2 CRNs as Markov processes

In order to properly analyse stochastic models of CRNs, we need some theory on **continuous-time Markov chains**, which we will abbreviate as **CTMCs**. Much of this theory takes us well beyond the scope of this course, so we will only outline the most essential parts of it, and in a form most suited to the application to CRNs.

Throughout this chapter, we consider a system of m chemical reactions on n species $A_1, \dots, A_n$ in a well-mixed compartment so we can ignore spatial dynamics. Let $X_i(t)$ be the random variable denoting the number of molecules of species A_i at time $t \geq 0$, so that $\mathbf{X}(t) := (X_1(t), \dots, X_n(t))^t$ is a stochastic process with outputs in (some subset of) $\mathbb{Z}^n_{\geq 0}$.

Recall our basic assumption about a well-mixed system: that the intensity of any reaction *depends only on the current state of the system* i.e., the current species numbers. We can make this precise by stipulating that our models of CRNs define time-homogeneous Markov chains. In particular, the stochastic process $\mathbf{X}(t)$:

- has the **Markov property**: given $k \in \mathbb{N}$, $0 \leq t_1 < t_2 < \cdots < t_{k+1} < \infty$, $A \subseteq \mathbb{Z}^n_{\geq 0}$ and $s_1, \dots, s_k \in \mathbb{Z}^n_{\geq 0}$, then

$$\mathbb{P}(\mathbf{X}(t_{k+1}) \in A \,|\, \mathbf{X}(t_k) = s_k, \dots, \mathbf{X}(t_1) = s_1) = \mathbb{P}(\mathbf{X}(t_{k+1}) \in A \,|\, \mathbf{X}(t_k) = s_k)\,;$$

- is **time-homogeneous**: given $t, u > 0$, $A \subseteq \mathbb{Z}^n_{\geq 0}$ and $s \in \mathbb{Z}^n_{\geq 0}$, then

$$\mathbb{P}(\mathbf{X}(t+u) \in A \,|\, \mathbf{X}(u) = s) = \mathbb{P}(\mathbf{X}(t) \in A \,|\, \mathbf{X}(0) = s)\,.$$

It is also very reasonable, and technically useful, to rule out "instantaneous states": if the system is in a state x, then with probability 1 the transition time after which it leaves this state will be positive (possibly infinite), but *not* zero. Consequently the paths $t \mapsto \mathbf{X}(t)$ are **right-continuous** step functions, i.e., $\mathbf{X}(t) = \lim_{u \searrow t} \mathbf{X}(u)$. This property is automatic when we consider systems with only a finite number of distinct reactions. It can also hold for systems with infinitely many reactions, but then we need to think more carefully about the intensities.

2.1 Transition probabilities

In order to understand the evolution of a CRN, we would like to know the probability of finding the system in a state x at time t, given that it was in state y at time $s \leq t$. Let $S \subseteq \mathbb{Z}^n_{\geq 0}$ be the set of possible states of the system; then we would like to know

$$\mathbb{P}(\mathbf{X}(t) = x \,|\, \mathbf{X}(s) = y)\,.$$

for all $x, y \in S$ and all times $0 \leq s \leq t$. These probabilities are called **transition probabilities**. They can be regarded as the "solution" of a stochastic model of a CRN. By analogy with the deterministic case, we would hope that knowing the reaction intensities would allow us to calculate the transition probabilities uniquely. As we might expect, there are some practical and theoretical difficulties we face trying to work out transition probabilities.

The transition matrix. The Markov assumptions mean that given $t \geq s$,

$$\mathbb{P}(\mathbf{X}(t) = x \,|\, \mathbf{X}(s) = y) = \mathbb{P}(\mathbf{X}(t-s) = x \,|\, \mathbf{X}(0) = y)$$

and so, in order to know all transition probabilities, it is sufficient if we can find

$$p_t(y, x) := \mathbb{P}(\mathbf{X}(t) = x \,|\, \mathbf{X}(0) = y\}\,.$$

for all states $x, y \in S$, and times $t \geq 0$. Since the state space S is discrete, we can put these probabilities together into a (possibly infinite) time-dependent matrix, say P_t. We will often also write this transition matrix as $P(t)$. Clearly P_0 is the identity matrix, denoted $\mathbb{I}$. Moreover, by a familiar application of the law of total probability, the transition probabilities satisfy the **Chapman-Kolmogorov equations**:

$$p_{t+u}(x, y) = \sum_{z \in S} p_t(x, z) p_u(z, y)\,. \tag{5}$$

The Chapman-Kolmogorov equations are equivalent to the **semigroup property** $P_{t+u} = P_t P_u$ for any $t, u \geq 0$.

If we have managed to calculate P_t for all t, then given any initial measure, say ϕ_0, on S, we can calculate $\phi_t = \phi_0 P_t$ which gives us a measure on S at time t. Occasionally it is quite straightforward to work out P_t; but often it can be challenging. We will see that in principle we can calculate P_t by solving a set of differential equations. However, even in cases where we cannot write down P_t we may be able to work out some useful limit such as $\lim_{t\to\infty} P_t$.

A very special kind of state of a CTMC is an **absorbing state**. This is a state x such that $p_t(x, x) = 1$ for some $t > 0$, and hence for all $t > 0$.

Example 2.1 (**Lotka reactions)**. Consider the Lotka reactions:

$$A \to 2A, \quad B \to 0, \quad A + B \to 2B\,.$$

With intensities satisfying our basic assumptions (the intensity of a reaction is positive if and only if there are sufficient molecules of all the species in its reactant complex for it to proceed), the state $(0, 0)$ is an absorbing state: without any molecules of A or B present, none of the reactions can proceed, i.e., all the intensities must be zero. We thus remain at $(0, 0)$ for all future time. Systems with absorbing states highlight a fundamental difference between deterministic and stochastic models: in a stochastic model, random fluctuations may take us to an absorbing state with positive probability, and after this we are stuck there; this cannot occur in the corresponding deterministic model.

2.2 Waiting times are exponentially distributed

Let us consider a CRN in some given state. How long do we have to wait until some reaction "fires" and the state of the system changes? Understanding how to compute this (random) time,

is key both to understanding and to simulating stochastic models. It is remarkable that we can infer some very specific and useful information from the Markov assumptions alone.

From now on, $\{\mathbf{X}(t)\}_{t\in T}$ will be a right-continuous, time-homogeneous, Markov process taking values in $S \subseteq \mathbb{Z}^n_{\geq 0}$ and describing the state of a CRN. Given $\mathbf{X}(0) = x$ we can define the **waiting time** associated with the state x,

$$T_x := \inf\{t > 0 \,|\, \mathbf{X}(t) \neq x\}\,.$$

We may also refer to T_x as the **first jump time**. In the context of CRNs, the random variable T_x, associated with each state x, tells us how long we have to wait until the next reaction fires and takes us to a new state. Our assumption about right-continuity means that T_x is positive or ∞. From the Markov properties, we can see that T_x has the memoryless property:

$$\mathbb{P}(T_x > t+u) = \mathbb{P}(T_x > t+u \,|\, T_x > t)\mathbb{P}(T_x > t) = \mathbb{P}(T_x > u)\mathbb{P}(T_x > t)\,.$$

It is well known that memorylessness for a positive continuous random variable implies that it is exponentially distributed. The proof is left as an exercise.

Exercise 2.2 (**Memorylessness implies the exponential distribution)**. Show that if a positive, continuous random variable T is memoryless then it has the exponential distribution and find its parameter. Here "memoryless" means that $\mathbb{P}(T > a+b) = \mathbb{P}(T > a)\mathbb{P}(T > b)$ for any positive a and b. [Hint: set $G(u) = \mathbb{P}(T > u)$ so that G satisfies the functional equation $G(t+u) = G(t)G(u)$. First work out G for all positive rational inputs, and then complete the argument for all positive inputs by noting that G is clearly a decreasing function.]

We thus have that T_x is either infinite (if x is an absorbing state) or exponentially distributed with parameter, say, $\gamma_x > 0$. Consequently

$$\mathbb{P}(\mathbf{X}(t+\Delta t) = x \,|\, \mathbf{X}(t) = x) = 1 - \gamma_x \Delta t + o(\Delta t) \quad \text{as } \Delta t \searrow 0\,.$$

Memorylessness: probabilities scale linearly with time for small times. The following properties are equivalent for a memoryless waiting time T. Let $\lambda > 0$.

