

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 T . We divide the spectrum T into m_1 frequency slices of width $F = T/m_1$ 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**.

Detailed state description

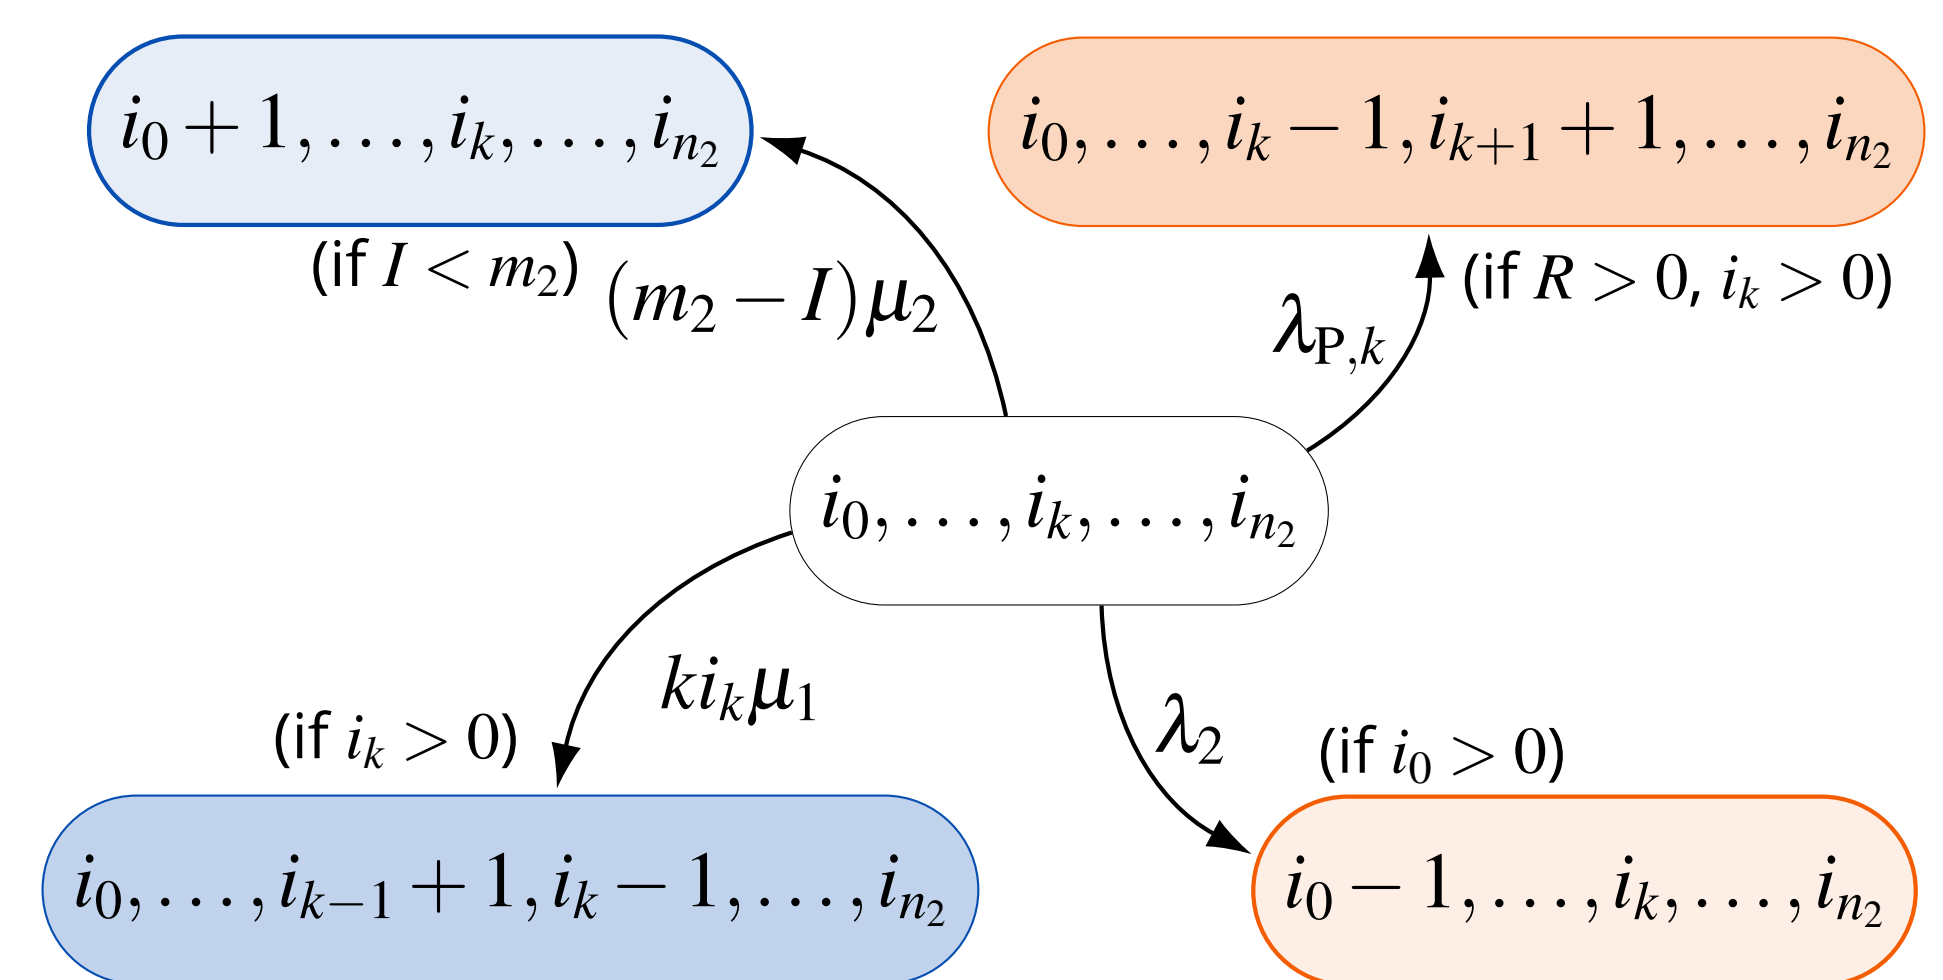
For our purposes, a sufficiently detailed state space is

