STOCHASTIC NETWORKS EXAMPLE SHEET 2 SOLUTIONS

ELENA YUDOVINA

Exercise 1. Show that the traffic equations for an open migration process have a unique solution, and that this solution is positive. [*Hint:* From the irreducibility of the open migration process deduce the irreducibility of a certain Markov process on $J+1$ states, and then use the fact that the equilibrium distribution for this process is unique.]

Deduce that in an open migration process $\alpha_j \lambda_j$ is the mean arrival rate at colony j, counting arrivals from outside the system and from other colonies.

Proof. Consider the Markov process on $J+1$ states corresponding, loosely, to the case of $N = 1$ person in the open migration network (state $J+1$ corresponds to "rest-of-the-universe"). The transition rates will be

$$\begin{aligned} q(j,k) &= \lambda_{jk} \quad j,k = 1,\dots,J \\ q(j, J+1) &= \mu_j \quad j = 1,\dots,J \\ q(J+1, j) &= \nu_j \quad j = 1,\dots,J \end{aligned}$$

I claim that this network is irreducible. Indeed, note that in the original network the states e_j = a single individual in colony j along with the state 0 = nobody in any colony all communicate – possibly going through states with multiple individuals in the system. In particular, everything communicates with state 0. (State e_j corresponds to j, and 0 corresponds to $J+1$.)

Now, to get from e_j to 0 we need the individual in j to leave somewhere. Let us simply trace his path through the network: it must be along edges with positive rates, and it must terminate with a transition out of the network. If we simply look at the corresponding transitions in the modified Markov process, they will have positive rates too (potentially different ones – if we had a few arrivals into state j first, the rate of leaving it may have greatly increased or decreased; however, it couldn't have become non-zero if the transition rate in our process is zero). Thus, we have shown that it is possible to get from j to $J+1$ for all $j = 1,\dots,J+1$. The reverse direction follows by tracing back the path along which the individual who is in colony j when the system state is e_j came into the network.

Now, the equilibrium distribution $(\tilde{\alpha}_j)$ of the modified Markov process satisfies

$$\begin{aligned} &\tilde{\alpha}_j\Big(\sum_{k=1}^{J} \lambda_{jk} + \mu_j\Big) = \sum_{k=1}^{J} \tilde{\alpha}_j \lambda_{kj} + \tilde{\alpha}_{J+1}\nu_j, \quad j = 1,\dots,J \\ &\tilde{\alpha}_{J+1}\Big(\sum_{k=1}^{J} \nu_j\Big) = \sum_{k=1}^{J} \tilde{\alpha}_j \mu_j \\ &\sum_{j=1}^{J+1} \tilde{\alpha}_j = 1 \end{aligned}$$

Since the Markov process is irreducible on a finite state space, there exists a unique distribution α_j with these properties, and it will then assign positive probabilities to all states (as otherwise the chain would have to be reducible).

Now, the solution to the traffic equations (α_j) corresponds to the first two equations above, but with a different normalisation, namely, $\alpha_{J+1} = 1$. Clearly, we can go between the two: given α_j we can produce $\tilde{\alpha}_j = \alpha_j/(1+\sum_j \alpha_j)$ (with $\tilde{\alpha}_{J+1} = 1/(1+\sum_j \alpha_j)$), and conversely given $\tilde{\alpha}_j$ we can define $\alpha_j = \tilde{\alpha}_j/\alpha_{J+1}$.

Thus, we conclude existence, uniqueness, and positivity of the solution to the traffic equations.

Finally, the mean arrival rate into colony j is

$$\sum_{\mathbf{n}} \pi(\mathbf{n})(\sum_k q(T_{jk}\mathbf{n}, \mathbf{n})+\nu_j) = \sum_{\mathbf{n}} \pi(\mathbf{n})(\sum_k \lambda_{kj}\phi_k(n_k+1))+\nu_j = \sum_k \lambda_{kj} \sum_{\mathbf{n}} \pi(\mathbf{n})\phi_k(n_k+1))+\nu_j$$
$$= \sum_k \alpha_k\lambda_{kj} + \nu_j = \alpha_j\lambda_j.$$

□

Exercise 2. Show that the reversed process obtained from a stationary closed migration process is also a closed migration process, and determine its transition rates.