1. T has $\mathrm{Exp}(\lambda)$ distribution.
2. $\mathbb{P}(T \leq \Delta t) = \lambda\,\Delta t + o(\Delta t) \quad$ as $\Delta t \searrow 0\,.$

In one direction this is obvious as $1 - e^{-\lambda\,\Delta t} = \lambda\,\Delta t + o(\Delta t)$. In the other direction, let $G(t) = \mathbb{P}(T > t)$. Clearly $G(0) = 1$. Then, by memorylessness, and our assumption,

$$G(t+\Delta t) = G(t)G(\Delta t) = G(t)(1 - \lambda\Delta t + o(\Delta t))\,.$$

I.e.

$$\frac{G(t+\Delta t) - G(t)}{\Delta t} = -G(t)\lambda + G(t)\frac{o(\Delta t)}{\Delta t}\,.$$

Taking the limit $\Delta t \to 0$ gives the ODE

$$G'(t) = -\lambda G(t) \quad \text{with solution} \quad G(t) = e^{-\lambda t},$$

where we have used the initial condition $G(0) = 1$.

Thus very general assumptions lead us to the conclusion that given any initial state, the waiting time until the state of the system changes (i.e., the first reaction fires) is exponentially distributed. We should bear in mind the special case that it may be infinite, if the current state of the system is an absorbing state.

The waiting parameter as a reaction intensity. Recall that given a CRN in some state x, we defined the intensity $\nu_j(x)$ of a reaction, by saying that the probability the reaction occurs in time $[t, t+\Delta t)$ should be $\nu_j(\mathbf{x})\Delta t + o(\Delta t)$ as $\Delta t \searrow 0$. We thus see that if we have *only a single reaction* in a CRN, then the parameter γ_x of the exponentially-distributed waiting time until the next event is precisely the intensity of this reaction, namely $\nu_j(x)$.

Next, in a system with several reactions we would like to know *which* reaction will fire first, and how we can relate the parameter γ_x to the reaction intensities.

Waiting times for each reaction. Suppose we could "switch off" all but one reaction in the system, say the jth reaction, and assume that the current state x of the system is such that the jth reaction can proceed, i.e., its intensity at x is positive. Then, by the argument in Exercise 2.2, the waiting time until the reaction fires would be exponentially distributed with parameter $\nu_j(x)$, its intensity. We thus have:

1. the waiting time before *some* reaction occurs is exponentially distributed with parameter γ_x; and
2. if we had only the jth reaction reaction, the waiting time until it occurs would be exponentially distributed with parameter $\nu_j(x)$.

We need to check that these claims are consistent and find the relationship between the parameters $\nu_j(x)$ and γ_x. Given current state x, we can think of each reaction as having its own random alarm clock: e.g., reaction j has an alarm clock set to go off at some exponentially distributed random time with parameter $\nu_j(x)$ (the intensity of the reaction). All of the clocks associated with the different reactions are independent. Of course, when any alarm clock goes off, the state of the system changes, and the intensities change, and the parameters for all the alarm clocks change.

To find the relationship between γ_x and $\nu_j(x)$, we merely need to observe that the minimum of a set of independent exponentially distributed random variables with parameters $\lambda_1, \ldots, \lambda_m$ is exponentially distributed with parameter $\lambda_1 + \cdots + \lambda_m$. We can also then find the probability that the jth reaction is the first to fire. This leads us to a very important exercise:

Exercise 2.3 (**Important exercise: independent waiting times)**. Let $T_1, \ldots T_m$ be a set of independent, exponentially distributed, random variables with parameters $\lambda_1, \ldots, \lambda_m$. (We can think of these as waiting times for a set of m reactions given some current state of a CRN.) Show that:

1. $T := \min_i\{T_i\}$ is an exponentially distributed random variable with parameter $\sum_{j=1}^m \lambda_j$;
2. The probability that the minimum T is achieved by T_i is $\lambda_i/(\sum_{j=1}^m \lambda_j)$.
3. Let K be the random variable defined by $T = T_K$ (i.e., the Kth reaction is the first to fire). Note that, by the previous part, K is defined with probability 1. Show that T and K are independent. Interpret this result intuitively.

Remark 2.4 (**Countably infinite reactions)**. The result of Exercise 2.3 can be extended to the case of infinitely many random times provided the parameters are summable, namely $\sum_{k=1}^\infty \lambda_i < \infty$.

The observations in Exercise 2.3 mean that the parameter γ_x for the distribution of the waiting time before *some* reaction fires is just equal to $\sum_{j=1}^m \nu_j(x)$, the sum of the reaction intensities. The exercise also tells us that the probability that the first reaction to fire is the jth reaction is the ratio $\nu_j(x)/\sum_{j=1}^m \nu_j(x)$. These observations are key to understanding the **Gillespie algorithm**, which we will present shortly.

A naive simulation algorithm. Before we come to this Gillespie algorithm, it is useful to consider the following "naive" approach to simulating a system of chemical reactions, based on our observations so far. We have seen that given a CRN with m reactions, currently in state x, the probability that some reaction occurs in a time Δt is $\gamma_x \Delta t + o(\Delta t)$ as $\Delta t \searrow 0$, where $\gamma_x = \sum_{j=1}^m \nu_j(\mathbf{x})$. If we choose the time-step Δt small enough, so that $\gamma_x \Delta t \ll 1$, then we can, at each time-step, just ask the "yes-no" question: "*Did some reaction occur during* Δt*?*" by choosing a random number uniformly distributed on $[0,1]$ and checking if its value is $\leq \gamma_x \Delta t$. If yes, then some reaction fired, and otherwise it didn't. We can then decide *which* reaction fired according to the probabilities $\nu_j(x)/\gamma_x$, e.g., in order to choose which reaction fired, we divide $[0,1]$ into m subintervals of lengths $\nu_j(x)/\gamma_x$ $(j = 1,\ldots,m)$ and see where a random number uniformly distributed on $[0,1]$ lands. Of course, this approach would be approximate, and its accuracy clearly depends on the step-size Δt, as we have ignored the $o(\Delta t)$ term. We will shortly see a better approach; but it is nevertheless useful to bear this naive algorithm in mind: we will see similar ideas at several points.

2.3 The Gillespie algorithm

We are ready to discuss the most important algorithm for the stochastic simulation of CRNs: the so-called Gillespie algorithm.

Suppose that our reaction network is currently in state x and the corresponding reaction intensities are $\nu_1(x),\ldots,\nu_m(x)$. Consequently the waiting times, say $T_1,\ldots,T_m$, for the m reactions are a set of independent exponentially distributed random variables with parameters $\nu_1(x),\ldots,\nu_m(x)$, whose minimum, say T, is exponentially distributed with parameter $\nu_1(x)+\cdots+\nu_m(x)$ (Exercise 2.3). Moreover (Exercise 2.3 again), the minimum is almost surely achieved by a unique T_K, and $\mathbb{P}(K = k) = \nu_k(x)/(\sum_{j=1}^m \nu_j(x))$. And, crucially (Exercise 2.3 again!), the random variable

T which tells us the firing time, and the random variable K which tells us which reaction fires first, are *independent*.

With these observations, we can then simulate the system as follows:

1. Sample from an exponential distribution with parameter $\nu_1(x) + \cdots + \nu_m(x)$ to get the time T of the next reaction;
2. Choose which reaction fires according to the probabilities $\nu_k(x)/(\sum_{j=1}^m \nu_j(x))$. I.e., if the random variable K tells us the next reaction to occur, then $\mathbb{P}(K = k) = \nu_k(x)/(\sum_{j=1}^m \nu_j(x))$;
3. At time T, update the species numbers $x \mapsto x+\zeta_K$, recalculate intensities $\nu(x) \mapsto \nu(x+\zeta_K)$, and return to step 1.

This is the Gillespie algorithm. Note that the Gillespie algorithm is *exact*: unlike in the naive algorithm presented above, we make no approximations, and our ability to generate exact sample paths is limited only by the precision of the machine we are working on (including our ability to sample from distributions). Of course, our model itself may not be perfect, but that is another matter.

2.4 The jump chain: irreducibility, recurrence, and transience

Let $x, y \in S$ be states of a CRN with $y \neq x$. Consider the probability that given *current* state x, the *next* state of the system will be y. If k reactions $j_1, \ldots, j_k$ take us from state x to state y, i.e., $y = x + \zeta_{j_i}$ for $i = 1, \ldots, k$ (recall: ζ_j is the reaction vector of reaction j), then the "jump" probability from state x to y is $\hat{p}(x, y) := (\nu_{j_1}(x) + \cdots + \nu_{j_k}(x))/\sum_{i=1}^m \nu_i(x)$. If, for some fixed state x, $\hat{p}(x, y) = 0$ for all $y \neq x$, then x is an absorbing state and we set $\hat{p}(x, x) = 1$. We can put these jump probabilities $\hat{p}(x, y)$ into a matrix, say $\widehat{P}$, to define a discrete-time Markov chain, termed the **jump chain** of the system, with transition matrix $\widehat{P} = (\hat{p}(x, y))$.

