Computational Modelling
Continuous Time Markov Chains

Luca Bortolussi[1]

[1]Dipartimento di Matematica e Geoscienze
Università degli studi di Trieste

Office 328, third floor, H2bis
lbortolussi@units.it

DSSC, Trieste

Outline

Outline

1. Preliminaries
   - Exponential Distribution
2. Continuous Time Markov Chains
   - Main concepts
   - Poisson Process
   - Time-inhomogeneous rates
3. Population Continuous Time Markov Chains
4. Simulation
   - SSA
   - Next Reaction Method
   - τ-leaping

EXPONENTIAL DISTRIBUTION

DEFINITION

A random variable $T : (\Omega, \mathcal{S}) \to [0, \infty]$ is $Exp(\lambda)$ iff

- Cdf is $\mathbb{P}(T < t) = 1 - e^{-\lambda t}$
- Survival probability is $\mathbb{P}(T > t) = e^{-\lambda t}$
- Density is $f_T(t) = \lambda e^{-\lambda t}$, $t \geq 0$.

The expected value of T is $\mathbb{E}(T) = \int_0^\infty \mathbb{P}(T > t)dt = \frac{1}{\lambda}$.

MEMORYLESS PROPERTY

$T \sim Exp(\lambda)$ if and only if the following memoryless property holds:

$$\mathbb{P}(T > s + t | T > s) = \mathbb{P}(T > t) \text{ for all } s, t \geq 0.$$

In fact
$\mathbb{P}(T > s + t | T > s) = \mathbb{P}(T > s + t) / \mathbb{P}(T > s) = e^{-\lambda(t+s)} e^{\lambda s} = e^{-\lambda t}$.

EXPONENTIAL DISTRIBUTION

DEFINITION

A random variable $T : (\Omega, \mathcal{S}) \to [0, \infty]$ is *Exp*(λ) iff

- Cdf is $\mathbb{P}(T < t) = 1 - e^{-\lambda t}$
- Survival probability is $\mathbb{P}(T > t) = e^{-\lambda t}$
- Density is $f_T(t) = \lambda e^{-\lambda t}$, $t \geq 0$.

The expected value of T is $\mathbb{E}(T) = \int_0^\infty \mathbb{P}(T > t)dt = \frac{1}{\lambda}$.

INSTANTANEOUS FIRING PROBABILITY

An exponential distribution with rate λ can be seen as the firing time of an event who has probability of firing between time t and $t + dt$ equal to λdt.

Call $p(t) = \mathbb{P}\{T \geq t\}$. Then $p(t + dt) = p(t) \cdot (1 - \lambda dt)$,

from which $\frac{dp(t)}{dt} = -\lambda p(t)$, that has solution $p(t) = e^{-\lambda t}$ (as $p(0) = 1$).

EXPONENTIAL DISTRIBUTION: RACE CONDITION

THEOREM

Let I be a countable set and let T_k, $k \in I$, be independent random variables, with $T_k \sim Exp(q_k)$ and $q = \sum_{k \in I} q_k < \infty$. Set $T = \inf_k T_k$. Then this infimum is obtained at a unique random value K of I, with probability 1. Moreover, T and K are independent, $T \sim Exp(q)$ and $\mathbb{P}(K = k) = q_k / q$.

PROOF

Set $K = k$ if $T_k < T_j$ for all $j \neq k$, K is undefined otherwise. Then

$$\begin{aligned}\mathbb{P}(K = k \text{ and } T \geq t) &= \mathbb{P}(T_k \geq t \text{ and } T_j > T_k \text{ for all } j \neq k) \\ &= \int_t^\infty q_k e^{-q_k s} \mathbb{P}(T_j > s \text{ for all } j \neq k) ds \\ &= \int_t^\infty q_k e^{-q_k s} \prod_{j \neq k} e^{-q_j s} ds \\ &= \int_t^\infty q_k e^{-qs} ds = \frac{q_k}{q} e^{-qt}\end{aligned}$$

Computing the marginal distributions for K and T, we obtain the claimed results. Moreover, their joint distribution turns out to be the product of the marginals, thus showing that K and T are independent and that $\mathbb{P}(K = k \text{ for some } k) = 1$.

Outline

CTMC: DEFINITION

S-VALUED STOCHASTIC PROCESS

Let S be finite or countable. A continuous-time random process $(X_t)_{t\geq 0} = \{X_t \mid t \geq 0\}$, with values in S, is a family of random variables $X_t : (\Omega, \mathcal{S}) \to (S, 2^S)$ that are *right-continuous* w.r.t. t.
Therefore, X_t (or $X(t)$) has *cadlag* sample paths.
Right continuous processes are determined by their *finite-dimensional distributions*.

CONTINUOUS TIME MARKOV CHAIN

A Continuous Time Markov Chain is a right-continuous continuous-time random process satisfying the memoryless condition: for each n, t_i and s_i:

$$\mathbb{P}(X_{t_n} = s_n \mid X_{t_0} = s_0, \ldots, X_{t_{n-1}} = s_{n-1}) = \mathbb{P}(X_{t_n} = s_n \mid X_{t_{n-1}} = s_{n-1}).$$

CTMC: RACE CONDITION

CTMC AS A GRAPH

A CTMC on a state space S can be seen as a labelled graph. Each edge takes some time to be crossed, exponentially distributed with the rate labelling the edge.
In each state, there is a race condition between the different exiting edges: the fastest is traversed.
The memoryless property follows from that of the exponential distribution.

Q-MATRIX

A Q-matrix is the $|S| \times |S|$ matrix such that:

1. $q_{ij} \geq 0$, $i \neq j$ is the rate of the exponential distribution giving the time needed to go from state s_i to state s_j
2. $q_{ii} = -\sum_{j \neq i} q_{ij}$ is the opposite of the exit rate from state i.

Therefore, each row of the Q-matrix sums up to zero.

