It is readily checked that the analogue of Theorem 10.1 holds: (Z^{**}, Γ^{**}) is periodically taboo stationary with period d *if and only if* there is a pair (Z, Γ) such that

$$\mathbf{P}((\theta_{nd-h}Z, \Gamma - nd) \in \cdot | \Gamma > nd) \stackrel{tv}{\to} \mathbf{P}((\theta_{-h}Z^*, \Gamma^*) \in \cdot), \quad n \to \infty.$$

Also, it is readily checked that the analogue of Theorem 10.2 holds: (Z^{**}, Γ^{**}) is periodically taboo stationary with period d *if and only if* Γ^{**}/d is geometric and independent of $\overset{\leftrightarrow}{\theta}_{\Gamma^{**}} Z^{**}$. In the above we can replace the pair (Z^{**}, Γ^{**}) by a triple $(Z^{**}, S^{**}, \Gamma^{**})$; see Remark 10.1.

Let (Z, S, Γ) be taboo regenerative (in the wide sense or not) and assume that $\mathbf{P}^\circ(X_1 \in \cdot | \Gamma^\circ \geqslant X_1)$ is lattice with span d and that

$$\mathbf{P}^\circ(S_0 \in d\mathbb{Z}) = 1 \quad \text{and} \quad \mathbf{P}^\circ(\Gamma^\circ \in d\mathbb{Z} | \Gamma^\circ < X_1) = 1.$$

Let $(\tilde{Z}, \tilde{S}, \tilde{\Gamma})$ be as in Theorem 10.3 and let (Z^*, S^*, Γ^*) be as at $(10.7a)$. Put

$$\Gamma^{**} = d + [\Gamma^*/d]d \quad \text{and} \quad (Z^{**}, S^{**}) := \overset{\leftrightarrow}{\theta}_{-\Gamma^{**}} \overset{\leftrightarrow}{\theta}_{\tilde{\Gamma}}(\tilde{Z}, \tilde{S}).$$

Then Γ^{**}/d is geometric and independent of $(\tilde{Z}, \tilde{S}, \tilde{\Gamma})$, and thus we have that $(Z^{**}, S^{**}, \Gamma^{**})$ is periodically taboo stationary with period d.

Observe that $\Gamma^{**} - \Gamma^*$ is $[0, d)$ valued and independent of $(Z^{**}, S^{**}, \Gamma^{**})$ and that

$$(Z^{**}, S^{**}, \Gamma^{**}) = \overset{\leftrightarrow}{\theta}_{\Gamma^* - \Gamma^{**}}(Z^*, S^*, \Gamma^*).$$

From this, the lattice assumption, and Theorems 10.5 and 10.6 we obtain easily that Theorems 10.5 and 10.6 hold with (Z^*, S^*, Γ^*) replaced by $(Z^{**}, S^{**}, \Gamma^{**})$, that is, $(Z^{**}, S^{**}, \Gamma^{**})$ is a periodically taboo stationary *version* of (Z, S, Γ). Now the proof of Theorem 10.7 (with obvious modifications) yields the following result: if the conditions $(9.18a)$ through $(9.18e)$ in Theorem 9.4 hold, then for $h \in [0, \infty)$,

$$\mathbf{P}(\theta_{nd-h}(Z, S, \Gamma) \in \cdot | \Gamma > nd) \stackrel{tv}{\to} \mathbf{P}(\theta_{-h}(Z^{**}, S^{**}, \Gamma^{**}) \in \cdot)$$

as $n \to \infty$.

11 Perfect Simulation – Coupling From-the-Past

We shall end this final chapter by considering the simulation aspects of the above theory. Coupling from-the-past (Section 8) can be applied to finite state space Markov chains to generate the *stationary* version of a time-homogeneous chain, the *two-sided* version of a time-inhomogeneous chain, and the *taboo stationary* version of a time-homogeneous chain (in particular, the so-called *quasi-stationary* distribution).

In Section 6.1 of Chapter 8 we discussed briefly the general problem of generating a stationary version of a given stochastic process, and the same discussion applies with obvious modification to the two-sided time-inhomogeneous case and the taboo case. We then gave a solution for Palm duals with bounded cycle-lengths using the *acceptance–rejection* algorithm. At the end of this section this algorithm is applied together with the structural results of Section 10 to generate the taboo stationary version of a *taboo regenerative* process when the minimum of the cycle-length and the taboo time is bounded and the exponential parameter α is known.

An important distinction between the acceptance–rejection algorithm and the coupling from-the-past algorithm is that the coupling algorithm works without knowledge of α. An interesting common feature of the algorithms is that the transition probabilities (or cycle distribution) of the processes need not be known. The processes could, for instance, be the output of another simulation.

11.1 Generating a Stationary Finite-State Markov Chain

Consider a Markov chain in discrete or continuous time

$$Z = (Z_k)_0^\infty \quad \text{or} \quad Z = (Z_s)_{s\in[0,\infty)}$$

with a finite state space E. Assume that Z is irreducible and, in the discrete-time case, that Z is aperiodic. Suppose it is known how to generate Z starting from any given initial state i. Note that the problem of generating the stationary version Z^* of Z can be reduced to that of generating the stationary initial state Z_0^*. The stationary initial state Z_0^* can be generated by coupling from-the-past as follows (see Figure 11.1).

INITIAL STEP. Start a family of independent versions of Z in all states at time -1 and run them up to time 0 (that is, one transition in the discrete-time case).

RECURSIVE STEPS. For each $n \geqslant 2$, start a new family of independent versions of Z in all states at time $-n$ and run them up to time $-n+1$. From time $-n+1$ let the chains continue up to time 0 as follows: if the chain starting in state i at time $-n$ is in state j at time $-n+1$, let it continue up to time 0 along the path of the chain that starts at time $-n+1$ in state j.

