Clustering in particle systems of fixed size

Oliver Tough
Department of Mathematics
University of Bath

Clustering in finite domains

- Phenomenon where the particles in a critical branching particle system are not evenly distributed, but rather cluster.
- Number of particles in a physical system is much larger than in the simulation. This causes clustering to occur in the simulation but not reality.
- In finite domains this depends on the ratio of two timescales:
  1. The mixing time τ_D. This represents how quickly the distribution of a given particle will converge to a limiting distribution. Roughly, it corresponds to how quickly a particle will explore the space.
  2. The extinction time τ_E, representing the time over which random fluctuations in the number of births and deaths lead to the population becoming extinct.
  3. The ratio of these times determines whether clustering will be observed:

$$\tau_D \gg \tau_E \Rightarrow \text{substantial clustering}$$

$$\tau_D \ll \tau_E \Rightarrow \text{negligible clustering}$$

Clustering in finite domains with population control

- We can link birth and death events to keep the population constant. For instance, de Mulatier, Dumonteil, Rosso and Zoir[a] consider a system such that whenever a neutron disappears and branches into m particles, we simultaneously remove $m-1$ particles chosen uniformly at random.
- The particles evolve as Brownian motions with diffusivity D (how quickly they randomly move) in a domain of volume V. Birth-death events occur at rate λ.
- Here the extinction time τ_E becomes the time over which all individuals descend from a single common ancestor.

[a]de Mulatier, C., Dumonteil, E., Rosso, A., Zoia, A. (2015). The critical catastrophe revisited. Journal of Statistical Mechanics: Theory and Experiment, 2015 (8), P08021.

Clustering in finite domains with population control

- The particles evolve as Brownian motions with diffusivity D (how quickly they randomly move) in a domain of volume V. Birth-death events occur at rate λ.
- De Mulatier, Dumonteil, Rosso and Zoir[a] argued that the mixing time τ_D and extinction time τ_E should be given by

$$\tau_D \propto \frac{V^{\frac{2}{d}}}{D}, \quad \tau_E \propto \frac{N}{\lambda}.$$

- We will consider a similar particle system, called the Fleming-Viot particle system. We will obtain an asymptotic description of τ_E for this particle system.
- We see that τ_E is, in general, smaller than $\frac{N}{\lambda}$. This means clustering will occur more easily than is predicted by only considering N and λ.

[a]de Mulatier, C., Dumonteil, E., Rosso, A., Zoia, A. (2015). The critical catastrophe revisited. Journal of Statistical Mechanics: Theory and Experiment, 2015 (8), P08021.

τ_E in population genetics

- Classical models of population genetics consider only a population of N individuals, with space not taken into account.
- Predictions are then made of the time to the most recent common ancestor, which is the extinction time τ_E here. These predictions are based only on the population size N and rate of births/deaths, λ.
- For a variety of reasons, it is observed that τ_E is typically much smaller than is predicted by these classical models.
- This is formulated in terms of an "effective population size".

τ_E in population genetics

- It is observed in population genetics that τ_E is typically much smaller than is predicted by considering only the population size N and rate of branching/killing λ.
- We will see that this is indeed the case with branching/killing particle systems with population control. In particular, τ_E is smaller than is predicted by simply counting the number of individuals and the branching/killing rate.
- We will quantify this.

Background

- In [a], Brown, Jenkins, Johansen and Koskela considered the genealogies of a sequential Monte Carlo algorithm,
- They obtained the coalescent time of finite samples of the population.
- Their time rescaling is rather different from what we obtain (the setup is different).

[a] Simple conditions for convergence of sequential Monte Carlo genealogies with applications (2021). Electron. J. Probab. 26 1-22, 2021.

