

Imprecise Markov Models for Scalable and Robust Performance Evaluation of Flexi-Grid Spectrum Allocation Policies

Alexander Erreygers, Cristina Rottondi, Giacomo Verticale and Jasper De Bock

Abstract—The possibility of flexibly assigning spectrum resources with channels of different sizes greatly improves the spectral efficiency of optical networks, but can also lead to unwanted spectrum fragmentation. We study this problem in a scenario where traffic demands are categorised in two types (low or high bit-rate) by assessing the performance of three distinct allocation policies. Our first contribution consists of exact Markov chain models for the allocation policies, which allow us to numerically compute the relevant performance measures. However, these exact models do not scale to large systems, in the sense that the computations required to determine the blocking probabilities—which measure the performance of the allocation policies—become intractable. In order to address this, we first extend an approximate reduced-state Markov chain model that is available in the literature to the three considered allocation policies. These reduced Markov chain models allow us to tractably compute approximations of the blocking probabilities, but the accuracy of these approximations cannot be easily verified. Our main contribution then is the introduction of reduced-state imprecise Markov chain models that allow us to derive guaranteed lower and upper bounds on blocking probabilities, for the three allocation policies separately or for all possible allocation policies.

Index Terms—Two-Service Optical Link; Blocking Probability; Imprecise Continuous-Time Markov Chain;

I. INTRODUCTION

Flexi-grid optical networks [1] have been envisioned as a novel paradigm to cope with the ever-growing Internet traffic: spectral resources are divided in small frequency slices (e.g., 12.5 GHz width, according to the ITU-T standard [2]) and groups of contiguous slices are adaptively assigned to different traffic requests—forming the so-called *superchannels*—according to their volume, the optical bandwidth of the transceivers in use and the adopted modulation format. Adjacent superchannels are separated by guardbands, constituted by one or multiple contiguous slices that are left unused. The advantages of flexi-grid networks have been quantified in terms of spectrum utilization reductions up to 30% with respect to traditional Wavelength Division Multiplexing (WDM)

systems [3]. However, the flexi-grid approach requires more advanced and costly optical devices such as ROADMs with colourless/directionless/contentionless add/drop functionalities equipped with dedicated tunable filters supporting coherent detection [4]. Moreover, the flexible spectrum allocation techniques enabled by flexi-grid networks typically induce spectrum fragmentation, as often groups of contiguous slices are formed that cannot be assigned to incoming traffic flows because the superchannel they would form is too narrow. This issue is further increased by the spectrum continuity constraint, which requires that the same spectrum portion is allocated to a traffic flow along all the physical links it traverses. Therefore, several studies have identified fragmentation-aware spectrum allocation policies and assessed their performance by determining the blocking probabilities for traffic requests of different sizes [5]. Analytical models based on Markov Chains (MCs) have been proposed to determine these blocking probabilities [6]. However, unfortunately, for realistic scenarios these models require unaffordable computational resources due to the high number of states that are required to correctly capture the degrees of freedom offered by the flexible grid, thus introducing scalability limitations.

In order to alleviate the spectrum fragmentation issue and limit the costs of equipment installation without renouncing the benefits of flexi-grid networks in terms of spectrum occupation reduction, alternative semi-flexible approaches have been proposed. One approach is to group traffic requests according to the number of slices required for transmission, and to place requests belonging to the same group along a dedicated fixed grid, with one edge of their superchannel anchored at a specific frequency [7]. Alternatively, a small set of predefined superchannel widths is defined. Traffic flows are then allocated in the smallest superchannel they fit in, at the price of leaving some spectrum slices unused [8]. Such scenarios allow for more scalable MC models, as for example is the case in [9] and [10]. More specifically, Kim et al. [9, 10] use an approximate MC to obtain blocking probabilities in a two-service semi-flexible optical link with a random spectrum allocation policy under the assumption that traffic demands are categorised in two types according to their bit-rate—high or low, respectively.