A SIMPLE EXAMPLE: THE MOOD CHAIN

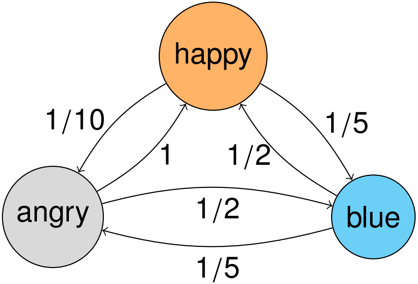

$$S = \{happy, blue, angry\}$$

$$Q = \begin{pmatrix} -\frac{3}{10} & \frac{1}{5} & \frac{1}{10} \\ \frac{1}{2} & -\frac{7}{10} & \frac{1}{5} \\ 1 & \frac{1}{2} & -\frac{3}{2} \end{pmatrix}$$

JUMP CHAIN AND HOLDING TIMES

FACTORIZING EACH JUMP

In each state i, we have a race condition between k transitions, each exponentially distributed with rate q_{ij}. Hence, the time spent is $T = \inf T_{ij}$. By the properties of the exponential distribution, we know that T has rate $q_i = \sum_j q_{ij}$, and that the transition that fires is independent from T and the next state j is chosen with probability q_{ij}/q_i.

JUMP CHAIN AND HOLDING TIMES

We can therefore factorize $X(t)$ into

- a DTMC Y_n, with probability matrix Π, defined by $\pi_{ij} = \frac{q_{ij}}{-q_{ii}}$, if $i \neq j$, and $\pi_{ii} = 0$;
- a sequence of jump times τ_n, where τ_n is the time of the n-th jump. Letting q_i the jump rate from state s_i, then $T_n = \tau_n - \tau_{n-1}$, the n-th holding time, is distributed exponentially with rate q_{Y_n}.
- Y_n and each T_i are independent.
- Hence $X(t) = Y_n$ for $\tau_n \leq t < \tau_{n+1}$.

A SIMPLE EXAMPLE: THE MOOD CHAIN

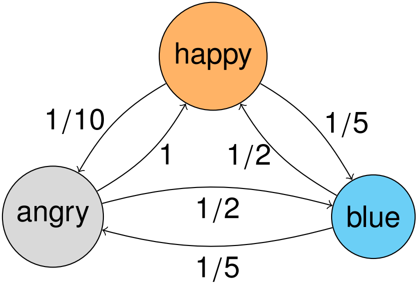

$$S = \{happy, blue, angry\}$$

Jump chain

$$\Pi = \begin{pmatrix} 0 & \frac{2}{3} & \frac{1}{3} \\ \frac{5}{7} & 0 & \frac{2}{7} \\ \frac{2}{3} & \frac{1}{3} & 0 \end{pmatrix}$$

Exit rates

$$q = \left(\frac{3}{10}, \frac{7}{10}, \frac{3}{2}\right)$$

CHAPMAN-KOLMOGOROV EQUATIONS

Let $P_{ij}(t) = \mathbb{P}\{X(t) = s_j \mid X(0) = s_i\}$. Then

$$\begin{aligned}
P_{ij}(t+s) &= \mathbb{P}\{X(t+s) = s_j \mid X(0) = s_i\} \\
&= \sum_k \mathbb{P}\{X(t+s) = s_j, X(t) = s_k \mid X(0) = s_i\} \\
&= \sum_k \mathbb{P}\{X(s) = s_j \mid X(0) = s_k\}\mathbb{P}\{X(t) = s_k \mid X(0) = s_i\} \\
&= \sum_k P_{ik}(s)P_{kj}(t).
\end{aligned}$$

Hence $P(t)$, as a matrix, satisfies

$$P(t+s) = P(t)P(s) = P(s)P(t),$$

which is the semigroup property, also known as Chapman-Kolmogorov equations.

KOLMOGOROV EQUATIONS

Using properties of the exponential, we can compute $P(dt)$:

- $P_{ij}(dt) = q_{ij}dt$, for $i \neq j$;
- $P_{ii}(dt) = 1 - \sum_{j \neq i} q_{ij}dt = 1 + q_{ii}dt$

Hence $P(dt) = I + Qdt$
From the CK equations: $P(t + dt) = P(t) + P(t)Qdt$, from which

$$\frac{dP(t)}{dt} = P(t)Q,$$

which is the forward Kolmogorov equation.
Using CK the other way round: $P(t + dt) = P(t) + QP(t)dt$, so

$$\frac{dP(t)}{dt} = QP(t),$$

which is the backward Kolmogorov equation.

A SIMPLE EXAMPLE: THE MOOD CHAIN

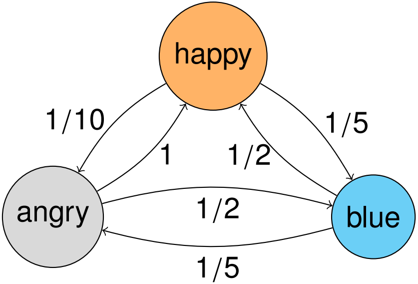

$$S = \{happy, blue, angry\}$$

$$p_0 = (0, 1, 0) \quad p = p_0 P$$

$$\frac{d}{dt} p_0 P = p_0 P Q \Rightarrow \frac{d}{dt} p = pQ$$

A simple example: the mood chain

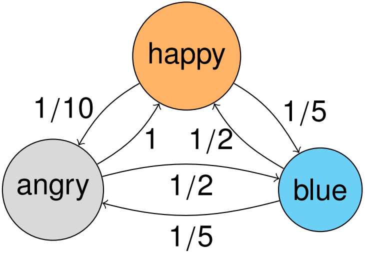

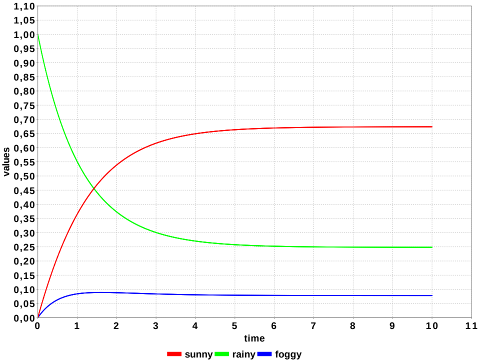

Outline

POISSON PROCESS: DEFINITION

A Poisson process $\mathcal{N}_\lambda(0, t)$ with rate λ is a process that counts how many times an exponential distribution with rate λ has fired from time 0 to time t.

