Exact methods for simulating stationary distributions

Alexander Cox
17th May, 2023

University of Bath

Convergence to Equilibrium

- Many problems, e.g. Fission Source Convergence can be thought of in terms of convergence of stochastic processes to equilibrium.
- Simple case: consider a Markov Chain X_n with initial law π_0, transition matrix P, i.e.

$$\mathbb{P}(X_0 = i) = (\pi_0)_i, \mathbb{P}(X_n = j) = \sum_i \mathbb{P}(X_{n-1} = i) P_{i,j}$$

  then we have $\mathbb{P}(X_n = i) =: (\pi_n)_i = (\pi_{n-1} P)_i$.
- Under reasonable conditions on the Markov Chain, we expect $\pi_n \to \pi_\infty$, where $\pi_\infty P = \pi_\infty$.
- In many numerical methods, we want to sample from π_∞. One solution is to just choose n large, and hope!

Are there other approaches to allow us to get guaranteed convergence?

Coupling from the Past

Based on Propp & Wilson ('96). Suppose we have a finite state Markov Chain. Want to simulate from the stationary distribution.

Key idea: start coupled evolutions from every state in the state space. Couple evolution using the *same* random noise.

Example: Consider a Markov Chain on $\{1, 2, \ldots, M\}$. Let U be a uniform $[0, 1]$ random variable. Can simulate steps of the Markov Chain using:

$$X_{n+1} = \min\{k : \sum_{j=1}^{k} P_{X_n,j} > U\} =: \Psi(X_n, U)$$

Let $\{U_{-n}, U_{-(n-1)}, \ldots, U_{-1}\}$ be a sequence of independent $U[0, 1]$ random variables, then

$$\bar{X}_{-n}(x) := (x, \Phi(x, U_{-n}), \Phi(\Phi(x, U_{-n}), U_{-(n-1)}), \ldots, \Phi(\ldots, U_{-1}))$$

determines a MC started at $-n$, run to 0.

