Markov chains

An introduction

Consider a generic continuous-time stochastic process $(X_t)_{t\in\mathbb{R}_{\geq 0}}$, where for all $t \in \mathbb{R}_{\geq 0}$ the state X_t is a random variable that takes values x in the finite state space $\mathscr{X}$. We provide $\mathscr{X}$ with some ordering, such that any real-valued function f on $\mathscr{X}$ can be identified with a row vector. We furthermore let $\mathscr{L}(\mathscr{X})$ denote the set of all real-valued functions on $\mathscr{X}$. Then any linear operator $T\colon \mathscr{L}(\mathscr{X}) \to \mathscr{L}(\mathscr{X})$ can be identified with a matrix.

Precise Markov chains

The stochastic process $(X_t)_{t\in\mathbb{R}_{\geq 0}}$ is a *precise (continuous-time) Markov chain* (pMC) if it satisfies the *Markov property*: where $n \geq 0$ is an integer and $\{t_1, \dots, t_n, s, t\}$ is a strictly increasing sequence of non-negative time points. The *transition matrix* T_s^t thus defined satisfies

$$\begin{aligned}[T_s^t f](x_s) &= \mathrm{E}(f(X_t)|X_s = x_s) \\ &= \mathrm{E}(f(X_t)|X_{t_1} = x_1, \dots, X_{t_n} = x_n, X_s = x_s).\end{aligned} \tag{P1}$$

A pMC is called *stationary* if it satisfies $T_t^{t+\Delta} = T_0^\Delta =: T_\Delta$ for all $t, \Delta \in \mathbb{R}_{\geq 0}$. In this case, there is a unique *transition rate matrix* Q—a matrix with non-negative off-diagonal elements and rows that sum up to zero—such that

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_\Delta = T_t^{t+\Delta} \approx I + \Delta Q \quad \text{for } \Delta \text{ suff. small.}$$

Furthermore, T_t then satisfies the differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t} T_t = Q T_t, \qquad \text{with } T_0 = I. \tag{P2}$$

Similarly, for any non-stationary pMC there is a time-dependent transition rate matrix Q_t such that

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_t^{t+\Delta} \approx I + \Delta Q_t \quad \text{for } \Delta \text{ suff. small.}$$

Imprecise Markov chains

It is often infeasible to precisely specify the transition rate matrix Q of a stationary pMC. Furthermore, assuming stationarity is not always justified. Therefore, we here consider the case where the (time-dependent) transition rate matrix Q_t of a (non-stationary) pMC is only known to be contained in some (non-empty and bounded) set $\mathscr{Q}$. In other words, we consider the set $\mathbb{P}_{\mathscr{Q}}$ of all pMCs that are consistent with $\mathscr{Q}$, in the sense that

$$(\forall t \in \mathbb{R}_{\geq 0})(\exists Q_t \in \mathscr{Q})\; T_t^{t+\Delta} \approx I + \Delta Q_t \quad \text{for } \Delta \text{ suff. small.}$$

This set $\mathbb{P}_{\mathscr{Q}}$ characterises an *imprecise (continuous-time) Markov chain* (iMC) as follows. Analogous to (P1), we define a *lower transition operator* $\underline{T}_s^t$ as

$$\begin{aligned}[\underline{T}_s^t f](x_s) &:= \underline{\mathrm{E}}(f(X_t)|X_s = x_s) \\ &= \underline{\mathrm{E}}(f(X_t)|X_{t_1} = x_1, \dots, X_{t_n} = x_n, X_s = x_s),\end{aligned} \tag{I1}$$

where $\underline{\mathrm{E}}(\cdot|\cdot)$ is the minimum of the conditional expectations that are induced by the set of consistent processes.

In case $\mathscr{Q}$ has separately specified rows, Krak et al. (2017) show that $\underline{T}_t^{t+\Delta} = \underline{T}_0^\Delta =: \underline{T}_\Delta$ for all $t, \Delta \in \mathbb{R}_{\geq 0}$. Moreover, they show that $\underline{T}_\Delta$ is the unique operator that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t} \underline{T}_t = \underline{Q}\,\underline{T}_t, \qquad \text{with } \underline{T}_0 = I. \tag{I2}$$