Thus, associated with a continuous-time CRN model is a discrete-time Markov chain, the jump chain, which tells us which states we can move directly to from the current state; and with what probabilities. The jump chain *does not* tell us anything about how fast or slowly these transitions might occur; nevertheless, many important properties of a continuous-time model can be inferred from properties of the jump chain.

Communicating classes. As in the case of discrete-time Markov chains, we can say that two states x and y **communicate** if it is possible to go from x to y via a sequence of jumps (reactions); and also back again from y to x via a sequence of jumps. As usual, we can divide the states into communicating classes. Clearly we can identify the **communicating classes** from the jump chain alone. Moreover, we can also define communicating classes to be **closed** or not in the usual way, again from examination of the jump chain alone.

Irreducibility. As in the discrete-time case, a CTMC is **irreducible** if for every pair of states x, y, there exists $t > 0$ s.t. $p_t(x, y) > 0$, in other words, all states belong to a single communicating class. However, now we have an important difference between continuous-time and discrete-time chains: if $p_t(x, y) > 0$ for some $t > 0$, then $p_t(x, y) > 0$ for *all* $t > 0$. We can see this as

follows. If $p_t(x, y) > 0$ for some $t > 0$ there must be a finite sequence of states connecting x and y, say $x = x_0, x_1, \ldots, x_n = y$, and there must exist (t_k) such that $p_{t_k}(x_{k-1}, x_k) > 0$ for $k = 1, \ldots n$. But, waiting times are exponentially distributed; and so, for each k and any $r > 0$, $p_{t_k/r}(x_{k-1}, x_k) > 0$. The claim follows.

Thus, if a continuous-time chain is irreducible then all entries of each transition matrix P_t $(t > 0)$ are strictly positive. Consequently, there is no analogue of **periodicity** in continuous-time Markov chains.

A rather trivial, but useful, observation is that a system with an absorbing state cannot be irreducible (assuming that it has more than one state).

Example 2.5 (**Lotka reactions: the chain is not irreducible)**. Consider again a stochastic model of the Lotka reactions which we saw in Example 2.1:

$$\mathrm{A} \to 2\mathrm{A}, \quad \mathrm{B} \to 0, \quad \mathrm{A} + \mathrm{B} \to 2\mathrm{B}\,.$$

We may take the state space to be $\mathbb{Z}^2_{\geq 0}$. Recall that $(0, 0)$ is an absorbing state of the system and thus the system is not irreducible. It is not hard to see that the only closed communicating class of the system is the absorbing state $\{(0, 0)\}$.

In fact we can divide $\mathbb{Z}^2_{\geq 0}$ into four regions: the absorbing state $(0, 0)$; states of the form $(n, 0)$ with $n > 0$ (each such state is an open class); states of the form $(0, n)$ with $n > 0$ (each such state is an open class); and the remaining states of the form (n, m) with $n > 0$, $m > 0$ which together form an open class.

Based on this we can guess that given any initial condition, with finite probability we will end up on one of the axes (i.e., in a state of the form $(0, n)$ or $(n, 0)$). We will explore this example further in an exercise.

Chains where all communicating classes are closed. Another observation is that if the state space of a CRN model consists of a set of closed communicating classes, then if we know the initial state we may restrict attention to exactly one of these; and thus assume the CRN gives rise to an irreducible chain (see, e.g., Example 1.13). Thus a CRN where all the communicating classes are closed is as good as an irreducible one in practice.

Recurrence and transience. Just as in the discrete-time case a state x of a CTMC is **recurrent** if, given $X(0) = x$, with probability one, we can find a sequence of times $t_k \to \infty$ such that $X(t_k) = x$ for each t_k. Clearly an absorbing state is recurrent. A CTMC is recurrent if all its states are recurrent. A state which is not recurrent is **transient**. For example, in a CRN with the single reaction $\mathrm{A} \to 2\mathrm{A}$, the state with one molecule of A is clearly transient.

Just as irreducibility can be inferred by looking at the jump chain, the same holds for recurrence and transience. From the equivalent discrete-time result, recurrence and transience are *class properties*: either every state in a communicating class is recurrent; or every state in the class is transient. The claim that we can simply look at the jump chain to decide on recurrence and transience is almost obvious; but the possibility of **explosion**, discussed in Section 2.6, slightly complicates the proof, and so we do not present the proof.

2.5 Mass action kinetics

So far, we have not considered how we might choose the reaction intensities in a stochastic model of a CRN. Our most basic assumption is that the intensity of a reaction is positive provided there are sufficient molecules of all the species in the reactant complex of a reaction for the reaction to proceed; and zero otherwise. We might expect a little more: that the intensity of a given reaction should increase if there are more ways for the molecules in its reactant complex to meet in the right combinations.

One common choice of reaction intensities satisfying this natural requirement is given by **stochastic mass action kinetics**. It is easiest to introduce by example.

Example 2.6 (**Mass action kinetics for a bimolecular reactant complex)**. Consider a reaction with reactant complex

$$\mathrm{A} + \mathrm{B}\,.$$

This reaction occurs with some nonzero probability when a molecule of A "meets" a molecule of B. Let us write n_A for the number of molecules of A, and n_B for the number of molecules of B at some moment in time. The mass action assumption is that the probability of a molecule of A meeting a molecule of B is proportional to the number of (A, B) pairs we can choose from n_A molecules of A and n_B molecules of B, namely $\mathrm{n}_\mathrm{A}\mathrm{n}_\mathrm{B}$. We would thus set the intensity to be

$$\alpha \mathrm{n}_\mathrm{A}\mathrm{n}_\mathrm{B}\,,$$

where the constant α is the **stochastic mass action rate constant** of the reaction. Note that if the reaction is unfolding in a chamber of volume V, we would also expect the rate to decrease as V increases (provided molecule numbers are held constant), reducing the likelihood of a molecule of A meeting one of B; thus we expect α to have an inverse relationship with the compartment volume V.

We can extract the key features of the previous example to define stochastic mass action kinetics.

Definition 2.7 (**Stochastic mass action kinetics).** *The intensity of a reaction is proportional to the number of available combinations of the molecules in its reactant complex. The constant of proportionality is the* **stochastic mass action rate constant** *of the reaction. Further, this rate constant is proportional to* $V^{1-\mathrm{k}}$*, where* V *is the compartment volume, and* k *is the order of the reaction.*

For example:

1. For a zeroth order reaction, i.e., where the reactant complex is 0, the rate is just a constant, say α, which is proportional to V. We could think of the production of some species from a "hidden species" whose number scales positively with V.
2. For a first order reaction with, say, reactant complex A, we would expect the rate to be $\alpha \mathrm{n}_\mathrm{A}$, where α is independent of V.

3. For a second order reaction with reaction complex $\mathrm{A}+\mathrm{B}$ we expect the rate to be of the form $\alpha n_A n_B$, where α depends inversely on V. If the reactant complex were $2\mathrm{A}$, we would expect the rate to have the form $\alpha n_A(n_A-1)/2$ where again we expect α to depend inversely on V.

The principle for higher order reactions is similar: we count the number of ways we can create the reactant complex out of available molecules, and multiply by a rate constant which we expect to be proportional to V^{1-k}, where k is the order of the reaction, and V is the compartment volume.

2.5.1 Relating stochastic and deterministic mass action kinetics

Often we are given are determinstic mass action rate constants, and we want to infer the mass action intensities from these. How do we do this? Let us start again with an example.

Example 2.8 (**Relating stochastic and deterministic mass action rate constants)**. Consider a reaction with reactant complex $2\mathrm{A}$ proceeding in a chamber with volume V. Let n_A be the number of molecules of A. In the derministic case, we can interpret the concentration of A as *number of molecules per unit volume*, i.e. n_A/V. Let the deterministic mass action rate constant be k so that the deterministic rate is kn_A^2/V^2. This is a rate of change of concentration; so the corresponding rate of change of number of molecules (regarded as a real variable!) would be kn_A^2/V. On the other hand, let the stochastic mass action rate constant be α, so the mass action intensity is $\alpha n_A(n_A-1)/2$. For large n_A, we may assume that $n_A(n_A-1)\simeq n_A^2$; then comparing kn_A^2/V (deterministic rate) to $\alpha n_A(n_A-1)/(2V)$ (stochastic intensity), we get $\alpha\simeq 2k/V$. We thus get the stochastic mass action intensity in terms of the deterministic rate constant k to be $kn_A(n_A-1)/V$.

With the above example in mind, it becomes easy to write down intensities in terms of deterministic mass action rate constants. The general rule is that a reaction with reactant complex $\alpha_1 A_1+\cdots+\alpha_n A_n$, with $|\alpha|:=\sum_i \alpha_i$ and deterministic mass action rate constant k, has stochastic mass action intensity

$$\frac{k}{V^{|\alpha|-1}}\frac{(A_1)!}{(A_1-\alpha_1)!}\frac{(A_2)!}{(A_2-\alpha_2)!}\cdots\frac{(A_n)!}{(A_n-\alpha_n)!},$$

where A_k is the number of molecules of species A_k.