The Fleming-Viot particle System

- Method for sampling quasi-stationary distributions

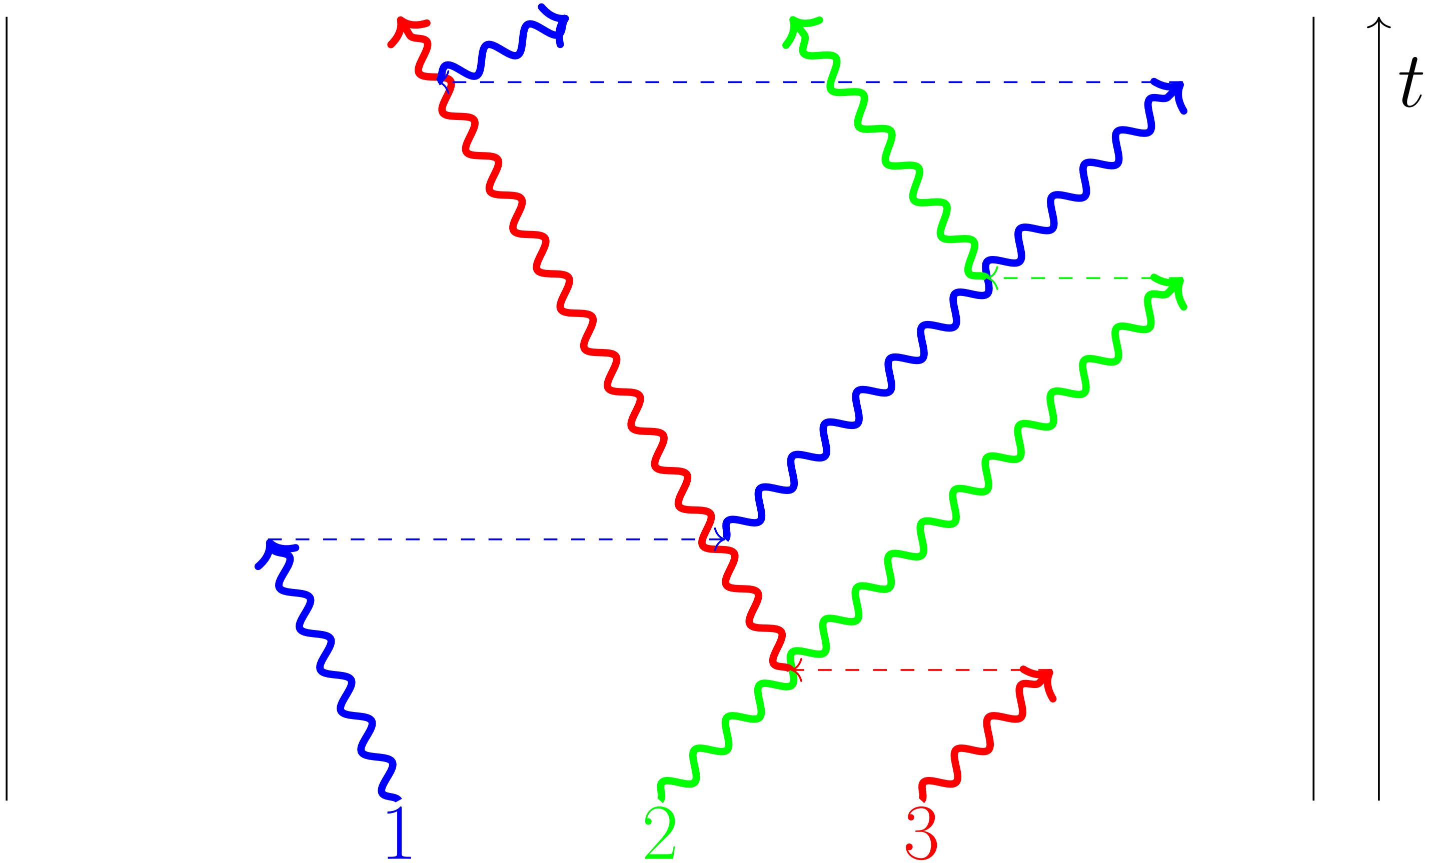

Figure 1: Each particle evolves randomly (as a Brownian motion) until being killed, then jumps onto one of the other particles chosen uniformly at random.

The Fleming-Viot particle system

- The Fleming-Viot particle system was introduced by Burdzy, Holyst and March [a].
- Each particle evolves randomly until being killed, then jumps onto one of the other particles chosen uniformly at random.
- Note that there is a deep connection between killing and branching/killing systems.
- Recall that de Mulatier, Dumonteil, Rosso and Zoir[b] consider a system such that whenever a neutron disappears and branches into m particles, we simultaneously remove $m-1$ particles chosen uniformly at random. Our results would follow by the same argument if we were also to allow such population controlled branching.

[a] Burdzy, K., Holyst, R., March, P. (2000). A Fleming-Viot Particle Representation of the Dirichlet Laplacian. Communications in Mathematical Physics, 214 (3), 679–703.

[b] de Mulatier, C., Dumonteil, E., Rosso, A., Zoia, A. (2015). The critical catastrophe revisited. Journal of Statistical Mechanics: Theory and Experiment, 2015 (8), P08021.