In this contribution, we first introduce exact MC models for three spectrum allocation policies. These models have a state space that grows exponentially with the dimensions of the system, making their use infeasible for realistic scenarios. Therefore, we also provide reduced-state MC models for the

A. Erreygers is a member of IDLab, Ghent University, 9052 Ghent, Belgium (email: alexander.erreygers@ugent.be).

C. Rottondi is with the Dalle Molle Institute for Artificial Intelligence, University of Lugano and University of Applied Science and Arts of Southern Switzerland, 6928 Manno, Switzerland (email: cristina.rotttondi@supsi.ch).

G. Verticale is with the Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, 20133 Milan, Italy (email: giacomo.verticale@polimi.it).

J. De Bock is a member of IDLab, Ghent University, 9052 Ghent, Belgium (email: jasper.debock@ugent.be).

same allocation policies. The downside of reducing the number of states is that the rates of some of the transitions cannot be precisely determined any more; instead, one can only determine lower and upper bounds. We compare two approaches to dealing with this indeterminacy. The first approach simply uses a precise but approximate value for each of the indeterminate transition rates. This approach was used in Kim et al. [9] for one of the policies, and we construct similar models for the two other allocation policies we consider. For the second approach, we make use of imprecise MC models to implicitly take into account the partially specified character of the transition rates.

Similar approaches were already put forward by us in [11], where imprecise MC models were applied to optical networks for the first time. The present contribution—an extended version of [12], a paper that is currently under peer-review in IEEE Transactions on Communications—extends [11]. Whereas in [11] we only considered a policy-independent imprecise MC model, we here also construct an imprecise MC model for each of the three policies. Furthermore, we provide a more thorough explanation on how to compute lower and upper bounds on the blocking probabilities. Finally, we illustrate the benefits of our imprecise MC models over the exact or approximate MC models by means of extensive numerical experiments.

The remainder of this contribution is organised as follows. We start in Section II by presenting an overview of the related scientific literature. Next, we provide some basic background on imprecise probabilities and imprecise continuous-time MCs in Section III and introduce the two-service optical link scenario under study in Section IV. We then go on in Section V to construct two precise MC models—an exact and an approximate one—for each of the three considered allocation policies. Section VI presents our main contribution, which is to introduce four imprecise MC models—one for each of the three allocation policies and a common policy-independent one. Finally, Section VII numerically compares the scalability of our (exact and approximate) precise MC models and our (policy-dependent and policy-independent) imprecise MC models. We conclude this contribution in Section VIII. For the sake of clarity, all proofs have been relegated to the appendix.

II. RELATED WORK

A rich body of research work on flexi-grid optical networks has appeared in the last few years. The reader is referred to [13] and [14] for a thorough overview. More specifically, both the static and the dynamic Routing and Spectrum Assignment (RSA) problem in flexi-grid networks has been extensively addressed; for a survey of the most relevant literature we refer to [15]. In the present contribution, we focus on a single-link dynamic scenario. The arrival and departure processes of traffic requests are random, and requests are either accommodated in a spectrum portion of the link or rejected in real-time.

Different semi-flexible approaches have been proposed to mitigate the issue of spectrum fragmentation. Comellas and Junyent [8] propose to partition the flexible grid in blocks

of a fixed number of slices and to assign one block to each traffic request, regardless of its size, possibly leaving some unused slices within the block. If a single block is not sufficient to accommodate the whole traffic demand, multiple contiguous blocks are assigned to it. Alternatively, Shen et al. [16] propose to reserve a dedicated spectrum portion to high bit-rate signals, whereas in [7] each specific bit-rate signal uses its own dedicated fixed grid and starts from predefined anchor frequencies. In this contribution, we adopt the latter approach.