Proof. Recall that the form of the equilibrium distribution for a closed migration process is

$$\pi(\mathbf{n}) = B_N \prod_{j=1}^{J} \frac{\alpha_j^{n_j}}{\prod_{r=1}^{n_j} \phi_j(r)}.$$

In that case, the transition probabilities for the reversed process are

$$q'(\mathbf{n}, T_{jk}(\mathbf{n})) = \frac{\pi(T_{jk}(\mathbf{n}))}{\pi(\mathbf{n})} q(T_{jk}(\mathbf{n}), \mathbf{n}) = \phi_k(n_k+1)/\pi_j(n_j)\lambda_{kj}\phi_k(n_k+1) = \phi_j(n_j)\frac{\alpha_k/\alpha_j}{\lambda \quad kj}.$$

Note that this has the required form, with $\phi'_j(n_j) = \phi_j(n_j)$ and $\lambda'_{jk} = \frac{\alpha_k}{\alpha_j}\lambda_{kj}$. □

Exercise 3. A restaurant has N tables, with a customer seated at each table. Two waiters are serving them. One of the waiters moves from table to table taking orders for food. The time that he spends taking orders at each table is exponentially distributed with parameter μ_1. He is followed by the wine waiter who spends an exponentially distributed time with parameter μ_2 taking orders at each table. Customers always order food first and then wine, and orders cannot be taken concurrently by both waiters from the same customer. All times taken to order are independent of each other. A customer, after having placed her two orders, completes her meal at rate ν, independently of the other customers. As soon as she finishes her meal, she departs and a new customer takes her place and waits to order. Model this as a closed migration process. Show that the stationary probability that both waiters are busy can be written in the form

$$\frac{G(N-2)}{G(N)} \cdot \frac{\nu^2}{\mu_1\mu_2},$$

for a function $G(\cdot)$, which may also depend on ν, μ_1, μ_2, to be determined.

In the above model it is assumed that the restaurant is always full. Develop a model in which this assumption is relaxed: for example, assume that customers enter the restaurant at rate λ while there are tables empty. Again obtain an expression for the probability that both waiters are busy.

Proof. The closed migration model will have three colonies: colony 1 "not-yet-ordered", colony 2 "ordered food but not wine", and colony 3 "ordered both". The transition rates are

$$\begin{aligned} \lambda_{12} &= \mu_1 & \phi_1 &= I[n_1 > 0] \\ \lambda_{23} &= \mu_2 & \phi_2 &= I[n_2 > 0] \\ \lambda_{31} &= \nu & \phi_3 &= n_3. \end{aligned}$$

(You may notice that this looks strikingly familiar from the lectures.) We conclude that $\alpha_0 : \alpha_1 : \alpha_2 = \mu_1^{-1} : \mu_2^{-1} : \nu^{-1}$, and

$$\pi(n_1, n_2, n_3) = C(N)^{-1} \frac{1}{\mu_1^{n_1} \mu_2^{n_2} \nu^{n_3} n_3!}$$

for the appropriate normalising constant $C(N)$. The probability in which we are interested is

$$\begin{aligned} \mathbb{P}(\text{both waiters busy}) &= \sum_{\substack{n_1>0, n_2>0, n_3 \\ n_1+n_2+n_3=N}} \pi(n_1, n_2, n_3) = C(N) \sum_{\substack{n_1>0, n_2>0, n_3 \\ n_1+n_2+n_3=N}} \frac{1}{\mu_1^{n_1} \mu_2^{n_2} \nu^{n_3} n_3!} \\ &= C(N)^{-1} \sum_{\substack{n_1', n_2', n_3 \\ n_1'+n_2'+n_3=N-2}} \frac{1/(\mu_1\mu_2)}{\mu_1^{n_1'} \mu_2^{n_2'} \nu_2^{n_3} n_3!} = \frac{1}{\mu_1\mu_2} \frac{C(N-2)}{C(N)}. \end{aligned}$$

Now, this has almost the right shape, but it would be nicer to express the answers as dimensionless – for example, if we were to try and scale the quantities in the problem by a constant factor, it would be nice to be able to easily see that this probability is unchanged. In light of this, we define

$$G(N) = \sum_{n_1+n_2+n_3=N} \frac{1}{n_3!} \left(\frac{\nu}{\mu_1}\right)^{n_1} \left(\frac{\nu}{\mu_2}\right)^{n_2} = \nu^N C(N),$$

which is obviously dimensionless, and for which the desired form holds. (Quantities of the form $\rho = \nu/\mu$, the load, are very natural in queueing theory.)