In (I2), $\underline{Q}$ is the so-called *lower transition rate operator* of $\mathscr{Q}$, which, for any $f \in \mathscr{L}(\mathscr{X})$ and $x \in \mathscr{X}$, is defined as

$$[\underline{Q} f](x) := \min\{[Qf](x)\colon Q \in \mathscr{Q}\}. \tag{I3}$$

Ergodicity

We are often interested in the long-term limit behaviour of stationary pMCs and iMCs. For iMCs, a special case is when

$$\lim_{t\to+\infty} [\underline{T}_t f](x) = \underline{\mathrm{E}}_\infty(f) \quad \text{for all } f \in \mathscr{L}(\mathscr{X}) \text{ and } x \in \mathscr{X}.$$

If this is the case, then the iMC is said to be *ergodic* and $\underline{\mathrm{E}}_\infty(f)$ is called the *limit lower expectation*. Similarly, a stationary pMC is ergodic if

$$\lim_{t\to+\infty} [T_t f](x) = \mathrm{E}_\infty(f) \quad \text{for all } f \in \mathscr{L}(\mathscr{X}) \text{ and } x \in \mathscr{X},$$

where E_∞ is now called the limit expectation.

Modelling spectrum assignment in a two-service flexi-grid optical link

Cristina Rottondi IDSIA
Alexander Erreygers Ghent University
Giacomo Verticale Politecnico di Milano
Jasper De Bock Ghent University

Two-service flexi-grid optical link

Consider a single optical link with total spectrum availability S. We divide the spectrum S into m_1 frequency slices of width $F = S/m_1$ called *channels*. Arriving messages are assigned to a number of contiguous channels according to the bandwidth they require. We assume that the link is used to send two types of messages: type 1 messages require 1 channel and type 2 messages require n_2 channels.

In order to limit spectrum fragmentation, we assign arriving type 2 messages to one of the fixed *superchannels*, which are formed by combining n_2 contiguous channels:

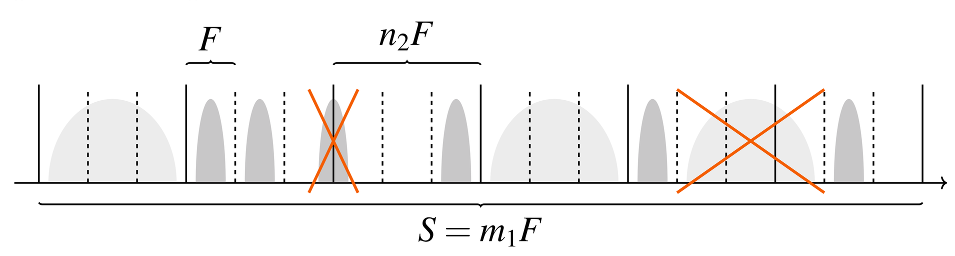

This way, we obtain $m_2 := m_1/n_2$ superchannels, where we assume that m_1 is an integer multiple of n_2.

We model the arrival of type 1 messages as a Poisson process with rate λ_1, and the arrival of type 2 messages as a Poisson process with rate λ_2. Furthermore, we assume that the sojourn time of a message is exponentially distributed, with rate μ_1 for the departure of a type 1 message and rate μ_2 for the departure of a type 2 message.

Spectrum assignment

Arriving type 1 message

If there is at least one free channel, the message is assigned to one of the free channels according to an *assignment policy*. We consider three such policies:

- **R** *randomly* select one of the free channels;
- **L** assign the arriving message to one of the free channels in the *least-filled* superchannels;
- **M** assign the arriving message to one of the free channels in the *most-filled* superchannels.

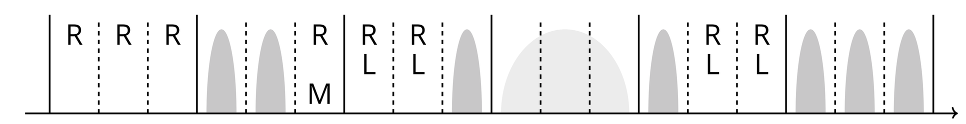

If there is no free channel, then the message is *blocked*.

Arriving type 2 message

If there is at least one free superchannel, the message is assigned to one of the free superchannels at random. If there is no free superchannel, then the message is *blocked*.

One measure that quantifies the performance of the assignment policies is the ***blocking probability*** *BP1 (BP2), which is the probability that a randomly selected type 1 (type 2) message is blocked.* We determine these blocking probabilities using **(im)precise (continuous-time) Markov chain** models, as introduced in **Markov chains: An introduction**.