Coupling from the Past

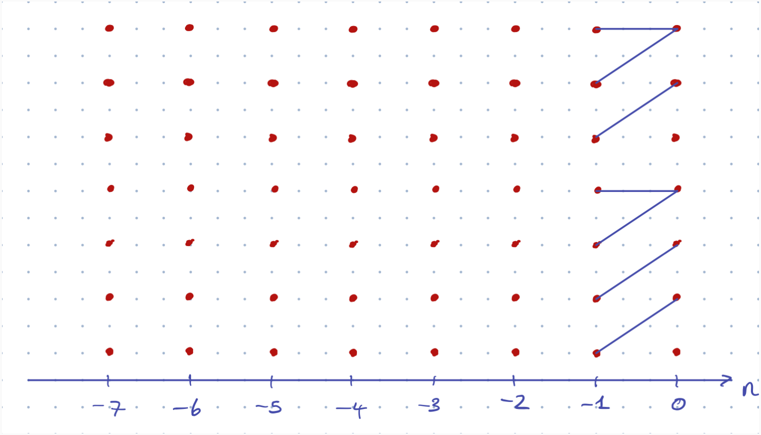

Figure 1: $\bar{X}_{-1}(x)$

Coupling from the Past

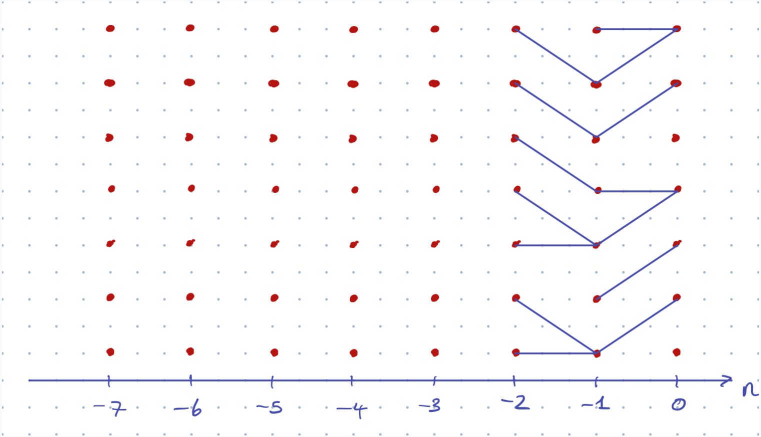

Figure 2: $\bar{X}_{-2}(x)$

Coupling from the Past

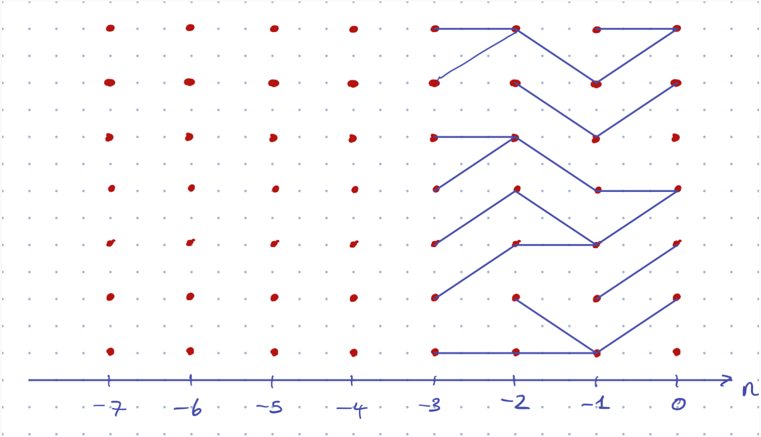

Figure 3: $\bar{X}_{-3}(x)$

Coupling from the Past

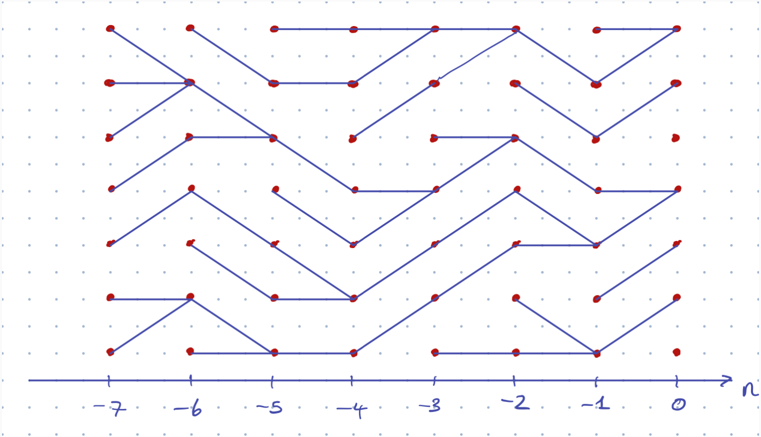

Figure 4: $\bar{X}_{-7}(x)$

Coupling from the Past

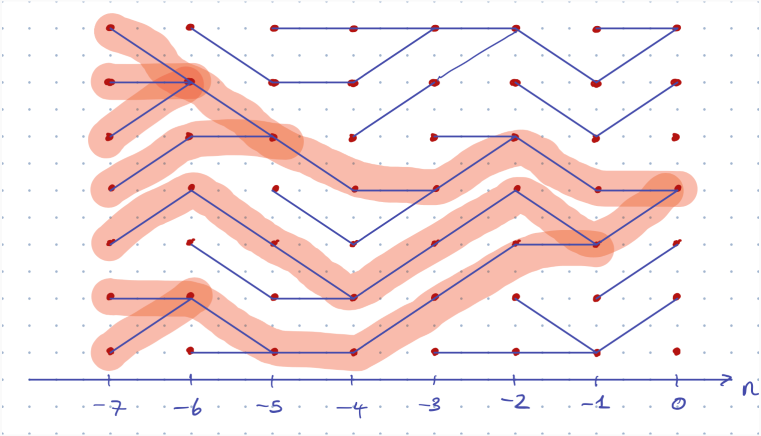

Figure 5: Stop when all starting points give the same value at time 0.

Coupling from the Past

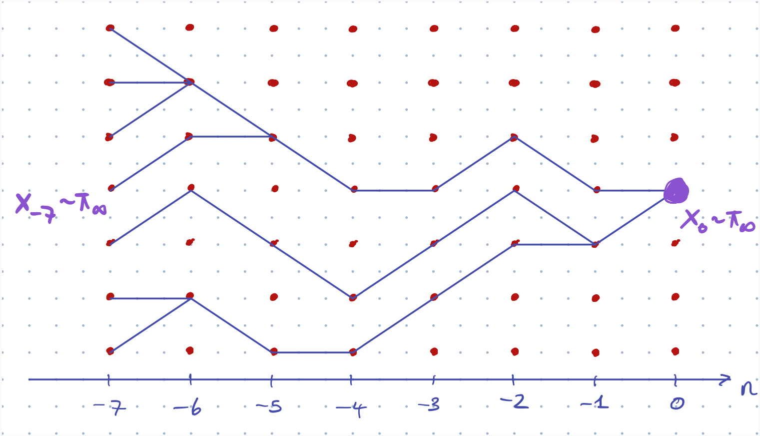

Figure 6: Choose $X_{-7} \sim \pi_\infty$. Then $X_0 \sim \pi_\infty$. But X_0 is the same for all X_{-7}

Coupling From the Past: Comments

- Coupling from the Past guarantees a sample from the stationary distribution...!
- ... but no free lunch:
    - We still have a random number of time steps.
    - Only gives one sample - we typically want many samples.
    - Does not avoid ‘traps’ in the underlying system.
    - Works better when there is a natural ordering on the state space.
- Ideas can be extended to much more general settings (see e.g. Wang et. al. ’21)

Exact estimation for Markov chain equilibrium expectations

Based on Glynn & Rhee ('14).