$$\mathcal{X}_{\text{det}} := \left\{ (i_0, \dots, i_{n_2}) \in \mathbb{N}^{(n_2+1)} : \sum_{k=0}^{n_2} i_k \leq m_2 \right\},$$

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 $\lambda_{p,k}$ of the transitions corresponding to the **arrival of a type 1 message** are policy and state-dependent, and satisfy $\sum_{k=0}^{n_2-1} \lambda_{p,k} = \lambda_1$.

Unfortunately, as $|\mathcal{X}_{\text{det}}| \sim \mathcal{O}((m_1/n_2)^{n_2})$, 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

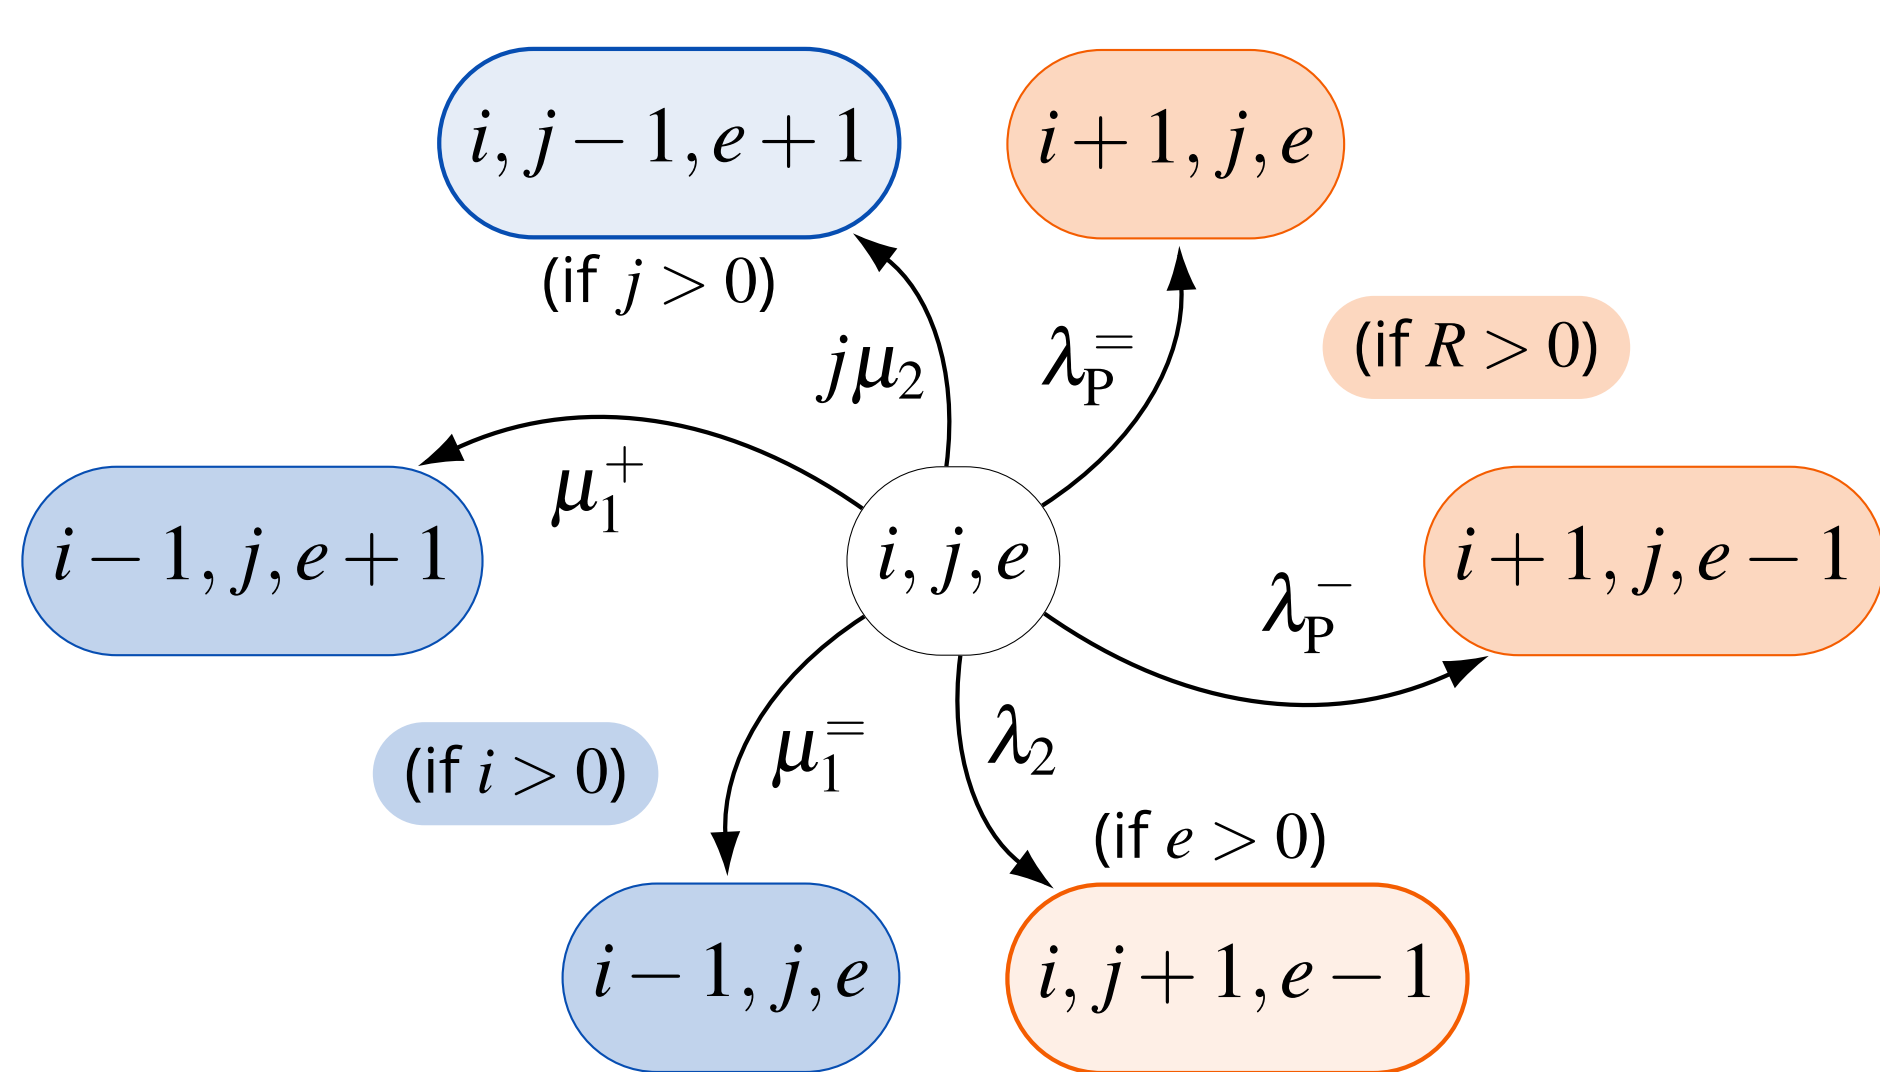
Reduced state description

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

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

where i (j) counts the number of assigned type 1 (type 2) messages and e counts the number of empty superchannels. As $|\mathcal{X}_{\text{red}}| \sim \mathcal{O}(m_1(m_1/n_2)^2)$, 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 \mathcal{X}_{red} , the transition diagram of which is depicted on the left. The rates λ_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 μ_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.



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

Kim et al. (2015) alleviate this problem by replacing the time and state-dependent transition rates μ_1^+ and μ_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

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

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.

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**, which is the probability that a randomly selected type 1 (or 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**.

$I := \sum_{k=0}^{n_2} i_k$ 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

Numerical example

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

