APTS Applied Stochastic Processes, Warwick, July 2013
Exercise Sheet for Assessment

The work here is "light-touch assessment", intended to take students up to half a week to complete. Students should talk to their supervisors to find out whether or not their department requires this work as part of any formal accreditation process (APTS itself has no resources to assess or certify students). It is anticipated that departments will decide the appropriate level *of assessment locally, and may choose to drop some (or indeed all) of the parts, accordingly.*

Students are recommended to read through the relevant portion of the lecture notes before attempting each question. It may be helpful to ensure you are using a version of the notes put on the web *after* the APTS week concluded.

1 Markov chains, reversibility and Poisson processes

Suppose that visitors to a museum arrive as a Poisson process of rate λ. Each visitor spends an $\mathrm{Exp}(\mu)$ time looking around, independently of the other visitors, and then leaves. Let X_t be the number of customers inside the museum at time $t \geq 0$. We assume that $\lambda, \mu > 0$.

(a) Convince yourself that $(X_t)_{t\geq 0}$ is a continuous-time Markov chain, and write down its transition rates.

(b) Using detailed balance, find the stationary distribution for $(X_t)_{t\geq 0}$. Do you need any constraints on the parameters λ and μ?

(c) Describe the process of departures from the museum in stationarity.

2 Martingale convergence

Imagine that an urn initially contains α red balls and β blue balls. At each step $n \geq 1$, you pick a ball uniformly at random from the urn and put it back in together with another ball of the same colour. For $n \geq 1$, let

$$X_n = \begin{cases} 1 & \text{if a red ball is picked at step n} \\ 0 & \text{otherwise} \end{cases}$$

and let M_n be the proportion of red balls in the urn, i.e. $M_0 = \frac{\alpha}{\alpha+\beta}$ and, for $n \geq 1$,

$$M_n = \frac{\alpha + \sum_{i=1}^n X_i}{\alpha + \beta + n}.$$

(a) Show that M is a martingale.

(b) Explain why M_n converges almost surely as $n \to \infty$.

(c) Show that for any sequence $x_1, x_2, \ldots, x_n \in \{0, 1\}$ such that $\sum_{i=1}^n x_i = k$, we have

$$\mathbb{P}\left[X_1 = x_1, \ldots, X_n = x_n\right] = \frac{\Gamma(\alpha+k)\Gamma(\beta+n-k)\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(\alpha+\beta+n)},$$

where Γ denotes the usual gamma function.

(d) Suppose now that $Y_1, Y_2, \ldots$ is a sequence of i.i.d. Bernoulli random variables with unknown parameter Θ. Put prior distribution Beta(α, β) on Θ. Show that the posterior distribution of $Y_1, Y_2, \ldots, Y_n$ is the same as the distribution of $X_1, X_2, \ldots, X_n$ for all $n \geq 1$.

(e) Conditionally on $\Theta = \theta$, the strong law of large numbers implies that

$$\frac{1}{n}\sum_{i=1}^{n} Y_i \to \theta \quad \text{almost surely as } n \to \infty.$$

So, unconditionally,

$$\frac{1}{n}\sum_{i=1}^{n} Y_i \to \Theta \quad \text{almost surely as } n \to \infty.$$

Deduce that $\Xi = \lim_{n\to\infty} M_n$ has Beta(α, β) distribution.

Note that α and β don't need to be integer-valued for all of this to work; $\alpha, \beta > 0$ would suffice (although we would have to come up with an interpretation for a non-integer number of balls!). The result in this question is an aspect of the theory of exchangeable random variables. *In particular, the last part is a special case of* de Finetti's theorem. *If you'd like to know more about the application of these concepts in Statistics, take a look at* Theory of Statistics *by M.J. Schervish (Springer, 1996).*

3 Recurrence and small sets

Define a Markov chain X taking values in $[0, \infty)$ as follows: for $n \geq 0$,

$$X_{n+1} = U_{n+1} X_n + E_{n+1},$$

where X_n, U_{n+1} and E_{n+1} are mutually independent, $U_{n+1} \sim \mathrm{U}[0,1]$ and $E_{n+1} \sim \mathrm{Exp}(1)$. It can be shown that X has transition density

$$p(x, y) = \frac{1}{x} e^{-y}(e^{x\wedge y} - 1), \quad x, y \geq 0,$$

where $x \wedge y = \min\{x, y\}$. *(If you're feeling keen, you might want to check this, but it's a rather tedious calculation!)*

(a) Argue that X is μ-irreducible, where μ denotes the Lebesgue measure on $[0, \infty)$.

(b) Show that any set of the form $[0, a]$ with $a > 0$ is small of lag 1. *(Hint: you may find it useful to observe that $\frac{1}{x}(e^{x\wedge y} - 1) \geq 1 \wedge (y/x)$.)*

(c) Show that X satisfies the Foster-Lyapunov criterion for positive recurrence. *(Hint: take $\Lambda(x) = x$.)*

(d) Show, using detailed balance, that the Gamma$(2, 1)$ distribution (which has density $\pi(x) = xe^{-x}$ for $x \geq 0$) is stationary for X.
(Another way to see this is to observe first that if $U \sim \mathrm{U}[0,1]$ and $\Gamma \sim \mathrm{Gamma}(2,1)$ then $U\Gamma \sim \mathrm{Exp}(1)$ and then to use the fact that the sum of two independent Exp(1) *random variables is* Gamma(2, 1).*)*

(e) Show that X satisfies the Foster-Lyapunov criterion for geometric ergodicity. *(Hint: take $\Lambda(x) = e^{x/2}$ and $C = \{x : \Lambda(x) \leq e^{2/\gamma}\}$ for any $\gamma \in (0, 1)$. You may find it helpful to note that $\mathbb{E}\left[e^{tU_1}\right] = \frac{1}{t}(e^t - 1)$ and that $\mathbb{E}\left[e^{tE_1}\right] = \frac{1}{1-t}$ for $t < 1$.)*