It is important to remember that for reactions of order 2 or more, the correspondence only holds approximately, becoming more accurate for large numbers of molecules.

Exercise 2.9. Write down the deterministic and stochastic mass action rates for a reaction with reactant complex

$$\mathrm{A}+2\mathrm{B}$$

and hence relate the deterministic and stochastic rate constants.

2.6 Explosion (blow-up)

Before we can go further, we need to consider a feature of CTMCs which cannot occur in the discrete-time case: "blow-up" or **explosion**. Just as for deterministic ODE models, we cannot assume that given a stochastic model of a CRN, and an initial condition, we can determine the state of the system for all time: it is possible that with positive probability some species numbers approach ∞ at some finite time, say T_∞. Note that T_∞, termed the **explosion time**, is a *random* time. If $\mathbb{P}(T_\infty < \infty | \mathbf{X}(0) = x) > 0$, then we say that the initial state x is **explosive**. A model with some explosive states will be called explosive.

The situation for explosive models is even more dramatic than in the case of ODEs with blow-up. In the ODE case, for any given initial condition, there is a finite time interval over which the solution exists. On the other hand, in a stochastic model with an explosive initial state, there is a nonzero probability of explosion by *any* time $t > 0$, however small. This follows from the remarks on irreducibility in Section 2.4 above.

When we have an explosive system, the evolution after the explosion time, say T, is not uniquely determined by the evolution prior to T_∞; we can set the system to behave in many different ways consistent with the memoryless assumption after explosion – for example we could choose it to return to zero. However, most natural is to stop our process at the explosion time T_∞ and formally set the state of the system to be "∞" after that time. In other words, we add a new absorbing state to the system which we call ∞. In this way, the transition matrices $P(t)$ can be defined for all time. This is the assumption we will make from now on.

Explosion is a general feature of continuous-time Markov processes, and is unavoidable in CRN models. Luckily, there are some general theorems to help us determine whether it happens or not for a given process. Here we will take an informal approach and check directly whether or not blow-up happens in particular CRN models.

Example 2.10 (**Blow-up**). Consider the following two CRNs. Both are examples of *birth processes*. In both cases n is the number of molecules of the species A, and $k > 0$ is a constant.

1. Consider the CRN

   $$A \to 2A$$

   with one or more initial molecules of A and intensity kn (i.e., mass action kinetics with rate constant k). Here kn can be regarded as a "birth rate". We can think of this as a model of a bacterial colony, where bacteria reproduce by splitting into two (and none ever die). Although $n \to \infty$ with probability 1 as time increases, this CRN model is *not* explosive. This can be proved using some theory (which is straightforward, but we will not present) from the fact that

   $$\sum_{n=1}^{\infty} \frac{1}{nk} = \infty .$$

2. Consider the CRN

   $$2A \to 3A$$

with two or more initial molecules of A and (mass action) intensity $kn(n-1)$ for some constant $k > 0$. We can think of this as a model of a colony where, with positive probability, whenever two individuals meet they produce a third (and none ever die). This model *is* explosive, which can be proven from the fact that

$$\sum_{n=2}^{\infty} \frac{1}{kn(n-1)} < \infty .$$

Thus given any $T > 0$, there is a positive probability that the state of the system will be ∞ at time T.

Note that if a model is non-explosive, then for any $t \geq 0$ the entries in each row of $P(t)$ must sum to 1 by the law of total probability. On the other hand, if a state is explosive, then the entries in the corresponding row of $P(t)$ sum to less than 1 for any $t > 0$. However, if we add in a new state ∞ as above, the entries in each row of the expanded matrix, say $\tilde{P}(t)$ (with one extra row and column corresponding to the state ∞), still sum to one.

Clearly, if a state is explosive, then it cannot be recurrent: with nonzero probability, we have explosion by any time $t > 0$, and so the probability of returning eventually to our initial state must be less than 1.

2.7 The reaction counting process

Up to now, we have been interested in the the stochastic process $\{\mathbf{X}(t)\}_{t\geq 0}$ where the ith component of $\mathbf{X}(t)$ is the number of molecules of the ith chemical species at time t. But it is useful also to consider another stochastic process which counts how many times each reaction has fired by time t.

Let $\mathbf{N}(t)$ be a stochastic process whose ith component, $N_i(t)$, counts how many times the ith reaction has fired up to time t. A process such as $N_i(t)$ which takes only nonnegative integer values and can only increase in steps of size 1 is called a **counting process**; and so $\mathbf{N}(t)$ is a vector of counting processes taking values in $\mathbb{Z}^m_{\geq 0}$ (recall: we are assuming there are m reactions in the CRN).

If we consider the vector $\mathbf{N}(t)$ of counting processes, then we can write down for the evolution of $\mathbf{X}(t)$ an equation which is entirely analogous to the deterministic evolution equation (2), namely:

$$\mathbf{X}(t) = \mathbf{X}(0) + \Gamma \mathbf{N}(t) , \tag{6}$$

where Γ is the stoichiometric matrix of the system. In order to see the contribution from each reaction, it is helpful to write (6) out in full:

$$\mathbf{X}(t) = \mathbf{X}(0) + N_1(t)\zeta_1 + N_2(t)\zeta_2 + \cdots + N_m(t)\zeta_m ,$$

where we recall that ζ_j is the reaction vector of the jth reaction, namely, the jth column of Γ.

An easy observation is that the discussion of conservation laws in the deterministic case remains valid in the stochastic case: if $\mathbf{p}$ is any vector in $\ker\Gamma^t$, then $\mathbf{p}^t\mathbf{X}(t)$ is in fact constant at the value $\mathbf{p}^t\mathbf{X}(0)$. Moreover, if there exists a vector $\mathbf{p} \in \ker\Gamma^t$ whose components are all *strictly positive*, then in fact the state space of the CRN is *finite*, being a bounded subset of $\mathbb{Z}^m_{\geq 0}$ (see Exercise 1.10).

What can we say about the reaction-counting process $\mathbf{N}(t)$? If the reaction intensities were not dependent on the state of the system, then the $N_i(t)$ would be independent Poisson processes with some fixed rates. (This is the case for zeroth order reactions with mass action kinetics.) However, whenever *any* reaction fires the molecule numbers change and, potentially, some or all of the rates associated with these Poisson processes need to be updated. Thus, we can think of $\mathbf{N}(t)$ as a vector of Poisson-like processes, but with variable rates $\nu_j(x)$, which depend on the state of the system.

Note that if $Y(t)$ is a unit Poisson process, then a Poisson process with rate λ can be written $Y(\lambda t)$. So, how is $N_j(t)$ distributed, namely, how many times does reaction j occur during the time interval $[0, t]$? Assuming the CRN is not explosive, with probability 1 we can partition the interval $[0, t]$ into $0 = t_0 \leq t_1 \leq \cdots \leq t_r = t$ such that $\mathbf{X}(t)$ is constant on each $[t_k, t_{k+1})$. Note that both the times t_k and the total number of jumps r are random variables. Then associated with each trajectory is an average intensity of the jth reaction, namely $\frac{1}{t}\sum_{k=0}^{r-1}\nu_j(\mathbf{X}(t_k))(t_{k+1} - t_k) = \frac{1}{t}\int_0^t \nu_j(\mathbf{X}(s))\,ds)$, and $N_j(t)$ can be identified with the Poisson process with (random) rate $\frac{1}{t}\int_0^t \nu_j(\mathbf{X}(s))\,ds$, namely,

$$N_j(t) = Y(\int_0^t \nu_j(\mathbf{X}(s))\,ds)\,.$$

Thus $N_j(t)$ is indeed Poisson-distributed; but with parameter $\frac{1}{t}\int_0^t \nu_j(\mathbf{X}(s))\,ds$ which depends on the history of the process. This perspective is quite useful: $\int_0^t \nu_j(\mathbf{X}(s))ds$ can be interpreted as a *rescaled time*; and now we can think of the counting processes $N_j(t)$ as a set of ordinary, independent, Poisson processes, evaluated at the (random) times $\tau_j(t) := \int_0^t \nu_j(\mathbf{X}(s))\,ds$.

It may be helpful to visualise this change of perspective in terms of our alarm clocks associated with each reaction. As before we have a set of m independent Poisson processes Y_j, each with its own clock. Before we had random alarms with different parameters for the alarms, but now the alarms all have the same parameter, *but the clocks run at different speeds*. So, at a particular "universal" time, we have reached different times in each counting process. Moreover, the speeds of the clocks are updated every time an alarm goes off. (Indeed some of the clocks can stop if, for example, one of the species in the corresponding reactant complex runs out.)