If we allow customers to leave the system after eating and to arrive to empty tables, then we will need to have an extra colony, 0 "table empty" in our system. The transition rates become

$$\begin{aligned} \lambda_{01} &= \lambda & \phi_0 &= I[n_0 > 0] \\ \lambda_{12} &= \mu_1 & \phi_1 &= I[n_1 > 0] \\ \lambda_{23} &= \mu_2 & \phi_2 &= I[n_2 > 0] \\ \lambda_{30} &= \nu & \phi_3 &= n_3. \end{aligned}$$

This gives

$$\alpha_0 : \alpha_1 : \alpha_2 : \alpha_3 = \frac{1}{\lambda} : \frac{1}{\mu_1} : \frac{1}{\mu_2} : \frac{1}{\nu}$$

and therefore

$$\pi(n_0, n_1, n_2, n_3) = C'(N)^{-1} \frac{1}{\lambda^{n_0} \mu_1^{n_1} \mu_2^{n_2} \nu^{n_3} n_3!}$$

We are still interested in summing over the states with $n_1 > 0$ and $n_2 > 0$, which will still give an expression of the form $\frac{C(N-2)}{C(N)} \frac{1}{\mu_1\mu_2}$, or the same terminal form as before, for

$$G'(N) = \nu^N C'(N) = \sum_{n_0+n_1+n_2+n_3=N} \frac{1}{n_3!} \left(\frac{\nu}{\lambda}\right)^{n_0} \left(\frac{\nu}{\mu_1}\right)^{n_1} \left(\frac{\nu}{\mu_2}\right)^{n_2}$$

Note that we can also define $G'(N) = \lambda^N C'(N)$ and get $\lambda^2/(\mu_1\mu_2)$ as the second factor. □

Exercise 4. Show that if the parameters of a stationary open migration process are such that there is no path by which an individual leaving colony k can later reach colony j, then the stream of individuals moving directly from j to k forms a Poisson process.

Proof. Note first of all that this does not show the system is reducible – all states can still communicate with the state "system empty".

Consider the communicating class of colony j, i.e. the set of colonies j' such that an individual can get from j to j' and back again without leaving the system. Observe that if we look only at the individuals in these colonies, then we have a Markov process, and its transition rates mean that it is itself an open migration process. Consequently, the exit stream from colony j in this process is Poisson. Now, the stream of individuals moving directly from j to k is a probabilistic thinning of this exit stream, and therefore also Poisson. □

Exercise 5. Airline passengers arrive at a passport control desk in accordance with a Poisson process of rate ν. The desk operates as a single-server queue at which service times are independent and exponentially distributed with mean μ ($< \nu^{-1}$) and are independent of the arrival process. After leaving the passport control desk a passenger must pass through a security check. This also operates as a single-server queue, but one at which service times are all equal to τ ($< \nu^{-1}$). Show that in equilibrium the probability that both queues are empty is

$$(1 - \nu\mu)(1 - \nu\tau).$$

It if takes a time σ to walk form the first queue to the second, what is the equilibrium probability that both queues are empty and there is no passenger walking between them?

Proof. To solve the first part of the problem, recall the result that queues in a series of $M/M/1$ queues are independent, and note that we can make the last queue in that a $\cdot/D/1$ queue without changing anything. Thus, the probability that both queues are empty is

$$\mathbb{P}(\text{both empty}) = \mathbb{P}(\text{queue 1 empty})\mathbb{P}(\text{queue 2 empty})$$

(Since the arrivals to both queues are Poisson, the probability that the queue is empty when an arrival happens is the same as the stationary probability of its being empty.) For queue 1 we know this probability exactly to be $(1-\nu\mu)$, but in fact it is insensitive to the service time distribution. Indeed, we can compute the stationary probability that a queue with mean service time τ is empty from Little's law: let the "number in the system" be the number of customers in service (0 or 1), in which case the probability that the system is occupied is the expected number in system L. The sojourn time is the service time, so $W = \tau$; the arrival rate is still ν, assuming the queue is stable (the customers must come up for service about as often as they come into the queue). Thus, $L = \nu\tau$, and the probability of the server being idle is $(1 - \nu\tau)$.

Adding a walking time is like adding an infinite-server queue with deterministic service time into the series queue. The event that the system is empty corresponds to the following: firstly, we need queues 1 and 2 to be empty at time 0, and secondly, we need no departures from queue 1 between $-\sigma$ until 0. Now that the input stream for queue 2 is simply the output stream from queue 1 shifted by σ; so the state of queue 2 at time 0 is determined by the departures from queue 1 during the time $(-\infty, -\sigma]$. The number of people walking is determined by the departures from queue 1 during the time $(-\sigma, 0)$. These are independent of each other and of the state of queue 1 at time 0. Therefore, the desired probability is

$$(1-\nu\mu)(1-\nu\tau)e^{-\nu\sigma}.$$

□

Exercise 6. Cars arrive at the beginning of a long road in a Poisson stream of rate ν from time $t = 0$ onwards. A car has a fixed velocity $V > 0$ which is a random variable. The velocities of cars are independent and identically distributed, and independent of the arrival process. Cars can overtake each other freely. Show that the number of cars on the first x miles of road at time t has a Poisson distribution with mean

$$\nu\mathbb{E}\left[\frac{x}{V} \wedge t\right].$$

Proof. By the theorem from lecture, we know that the number of cars on the first x miles of road at time t has a Poisson distribution with mean

$$M(t, [0, x]) = \nu \int_0^t P(u, [0, x])du,$$

where $P(u, [0, x])$ is the probability that a car is in the set $[0, x]$ a time u after entering the system.