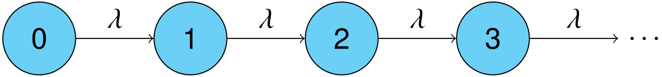

It can be seen as a CTMC on the state space $S = \mathbb{N}$, with rate matrix Q given by $q_{i,i+1} = \lambda$, and zero elsewhere.
It's a very common process. For instance, it is the simplest model of job arrivals in a queue.

POISSON PROCESS: BASIC PROPERTIES

POISSON RANDOM VARIABLE

A Poisson r.v. $\mathcal{Y}(\lambda)$ with rate λ ($\mathcal{Y}(\lambda) \sim Poisson(\lambda)$) is a r.v. on $\mathbb{N}$ with probability distribution $\mathbb{P}\{\mathcal{Y}(\lambda) = n\} = \frac{e^{-\lambda}\lambda^n}{n!}$.
Its generating function is $G(z) = \mathbb{E}[z^{\mathcal{Y}(\lambda)}] = e^{\lambda(z-1)}$.

POISSON PROCESS DISTRIBUTION

The distribution of $\mathcal{N}_\lambda(0,t)$ is $Poisson(\lambda t)$.

We show that $G_t(z) = \mathbb{E}[z^{\mathcal{N}(0,t)}] = e^{\lambda t(z-1)}$.
By the Markov property, $\mathcal{N}(0, t+s) = \mathcal{N}(0,t) + \mathcal{N}(t,s)$, and the two processes on the right are independent.
Then $G_{t+dt}(z) = \mathbb{E}[z^{\mathcal{N}(0,t)}]\mathbb{E}[z^{\mathcal{N}(t,t+dt)}]$. But $\mathbb{E}[z^{\mathcal{N}(t,t+dt)}] = (1-\lambda dt)z^0 + \lambda dt z^1$, hence $G_{t+dt}(z) = G_t(z) + \lambda(z-1)G_t(z)dt$, and so

$$\frac{dG_t(z)}{dt} = \lambda(z-1)G_t(z),$$

which has solution $G_t(z) = e^{\lambda t(z-1)}$, as $\mathcal{N}_\lambda(0,0) = 0$ with probability 1.

INVARIANT MEASURES AND STEADY STATE

INVARIANT MEASURE

Consider a CTMC with rate matrix Q and finite state space S. An invariant measure for the CTMC is a probability distribution π satisfying

$$\pi Q = 0.$$

If Q is irreducible (has a strongly connected graph), then it has a unique invariant measure.

STEADY STATE BEHAVIOUR

Consider an irreducible CTMC with rate matrix Q and finite state space S, and let π be its invariant probability distribution. Then, for each $s_i, s_j \in S$,

$$\lim_{t\to\infty} P_{ij}(t) = \pi_j.$$

Notice that aperiodicity is not required. Why?

EXAMPLE: BIRTH-DEATH PROCESS

A birth-death process is a CTMC on $S = \mathbb{N}$ with birth rate λ_i (from i to $i+1$) and death rate μ_i (from i to $i-1$).

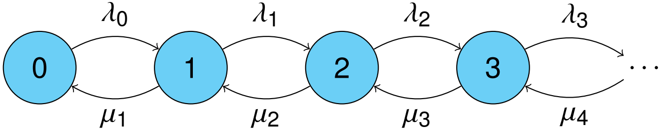

Example: birth-death process

A birth-death process is a CTMC on $S = \mathbb{N}$ with birth rate λ_i (from i to $i+1$) and death rate μ_i (from i to $i-1$).

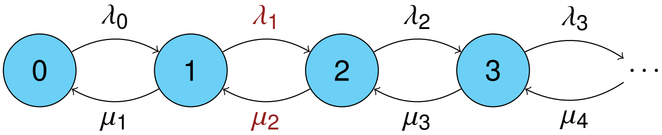

To derive the steady state π, we can use the fact that the net flow along each cut must be zero (why?):

$$\pi_i \lambda_i = \pi_{i+1} \mu_{i+1}$$

Hence we get:

$$\pi_k = \prod_{i=0}^{k-1} \frac{\lambda_i}{\mu_{i+1}} \pi_0; \qquad \pi_0 = \left(1 + \sum_{k=1}^{\infty} \prod_{i=0}^{k-1} \frac{\lambda_i}{\mu_{i+1}} \right)^{-1}$$

EXAMPLE: BIRTH-DEATH PROCESS

Consider a birth-death process with constant birth rate λ and constant death rate μ. It is the model of an M/M/∞ queue.

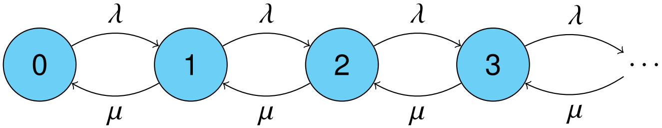

$$\pi_k = \left(\frac{\lambda}{\mu}\right)^k \pi_0; \qquad \pi_0 = \left(1 + \sum_{k=1}^{\infty}\left(\frac{\lambda}{\mu}\right)^k\right)^{-1}$$

- If $\lambda \geq \mu$, then $\pi_0 = 0 = \pi_k$. No state is positive recurrent, there is no invariant measure. The chain escapes to infinity.
- If $\lambda < \mu$, then $\pi_0 = \frac{1-\lambda/\mu}{2-\lambda/\mu}$ and $\pi_k = \left(\frac{\lambda}{\mu}\right)^k \frac{1-\lambda/\mu}{2-\lambda/\mu}$

EXAMPLE: BIRTH-DEATH PROCESS

If $\lambda < \mu$, then $\pi_0 = \frac{1-\lambda/\mu}{2-\lambda/\mu}$ and $\pi_k = \left(\frac{\lambda}{\mu}\right)^k \frac{1-\lambda/\mu}{2-\lambda/\mu}$

Assume $\lambda = 1$, $\mu = 2$.