Writing out the solution (6) in full we now have

$$\mathbf{X}(t) = \mathbf{X}(0) + \zeta_1 Y_1(\int_0^t \nu_1(\mathbf{X}(s))ds) + \cdots + \zeta_m Y_m(\int_0^t \nu_m(\mathbf{X}(s))ds)\,. \tag{7}$$

We have thus found a way of writing the stochastic evolution of a CRN in terms of a set of independent Poisson-like processes; and in a form which looks, formally, very similar to the deterministic solution (3).

Although our presentation has been heuristic, this approach can be made precise. We will see later that the representation in terms of Poisson-like processes is the natural starting point for writing down the **Chemical Langevin equation** and the **Chemical Fokker-Planck equation**.

2.7.1 τ-leaping

While the Gillespie algorithm is exact, (7) also suggests another, more approximate, approach to simulating a stochastic CRN, termed **tau-leaping**. If we have many reactions in a network, and large numbers of molecules, then we may not wish to track every reaction occurrence. Instead, we might fix a small time-interval, say Δt, and hope to estimate how many times each reaction occurs during this time.

Let us suppose the system is currently in state x. At this moment, the jth reaction can be regarded as a Poisson process with rate $\nu_j(x)$. Of course, this is only true until some reaction occurs; but let us pretend it is true over the whole time interval Δt. Then, we expect the jth reaction to occur n_j times over this interval, where n_j is Poisson distributed with parameter $\nu_j(x)\,\Delta t$. And similarly for all the other reactions: hence, at the end of the time-interval, we would expect the system to be in the new state

$$x + \sum_{j=1}^{m} n_j \zeta_j \,,$$

obtained by taking the current state and allowing each of the m reactions to occur n_j times. For each j, we choose n_j from (independent) Poisson distributions with parameters $\nu_j(x)\,\Delta t$.

We can now update the state of the system, and continue again. Of course, we have to be careful not to allow reactions to fire in such a way that some species numbers become negative. Indeed, if some species numbers are small, then the approximation is almost certainly a bad one! Moreover, the outcome can only be approximate, because the rates were not truly constant over the interval Δt. In particular, we expect this approximation to be better if

- Δt is small, and
- the species numbers are large enough that the fractional changes in intensities each time a reaction fires are small.

Despite not being exact, τ-leaping can be useful for simulating large systems where the usual Gillespie algorithm can become prohibitively slow. It is also conceptually helpful when we write down the chemical Langevin and Fokker-Planck equations after we have discussed stochastic differential equations (SDEs).

2.8 The generator matrix for a CRN

We now return to the question of how to compute transition probabilities, and thus the evolution of a stochastic CRN model. Given such a model, we will first define a matrix $Q = (q(x, y))$, termed the the **generator matrix** of the CRN; and we will then try to understand how Q arises and how it can be used. We will also refer to this matrix, informally, as the **Q-matrix** of the CRN.

As usual, let us suppose that the CRN has m reactions and that the jth reaction has reaction vector ζ_j. Let us suppose that at some particular time, the system is in state x. We have seen

that waiting times for each reaction are independent and exponentially distributed; and so, the probability that the jth reaction occurs in a time interval Δt is $\nu_j(x)\Delta t + o(\Delta t)$. Moreover, the waiting times for reactions to fire are *independent* random variables. Define $\nu_{tot}(x,y) := \sum_{j \text{ s.t. } y=x+\zeta_j} \nu_j(x)$, i.e., $\nu_{tot}(x,y)$ is the total intensity of all reactions which take us from state x to state y. Then, as $\Delta t \searrow 0$, we have, for the transition probabilities

$$P_{\Delta t}(x,y) = \nu_{tot}(x,y)\,\Delta t + o(\Delta t) \ (j = 1,\dots m) \quad \text{and} \quad P_{\Delta t}(x,x) = 1 - \sum_{j=1}^{m} \nu_j(x)\Delta t + o(\Delta t)\,.$$

with all other transition probabilities being $o(\Delta t)$. Since $P(0) = \mathbb{I}$, the identity matrix, we quite naturally define Q as the (right) derivative of P at 0, namely,

$$Q := P'(0) := \lim_{\Delta t \searrow 0} \frac{P(\Delta t) - \mathbb{I}}{\Delta t}\,. \tag{8}$$

Thus we have, for $x \neq y$, $q(x,y) = \nu_{tot}(x,y)$; and $q(x,x) = -\sum_{j=1}^{m} \nu_j(x) = -\sum_{j=1}^{m} q(x, x+\zeta_j)$. Thus each row of Q has at most m nonzero off-diagonal entries, corresponding to each of the reactions firing. And the diagonal entry is set by ensuring that sum of all entries in a row is zero. Of course, it is possible that some of the $\nu_j(x)$ will be zero, or different reactions have the same reaction vector, in which case the corresponding row(s) of Q will have fewer than m nonzero off-diagonal entries.

Remark 2.11 (**Generator matrices for CTMCs)**. We remark that the right differentiability of $P(t)$ at 0, and hence the existence of the limit in (8), follows from very general theory on continuous time Markov chains with right-continuous paths, and is not specific to CRN models. Thus all such chains are associated with a generator matrix Q.

The transition matrix for the jump chain from the Q-matrix. For a CRN, calculating the Q-matrix is simply a matter of knowing the intensities associated with each state. Thus the relative sizes of off-diagonal entries in the Q-matrix tells us about the relative probabilities of different reactions firing. For example, if x is not an absorbing state, then given state x the probability that the first reaction to fire is reaction j is just

$$\frac{\nu_j(x)}{\sum_{j=0}^{m} \nu_j(x)} = \frac{q(x, x+\zeta_j)}{-q(x,x)}\,.$$

Thus, to get the transition matrix $\widehat{P}$ of the jump chain of a CRN from the Q-matrix is simple.

- Given an absorbing state x, we simply set the corresponding diagonal entry $\widehat{P}(x,x) = 1$, and $\widehat{P}(x,y) = 0$ for all $y \neq x$.
- Given a state x which is not absorbing, we set $\widehat{P}(x,x) = 0$ and $\widehat{P}(x,y) = \frac{q(x,y)}{-q(x,x)}$ for $y \neq x$.

Example 2.12 (Q **matrix and jump chain of a CRN with finite state space)**. [corrected.] Consider the CRN with four reactions

$$A \xrightarrow{1} B + C, \quad B + C \xrightarrow{1} A, \quad B \xrightarrow{1/2} C, \quad C \xrightarrow{1/2} B\,.$$

Let us assume that all the reactions occur in a cell of volume $V = 1$ and have deterministic mass action rate constants as shown. There is a strictly positive vector in $\ker \Gamma^t$, where Γ is the stoichiometric matrix (check this!), and so, given any initial state, this CRN has finite state space.

Suppose we initially have one molecule of A. The system then has four states it can reach from this initial state:

$$\{\{A\},\ \{B, C\},\ \{B, B\},\ \{C, C\}\},$$

(or, as vectors, $(1,0,0), (0,1,1), (0,2,0)$ and $(0,0,2)$.) The Q-matrix and transition matrix of the corresponding jump-chain are:

$$Q = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 1 & -2 & 1/2 & 1/2 \\ 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix} \quad \text{and} \quad \widehat{P} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1/2 & 0 & 1/4 & 1/4 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

Note that the chain is irreducible and recurrent because the corresponding jump-chain is.

We will return to this example below.

2.9 The forward Kolmogorov and backward equations

Recall that for discrete-time Markov chains we can calculate n-step transition probabilities from one-step transition probabilities. We can thus think of the one-step transition matrix as *generating* all the possible transitions. We expect something similar to occur in continuous time: we hope to be able to infer $P(t)$ for any given t from the generator Q.

We will show how to derive two systems of ODEs satisfied by the transition probabilities in the case of a *finite* state space. But in fact the systems are valid in the infinite case too.

Recall that the semigroup property implies that $P(t + \Delta t) = P(t)P(\Delta t)$. In the case that the state space S is finite, we can subtract $P(t)$ from both sides, divide through by Δt, and take the limit $\Delta t \searrow 0$ to get

$$P'(t) = P(t)Q, \quad \text{or in components} \quad \frac{d}{dt}p_t(x, y) = \sum_{z \in S} p_t(x, z)q(z, y). \tag{9}$$

This is called the the **forward Kolmogorov equation**. Similarly, if we start by writing $P(t + \Delta t) = P(\Delta t)P(t)$ and carry out a similar procedure, we obtain the equation

$$P'(t) = QP(t), \quad \text{or in components} \quad \frac{d}{dt}p_t(x, y) = \sum_{z \in S} q(x, z)p_t(z, y), \tag{10}$$

which is the **backward Kolmogorov equation**. Thus the transition matrices $P(t)$ satisfy the two differential equations (9) and (10). Moreover, still assuming a finite state space, it is easy to find the unique solution to these two equations with initial condition $P(0) = \mathbb{I}$, namely

$$P(t) = \exp(tQ) = \mathbb{I} + tQ + t^2Q^2/2 + \cdots.$$

Thus in the case of a finite state space we can, in theory, always solve to find the transition matrices. From this, we can infer the long-term evolution of any distribution on the state space. Let us see this by example.