Several authors have investigated the unfairness of blocking probabilities in semi-flexible optical networks in the two-service scenario. Kim et al. [9] propose an approximate MC model for the calculation of blocking probabilities of the two types of traffic requests along a single optical link, assuming that spectrum allocation is performed randomly. This model has been refined in [10, 17, 18] to include spectrum portions exclusively reserved for each of the two connection types. In this contribution, we generalise the original model by means of an imprecise MC, which is an MC that allows for imprecise—partially specified—transition rates. In this way, we can provide guaranteed lower and upper bounds on the performance of a wide class of spectrum allocation policies. As exemplificative cases, we consider three such cases: the random, least-filled and most-filled allocation policies. Furthermore, we also use our partially specified transition rates to provide guaranteed lower and upper bounds on the performance of *any* spectrum allocation policy.

A general introduction to the theory of imprecise probabilities can be found in [19]. This theory has been used to generalise both discrete-time and continuous-time MCs, the latter more recently than the former. A basic treatment of imprecise continuous-time MCs can be found in [20–24]. For their discrete-time counterpart, see for example [25, 26].

III. BACKGROUND

A. Imprecise Probabilities

Whenever it is impossible or impractical to provide, obtain or compute exact values for the probability of some event, or the expectation of some function, the theory of imprecise probabilities allows for these quantities to be described ‘imprecisely’, using lower and upper bounds. For a detailed treatment of this theory, the interested reader is referred to [27] and [19]. For our present purposes, and in the context of stochastic processes, it suffices to understand the following basic concepts.

In the remainder of this paper, we will use $(X_t)_{t \in \mathbb{R}_{\geq 0}}$ ¹ to denote a generic continuous-time stochastic process, where for all $t \in \mathbb{R}_{\geq 0}$ the state X_t is a random variable that takes values x in a finite state space \mathcal{X} . Moreover, we let \mathbb{I}_A be the indicator $\mathbb{I}_A: \mathcal{X} \rightarrow \{0, 1\}$ of the event $A \subseteq \mathcal{X}$, defined by $\mathbb{I}_A(x) := 1$ if $x \in A$ and $\mathbb{I}_A(x) := 0$ otherwise.

For our present purposes, the most important notion is that of a conditional lower expectation, which can be interpreted as

¹ \mathbb{R} denotes the set of real numbers, while the set of all non-negative real numbers is denoted by $\mathbb{R}_{\geq 0}$. Furthermore, we use \mathbb{N} to denote the set of all natural numbers (including zero), and $\mathbb{N}_{>0}$ to denote the set of all strictly positive natural numbers.

a lower bound on a conditional expectation. More formally, for any $s, t \in \mathbb{R}_{\geq 0}$ such that $t \geq s$, any $x_s \in \mathcal{X}$ and any function $f: \mathcal{X} \rightarrow \mathbb{R}$, the lower expectation of f at time t , conditional on $X_s = x_s$, is

$$\underline{E}(f(X_t)|X_s = x_s) := \inf_{\mathbb{E} \in \mathcal{E}} \mathbb{E}(f(X_t)|X_s = x_s),$$

where \mathcal{E} is the set of conditional expectations that corresponds to some set of stochastic processes \mathbb{P} —see Section III-B for a concrete example of such a set of processes \mathbb{P} .

The reason why we can focus on lower expectations, and not on upper expectations or lower/upper probabilities, is because the latter can all be obtained as special cases. On the one hand, upper expectations are conjugate to lower expectations, in the sense that

$$\bar{E}(f(X_t)|X_s = x_s) = -\underline{E}(-f(X_t)|X_s = x_s). \quad (1)$$

On the other hand, lower and upper probabilities are special cases of lower and upper expectations, in the sense that for any event $A \subseteq \mathcal{X}$:

$$\underline{P}(X_t \in A|X_s = x_s) = \underline{E}(\mathbb{I}_A(X_t)|X_s = x_s) \quad (2)$$

and

$$\bar{P}(X_t \in A|X_s = x_s) = \bar{E}(\mathbb{I}_A(X_t)|X_s = x_s). \quad (3)$$

Lower and upper probabilities are also conjugate, in the sense that

$$\bar{P}(X_t \in A|X_s = x_s) = 1 - \underline{P}(X_t \notin A|X_s = x_s).$$

B. Imprecise Continuous-Time Markov Chains

Recently, several authors have used the theory of imprecise probabilities to develop the notion of an imprecise continuous-time Markov chain, which we will here abbreviate as ‘imprecise MC’ [20, 21, 23, 24, 28].