The Fleming-Viot particle system

- The particles in the Fleming-Viot particle system follow the dynamics of a process B_t with killing time τ_∂.
- The long-time limits

$$\mathcal{L}(B_t|\tau_\partial) \to \varphi(x)dx$$

  are known as quasi-stationary distributions (QSDs). This corresponds to the flux or importance map.
- The Fleming-Viot particle system provides an approximation method for QSDs.

The Fleming-Viot particle system

- We attach genetic information to the particles in the Fleming-Viot particle system.
- We obtain a scaling limit for this genetic information as time is rescaled by $t \mapsto Nt$ [a].

[a] T. Scaling limit of the Fleming-Viot multi-colour process. Arxiv:2110.05049.

The Wright-Fisher process

- The limit we obtain is called the Wright-Fisher process.
- For a survey of the Wright-Fisher process see Ethier and Kurtz's survey [a]
- The Wright-Fisher process was introduced by Fleming and Viot in [b]. In the population genetics literature it is typically called a "Fleming-Viot process", but we obviously can't use that name here.
- In contrast, the Fleming-Viot particle system has nothing to do with Fleming or Viot. The two processes are very different.

---

[a] Ethier, S. N., Kurtz, T. G. (1993). Fleming-Viot processes in population genetics. SIAM Journal on Control and Optimization, 31(2), 345–386.

[b] Fleming, W., Viot, M. (1979). Some Measure-Valued Markov Processes in Population Genetics Theory. Indiana University Mathematics Journal, 28 (5), 817–843.

The Wright-Fisher process

- The Wright-Fisher process arises as the scaling limit of a number of classical models in population genetics.
- The construction is not straightforward.
- The Wright-Fisher process describes forwards-in-time dynamics - i.e. how do the proportions of different genes evolve forwards in time.
- The Wright-Fisher process corresponds to Kingman's coalescent, a canonical coalescent process from population genetics.
- The difference is that Kingman's coalescent describes backwards-in-time dynamics - i.e. how do ancestral lineages coalesce as you look backwards in time.
- This connection will give us the distribution of τ_E.

Recap

- The Fleming-Viot particle system is a system of branching and killing particles, with population control.
- The Wright-Fisher process is a process from population genetics that arises as the large population scaling limit of many models.
- We will obtain the Wright-Fisher process as a scaling limit for the Fleming-Viot particle system.
- This will give us the extinction time, τ_E, of the Fleming-Viot particle system.
- We will find this extinction time to be less than is predicted by the population N and branching/killing rate. This indicates more clustering.

Main result

- We take the Fleming-Viot particle system with N particles, and attach genetic information to the particles.
- We write $\varphi(x)dx$ for the flux/importance map and $\tilde{\varphi}(x)$ for the adjoint flux. These represent the principal left (respectively right) eigenfunctions of the generator of the associated semigroup.
- We rescale time by

$$t \mapsto 2\frac{N}{\lambda} \times \frac{\int [\tilde{\varphi}(x)]^2 \varphi(x) dx}{\left(\int \tilde{\varphi}(x)\varphi(x) dx \right)^2} t.$$

- Then the genetic information converges to a Wright-Fisher process.
- This time rescaling therefore tells you how quickly the merging of ancestral lineages occurs.

τ_E in the Fleming-Viot particle system

- It follows from our time rescaling that the extinction time is given by [a]

$$\mathbb{E}[\tau_E] \sim \frac{N}{\lambda} \times \frac{\left(\int \tilde{\varphi}(x)\varphi(x)dx\right)^2}{\int [\tilde{\varphi}(x)]^2\varphi(x)dx}.$$

- We would expect $\frac{N}{\lambda}$ by simply counting the number of births and deaths.

[a] T. Scaling limit of the Fleming-Viot multi-colour process. Arxiv:2110.05049

τ_E in the Fleming-Viot particle system

- We normalise $\tilde{\varphi}$ so that $\int \tilde{\varphi}(x)\varphi(dx) = 1$. Then the correction term can be written

$$\frac{\left(\int \tilde{\varphi}(x)\varphi(x)dx\right)^2}{\int [\tilde{\varphi}(x)]^2\varphi(x)dx} = \frac{1}{1 + \mathrm{Var}_\varphi(\tilde{\varphi})}.$$

- When the dynamics are very inhomogeneous, $\tilde{\varphi}$ will be inhomogeneous, hence $\mathrm{Var}_\varphi(\tilde{\varphi})$ will be large, so the correction term will be much smaller than 1.

