 < \cdots < \tau_n < \cdots$ of resampling/selection times.

4. If $\tau_1 \in \{\tau_\partial^{i_0}, \tau_\kappa^{i_0}\}$, we say a **killing** event occurred.
   $\rightsquigarrow$ Remove particle i_0 from the system
   $\rightsquigarrow$ With probability $p^{i_0}(X^i_{\tau_1}, i \in S_0)$ a **resampling** event occurs: another particle is chosen uniformly at random and is duplicated.

5. If $\tau_1 = \tau_b^{i_0}$, we say a **branching** event occurred.
   $\rightsquigarrow$ Duplicate particle i_0.
   $\rightsquigarrow$ With probability $q_{i_0}(X^i_{\tau_1}, i \in S_0)$ a **selection** event occurs: one of the particles is selected uniformly at random and removed from the system.

Repeat to generate a sequence $\tau_0 := 0 < \tau_1 < \tau_2 < \cdots < \tau_n < \cdots$ of resampling/selection times.

BBMMI: Evolution of the system

4. If $\tau_1 \in \{\tau_\partial^{i_0}, \tau_\kappa^{i_0}\}$, we say a **killing** event occurred.
   $\rightsquigarrow$ Remove particle i_0 from the system
   $\rightsquigarrow$ With probability $p^{i_0}(X_{\tau_1}^i, i \in S_0)$ a **resampling** event occurs: another particle is chosen uniformly at random and is duplicated.

5. If $\tau_1 = \tau_b^{i_0}$, we say a **branching** event occurred.
   $\rightsquigarrow$ Duplicate particle i_0.
   $\rightsquigarrow$ With probability $q_{i_0}(X_{\tau_1}^i, i \in S_0)$ a **selection** event occurs: one of the particles is selected uniformly at random and removed from the system.

Repeat to generate a sequence $\tau_0 := 0 < \tau_1 < \tau_2 < \cdots < \tau_n < \cdots$ of resampling/selection times.

BBMMI: Example

Example: branching Brownian motion killed at 0 and 1.

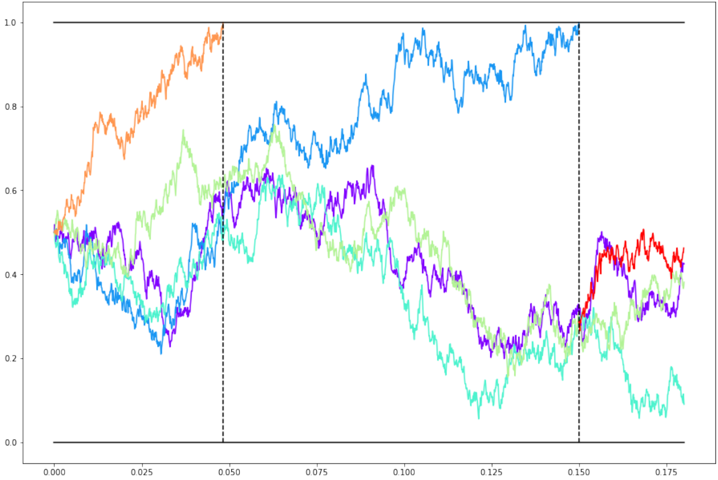

BBMMI: Example

Example: branching Brownian motion killed at 0 and 1.