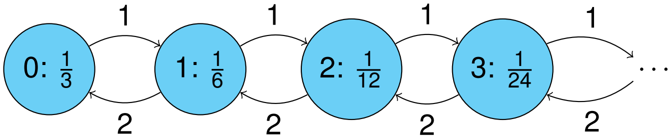

MATRIX EXPONENTIAL

The solution of the forward Kolmogorov equation $\frac{dP(t)}{dt} = P(t)Q$, for a generic CTMC, can be given in terms of the matrix exponential

$$P(t) = e^{Qt} = \sum_{n=0}^{\infty} \frac{t^n Q^n}{n!}.$$

However, numerical computation of the series expansion is numerically unstable.

UNIFORMIZATION

A more efficient strategy is to solve the uniformized CTMC.
Let $\lambda \geq \max_i\{-q_{ii}\}$.
Then one considers a CTMC with jump chain $Y(n)$ with matrix

$$\Pi = I + \frac{1}{\lambda}Q,$$

and uniform exit rate λ.
The number of fires of this CTMC up to time t is a Poisson process $N_\lambda(0,t)$, and so

$$X(t) = Y_{N(0,t)} = Y_{\mathcal{Y}(\lambda t)}.$$

It follows that

$$P(t) = \sum_{n=0}^{\infty} \frac{e^{-\lambda t}(\lambda t)^n}{n!}\Pi^n,$$

which can be truncated above (and below) by bounding the Poisson r.v.

A SIMPLE EXAMPLE: THE MOOD CHAIN

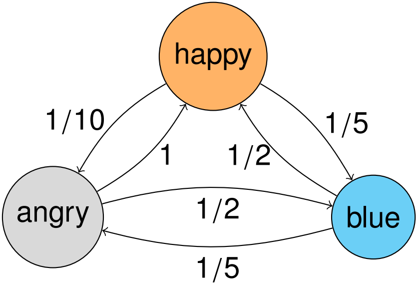

Upper bound on exit rate: 2

$$P(t) = \sum_{n=0}^{\infty} \frac{e^{-2t}(2t)^n}{n!} \Pi^n$$

$$\Pi = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} -\frac{3}{10} & \frac{1}{5} & \frac{1}{10} \\ \frac{1}{2} & -\frac{7}{10} & \frac{1}{5} \\ 1 & \frac{1}{2} & -\frac{3}{2} \end{pmatrix} = \begin{pmatrix} \frac{17}{20} & \frac{2}{20} & \frac{1}{20} \\ \frac{5}{20} & \frac{13}{20} & \frac{2}{20} \\ \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}$$

Outline

TIME-INHOMOGENEOUS EXPONENTIAL

DEFINITION

A exponential random variable $T \sim Exp(\lambda)$ has time inhomogeneous rate iff $\lambda = \lambda(t)$ is a function $\lambda : [0, \infty[\rightarrow \mathbb{R}^+$.

- Cumulative rate is $\Lambda(t) = \int_0^t \lambda(s)ds$
- Cdf is $\mathbb{P}(T < t) = 1 - e^{-\Lambda(t)}$
- Survival probability is $\mathbb{P}(T > t) = e^{-\Lambda(t)}$

INVERSION METHOD

One can simulate unidimensional random variables by sampling a uniform r.v. $U \in [0, 1]$, and then finding t^* such that $t^* = \inf_t \mathbb{P}(T \leq t) = U$.
For a time-inhomogeneous $Exp(\lambda(t))$, one has to solve $e^{-\Lambda(t)} = U$, iff $\Lambda(t) = -\log U = \xi$, with $\xi \sim Exp(1)$.
If λ is constant, then $\Lambda(t) = \lambda t$, and one has $t = -\frac{1}{\lambda}\log(U)$.
In general, one can either integrate $\lambda(t)$ or the equivalent ODE $\frac{d\Lambda(t)}{dt} = \lambda(t)$, and check for the root of $\Lambda(t) + \log(U)$ along the solution.

TIME-INHOMOGENEOUS POISSON PROCESS

A time-inhomogeneous Poisson process $\mathcal{N}_\lambda(0,t)$, $\lambda = \lambda(t)$, is a Poisson process with time-varying rate.

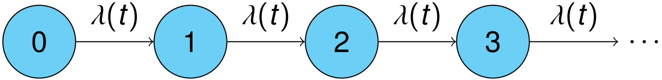

It can be shown (same generating function argument as above) that the distribution of $\mathcal{N}_\lambda(0,t)$ is *Poisson*$(\Lambda(t))$, i.e. it is the r.v.

$$\mathcal{Y}(\Lambda(t)) = \mathcal{Y}\left(\int_0^t \lambda(s)ds\right).$$

TIME-INHOMOGENEOUS CTMC

TIME-INHOMOGENEOUS CTMC

In general, if the rate matrix Q of a CTMC depends on time, $Q = Q(t)$, then the CTMC is time inhomogeneous.
The probability semigroup depends now also on the initial time: $P_{ij}(t_1, t_2) = \mathbb{P}\{X(t_2) = s_j \mid X(t_1) = s_i\}$.