Detailed state description

For our purposes, a sufficiently detailed state space is

$$\mathscr{X}_{\mathrm{det}} := \Big\{(i_0, \dots, i_{n_2}) \in \mathbb{N}^{(n_2+1)}\colon \sum_{k=0}^{n_2} i_k \leq m_2\Big\},$$

where, for $k \in \{0, \dots, n_2\}$, i_k counts the number of superchannels that are currently assigned k type 1 messages and not a type 2 message.

Using this state space, for each of the three policies we construct a stationary **precise Markov chain** (pMC) model of the optical link that *exactly* models the dynamics of the system—at least probabilistically speaking. The transition diagram of these pMCs is depicted on the right. The rates $\lambda_{\mathrm{P},k}$ of the transitions corresponding to the arrival of a type 1 message are policy- and state-dependent, and satisfy $\sum_{k=0}^{n_2-1} \lambda_{\mathrm{P},k} = \lambda_1$.

Unfortunately, as $|\mathscr{X}_{\mathrm{det}}| \sim \mathscr{O}((m_1/n_2)^{n_2})$, using this stationary pMC model to exactly determine the blocking probabilities of large systems (i.e., m_1 large and $n_2 > 2$) is infeasible.

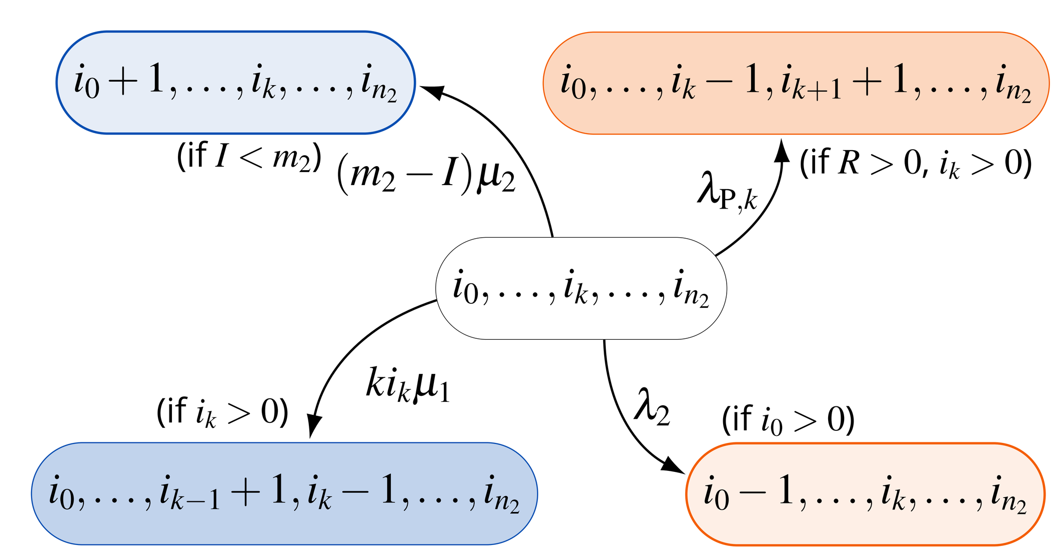

$I := \sum_{k=0}^{n_2} i_k$ denotes the number of superchannels not occupied by a type 2 message