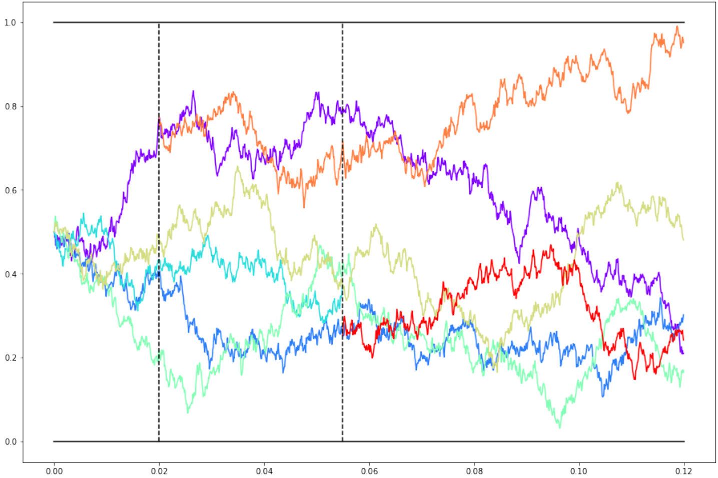

BBMMI: Notation

- Define
  - $\rho_n = n$-th resampling time.
  - $\sigma_n = n$-th selection time.

- Define
  - $A_t := \sup\{n \geq 0 : \rho_n \leq t\}$, the number of resampling events up to time t,
  - $B_t := \sup\{n \geq 0 : \sigma_n \leq t\}$, the number of selection events up to time t,

- Define $\Pi_t^A := \prod_{i=1}^{A_t} \left(\frac{N_{\rho_n} - 1}{N_{\rho_n}} \right)$ and $\Pi_t^B := \prod_{i=1}^{B_t} \left(\frac{N_{\sigma_n} + 1}{N_{\sigma_n}} \right)$

BBMMI: Main result

- ψ_t = semigroup of the branching process/many-to-one
- m_t = empirical distribution of the BBMMI at time t.

Theorem (Cox, H. & Villemonais)

Under certain assumptions, for all time $T \geq 0$ and all bounded measurable functions $f : E \to \mathbb{R}$, the BBMMI satisfies

$$\psi_T f(x) = \mathbf{E}_x \left[\Pi_T^A \Pi_T^B m_T(f) \right]. \quad (1)$$

Moreover,

$$\left\| \frac{\psi_T f(x)}{\psi_T \mathbf{1}_E(x)} - \frac{m_T(f)}{m_T(\mathbf{1}_E)} \mathbf{1}_{m_T \neq 0} \right\|_2 \leq \frac{C_T}{\sqrt{N_0}} \frac{\|f\|_\infty}{m_0 \psi_T \mathbf{1}_E / N_0}, \quad (2)$$

where C_T is a constant depending on $\|b\|_\infty$ and T.

BBMMI: The assumptions

(B1) For any $x \in E$ and $t \geq 0$, $\mathbb{P}_x(\tau_\partial = t) = 0$ and $\mathbb{P}_x(\tau_\partial > t) > 0$.

(B2) For all $s \in \mathcal{P}_f(\mathbb{N})$, $(x_i)_{i \in s}$ and $i_0 \in s$,

$$b_{i_0}(x_i, i \in s) - \kappa_{i_0}(x_i, i \in s) = b(x_{i_0}) - \kappa(x_{i_0}).$$

(B3) For all $s \in \mathcal{P}_f(\mathbb{N})$, $(x_i)_{i \in s}$ and $i_0 \in s$,

$$p_{i_0}(x_i, i \in s) = 0 \text{ whenever } |s| = 1.$$

(B4) The sequence, $(\tau_n)_{n \geq 1}$, of event times satisfies $\tau_n \to \infty$ as $n \to \infty$.

Sketch proof

- Define
$$\nu_t^f = \Pi_t^A \Pi_t^B \sum_{i \in S_t} \psi_{T-t} f(X_t^i).$$
- Then the many-to-one formula (1) in the theorem can be written as
$$\nu_0^f = \psi_T f(x) = \mathbf{E}_x \left[\Pi_T^A \Pi_T^B m_T f \right] = \mathbf{E}_x[\nu_T^f].$$
- The idea behind the proof is to find a martingale decomposition for $\nu_T^f - \nu_0^f$.
  - (1) then follows by taking expectations.
  - (2) follows by studying the L^2 norm of the martingales.