Example 2.13 (**The transition matrix from the Kolmogorov equations**). [corrected.] Consider again the CRN in Example 2.12, namely,

$$A \xrightarrow{1} B + C, \quad B + C \xrightarrow{1} A, \quad B \xrightarrow{1/2} C, \quad C \xrightarrow{1/2} B,$$

giving

$$Q = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 1 & -2 & 1/2 & 1/2 \\ 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix}.$$

We can confirm that Q has eigenvalues $-3, -1, -1, 0$ and corresponding matrix of eigenvectors

$$E = \begin{pmatrix} 1 & 1 & 0 & 1 \\ -2 & 0 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & -2 & -1 & 1 \end{pmatrix}.$$

With $D = \mathrm{diag}\{-3, -1, -1, 0\}$, we calculate

$$e^{tQ} = Ee^{tD}E^{-1} = \frac{1}{12}\begin{pmatrix} 6e^{-t} + 2e^{-3t} + 4 & -4e^{-3t} + 4 & -3e^{-t} + e^{-3t} + 2 & -3e^{-t} + e^{-3t} + 2 \\ -4e^{-3t} + 4 & 8e^{-3t} + 4 & -2e^{-3t} + 2 & -2e^{-3t} + 2 \\ -6e^{-t} + 2e^{-3t} + 4 & -4e^{-3t} + 4 & 9e^{-t} + e^{-3t} + 2 & -3e^{-t} + e^{-3t} + 2 \\ -6e^{-t} + 2e^{-3t} + 4 & -4e^{-3t} + 4 & -3e^{-t} + e^{-3t} + 2 & 9e^{-t} + e^{-3t} + 2 \end{pmatrix}$$

We see that as $t \to \infty$, any initial probability distribution on the states $\{\{A\}, \{B, C\}, \{B, B\}, \{C, C\}\}$ approaches $(1/3, 1/3, 1/6, 1/6)$. In other words, in the long run, the system spends $1/3$ of the time in the state with one molecule of A, $1/3$ of the time in the state with one molecule of B and one of C, $1/6$ of the time in the state with two molecules of B, and $1/6$ of the time in the state with two molecules of C. If we increase the rate constant of the reaction $A \longrightarrow B + C$ to 2 (while keeping all other rate constants as before) you should find that any initial probability distribution on the states now approaches $(1/5, 2/5, 1/5, 1/5)$ as $t \to \infty$. Check this as an exercise.

The Kolmogorov equations for a countably infinite state space. In the case of a countably *infinite* state space, the derivations above of the forward and backward equations do not work. We assumed that a sum of $o(\Delta t)$ terms is again an $o(\Delta t)$ term; this assumption is not automatically justified in the case of an infinite sum. We have to think more carefully about interchanging limiting operations. However, using more involved arguments, we can indeed show that under the assumptions we have made, $P(t)$ satisfies the forward and backward equations, which are now infinite systems of coupled first order ODEs.

The meaning of "forward" and "backward". When we write $P(t + \Delta t) = P(t)P(\Delta t)$, we consider the evolution first in a step of size t, and then in a step of size Δt. I.e., we perturb the

system a little bit at the end of our time-interval to see how probabilities change. When we write $P(t+\Delta t) = P(\Delta t)P(t)$, the time-step Δt occurs at the start of the time-interval of interest. We can think of this as perturbing a little bit at the start. We will see more clearly later that forward equations correspond to initial value problems: we fix the initial state and see how probabilities evolve forward in time. Meanwhile backward equations correspond to "final value problems": we have some target final state and want to find a probability distribution on initial states which lead us to this state after a time t.

The Kolmogorov equations for an uncountable state space. Later on, when we study stochastic differerntial equations (SDEs), we will meet the Kolmogorov equations in the case of an uncountable state space, where they become partial differential equations describing the evolution of probability density functions, rather than systems of ODEs describing the evolution of probability mass functions. The forward Kolmogorov equation is, in this case, often called the **Fokker-Planck equation**.

2.10 Stationary measures and distributions

A function which assigns a nonnegative number to each element of the state space of the CRN is a **measure** on the state space. Transition matrices tell us how measures evolve. Given a continuous-time Markov chain with state space S and transition matrices $P(t)$, and an initial measure ϕ_0 on S, after time t we have the measure $\phi_t = \phi_0 P(t)$ on the state space. This is just a statement of the law of total probability.

If we are able to normalise a measure on S (i.e., if it has finite total mass), then we can define a **distribution** on S, namely a measure with total mass 1. If $P(t)$ are the transition matrices of a non-explosive Markov chain, and ϕ_0 is some initial distribution, then $\phi_t = \phi_0 P(t)$ again has total mass 1, and can be regarded as the probability distribution on the states of the chain given the initial distribution ϕ_0.

Definition 2.14 (**Stationary/invariant measure for a CRN).** *A measure ϕ for a CRN with transition matrices $P(t)$ is* **stationary** *(or* **invariant***) if*

$$\phi = \phi P(t) \tag{11}$$

for all $t \geq 0$.

Note that there is no reason in general for a CRN to have a stationary measure.

Example 2.15. Consider the CRN

$$0 \to A$$

with mass action kinetics and rate constant k. In this case, it is clear that the transition matrix $P(t)$ is upper triangular with nonzero diagonal elements, and so (11) can have no nonzero solution.

We have the following important result for stationary measures of irreducible and recurrent chains:

Theorem 2.16. *Suppose that a CRN model is irreducible and recurrent: then it has a unique invariant measure up to scalar multiples. Moreover, this measure is the unique (nonnegative) element of* $\ker Q^t$ *up to scalar multiples.*

We won't prove this theorem, but note that it follows by examining the corresponding jump chain, and using results on invariant measures for irreducible, recurrent, discrete-time Markov chains.

In the case of a finite state space, the claim in Theorem 2.16 that invariant measures correspond to elements of $\ker Q^t$ follows by differentiating both sides of (11) and using the backward equation to get

$$0 = \phi P'(t) = \phi Q P(t)$$

for all $t > 0$. Since $\lim_{t \searrow 0} P(t) = P(0) = \mathbb{I}$, we must have

$$0 = \phi Q\,.$$

Similarly, if ϕ satisfies $0 = \phi Q$, then multiplying on the right by $P(t)$, and again using the backward equation gives

$$0 = \phi Q P(t) = \phi P'(t) = \frac{d}{dt}(\phi P(t))\,,$$

implying that $\phi P(t)$ is a constant. Since $P(0) = \mathbb{I}$ we recover $\phi = \phi P(t)$. We need a more elaborate proof in the case of an infinite state space as we cannot automatically justify exchanging infinite summation and differentiation, but the result still holds.

Thus stationary measures for irreducible and recurrent chains correspond precisely to elements in $\ker Q^t$.

Remark 2.17 (**detailed balance)**. Note that the stationarity condition $0 = \phi Q$ can be written

$$0 = \phi Q = \sum_y \phi(y) q(y,x) = \sum_{y \neq x} (\phi(y) q(y,x) - \phi(x) q(x,y))\,.$$

A sufficient condition for stationarity of a measure is thus if the **detailed balance** condition:

$$\phi(y) q(y,x) = \phi(x) q(x,y) \quad \text{for all } x, y \in S$$

is satisfied.

Example 2.18 (**The invariant distribution for a simple reversible reaction)**. Consider the following production/degratdation system with mass action kinetics and both rate constants equal to 1:

$$0 \rightleftharpoons A\,.$$

We will write down the stationary distribution for this CRN. When there are n molecules of A, then the intensity of the forward reaction is 1, and if $n \geq 1$, the intensity of the reverse reaction

is n, giving the Q-matrix:

$$Q = \begin{pmatrix} -1 & 1 & 0 & 0 & 0 & \cdots \\ 1 & -2 & 1 & 0 & 0 & \cdots \\ 0 & 2 & -3 & 1 & 0 & \cdots \\ 0 & 0 & 3 & -4 & 1 & \cdots \\ 0 & 0 & 0 & 4 & -5 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

If ϕ is in the kernel of Q^t, we thus have the system of equations

$$\phi(1) = \phi(0), \quad 2\phi(2) = 2\phi(1) - \phi(0), \quad \dots \quad n\phi(n) = n\phi(n-1) - \phi(n-2).$$

Solving these recursively in terms of $\phi(0)$, we get

$$\phi = \phi(0)\left(1,\ 1,\ \frac{1}{2},\ \frac{1}{3!},\ \dots\right)$$

and we can easily prove by induction that the nth term is $\frac{1}{n!}$. Normalising, we get that $\phi(0) = e^{-1}$, so $\phi(n) = e^{-1}/n!$.