$R := \sum_{k=0}^{n_2-1} i_k(n_2 - k)$ denotes the number of free channels

Reduced state description

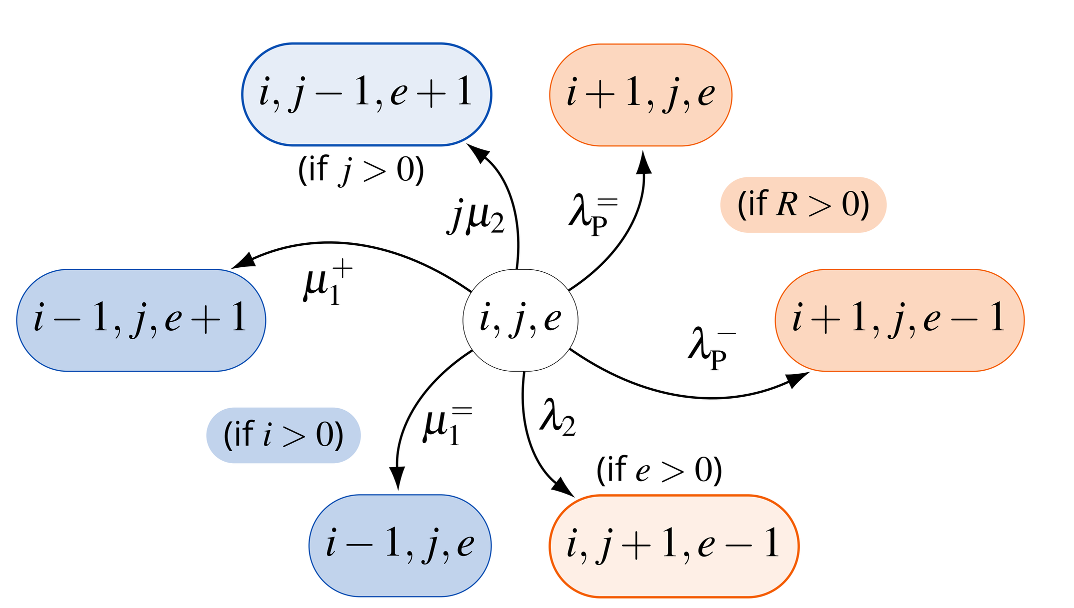

$R := m_1 - i - jn_2$ denotes the number of free channels

Kim et al. (2015) propose to use the reduced state space

$$\mathscr{X}_{\mathrm{red}} := \big\{(i, j, e) \in \mathbb{N}^3\colon m_2 \leq i + j + e, i + (j+e)n_2 \leq m_1\big\},$$

where i (j) counts the number of assigned type 1 (type 2) messages and e counts the number of empty superchannels. As $|\mathscr{X}_{\mathrm{red}}| \sim \mathscr{O}(m_1(m_1/n_2)^2)$, this reduced state description is better suited to model large systems (i.e., m_1 large and $n_2 > 2$).

By *lumping*—see **Handling state space explosion in Markov chains**—states in the exact pMCs, we obtain pMCs with state space $\mathscr{X}_{\mathrm{red}}$, the transition diagram of which is depicted on the left. The rates $\lambda_{\mathrm{P}}^{=}$ and λ_{P}^{-} of the transitions corresponding to the arrival of a type 1 message are state- and policy-dependent, and their sum equals λ_1. If $n_2 > 2$, the (state-dependent) rates μ_1^{+} and $\mu_1^{=}$ of the transitions corresponding to the departure of a type 1 message are time-dependent and indeterminable. Therefore, using these pMCs to determine the blocking probabilities is infeasible.

Kim et al. (2015) alleviate this problem by replacing the time and state-dependent transition rates μ_1^{+} and $\mu_1^{=}$ with constant state-dependent approximations $\tilde{\mu}_1^{+}$ and $\tilde{\mu}_1^{=}$, thus obtaining stationary but approximate pMCs. They then use these pMCs to *approximate* the blocking probabilities.

Instead of determining a single approximation without any sense of its accuracy, we opt for an approach that results in guaranteed lower and upper bounds. While precisely determining $\mu_1^{+}(i,j,e,t)$ and $\mu_1^{=}(i,j,e,t)$ is infeasible, we do know that

$$\mu_1^{+}(i,j,e,t) + \mu_1^{=}(i,j,e,t) = i\mu_1 \ \text{ and } \ i_{\min}(i,j,e)\mu_1 \leq \mu_1^{+}(i,j,e,t) \leq i_{\max}(i,j,e)\mu_1, \tag{1}$$

where $i_{\min}(i,j,e)$ ($i_{\max}(i,j,e)$) denotes the minimum (maximum) number of type 1 messages that are alone in their superchannel. Instead of considering a single transition rate matrix that, for every state, fixes values for $\mu_1^{+}(i,j,e,t)$ and $\mu_1^{=}(i,j,e,t)$ that satisfy (1), we consider the set of all transition rate matrices that satisfy (1). This way, we characterise policy-dependent **imprecise Markov chains** (iMCs). For every policy, we use this iMC to obtain *guaranteed lower and upper bounds* for the blocking probabilities. More generally, we also characterise a policy-independent iMC. Using this iMC, we can provide *policy-free bounds* for the blocking probabilities.

Numerical example

Below, we depict (the bounds on) the blocking probabilities for a system with $m_1 = 40$, $n_2 = 4$, $\mu_1 = \mu_2 = 1$ and $\lambda_1 = \lambda_2 = \lambda$.