In the previous example, we tried to simulate from π_∞. In many applications, we only want to compute an expectation:

$$\mathbb{E}[h(X_\infty)] = \lim_{n\to\infty} \mathbb{E}[h(X_n)]$$

Can rewrite this as:

$$\mathbb{E}[h(X_\infty)] = \mathbb{E}[h(X_0)] + \mathbb{E}\left[\sum_{k=1}^{\infty}(h(X_k) - h(X_{k-1}))\right]$$

Now consider a further process Y with the same law as X such that Y_k has the same distribution as X_{k-1}. In particular $Y_0 = X_{-1}$. Then:

$$\begin{aligned}\mathbb{E}[h(X_\infty)] &= \mathbb{E}[h(X_0)] + \mathbb{E}\left[\sum_{k=1}^{\infty}(h(X_k) - h(X_{k-1}))\right] \\ &= \mathbb{E}[h(X_0)] + \mathbb{E}\left[\sum_{k=1}^{\infty}(h(X_k) - h(Y_k))\right]\end{aligned}$$

Exact estimation for Markov chain equilibrium expectations

We want to compute:

$$\mathbb{E}\left[h(X_\infty)\right] = \mathbb{E}\left[h(X_0)\right] + \mathbb{E}\left[\sum_{k=1}^{\infty}(h(X_k) - h(Y_k))\right]$$

where X and Y are not necessarily independent Markov processes with $X_{-1} = Y_0 = x$. Note that if we can find a random time τ when $X_\tau = Y_\tau$ then:

$$\mathbb{E}\left[h(X_\infty)\right] = \mathbb{E}\left[h(X_0)\right] + \mathbb{E}\left[\sum_{k=1}^{\tau-1}(h(X_k) - h(Y_k))\right]$$

Need to choose a coupling of X and Y to ensure they hit in "good" time. E.g. 'mirror coupling'.

Mirror Coupling

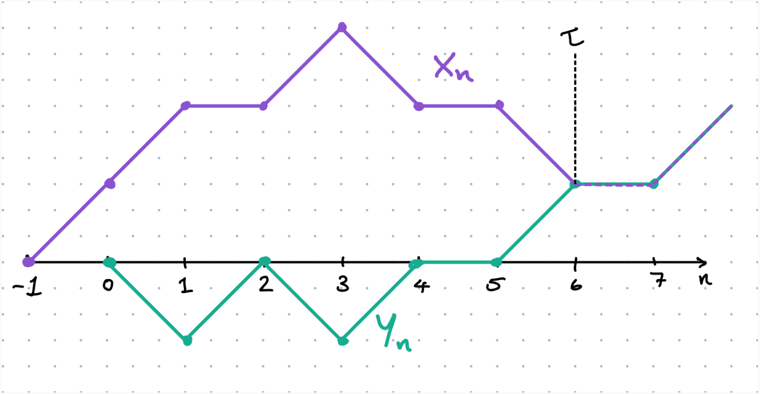

Figure 7: Mirror coupling of *X* and *Y*. Random walk with mild drift to the centre.

Exact estimation for Markov chain equilibrium expectations: comments

- Variance of estimator will grow if τ is large. Need fast coupling for low variance.
- As before, if the underlying space is complex, this could make coupling hard?
- Estimates may not always be physically meaningful: e.g. $h \geq 0$ does not mean $h(X_0) + \sum_{k=1}^{\tau-1}(h(X_k) - h(Y_k)) \geq 0$.
- Variance of the unbiased estimator seems to generally be higher.
- Active area of research for Markov Chain Monte Carlo (MCMC) and Hamiltonian Monte Carlo (HMC) (e.g. Jacob et al. '20, Heng et al. '19)

References

- Propp, J., and D. Wilson. ‘Exact Sampling with Coupled Markov Chains and Applications to Statistical Mechanics’. *Random Structures & Algorithms* (1996)
- Wang, A., M. Pollock, G. Roberts, and D. Steinsaltz. ‘Regeneration-Enriched Markov Processes with Application to Monte Carlo’. *The Annals of Applied Probability* (2021)
- Glynn, P., and C.-H. Rhee. ‘Exact Estimation for Markov Chain Equilibrium Expectations’. *Journal of Applied Probability* (2014)
- Jacob, P., F. Lindsten, and T. Schön. ‘Smoothing With Couplings of Conditional Particle Filters’. *Journal of the American Statistical Association* (2020)
- Jacob, P., J. O’Leary, and Y. Atchadé. ‘Unbiased Markov Chain Monte Carlo Methods with Couplings’. *Journal of the Royal Statistical Society: Series B* (2020)