FORWARD KOLMOGOROV EQUATION

$$\frac{\partial P(t_1, t_2)}{\partial t_2} = P(t_1, t_2)Q(t_2)$$

BACKWARD KOLMOGOROV EQUATION

$$\frac{\partial P(t_1, t_2)}{\partial t_1} = -Q(t_1)P(t_1, t_2)$$

Outline

POPULATION PROCESSES

SIR epidemics model
single individual

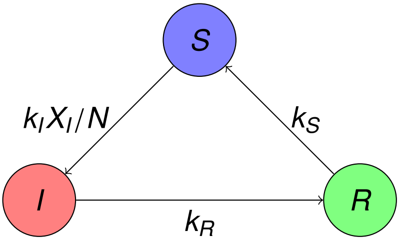

- Consider a CTMC model of a population epidemics in which each of N individuals can be in one of three states: susceptible (S), infected (I), and recovered (R);
- Infection rate depends on the density of infected individuals;
- The CTMC for N agents has 3^N states (if we distinguish the individuals) or $(N+1)^2$ states (if we just count them): *it's impossible to write down the Q matrix explicitly*.
- We need a better description of population CTMCs.

POPULATION CTMC

A population CTMC model is a tuple $\mathcal{X} = (\mathbf{X}, \mathcal{D}, \mathcal{T}, \mathbf{x_0})$, where:

1. $\mathbf{X}$ — vector of *variables* counting how many individuals in each state.
2. $\mathcal{D} = \prod_i \mathcal{D}_i$ — (countable) state space.
3. $\mathbf{x_0} \in \mathcal{D}$ —*initial state*.
4. $\eta_i \in \mathcal{T}$ — *global transitions*, $\eta_i = (a, \phi(\mathbf{X}), \mathbf{v}, r(\mathbf{X}))$
    1. a — event name (optional).
    2. $\phi(\mathbf{X})$ — guard.
    3. $\mathbf{v} \in \mathbb{R}^n$ — *update vector* (from $\mathbf{X}$ to $\mathbf{X} + \mathbf{v}$)
    4. $r : \mathcal{D} \to \mathbb{R}_{\geq 0}$ — rate function.

EXAMPLE: SIR EPIDEMICS

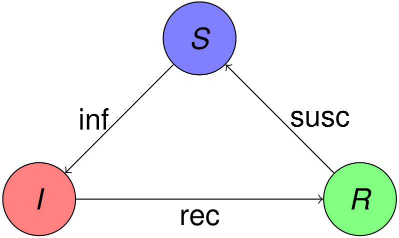

Three variables: X_S,X_I,X_R.
State space:
$\mathcal{D} = \{(n_1, n_2, n_3) \mid n_1 + n_2 + n_3 = N\} \subset \{0, \ldots, N\}^3$.

EXAMPLE: SIR EPIDEMICS

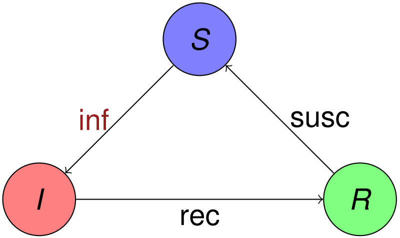

Three variables: X_S,X_I,X_R.
State space:
$\mathcal{D} = \{(n_1, n_2, n_3) \mid n_1 + n_2 + n_3 = N\} \subset \{0, \dots, N\}^3$.

Transitions:

- $(inf, \top, (-1, 1, 0) k_I \frac{X_I}{N} X_S)$

EXAMPLE: SIR EPIDEMICS

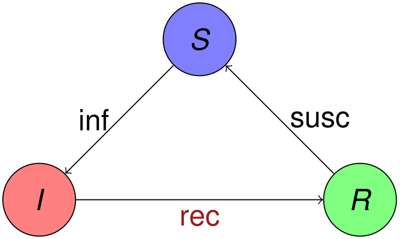

Three variables: X_S,X_I,X_R.
State space:
$\mathcal{D} = \{(n_1, n_2, n_3) \mid n_1 + n_2 + n_3 = N\} \subset \{0, \ldots, N\}^3$.

Transitions:

- $(inf, \top, (-1, 1, 0) k_I \frac{X_I}{N} X_S)$
- $(rec, \top, (0, -1, 1), k_R X_I)$

EXAMPLE: SIR EPIDEMICS

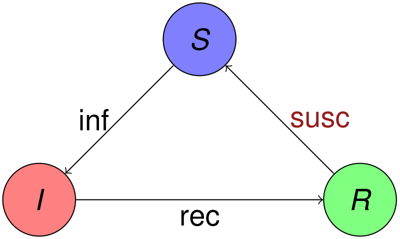

Three variables: X_S,X_I,X_R.
State space:
$\mathcal{D} = \{(n_1, n_2, n_3) \mid n_1 + n_2 + n_3 = N\} \subset \{0, \ldots, N\}^3$.

Transitions:

- $(inf, \top, (-1, 1, 0) k_I \frac{X_I}{N} X_S)$
- $(rec, \top, (0, -1, 1), k_R X_I)$
- $(susc, \top, (1, 0, -1), k_S X_R)$

EXAMPLE: SIR EPIDEMICS

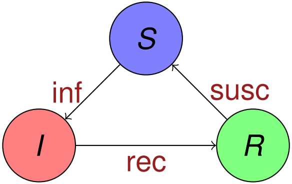

$N = 100,\ k_I = 1,\ k_R = 0.05,\ k_S = 0.01$

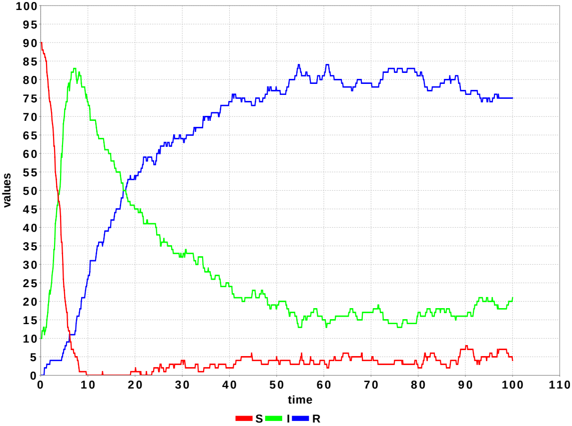

(1 run)