Exercise 2.19 (**The invariant distribution for a simple reversible reaction)**. Consider again the production-degration system with stochastic mass action kinetics, forward rate constant equal to k_1 and reverse rate constant equal to k_2:

$$0 \xrightarrow{k_1} A, \quad A \xrightarrow{k_2} 0.$$

Write down the Q-matrix for this system, and by examining $\ker Q^t$ find the unique invariant distribution ϕ for this CRN. [You should obtain a Poisson distribution, with parameter to be determined.]

Exercise 2.20 (**A stationary measure, but no stationary distribution)**. Consider the following CRN with mass action kinetics:

$$0 \to A, \quad A \to 0, \quad A \to 2A.$$

- Show that, with all rate constants set to 1, the system admits a stationary measure, but no stationary distribution.
- Now assume that the first and third (production) reactions have rate constants 1, while the second reaction (degradation) has rate constant 2. In this case, does the system admit a stationary measure? Does it admit a stationary distribution?

2.11 The chemical master equation

Recall that given any initial probability distribution, say ϕ_0, and the transition matrices $P(t)$, we can calculate the probability of finding the system in each state at time t simply as $\phi_t = \phi_0 P(t)$. Differentiating and applying the forward equation gives

$$\frac{d\phi_t}{dt} = \phi_t Q\,. \tag{12}$$

Let us consider each component of this equation: for each state x from the state space S, we write $p_x(t)$ for $\mathbb{P}(\mathbf{X}(t) = x)$ (i.e., $\phi_t(x)$), to get:

$$\frac{dp_x}{dt} = \sum_{y \in S} p_y(t) Q(y, x) = \sum_{j=1}^{m} \big(\nu_j(x - \zeta_j) p_{x-\zeta_j}(t) - \nu_j(x) p_x(t)\big) \quad (x \in S)\,. \tag{13}$$

In writing the second equality, we have abused notation a little: it is possible that $x - \zeta_j$ does not belong to the state-space S, e.g. if it has some negative components, in which case we take $\nu_j(x - \zeta_j) p_{x-\zeta_j}(t) = 0$. The system of ODEs (13), is termed the **chemical master equation** (or **CME** for short). It is simply the forward equation written out component by component.

Note that since we have assumed a finite number m of chemical reactions, each ODE in the CME has only finitely many terms. In fact, we have two kinds of contributions to the probability of finding the system in state x at time t: *positive terms* corresponding to arrival into state x via the firing of a single reaction; and *negative terms* corresponding to leaving state x via reactions which take us away from state x. Note that the negative terms can be aggregated into a single term $-p_x(t) \sum_{j=1}^{m} \nu_j(x)$.

Example 2.21 (**Writing down the CME of a simple CRN)**. Consider the following system of reactions, with mass action kinetics, and all rate constants set to 1.

$$0 \to A, \quad A \to 0, \quad A \to 2A\,.$$

We will write down the CME for this system. Let $p_n(t)$ be the probability that there are n molecules of A at time t. The three "reaction vectors" are, in this case, just the scalars $\zeta_1 = 1, \zeta_2 = -1$ and $\zeta_3 = 1$ respecitively. The mass action rates are $\nu_1(n) = 1$, $\nu_2(n) = n$ $(n \geq 1)$ and $\nu_3(n) = n$ $(n \geq 1)$. We thus have,

$$\frac{dp_0}{dt} = -p_0(t) + p_1(t),$$

and, for $n \geq 1$,

$$\begin{aligned}
\frac{dp_n}{dt} &= \sum_{j=1}^{3} \big(\nu_j(n - \zeta_j) p_{n-\zeta_j}(t) - \nu_j(n) p_n(t)\big) \\
&= p_{n-1}(t) - p_n(t) + (n+1) p_{n+1}(t) - n p_n(t) + (n-1) p_{n-1}(t) - n p_n(t) \\
&= n p_{n-1}(t) - (2n+1) p_n(t) + (n+1) p_{n+1}(t)\,.
\end{aligned}$$

In an earlier exercise (Exercise 2.20), we computed a stationary measure for this system and observed that it admits no stationary distribution. We can find the same measure by setting $\frac{dp_n}{dt} = 0$ for each n and solving the resulting recurrence relations. But note that in this case we can no longer regard the quantities $p_n(t)$ as probabilities.

Example 2.22 (**A model of enzyme kinetics)**. Consider the following CRN which represents a model of enzyme kinetics with inflow and outflow of the enzyme and main substrate:

$$\mathrm{S} + \mathrm{E} \rightleftharpoons \mathrm{C} \rightleftharpoons \mathrm{P} + \mathrm{E}, \quad \mathrm{S} \rightleftharpoons 0 \rightleftharpoons \mathrm{E}\,.$$

This model has four species $\{\mathrm{S}, \mathrm{E}, \mathrm{C}, \mathrm{P}\}$ which we take in this order, and eight reactions. The stoichiometric matrix can be checked to have rank 4 (check this!) and so there are no conservation laws. The state space consists of all of $\mathbb{Z}^4_{\geq 0}$. With mass action kinetics and any rate constants, the system has a unique stationary distribution. This time, however, it is not easy to find this distribution directly and we need to call on additional theory to prove its existence and find its form.

For simplicity we set all the rate constants to be 1, and the reactor volume to be 1. Let $x = (s, \eta, c, p) \in \mathbb{Z}^4_{\geq 0}$ denote the vector of species numbers of the four species involved. Then

$$\phi(x) := \frac{e^{-4}}{s!\,\eta!\,c!\,p!}, \quad x \in \mathbb{Z}^4_{\geq 0}$$

is a stationary distribution of the system. Seeing that this is a product of four independent Poisson distributions with parameter 1, allows us to confirm that $\sum_{x \in \mathbb{Z}^4_{\geq 0}} \phi(x) = 1$. We can check that the distribution is stationary by appealing directly to the CME (13). We find that each reversible pair of reactions contributes a pair of terms summing to zero in $\frac{\mathrm{d}p_x}{\mathrm{d}t}$. For example the pair $\mathrm{S} + \mathrm{E} \rightleftharpoons \mathrm{C}$ contribute the terms

$$\frac{e^{-4}(c - s\eta)}{s!\,\eta!\,c!\,p!} \quad \text{and} \quad \frac{e^{-4}(s\eta - c)}{s!\,\eta!\,c!\,p!}$$

to $\frac{\mathrm{d}p_x}{\mathrm{d}t}$, and these clearly add up to zero.

Exercise 2.23. Confirm the details in Example 2.22. You do not need to calculate every term in the CME; but make sure you can correctly write down terms in the CME for, say the first pair of reactions $\mathrm{S} + \mathrm{E} \rightleftharpoons \mathrm{C}$, and check that they are as given and add up to zero.

2.12 Evolution of the mean and moments

We can use the CME (13) to write down an equation for the evolution of the mean $\mathbb{E}\mathbf{X}(t)$. We multiply (13) through by x_i, and sum over all states $x \in S$:

$$\begin{aligned}\frac{d}{dt}\mathbb{E}X_i(t) = \sum_{x\in S} x_i \frac{dp_x}{dt} &= \sum_{j=1}^{m}\sum_{x\in S}\left(x_i\nu_j(x-\zeta_j)p_{x-\zeta_j}(t) - x_i\nu_j(x)p_x(t)\right)\\ &= \sum_{j=1}^{m}\sum_{x\in S}\left((x_i+\zeta_{ij})\nu_j(x)p_x(t) - x_i\nu_j(x)p_x(t)\right)\\ &= \sum_{j=1}^{m}\zeta_{ij}\mathbb{E}[\nu_j(X(t))]\,.\end{aligned}$$

To go from the first to the second line, we reindex the first part of the sum, bearing in mind that it is over *all* states of the system. I.e., we define the new (dummy) variable $y := x - \zeta_j$, and then rename it as x. We have assumed as usual that any reasonable intensity functions prevent X_i from becoming negative; also, we have chosen to write the ith element of reaction vector ζ_j as ζ_{ij} as this corresponds to the (i,j)th element of the stoichiometric matrix of the system. In vector notation, and calling the stoichiometric matrix Γ, we could thus write,

$$\frac{d}{dt}\mathbb{E}\mathbf{X}(t) = \Gamma\,\mathbb{E}[\boldsymbol{\nu}(\mathbf{X}(t))]\,. \tag{14}$$

Note the formal similarity between (14) and the deterministic system (1), namely $\dot{x} = \Gamma\boldsymbol{\nu}(x)$. In fact, if all components of $\boldsymbol{\nu}$, the vector of intensities, are linear or constant functions of $\mathbf{X}$, then $\mathbb{E}[\boldsymbol{\nu}(\mathbf{X}(t))] = \boldsymbol{\nu}(\mathbb{E}[\mathbf{X}(t)])$ and we get *exactly* the same equation as in the deterministic case: i.e., the mean evolves according to the usual deterministic equation. But otherwise this is not, in general, the case, as the following example illustrates:

