Models and Languages for Computational Systems Biology

Lecture 10: Stochastic Petri Nets and Markov Chains

Ian Stark

School of Informatics
The University of Edinburgh

Thursday 11 February 2010
Semester 2 Week 5

Stochastic Petri Nets

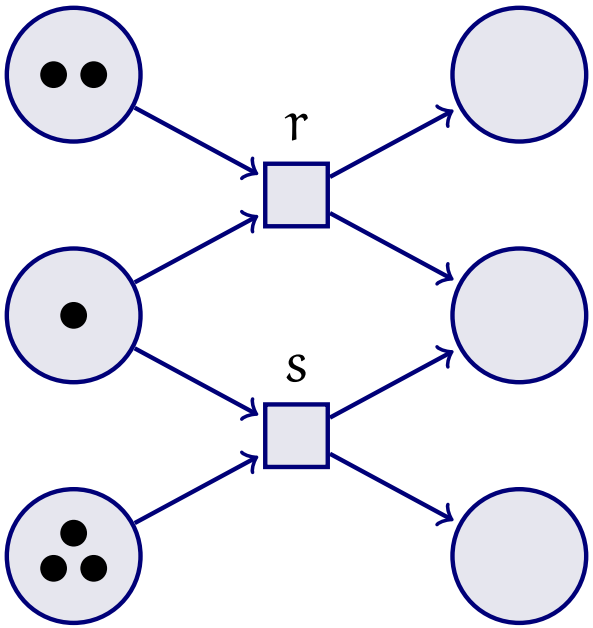

To introduce continuous time into a Petri nets, we label each transition with a rate.

When enabled, the transition fires at that rate (exponential distribution) until it is no longer enabled.

Where more than one transition is enabled, we have a race.

The combined rate of firing is $r + s$, with respective probabilities

$$\frac{r}{r+s} \quad \text{and} \quad \frac{s}{s+r}\,.$$

Simulation

Moving from a qualitative model (all possible actions are equal) to quantitative one (some are more equal than others), it is appropriate to consider simulation.

A single simulated run of a model only explores one possible outcome. Repeated simulations can explore many possible outcomes; and, moreover, by construction they will be clustered in regions of greater likelihood.

Such *ensemble* simulations can properly give information about the likely behaviour of the system being modelled.

They may involve either repeated runs with the same initial conditions, or with varying initial conditions according to some parameter scan.

For large, complex, systems, or large populations of small systems, simulation may the only feasible way to obtain any analysis at all.

In many cases, however, symbolic and numerical analysis can give results without simulation; or even results that cannot be obtained by simulation — so-called *beyond simulation*.

Markov chains

A (discrete-time) *Markov chain* is a transition system where transitions between states are labelled with probabilities.

For any given state, the exit transitions should have probabilities totalling 1; this may include a self-transition back to the same state.

$$M = (S, T : S \times S \to \mathbb{R})$$

A Markov chain can also be considered as sequence of random variables X_n drawn from S, for $n \in \mathbb{N}$:

$$P(X_{n+1} = t \mid X_n = s_n \ldots X_0 = s_0) = P(X_{n+1} = t \mid X_n = s_n) = T(s, t) \ .$$

Each transition is independent of all previous ones: the *Markov property*.

Markov chains as matrices

$$M = (S, T : S \times S \to \mathbb{R})$$

The description of a Markov chain can readily be given as a $|S| \times |S|$ matrix of probabilities.

Such a matrix will necessarily be a *stochastic matrix*: all its elements are non-negative, and all rows sum to 1. These have particular properties:

- The product of two stochastic matrices is also a stochastic matrix.
- Every eigenvalue λ of a stochastic matrix has $|\lambda| \leqslant 1$.
- Every stochastic matrix has at least one eigenvalue equal to exactly 1.

For a Markov chain: the powers of its stochastic matrix represent repeated steps of the chain; the eigenvalues of norm below 1 represent transient behaviour; and eigenvalues of 1 represent steady-state behaviour.

Stationary and limit distributions

The norm-1 eigenvectors of the stochastic matrix for a Markov chain correspond to *stationary distributions* of the chain: if probability mass is distributed over the states of the chain according to a stationary distribution, then this remains unchanged after every successive transition.

Such stationary distributions π are exactly the solutions of

$$\pi = \mathrm{T}\pi, \quad \text{equivalently} \quad (\mathrm{T} - \mathrm{I})\pi = 0 \; .$$

It may be that for some, or any, initial distribution π_0 the iterated distribution $\mathrm{T}^n\pi_0$ after n steps of the chain does converge as $n \to \infty$ to a stationary distribution, the *equilibrium distribution* of the chain.

A Markov chain is *ergodic* if it converges to a unique stationary distribution. Equivalently, if its stochastic matrix has exactly one norm-1 eigenvector.

Markov decision processes

A *Markov Decision Process* (MDP) is a labelled transition system, where the transitions are also assigned probabilities.

$$M = (S, L, T : S \times L \times S \to \mathbb{R})$$

Labels represent actions under external control, so for any state, the probabilities for all transitions leaving that state with a given label must total to 1.

$$\forall s \in S, a \in L \,.\, \left(\sum_{t \in S} T(s, a, t) = 1 \right)$$

Continuous Time Markov Chains

A *Continuous Time Markov Chain* (CTMC) is a regular Markov chain, but where the transitions arise from exponentially-distributed events occurring as time passes.

Thus each transition is labelled by a *rate*, rather than a probability; the *exit rate* of a state is the total of all these rates; and the relative probabilities of different transitions are determined by their relative rates, as in a race.

As the (qualitative behaviour) of a Petri Net gives rise to a transition system, so the (quantitative) behaviour of a Stochastic Petri Net is described by a Continuous Markov Chain.

Reading

Homework

Read §§1–3 of the following tutorial article.

Marta Kwiatkowska, Gethin Norman, and David Parker.
Probabilistic model checking for systems biology.
In *Symbolic Systems Biology: Theory and Methods*. Jones and Bartlett, 2010. To appear.

You can find more about the *Prism* model checker at http://www.prismmodelchecker.org.