EXAMPLE: SIR EPIDEMICS

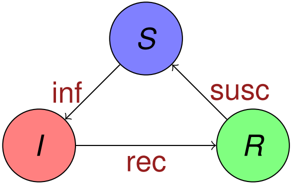

$N = 100,\ k_I = 1,\ k_R = 0.05,\ k_S = 0.01$

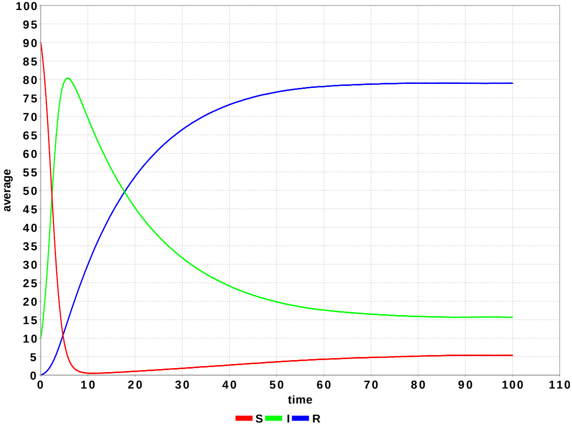

(average)

MASTER EQUATION

The Kolmogorov equation in the context of Population Processes is often know as master equation.

There is one equation per state $\mathbf{x} \in \mathcal{D}$, for the probability mass $P(\mathbf{x}, t)$, which considers the inflow and outflow of probability at time t.

$$\frac{dP(\mathbf{x}, t)}{dt} = \sum_{\eta \in \mathcal{T}} r_\eta(\mathbf{x} - \mathbf{v}_\eta) P(\mathbf{x} - \mathbf{v}_\eta, t) - \sum_{\eta \in \mathcal{T}} r_\eta(\mathbf{x}) P(\mathbf{x}, t)$$

POISSON REPRESENTATION

Population CTMC admit a simple description in terms of Poisson processes.
Essentially, we introduce variables $R_\eta(t)$ counting how many times each transition η has fired up to time t. Hence we can write:

$$X(t) = X(0) + \sum_{\eta \in \mathcal{T}} \mathbf{v}_\eta R_\eta(t).$$

It turns out that $R_\eta(t)$ is a time-inhomogeneous Poisson process with cumulative rate $\int_0^t r_\eta(X(s))ds$, independent from the other $R_{\eta'}$.
Hence, let $\mathcal{N}_\eta$ be independent Poisson processes. For each $t \geq 0$:

$$X(t) = X(0) + \sum_{\eta \in \mathcal{T}} \mathbf{v}_\eta \mathcal{N}_\eta\left(\int_0^t r_\eta(X(s))ds\right).$$

Equivalently, let $\mathcal{Y}_\eta$ be independent Poisson r.v. It holds:

$$X(t) = X(0) + \sum_{\eta \in \mathcal{T}} \mathbf{v}_\eta \mathcal{Y}_\eta\left(\int_0^t r_\eta(X(s))ds\right).$$

Outline

1. Preliminaries
   - Exponential Distribution
2. Continuous Time Markov Chains
   - Main concepts
   - Poisson Process
   - Time-inhomogeneous rates
3. Population Continuous Time Markov Chains
4. Simulation
   - SSA
   - Next Reaction Method
   - τ-leaping

SIMULATING A POPULATION CTMC

Population CTMC have generally a complex dynamics and state space which is too large for

1. Solving the CTMC analytically
2. Performing numerical computations like solution of the Kolmogorov equation, transient analysis by uniformization, or computation of steady state.

Therefore, one can resort to statistical tools.
One samples a (large) set of trajectories from the distribution induced by the CTMC in the space of traces (cadlag functions), and then uses statistical methods to extract information about the process from these samples.
We will review some simulation algorithms, exploiting the different characterizations of (population) CTMCs.

DIRECT METHOD

RACE CONDITION CHARACTERIZATION OF A PCTMC

In each state $\mathbf{x}$, the m transitions in $\mathcal{T}$ compete in a race condition: the fastest wins and is executed.

DIRECT METHOD

At each step, with current state $\mathbf{x}$ and current time t

1. sample m uniform r.v. U_η;
2. compute $T_\eta = -\frac{1}{r_\eta(\mathbf{x})} \log(U_\eta)$;
3. find $\bar{\eta} = \mathrm{argmin}_{\eta \in \mathcal{T}} T_\eta$;
4. execute transition $\bar{\eta}$ updating the current state from $\mathbf{x}$ to $\mathbf{x} + \mathbf{v}_\eta$ and current time to $t + T_\eta$.

Outline

STOCHASTIC SIMULATION ALGORITHM

JUMP CHAIN AND HOLDING TIMES

We can improve the previous simulation by using the characterization with Jump Chain and Holding Times, which for population CTMC reads:

HOLDING TIME $r(\mathbf{x}) = \sum_{\eta \in \mathcal{T}} r_\eta(\mathbf{x})$

JUMP CHAIN $P(\eta \mid \mathbf{x}) = \frac{r_\eta(\mathbf{x})}{r(\mathbf{x})}$

SSA

At each step, with current state $\mathbf{x}$ and current time t

1. sample the next transition η from the jump chain;
2. sample the holding time from an $Exp(r(\mathbf{x}))$;
3. update current state and current time.

This method in biochemistry and system biology is also known as Gillespie Algorithm.