Now, a car is in the set $[0, x]$ a time u after entering the system if and only if its speed V satisfies $V \leq x/u$, or equivalently $x/V \geq u$. Thus, we have

$$M(t, [0, x]) = \nu \int_0^t \mathbb{P}(\frac{x}{V} \geq u)du = \nu\mathbb{E}\left[\frac{x}{V} \wedge t\right].$$

(Here, $\wedge$ stands for minimum.) To see the equality, recall that for a random variable X we have $\mathbb{E}[X] = \int_0^\infty \mathbb{P}(X > x)dx$; here we simply use the fact that, for the random variable $X = \min(x/V, t)$ the probability that $X > t$ will be 0, so we can truncate the integral. □

Exercise 7. Recall the mathematical model for a loss network with fixed routing.

A network consists of three nodes, with each pair of nodes connected by a link. A call in progress between two nodes may be routed on the direct link between the nodes, on on the two link path through the third node. A call in progress can be rerouted if this will allow an additional arriving call to be accepted. Describing carefully the modelling assumptions you make, obtain an exact expression for the probability an arriving call is lost, and sketch a network with fixed routing which shares the same loss probabilities. Deduce an Erlang fixed point approximation for the loss probabilities in the original network, in terms of blocking probabilities across certain *cuts* in the network.

Proof. This problem is solved in the lecture notes. We will assume Poisson arrivals of calls (with rates ν_{ij} for calls going between i and j), exponential holding times of rate 1 for each call, and independence of lots of things. □

Exercise 8. Calls arrive as a Poisson process of rate ν at a link with capacity C circuits. A call is blocked and lost if all C circuits are busy; otherwise the call is accepted and occupies a single circuit for an exponentially distributed holding time with mean one. Holding times are independent of each other and of the arrival process. Calculate the mean number of circuits in use, $M(\nu, C)$. Show that, as $n \to \infty$,

$$\frac{M(\nu n, Cn)}{n} \to \nu \wedge C,$$

and the mean number of idle circuits satisfies

$$Cn - M(\nu n, Cn) \to \frac{C}{\nu \vee C - C}.$$

Proof. From the Erlang formula,

$$M(\nu, C) = \sum_{k=1}^{C} k \frac{\nu^k/k!}{\sum_{k=0}^{C} \nu^k/k!}$$

The first method of proof is to observe that

$$M(\nu, C) = \nu(1 - \mathcal{E}(\nu, C))$$

and then

$$\mathcal{E}(\nu n, Cn) = \frac{(\nu n)^{Cn}/(Cn)!}{\sum_{k=0}^{nC} \nu^k/k!} = \frac{1}{1 + \frac{Cn}{\nu n} + \frac{(Cn)(Cn-1)}{(\nu n)^2} + \ldots} \to \frac{1}{1 + \frac{C}{\nu} + \frac{C^2}{\nu^2} + \ldots} = \begin{cases} 0, & C \geq \nu \\ 1 - C/\nu, & C < \nu \end{cases}$$

(This is monotone convergence – the denominator converges to the geometric series term-by-term from below.) Then

$$\frac{M(\nu, C)}{n} \to \nu \wedge C$$

as required.

An alternative method of proof is to show that the distribution of the number of busy circuits is tightly centered on the smaller of Cn and νn. Suppose first that $\nu > C$, and let $C/\nu = \alpha < 1$. Consider k with $\frac{k}{n} < C$. Then

$$\pi(k-1) = \pi(k) \cdot \frac{k}{\nu n} \leq \pi(k) \cdot \frac{C}{\nu} = \pi(k) \cdot \alpha \leq \pi(Cn) \cdot \alpha^{Cn-k}.$$

Summing over all $k < (1-\epsilon)Cn$, we get

$$\sum_{k<(1-\epsilon)Cn} \pi(k) < \pi(C) \sum_{k=0}^{(1-\epsilon)Cn} \alpha^{Cn-k} < \pi(C)\alpha^{\epsilon n} \sum_{k'=0}^{\infty} \alpha^{k'} = \pi(C) \frac{\alpha^{\epsilon n}}{1-\alpha}.$$