The starting point are traditional precise (homogeneous) MCs, which are characterised by a transition rate matrix Q (a real-valued square matrix whose off-diagonal elements are non-negative and whose rows sum to zero). For any $x, y \in \mathcal{X}$ and any $s, t \in \mathbb{R}_{\geq 0}$ such that $s \leq t$, the (x, y) -component of the transition matrix T_s^t is then

$$T_s^t(x, y) := P(X_t = y|X_s = x) = [e^{(t-s)Q}](x, y),$$

where $e^{(t-s)Q}$ is the matrix exponential of $(t-s)Q$, defined as

$$e^{(t-s)Q} := \lim_{n \rightarrow +\infty} \sum_{k=0}^n \frac{(t-s)^k}{k!} Q^k = \lim_{n \rightarrow +\infty} \left(I + \frac{t-s}{n} Q \right)^n. \quad (4)$$

Instead of a single precisely specified transition rate matrix Q , an imprecise MC now considers a (non-empty and bounded) set \mathcal{Q} of transition rate matrices. In practice, this is typically useful in cases where the values of the transition rates $Q(x, y)$ cannot be determined exactly, as is the case in Section V-B2.

More formally, instead of a single homogeneous MC, an imprecise MC considers the set $\mathbb{P}_{\mathcal{Q}}$ of all MCs that are consistent with \mathcal{Q} , in the sense that at every point in time $t \in \mathbb{R}_{\geq 0}$, and for $\Delta \in \mathbb{R}_{\geq 0}$ sufficiently small, the transition matrix $T_t^{t+\Delta}$ is approximately equal to $I + \Delta Q_t$, for some

$Q_t \in \mathcal{Q}$. Note that the Markov chains in this set $\mathbb{P}_{\mathcal{Q}}$ are not assumed to be homogeneous, in the sense that Q_t is not required to be constant. The only thing that is assumed about Q_t is that it is a—possibly unknown—function of time that takes values in \mathcal{Q} .

Since we consider a set of MCs, the transition matrices T_s^t are no longer uniquely known—as was the case for precise MCs. Instead, an imprecise MC is characterised by a lower transition operator \underline{T}_s^t . For any $f: \mathcal{X} \rightarrow \mathbb{R}$ and any $s, t \in \mathbb{R}_{\geq 0}$ such that $t \geq s$, its value $\underline{T}_s^t f$ is again a real-valued function on \mathcal{X} , defined by

$$[\underline{T}_s^t f](x_s) := \underline{E}(f(X_t)|X_s = x_s) \text{ for all } x_s \in \mathcal{X}, \quad (5)$$

where $\underline{E}(f(X_t)|X_s = x_s)$ is the infimum of the conditional expectations that are induced by the set of consistent processes $\mathbb{P}_{\mathcal{Q}}$. Of course, determining the set of all consistent processes $\mathbb{P}_{\mathcal{Q}}$ explicitly and then computing the infimum of the corresponding expectations is infeasible, if not impossible. Fortunately, this is not always necessary because the lower transition operator \underline{T}_s^t can often also be characterised by a non-linear version of the matrix exponential (4):

$$\underline{T}_s^t = \lim_{n \rightarrow +\infty} \left(I + \frac{t-s}{n} \underline{Q} \right)^n, \quad (6)$$

where \underline{Q} is the so-called *lower transition rate operator* of \mathcal{Q} , which maps any $f: \mathcal{X} \rightarrow \mathbb{R}$ to $\underline{Q}f: \mathcal{X} \rightarrow \mathbb{R}$, defined by

$$[\underline{Q}f](x) := \inf \{ [Qf](x) : Q \in \mathcal{Q} \} \text{ for all } x \in \mathcal{X}. \quad (7)$$