EXAMPLE: SIR EPIDEMICS

$N = 10$, $k_I = 1$, $k_R = 0.05$, $k_S = 0.01$
$X_S(0) = 8$, $X_I(0) = 2$, $X_R(0) = 0$.

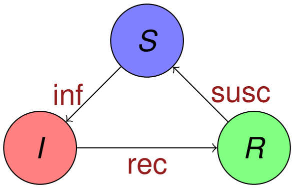

STEP 0: RATES OF TRANSITIONS

INFECTION: $\frac{1}{10} \cdot 8 \cdot 2 = 1.6$

RECOVERY: $0.05 \cdot 2 = 0.1$

IMMUNITY LOSS: 0

EXAMPLE: SIR EPIDEMICS

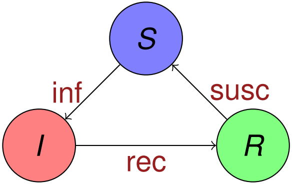

$N = 10$, $k_I = 1$, $k_R = 0.05$, $k_S = 0.01$
$X_S(0) = 8$, $X_I(0) = 2$, $X_R(0) = 0$.

STEP 0: RATES OF TRANSITIONS

INFECTION: $\frac{1}{10} \cdot 8 \cdot 2 = 1.6$

RECOVERY: $0.05 \cdot 2 = 0.1$

IMMUNITY LOSS: 0

TIME DELAY

Exponential with rate $1.6 + 0.1 = 1.7$.

NEXT STATE

- $X_S(0) = 7$, $X_I(0) = 3$, $X_R(0) = 0$ with prob. $\frac{1.6}{1.6+0.1} = 0.9412$
- $X_S(0) = 8$, $X_I(0) = 1$, $X_R(0) = 1$ with prob. $\frac{1.6}{1.6+0.1} = 0.0588$

OUTLINE

NEXT REACTION METHOD/GIBSON-BRUCK (SKETCH)

- Consider a single η transition in a time interval $[0, t]$ in which it never fires.
- As other transitions may fire, its rate $r_\eta(\mathbf{X}(s))$ is a time-dependent function.
- Therefore, we can sample the firing time of η using the inversion method for time-inhomogeneous exponential distribution, solving for t

$$\Lambda_\eta(t) = \int_0^t r_\eta(\mathbf{X}(s))ds = \xi \sim Exp(1).$$

NEXT REACTION METHOD/GIBSON-BRUCK (SKETCH)

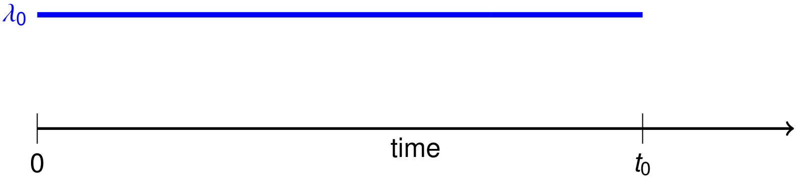

- Start at time 0, and suppose the rate of η is λ_0. Assuming it does not change in time, the firing time would be $t_0 = \frac{1}{\lambda_0}\xi \sim Exp(\lambda_0)$.

NEXT REACTION METHOD/GIBSON-BRUCK (SKETCH)

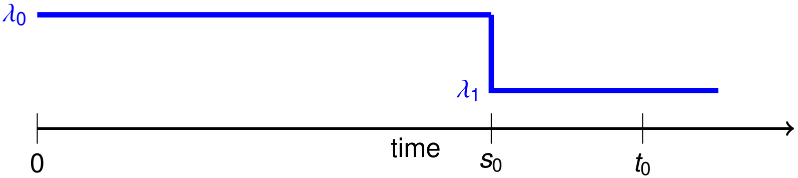

- Start at time 0, and suppose the rate of η is λ_0. Assuming it does not change in time, the firing time would be $t_0 = \frac{1}{\lambda_0}\xi \sim Exp(\lambda_0)$.
- Now, suppose at time s_0 another event η' fires, and this changes the rate of η to λ_1.

NEXT REACTION METHOD/GIBSON-BRUCK (SKETCH)

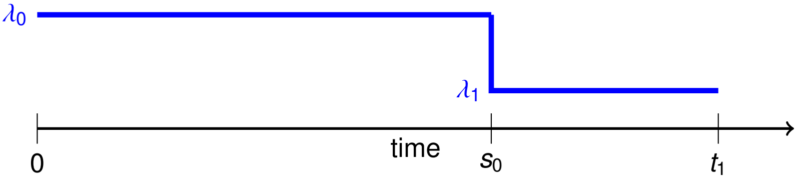

- Start at time 0, and suppose the rate of η is λ_0. Assuming it does not change in time, the firing time would be $t_0 = \frac{1}{\lambda_0}\xi \sim Exp(\lambda_0)$.
- Now, suppose at time s_0 another event η' fires, and this changes the rate of η to λ_1.
- Then the firing time of η would be found by solving $\lambda_0 s_0 + \lambda_1(t_1 - s_0) = \xi$, from which

$$t_1 = s_0 + \frac{\lambda_0}{\lambda_1}\left(\frac{1}{\lambda_0}\xi - s_0\right) = s_0 + \frac{\lambda_0}{\lambda_1}(t_0 - s_0).$$

NEXT REACTION METHOD/GIBSON-BRUCK (SKETCH)

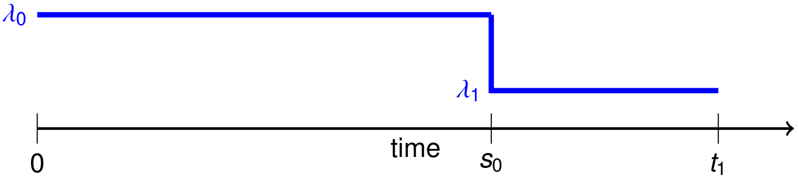

- Start at time 0, and suppose the rate of η is λ_0. Assuming it does not change in time, the firing time would be $t_0 = \frac{1}{\lambda_0}\xi \sim Exp(\lambda_0)$.
- Now, suppose at time s_0 another event η' fires, and this changes the rate of η to λ_1.
- Then the firing time of η would be found by solving $\lambda_0 s_0 + \lambda_1(t_1 - s_0) = \xi$, from which

$$t_1 = s_0 + \frac{\lambda_0}{\lambda_1}\left(\frac{1}{\lambda_0}\xi - s_0\right) = s_0 + \frac{\lambda_0}{\lambda_1}(t_0 - s_0).$$

- This is the update formula of Gibson-Bruck algorithm (can be easily generalized to n intermediate events by induction).