Sketch proof

We can find two martingales $\mathbb{M}$ and $\mathcal{M}$ such that

$$\nu_t^f - \nu_0^f = \Pi_t^A \Pi_t^B (\mathbb{M}_t - \mathbb{M}_{\tau_{C_t}}) + \sum_{n=1}^{C_t} \Pi_{\tau_n}^A \Pi_{\tau_n}^B (\mathbb{M}_{\tau_n} - \mathbb{M}_{\tau_{n-1}})$$

$$+ \sum_{n=1}^{A_t} \Pi_{\rho_n}^A \Pi_{\rho_n}^B (\mathcal{M}_{\rho_n} - \mathcal{M}_{\rho_{n-1}}) + \sum_{n=1}^{B_t} \Pi_{\sigma_n}^A \Pi_{\sigma_n}^B (\mathcal{M}_{\sigma_n} - \mathcal{M}_{\sigma_{n-1}}). \quad (3)$$

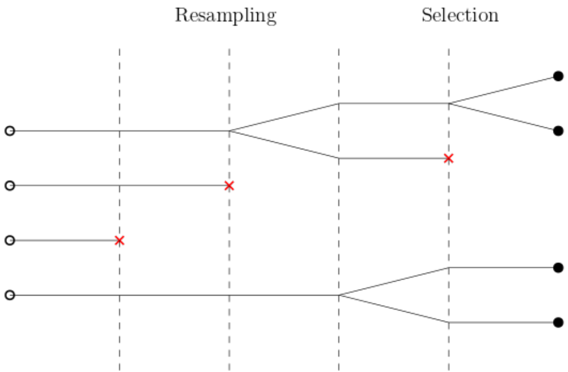

- $p = q = 0 \rightsquigarrow$ binary branching process.
  Then $\Pi_T^A = \Pi_T^B = 1$ a.s., and (1) is the classical many-to-one formula.
- When $p = q = 1$, we obtain an extension of the (F)FV particle system
- It is possible to let p and q depend on time.
- The above dynamics allow one to constrain the size of the process to remain between two bounds $0 \leq N_{min} \leq N_{max}$, $N_{min} \neq 1$, by choosing $p_{i_0}(x_i, i \in s) = 1$ whenever $|s| = N_{min}$, and $q_{i_0}(x_i, i \in s) = 1$ whenever $|s| = N_{max}$. We call this the N_{min}-N_{max} model.

Contents

A branching birth and death process

$E = \{1, \dots, M\}$ for some $M \geq 1$. When at $x \in E$, one of the following things may occur:

- the particle jumps with rate x^2 to state
  - $\max\{1, x-1\}$ with probability $x/(x+1)$,
  - $\min\{x+1, M\}$ with probability $1/(x+1)$.

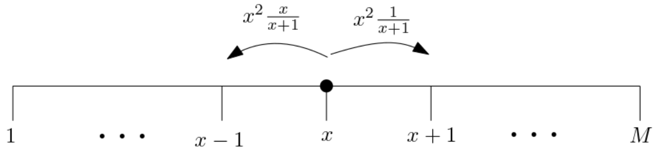

A branching birth and death process

$E = \{1, \ldots, M\}$ for some $M \geq 1$. When at $x \in E$, one of the following things may occur:

- the particle jumps with rate x^2 to state
    - $\max\{1, x-1\}$ with probability $x/(x+1)$,
    - $\min\{x+1, M\}$ with probability $1/(x+1)$.
- at rate $b(x) = x$, a new particle is produced at the same site, which will continue to evolve independently according to the same dynamics.

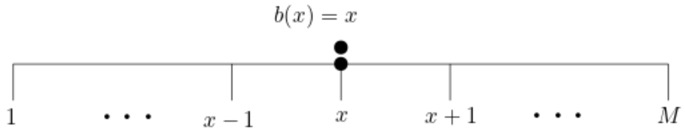

N_{min}-N_{max} vs (F)FV

- We now consider the problem of numerically approximating ν_M, the **left eigenmeasure** of the process.