As proved in [24], a sufficient condition for this to be possible is that \mathcal{Q} has separately specified rows, which more or less means that for every $f: \mathcal{X} \rightarrow \mathbb{R}$, there is some $Q \in \mathcal{Q}$ such that $[\underline{Q}f](x) = [Qf](x)$ for all $x \in \mathcal{X}$; see [21, Definition 7.3] for a formal definition. A lower transition rate operator \underline{Q} also has a norm, which is equal to

$$\|\underline{Q}\| := \max \{ \|[Q\mathbb{I}_x](x)\| : x \in \mathcal{X} \}.$$

That said, in this contribution, we will be solely interested in providing lower and upper bounds on the value of $P(X_t \in A|X_0 = x_0)$ as t approaches infinity, that is, we are interested in $\lim_{t \rightarrow +\infty} \underline{P}(X_t \in A|X_0 = x_0)$ and $\lim_{t \rightarrow +\infty} \bar{P}(X_t \in A|X_0 = x_0)$. As shown in [24], these limits always exist. If they furthermore do not depend on $x_0 \in \mathcal{X}$, then we call them the *lower and upper limit probability* of A , and denote them by $\underline{\pi}_A$ and $\bar{\pi}_A$, respectively. A sufficient condition for the existence of $\underline{\pi}_A$ and $\bar{\pi}_A$ is the ergodicity of the lower transition rate operator \underline{Q} , in the sense that for any real-valued function f on \mathcal{X} and any $s \in \mathbb{R}_{\geq 0}$, $\lim_{t \rightarrow +\infty} [\underline{T}_s^t f](x_s)$ does not depend on the initial state $x_s \in \mathcal{X}$ [24]. A simple necessary and sufficient condition for ergodicity is provided in [24, Theorem 19].

Provided that \mathcal{Q} has separately specified rows and \underline{Q} is ergodic, it follows from (1)–(3), (5) and (6) that

$$\underline{\pi}_A = \lim_{t \rightarrow +\infty} \lim_{n \rightarrow +\infty} \left[\left(I + \frac{t}{n} \underline{Q} \right)^n \mathbb{I}_A \right] (x_0) \quad (8)$$

and

$$\bar{\pi}_A = - \lim_{t \rightarrow +\infty} \lim_{n \rightarrow +\infty} \left[\left(I + \frac{t}{n} \underline{Q} \right)^n (-\mathbb{I}_A) \right] (x_0) \quad (9)$$

Algorithm 1: Approximating the lower limit probability

Data: An ergodic lower transition rate operator \underline{Q} , an event $A \subseteq \mathcal{X}$, a time step $\delta \in (0, 2/\|\underline{Q}\|)$, a relative tolerance $\phi \in \mathbb{R}_{>0}$ and a maximum number of iterations $N \in \mathbb{N}$.

Result: An approximation $\hat{\pi}_A = \underline{\pi}_A \pm \epsilon'$

```

1  $i \leftarrow 0$ 
2  $g_0 \leftarrow \mathbb{I}_A$ 
3 while  $\|g_i\|_v > \phi |\text{mid}(g_i)|$  and  $i < N$  do
4    $i \leftarrow i + 1$ 
5    $g_i \leftarrow g_{i-1} + \delta \underline{Q} g_{i-1}$ 
6  $\hat{\pi}_A \leftarrow \text{mid}(g_i)$ 
7 return  $\hat{\pi}_A$ 

```

for any $x_0 \in \mathcal{X}$, which allows us to compute them both by choosing t and n sufficiently large. Unlike for the case of precise MCs, where literature offers a plethora of methods to compute limit probabilities, this is—to the best of our knowledge—currently the only method to compute lower or upper limit probabilities for imprecise MCs.