- Since the expectation of the extinction time is given by

$$\mathbb{E}[\tau_E] \sim \frac{N}{\lambda} \times \frac{1}{1 + \mathrm{Var}_\varphi(\tilde{\varphi})},$$

  this would mean that the extinction time is much smaller than the $\frac{N}{\lambda}$ value you get by counting the number of births and deaths.

τ_E in the Fleming-Viot particle system

- Normalise $\tilde{\varphi}$ so that $\int \tilde{\varphi}(x)\varphi(dx) = 1$.
- The expectation of the extinction time is given by

$$\mathbb{E}[\tau_E] \sim \frac{N}{\lambda} \times \frac{1}{1 + \mathrm{Var}_{\varphi}(\tilde{\varphi})}.$$

- When the dynamics are inhomogeneous, the extinction time is much smaller than the $\frac{N}{\lambda}$ value you get by counting the number of births and deaths.
- Since clustering occurs when $\tau_E \ll \tau_D$, this means clustering occurs more easily.

τ_E in the Fleming-Viot particle system

- We can describe the distribution of τ_E.
- Let $\exp(k(k-1))$ for $k \geq 2$ be independent random variables whose distribution is the exponential distributions with parameter $k(k-1)$.
- Then the distribution of the extinction time of the Fleming-Viot process is given by

$$\tau_E \sim \underbrace{\frac{N\Big(\int \tilde{\varphi}(x)\varphi(x)dx\Big)^2}{\lambda \int [\tilde{\varphi}(x)]^2 \varphi(x)dx}}_{\text{non-random}} \quad \underbrace{\Big[\sum_{k\geq 2} \exp(k(k-1))\Big]}_{\text{sum of independent random variables}} \quad .$$

- Recall that the Wright-Fisher process corresponds to "Kingman's coalescent". The above random sum is simply the total coalescence time of Kingman's coalescent.

Less idealised systems?

- Question: can these results be extended to less idealised Monte Carlo algorithms?

τ_E in particle systems without population control

- We have focused on particle systems without population control. We investigate the effect of spatial inhomogeneity on τ_E.
- In population genetics it is observed that genetic bottlenecks caused by large fluctuations in the size of the population can also substantially reduce τ_E.
- This may explain the observation of de Mulatier, Dumonteil, Rosso and Zoir[a] that systems with population control seem to have less clustering.

[a]de Mulatier, C., Dumonteil, E., Rosso, A., Zoia, A. (2015). The critical catastrophe revisited. Journal of Statistical Mechanics: Theory and Experiment, 2015 (8), P08021.

Mixing time

- We recall that τ_D is the mixing time and τ_E the extinction time. The ratio of these times determines whether clustering will be observed in finite domains:

$$\tau_D \gg \tau_E \Rightarrow \text{substantial clustering}$$

$$\tau_D \ll \tau_E \Rightarrow \text{negligible clustering}$$

- What can we say about the mixing time?
- The mixing rate is given by a spectral gap (difference between first and second eigenvalue).
- Clearly, the mixing time is inversely proportional to the mixing rate.
- An upper bound on the mixing time/lower bound on the mixing rate tells you how large τ_E has to be to ensure no clustering.

Mixing time

- There is a very large literature on convergence/mixing rates for Markov processes without killing.
- Since we are dealing with killing, things become substantially harder. Results without killing don't apply, and much less is known.

Mixing time

- In [a], we introduce a criterion for the mixing of killed processes.
- This gives a rate of convergence/lower bound on the spectral gap.
- It is applied to piecewise-deterministic processes, so should be applicable in the case of neutron transport dynamics.

[a] T. (2022). L^∞-convergence to a quasi-stationary distribution. Arxiv:2210.13581

Mixing time

- Suppose that $\varphi(x)$ is the importance map.
- We assume that one can formulate an "adjoint process" $(\tilde{X}_t)_{0\le t<\tilde{\tau}_\partial}$.
- We assume that there is a time $t_0 > 0$, constant $c_0 > 0$, and probability measure $\nu(x)dx$ such that this adjoint process satisfies

$$\mathbb{P}_{x_0}(\tilde{X}_{t_0} \in dx | \tilde{\tau}_\partial > t_0) \ge c_0 \nu(x)dx \quad \text{for all initial conditions } x_0.$$

- Then the spectral gap is bounded below by

$$\frac{c_0 \int \varphi(x)\nu(x)dx}{\sup_x \varphi(x) t_0}.$$

- It follows that the mixing time is bounded by

$$\tau_D \lesssim \frac{\sup_x \varphi(x) t_0}{c_0 \int \varphi(x)\nu(x)dx}.$$

The End