TERMINATION CONDITION. Terminate the recursion at the first $n \geqslant 1$ such that all the chains that start at time $-n$ are in the same state at time 0. This state is a realization of the stationary state Z_0^*, according to the following theorem.

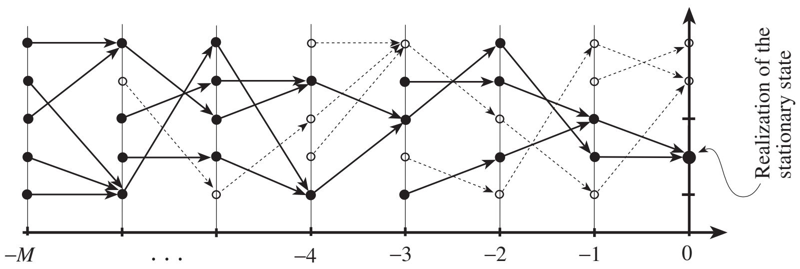

FIGURE 11.1. Coupling from-the-past.

Theorem 11.1. *Let M be the first $n \geqslant 1$ such that all the chains that start at time $-n$ are in the same state at time* 0*. Call this state Y. Then $\mathbf{P}(M < \infty) = 1$ and Y is a copy of Z_0^*.*

PROOF. Let $\mathbf{P}_i$ indicate that Z starts in the state i, that is, $\mathbf{P}_i(Z_0 = i) = 1$. According to Theorems 3.1 and 3.2 in Chapter 2, $\lim_{n\to\infty} \mathbf{P}_i(Z_n = j) = \mathbf{P}(Z_0^* = j) > 0$ for all states i and j. Thus there are j_0, n_0, and $p > 0$ such that

$$\mathbf{P}_i(Z_{n_0} = j_0) \geqslant p, \quad i \in E.$$

For each $n \geqslant n_0$ the probability that all the chains starting at time $-n$ are in the same state at time $-n + n_0$ is no less than the probability that independent chains starting from all states at time $-n$ are all in the state j_0 at time $-n + n_0$. Thus

$$\mathbf{P}(M > (k+1)n_0 | M > kn_0) \leqslant (1 - p^{\#E}), \quad k \geqslant 0.$$

This implies

$$\mathbf{P}(M > kn_0) \leqslant (1 - p^{\#E})^k \to 0 \text{ as } k \to \infty,$$

that is, $\mathbf{P}(M < \infty) = 1$.

Now fix an $i \in E$. Let $Z_0^{(-n,i)}$ be the state at time 0 of the chain that starts in i at time $-n$ and note that $Z_0^{(-n,i)} = Y$ on $\{M \leqslant n\}$. Thus, for each $j \in E$,

$$\begin{aligned}\mathbf{P}(Z_0^{(-n,i)} = j) &= \mathbf{P}(Z_0^{(-n,i)} = j, M > n) + \mathbf{P}(Y = j, M \leqslant n)\\ &\to \mathbf{P}(Y = j), \quad n \to \infty.\end{aligned}$$

But for each $j \in E$,

$$\mathbf{P}(Z_0^{(-n,i)} = j) = \mathbf{P}_i(Z_n = j) \to \mathbf{P}(Z_0^* = j) \text{ as } n \to \infty,$$

and thus $\mathbf{P}(Y = j) = \mathbf{P}(Z_0^* = j)$ as desired. $\square$

Remark 11.1. Why use coupling from-the-past and not ordinary coupling to-the-future? Simply because the latter does not work. The coupling time T, when the chains starting from all states at time zero merge, is also the state of the stationary chain. But T is random, and thus the stationary chain need not have the stationary distribution at time T. (In order to see this, consider a chain with state space $E = \{1, 2, 3\}$ which goes from 1 to 2 with probability $\frac{1}{2}$, from 1 to 3 with probability $\frac{1}{2}$, from 2 to 3 with probability 1, and from 3 to 1 with probability 1. The versions starting from all states at time zero merge in state 3. The nonrandom state 3 is not the stationary state.)

11.2 More Efficient Algorithm for Birth and Death Chains

The above algorithm shows that perfect simulation is theoretically possible, but the algorithm is not very efficient. Typically, in practical simulations the number of states is several thousands or millions, even trillions or more, and thus having to generate independent transitions from all states is astronomically expensive and time-consuming. However, in special cases there are efficient versions of the above algorithm.

For instance, if Z is a birth and death chain in continuous time, then we can (recursively for $n \geqslant 1$) generate two chains starting at time $-n$, one starting from top (from the highest state) and the other from bottom (from zero), and run each of them until time zero or until it merges with a chain starting at time $-k$ for some $k < n$. Repeat this until the first n such that the two chains starting at time $-n$ are in the same state at time zero. This common state is a realization of the stationary state because all chains coming in from the past (in particular the stationary chain) are captured by these chains and have to merge with them.

The same trick can be used for more complicated *monotone* chains, that is, chains having a partially ordered state space and transition probabilities that preserve this partial ordering. An example is the (finite) Ising model, a Markov chain with state space $\{-1, 1\}^{\{0,\ldots,k\}^d}$ (a state is a configuration of minus ones and ones indexed by a location in $\{0, \ldots, k\}^d$) and with the property that a state changes either by switching a single -1 to 1 or a single 1 to -1. The chains starting from top (with ones at all locations) and the chains starting from bottom (with zeros at all locations) will then capture all chains that come in from the past.

11.3 Generating a Two-Sided Time-Inhomogeneous Chain

The above coupling algorithm also works for time-inhomogeneous chains. To simplify the presentation we shall only consider the discrete-time case. Analogous results hold in the continuous-time case.