Note that as $n \to \infty$ this tends to 0, for any ϵ. Consequently, for any $\epsilon > 0$ and any $\delta > 0$, for large enough n, the probability that the number of busy circuits is at least $(1-\epsilon)Cn$ will be at least $1 - \delta$. Therefore,

$$\frac{1}{n} M(\nu n, Cn) \to C.$$

Remark. We do not need to worry about tightness of measures here – we are interested in the distribution of $\frac{k}{n}$, which lives on the compact set $[0, C]$. We've shown that the probability that it lives outside of $[C(1-\epsilon), C]$ can be made arbitrarily small, which tells us that its average will converge to C.

On the other hand, if $\nu < C$, we consider the probability that $\frac{k}{n} < (1-\epsilon)\nu$ or $\frac{k}{n} > (1-\epsilon)^{-1}\nu$. In the first case, the calculation will be the same, using the bound $\pi(k-1) = \pi(k) \cdot \frac{k}{\nu n} < \pi(k) \cdot (1-\epsilon)$ for $\frac{k}{n} \leq (1-\epsilon)\nu$. This gives

$$\pi(k) < \pi((1-\epsilon)\nu n) \cdot (1-\epsilon)^{(1-\epsilon)\nu n - k}$$

and the computation continues as above to give a total probability that tends to 0 as $n \to \infty$ for any fixed ϵ. For the upper bound, we note $\pi(k+1) = \pi(k) \cdot \frac{\nu n}{k+1} < \pi(k) \cdot (1-\epsilon)$ if $k > (1-\epsilon)^{-1}\nu n$, and repeat the calculation. Thus, if $\nu < C$ we have shown

$$\frac{1}{n} M(\nu n, Cn) \to \nu.$$

The case $\nu = C$ can be incorporated into either of these.

For the mean number of idle circuits, first note that if $\nu < C$, i.e. $\nu \vee C = \max(\nu, C) = C$, then $M(\nu n, Cn)/n \to \nu$ and therefore $Cn - M(\nu n, Cn) = Cn - \nu n + o(n) \to \infty$. Thus, we are only interested in the case $\nu \geq C$.

In this case, as we saw above, the probability that $k < (1-\epsilon)Cn$ circuits are occupied can be made arbitrarily small. Consequently, for the system as a whole the arrival rate is νn and the service rate is between $(1-\delta)(1-\epsilon)Cn$ and Cn. Thus, the free circuits in the system behave (almost) as an $M/M/1$ queue with arrival rate Cn and service rate νn. Their mean, therefore, converges to

$$\frac{Cn}{\nu n - Cn} = \frac{C}{\nu - C}.$$

(This could be made precise with a coupling argument: think of a queue with arrival rate $(1-\delta)(1-\epsilon)Cn$ for the lower bound, and with arrival rate Cn for the upper bound.) □

Exercise 9. Show that the solutions of the fixed point equation

$$B = \mathcal{E}(\nu[1 + 2B(1-B)], C)$$

locate stationary points of the potential function

$$\nu \left[e^{-y} + e^{-2y}(1 + \frac{2}{3}e^{-y}) \right] + \int_0^y U(z, C) dz.$$

Proof. Differentiate the potential function with respect to y, and set to 0:

$$0 = -\nu e^{-y} - 2\nu e^{-2y}(1 - \frac{2}{3}e^{-y}) + \frac{2}{3}\nu e^{-2y} e^{-y} + U(y, C)$$

Simplifying:

$$\nu e^{-y} + 2\nu e^{-2y} - 2\nu e^{-3y} = U(y, C).$$

Letting $B = 1 - e^{-y}$, the left-hand side becomes

$$\nu \left((1-B) + 2(1-B)^2 - 2(1-B)^3 \right) = \nu(1-B)(1 + 2B(1-B)),$$

while the right-hand side is $U(-\log(1-B), C)$. The statement that

$$\nu(1-B)(1 + 2B(1-B)) = U(-\log(1-B), C)$$

precisely expresses the fact that

$$B = \mathcal{E}(\underbrace{\nu[1 + 2B(1-B)]}_{\nu'}, C) :$$

indeed,

$$U(-\log(1-B), C) = U(-\log(1-\mathcal{E}(\nu', C)), C) = \underbrace{\nu[1+2B(1-B)]}_{\nu'}(1-\mathcal{E}(\nu')], C) = [1+2B(1-B)](1-B).$$

$\square$