C. Numerical Approximation Method

Assume we have some ergodic lower transition rate operator \underline{Q} and some event $A \subseteq \mathcal{X}$. We can then use Alg. 1—a slightly modified version of [21, Algorithm 1]—to compute the lower limit probability $\underline{\pi}_A$ of the event $A \subseteq \mathcal{X}$. This algorithm uses the variation norm, defined for all $f: \mathcal{X} \rightarrow \mathbb{R}$ as $\|f\|_v := (\max f - \min f)/2$, and the midpoint, defined as $\text{mid}(f) := (\max f + \min f)/2$. Note that to compute $\underline{\pi}_A$, we simply need to change line 2 in Alg. 1 to “ $g_0 \leftarrow -\mathbb{I}_A$ ” and line 6 to “ $\hat{\pi}_A \leftarrow \text{mid}(-g_n)$ ”.

It is worth noting that we need to specify a time step $\delta \in (0, 2/\|\underline{Q}\|)$, a maximal relative tolerance $\phi \in \mathbb{R}_{>0}$ and a maximum number of iterations $N \in \mathbb{N}$.

From [23, Proposition 11] and the theory of ergodic imprecise discrete-time Markov chains [25], it follows that if \underline{Q} is ergodic, then for any $\delta \in (0, 2/\|\underline{Q}\|)$, $\lim_{n \rightarrow +\infty} [(I + \delta \underline{Q})^n \mathbb{I}_A](x_0)$ is the same for all $x_0 \in \mathcal{X}$. However, this limit value is not necessarily the same for all allowable values of δ . Empirically, we observe that for the lower transition rate operators we introduce in Sections VI-A and VI-B this limit value increases (i.e., the obtained lower bound is tighter) for decreasing δ . One way to empirically confirm that the lower bound $\hat{\pi}_A$ obtained using Alg. 1 is sufficiently tight is proposed—albeit in a slightly different form—in [22]: run Alg. 1 for the step sizes δ and $\delta/2$, and compare the obtained bounds $\hat{\pi}_{A,1}$ and $\hat{\pi}_{A,2}$. If $|\hat{\pi}_{A,1} - \hat{\pi}_{A,2}| < \phi \hat{\pi}_{A,2}$ —for example, for $\phi = 10^{-m}$, if $\hat{\pi}_{A,1}$ and $\hat{\pi}_{A,2}$ differ only after the m -th significant digit—then we can be relatively confident that the obtained bound is sufficiently tight.

If one wants a theoretically guaranteed instead of empirically verified measure of the accuracy of the bound, things get a little more intricate. To the best of our knowledge, existing methods are based on the absolute error $\epsilon_a := |\underline{\pi}_A - \text{mid } g_n|$. Let g_0, g_1, \dots, g_n be obtained by running Alg. 1.

Then by [23, Theorem 5 and Proposition 12] a theoretically guaranteed upper bound for the absolute error ϵ_a is $\epsilon' := \max \left\{ 2\delta^2 \|\underline{Q}\|^2 \sum_{i=0}^{n-1} \|g_i\|_v, \|g_n\|_v \right\}$. However, as illustrated in [23, Example 6], this upper bound can be overly conservative. Therefore, and because in the current setting it makes more sense to consider the relative error instead of the absolute error, we would here opt to use the empirical way of assessing the accuracy. Nevertheless, in Section VII-C3 we simply use $\delta = 2/(4\|\underline{Q}\|)$, which—although we do not explicitly mention it there—turns out to yield the required accuracy in all cases, except for two cases where the computations preemptively stop because the maximum number of iterations is reached.

IV. THE SINGLE OPTICAL LINK UNDER STUDY

We consider a single optical-fiber link with overall spectrum availability T . The spectrum is partitioned in slices of width F , for a total number of $S := T/F$ slices. For the sake of convenience, we assume that T is an integer multiple of F . Slices are sequentially numbered from 1 to S .