- Fix $T > 0$ and $N_0 = N_{min} = N_{max}$.

- $\mathcal{X}_N^M$ = empirical distribution of the N_{min}–N_{max} particle system.
  $\mathcal{Y}_N^M$ = the empirical distribution FV particle systems.

- The estimator is $\hat{\theta}_N$, the empirical distribution of the relevant particle system.

- We compare three quantities:
  - the bias of the estimator, $|\mathbb{E}(\hat{\theta}_N(f)) - \nu_M(f)|$,
  - the standard deviation of the estimator, $\mathrm{Std}(\hat{\theta}_N)$,
  - the number of interaction events, $(A_T + B_T)/T$.

N_{min}-N_{max} vs (F)FV

$N = 10$		$\lvert\mathbb{E}(\hat{\theta}_N) - \nu_M(f)\rvert$	$\mathrm{Std}(\hat{\theta}_N)$	$(A_T + B_T)/T$
$M = 10$	N_{min}–N_{max}	0.08	0.30	14.0
	FV	0.10	0.41	87.2
$M = 100$	N_{min}–N_{max}	0.08	0.30	14.0
	FV	0.20	0.51	988
$M = 1000$	N_{min}–N_{max}	0.08	0.30	14.0
	FV	0.22	0.53	9989
$M = +\infty$	N_{min}–N_{max}	0.08	0.30	14.0
	FV	$*$	$*$	$*$

N_{min}-N_{max} vs (F)FV

$N = 100$		$\lvert\mathbb{E}(\hat{\theta}_N) - \nu_M(f)\rvert$	$\mathrm{Std}(\hat{\theta}_N)$	$(A_T + B_T)/T$
$M = 10$	N_{min}–N_{max}	0.01	0.12	144
	FV	0.02	0.18	857
$M = 100$	N_{min}–N_{max}	0.01	0.12	144
	FV	0.10	0.39	9866
$M = 1000$	N_{min}–N_{max}	0.01	0.12	144
	FV	0.20	0.50	99873
$M = +\infty$	N_{min}–N_{max}	0.01	0.12	144
	FV	$*$	$*$	$*$

N_{min}-N_{max} vs (F)FV

"Theorem"

The BBMMI wins!

Future work

- Extending the results to more general branching processes
- Scaling limits and genealogical structure
- Central limit theorem
- Branching processes with infinite branching rate
- Relation to adaptive multilevel splitting for rare event simulation
- ...

Thank you!

Define

$$\mathbb{M}_t := \sum_{i\in S_t} \psi_{T-t} f(X_t^i) - \sum_{i\in S_0} \psi_T f(X_0^i) - \sum_{n=1}^{A_t} \psi_{T-\rho_n} f(X_{\rho_n}^{i_n}) + \sum_{n=1}^{B_t} \psi_{T-\sigma_n} f(X_{\sigma_n}^{j_{n-1}}). \quad (4)$$

and

$$\mathcal{M}_{\rho_n} - \mathcal{M}_{\rho_n -} := \psi_{T-\rho_n} f(X_{\rho_n}^{i_{n-1}}) - \frac{1}{N_{\rho_n} - 1} \sum_{i\in S_{\rho_n -}\setminus\{i'_{n-1}\}} \psi_{T-\rho_n} f(X_{\rho_n}^i),$$

$$\mathcal{M}_{\sigma_n} - \mathcal{M}_{\sigma_n -} := \frac{1}{N_{\sigma_n} + 1} \left(\sum_{i\in S_{\sigma_n -}} \psi_{T-\sigma_n} f(X_{\sigma_n}^i) + \psi_{T-\sigma_n} f(X_{\sigma_n}^{j'_{n-1}}) \right) - \psi_{T-\sigma_n} f(X_{\sigma_n}^{j'_{n-1}}).$$