Example 2.24 (**Evolution of the mean of a simple CRN**). Consider the CRN

$$2A \to 0, \quad 0 \to A\,,$$

with mass action kinetics. Letting n denote the number of molecules of A, and fixing the reactor volume at 1, the intensity of $2A \to 0$ is $\nu_1 := k_1 n(n-1)$ (so $\mathbb{E}\nu_1 = k_1(\mathbb{E}[n^2] - \mathbb{E}n)$), while that of $0 \to A$ is $\nu_2 := k_2$ (so $\mathbb{E}\nu_2 = k_2$). Here k_1 and k_2 are the deterministic mass action rate constants of the reactions. Using (14) we find, for the evolution of the mean

$$\frac{d}{dt}\mathbb{E}n = k_2 - 2k_1(\mathbb{E}[n^2] - \mathbb{E}n)\,.$$

The corresponding deterministic equation is $\frac{da}{dt} = k_2 - 2k_1 a^2$, where a denotes the concentration of A. Only if $\mathbb{E}[n^2] - \mathbb{E}n = (\mathbb{E}n)^2$, i.e., the variance of n is equal to its expectation, do we get an equation formally identical to the deterministic equation; and we have no reason to expect that to hold in this case (it doesn't).

More importantly, note the main difficulty: the evolution of $\mathbb{E}n$ depends on $\mathbb{E}[n^2]$. In fact, if we continue and write down an ODE for the evolution of $\mathbb{E}[n^2]$ (try it!) we find that it depends on

$\mathbb{E}[n^3]$, and so forth. Thus we are back in the situation of an infinite system of coupled ODEs for the evolution of the moments.

This example appears in Section 1.4 of [Erban and Chapman] and we will see below how to partially solve this problem for this network and find the moments of the steady state distribution for this CRN.

In the same way that we can start with the chemical master equation and use it to write down a differential equation for the evolution of the mean $\mathbb{E}\mathbf{X}(t)$, we can do the same thing with higher moments of the random variables $\mathbf{X}(t)$. In fact, we can do this for all moments at once, via the probability generating function.

2.12.1 Evolution of the probability generating function

Recall that the probability generating function (PGF) of a random variable Y taking non-negative integer values is defined as

$$G(s) = \sum_{n=0}^{\infty} s^n \mathbb{P}(Y = n)\,.$$

The function G has radius of convergence at least 1 and $G(1)$ and $G(-1)$ are always well defined. Moreover, moments of the distribution of Y (assuming they exist) are easily computed as derivatives of G at 1; e.g., $\mathbb{E}[Y] = G'(1)$, and $\mathbb{E}[Y^2] = G''(1) + G'(1)$.

In the case of a CRN involving only one species, if we take $X(t)$ to be the number of molecules of this species, at time t, we get the PGF for $X(t)$

$$G(s,t) = \sum_{n=0}^{\infty} s^n \mathbb{P}(X(t) = n) = \sum_{n=0}^{\infty} s^n p_n(t)\,.$$

Differentiating with respect to time and using the chemical master equation gives

$$\begin{aligned}
\frac{\partial}{\partial t} G(s,t) &= \sum_{n=0}^{\infty} s^n \frac{dp_n}{dt} = \sum_{n=0}^{\infty} \sum_{j=1}^{m} \left(s^n \nu_j(n-\zeta_j) p_{n-\zeta_j}(t) - s^n \nu_j(n) p_n(t) \right) \\
&= \sum_{n=0}^{\infty} \sum_{j=1}^{m} \left(s^{n+\zeta_j} \nu_j(n) p_n(t) - s^n \nu_j(n) p_n(t) \right) \\
&= \sum_{j=1}^{m} (s^{\zeta_j} - 1) \sum_{n=0}^{\infty} s^n \nu_j(n) p_n(t)\,.
\end{aligned}$$

We can use this expression to write down a PDE for the evolution of the PGF, as we can see in the next example.

Example 2.25 (**A PDE for the evolution of the probability generating function)**. We consider again the CRN which appeared in Example 2.24, namely,

$$2A \to 0, \quad 0 \to A\,,$$

with n denoting the number of molecules of A, and intensities $\nu_1(n) := k_1 n(n-1)$ and $\nu_2(n) := k_2$ (we may consider k_1, k_2 as deterministic rate constants, and we set the reactor volume to be 1). Note that in this case we have the reaction "vectors" $\zeta_1 = -2,\ \zeta_2 = 1$. Applying the formula above, the PGF $G(s,t) := \sum_{n=0}^{\infty} s^n p_n(t)$ now evolves according to

$$\frac{\partial}{\partial t}G(s,t) = \sum_{n=0}^{\infty}\left((s^{-2}-1)s^n k_1 n(n-1)p_n(t) + (s-1)s^n k_2 p_n(t)\right)\,.$$

On the other hand

$$\frac{\partial}{\partial s}G(s,t) = \sum_{n=0}^{\infty} n s^{n-1} p_n(t)\,, \quad \frac{\partial^2}{\partial s^2}G(s,t) = \sum_{n=0}^{\infty} n(n-1)s^{n-2}p_n(t)$$

Comparing expressions we see that $G(s,t)$ satisfies the second-order linear PDE

$$\frac{\partial G}{\partial t} = k_1(1-s^2)\frac{\partial^2 G}{\partial s^2} + k_2(s-1)G\,.$$

As a consequence, assuming the existence of a stationary distribution ϕ, the stationary PGF, $\widehat{G}(s) = \sum_{n=0}^{\infty} s^n \phi(n)$, must satisfy

$$0 = k_1(1-s^2)\frac{d^2\widehat{G}}{ds^2} + k_2(s-1)\widehat{G}, \quad \text{or, equivalently,} \quad \frac{d^2\widehat{G}}{ds^2} = \frac{k_2}{k_1(1+s)}\widehat{G}\,.$$

We need two boundary conditions in order to solve this equation: one is simply that $\widehat{G}(1) = \sum_{n=0}^{\infty}\phi(n) = 1$; the second is that $\widehat{G}(-1) = 0$. With these boundary conditions, we can find the solution to this second order, linear, ODE in terms of Bessel functions and compute the mean and variance of the stationary distribution from its PGF (see Section 1.4 in [Erban and Chapman]). We find that, as predicted in Example 2.24, the steady state mean differs (although only slightly) from the mean predicted by the deterministic equation.

Remark on the multivariate case. In the case of a CRN involving n species we can use the multivariate PGF $G_X(s_1,\ldots,s_n) = \mathbb{E}[s_1^{X_1}s_2^{X_2}\cdots s_n^{X_n}]$. We can abbreviate $s_1^{X_1}s_2^{X_2}\cdots s_n^{X_n}$ as $\mathbf{s}^X$ and write the PGF as $G_X(\mathbf{s}) = \mathbb{E}[\mathbf{s}^X]$. We can write down an equation for its evolution similarly to the 1D case, using the chemical master equation:

$$\begin{aligned}
\frac{\partial}{\partial t}G_X(\mathbf{s},t) &= \sum_x\sum_{j=1}^{m}\mathbf{s}^x\frac{dp_x}{dt} = \sum_x\sum_{j=1}^{m}\left(\mathbf{s}^x\nu_j(x-\zeta_j)p_{x-\zeta_j}(t) - \mathbf{s}^x\nu_j(x)p_x(t)\right)\\
&= \sum_x\sum_{j=1}^{m}\left(\mathbf{s}^{x+\zeta_j}\nu_j(x)p_x(t) - s^x\nu_j(x)p_x(t)\right)\\
&= \sum_{j=1}^{m}(\mathbf{s}^{\zeta_j}-1)\sum_x \mathbf{s}^x\nu_j(x)p_x(t)\,.
\end{aligned}$$

Here the first summation on each line is over all states of the system. As before, we can hope to write down a PDE for the evolution of the PGF in this case and, in theory, use this to make claims about evolution of the distribution.

Exercise 2.26 (**Writing down a PDE for the evolution of a multivariate PGF)**. Write down a PDE for the evolution of the multivariate PGF of the CRN:

$$0 \xrightarrow{k_1} A \xrightarrow{k_2} B \xrightarrow{k_3} 0 .$$

You may assume that the rate constants are deterministic, and the reactor volume is 1. It simplifies notation to write $n := n_A$, $m := n_B$, so that $G(s_1, s_2, t) = \mathbb{E}[s_1^n s_2^m]$. You should obtain the PDE:

$$G_t = k_1(s_1 - 1)G + k_2(s_2 - s_1)G_{s_1} + k_3(1 - s_2)G_{s_2} .$$