NEXT REACTION METHOD/GIBSON-BRUCK (SKETCH)

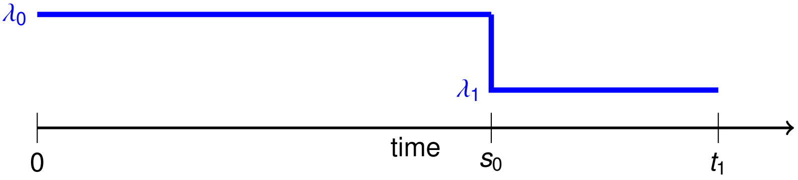

NEXT REACTION METHOD

At each step, with current state **x** and current time t

1. execute transition η with smallest time;
2. update rates and firing times of other transitions;
3. sample a new firing time for η.

the algorithm uses a priority queue and a dependency graph to speed up operations.

EXAMPLE: SIR EPIDEMICS

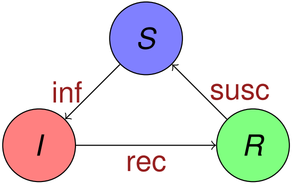

$N = 10$, $k_I = 1$, $k_R = 0.05$, $k_S = 0.01$
$X_S(0) = 8$, $X_I(0) = 2$, $X_R(0) = 0$.

STEP 1: RATES OF TRANSITIONS

INFECTION: $\frac{1}{10} \cdot 8 \cdot 2 = 1.6$

RECOVERY: $0.05 \cdot 2 = 0.1$

IMMUNITY LOSS: 0

STEP 2: COMPUTE FIRING TIMES

INFECTION: $\frac{1}{1.6} \cdot 0.2228 = 0.1392$

RECOVERY: $\frac{1}{0.1} \cdot 1.9527 = 19.5273$

IMMUNITY LOSS: $\frac{1}{0} \cdot 0 = \infty$

EXAMPLE: SIR EPIDEMICS

$N = 10$, $k_I = 1$, $k_R = 0.05$, $k_S = 0.01$
$X_S(0.1392) = 7$, $X_I(0.1392) = 3$,
$X_R(0.1392) = 0$.

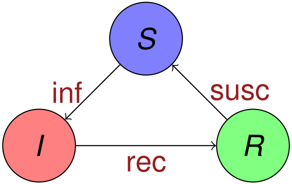

STEP 1: RATES OF TRANSITIONS

INFECTION: $\frac{1}{10} \cdot 7 \cdot 3 = 2.1$

RECOVERY: $0.05 \cdot 3 = 0.15$

IMMUNITY LOSS: 0

STEP 2: REEVALUATE FIRING TIMES

INFECTION: $\frac{1}{2.1} \cdot 3.3323 = 1.5868$

RECOVERY: $0.1392 + \frac{0.1}{0.15} \cdot (19.5273 - 0.1392) = 13.0646$

IMMUNITY LOSS: ∞

Outline

τ-LEAPING (SKETCH)

Consider the Poisson representation of a population CTMC at time τ

$$X(\tau) = X(0) + \sum_{\eta \in \mathcal{T}} \mathbf{v}_\eta \mathcal{Y}_\eta \left(\int_0^\tau r_\eta(X(s)) ds \right).$$

If τ is sufficiently small, we may assume that the rates $r_\eta(X(s))$ are approximately constant in $[0, \tau]$ and equal to a_η.
Then $\int_0^t r_\eta(X(s)) ds \approx a_\eta \tau$, hence

$$X(\tau) \approx X(0) + \sum_{\eta \in \mathcal{T}} \mathbf{v}_\eta \mathcal{Y}_\eta \left(a_\eta \tau \right).$$

τ-LEAPING (SKETCH)

τ-LEAPING

At each step, with current state $\mathbf{x}$ and current time t

1. choose τ;
2. for each η, sample n_η from the Poisson r.v. $\mathcal{Y}_\eta\left(a_\eta \tau\right)$;
3. update $\mathbf{x}$ to $\mathbf{x} + \sum_\eta \mathbf{v}_\eta n_\eta$ and time to $t + \tau$.

CHOICE OF τ: LEAPING CONDITION

The choice of τ is an art:

- it has to be small for rates to be approximately constant in $[t, t + \tau]$;
- it has to be as large as possible to make $\mathcal{Y}_\eta(a_\eta \tau)$ large to gain in computational efficiency;
- one has to avoid the generation of negative populations.

REFERENCES

- J.R. Norris. Markov Chains, Cambridge University Press, 1998.
- R. Durrett, Essentials of Stochastic Processes, Springer-Verlag, 1998.
- D.T. Gillespie (1976). A General Method for Numerically Simulating the Stochastic Time Evolution of Coupled Chemical Reactions. Journal of Computational Physics, 22(4): 403-434
- M.A. Gibson and J. Bruck (2000). Efficient Exact Stochastic Simulation of Chemical Systems with Many Species and Many Channels. Journal of Physical Chemistry A, 104(9).
- Y Cao, DT Gillespie, LR Petzold (2006). Efficient step size selection for the tau-leaping simulation method. J Chem Phys, 28;124(4).