The optical link is used to transmit two types of flows with distinct bandwidth demands: flows of type 1 require $n_1 > 0$ slices, whereas flows of type 2 require $n_2 n_1 > n_1$ slices. Type 1 (respectively type 2) flows arrive according to a Poisson process with arrival rate λ_1 (respectively λ_2) and, if assigned to some free slices in the spectrum, cease after a holding time that is exponentially distributed and has average service time $1/\mu_1$ (respectively $1/\mu_2$). We are not interested in edge cases, so throughout this contribution we assume that $\lambda_1, \lambda_2, \mu_1$ and μ_2 are all known positive real numbers.

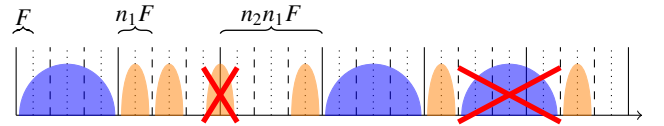


Fig. 1. Example of the (super)channel ordering in a single optical link with $n_1 = 2$, $n_2 = 3$ and $S = 36$.

As depicted in Fig. 1, we consider two overlapping grids of different granularity. The first consists of a sequence of adjacent channels of width $n_1 F$: the first channel comprises slices $1, \dots, n_1$, the second comprises slices $n_1 + 1, \dots, 2n_1$, and so on. Similarly, the second grid consists of a sequence of adjacent superchannels of width $n_2 n_1 F$. The number of channels of size $n_1 F$ is thus $m_1 := S/n_1$, whereas the number of superchannels of size $n_2 n_1 F$ is $m_2 := m_1/n_2$. For the sake of convenience, we assume that S is an integer multiple of $n_1 n_2$.

An incoming type 1 flow, in the remainder referred to as a *type 1 request*, must be assigned to a free channel of width $n_1 F$, whereas an incoming type 2 flow, in the remainder referred to as a *type 2 request*, must be assigned to a free superchannel of width $n_2 n_1 F$. If there is only a single free (super)channel, we always assign the request to this free (super)channel. If there are multiple free (super)channels, we assign the request to one of the free (super)channels according to some spectrum allocation policy. For type 2 requests, it will become evident in the remainder that the allocation policy is of

no importance to our analysis. For type 1 requests, we restrict ourselves to allocation policies that only depend on the number of type 1 flows currently occupying each superchannel and not on the specific ordering of the type 1 flows along the spectrum.

We consider three such allocation policies as exemplificative cases. The Random Allocation (RA) policy assigns a type 1 request to a randomly selected free channel if possible, where every free channel has the same probability of being selected [9, 10]. Alternatively, we can assign the type 1 request to one of the free channels in a partially occupied (i.e., non-empty and non-full) superchannel that contains either the lowest number of type 1 flows or the highest number of type 1 flows. The former is called the Least-Filled (LF) policy, while the latter is called the Most-Filled (MF) policy. For both of these policies, if all superchannels are empty or full, the type 1 request is assigned to one of the channels in an empty superchannel—provided they are not all full of course. Throughout the remainder, we use AP to denote a generic allocation policy.

If no (super)channel of the required size is available, the incoming traffic request is *blocked*. Determining the blocking probability of type 1 or type 2 traffic requests (or more specifically, the long-term limit of the fraction of incoming requests of type 1 or type 2 that are blocked), denoted by BP1 and BP2 respectively, is the main objective of this contribution.

V. PRECISE MARKOV CHAIN MODELS

A. Exact Model

1) *Detailed State Space Description*: To determine the blocking probabilities of type 1 and type 2 requests, we will construct policy-dependent MC models and determine the limit probabilities of the event that no type 1 (or type 2) flow can be allowed into the system.² The state description of these MC models should allow us to (i) determine whether or not an incoming flow is blocked or allowed into the system, and (ii) accurately model the allocation of the flows and the completion of their holding times. Due to the memorylessness of the exponential distribution and our requirement that the allocation policy for type 1 requests only depends on the number of type 1 flows currently occupying each superchannel, a sufficiently detailed state description is $(i_0, i_1, \dots, i_{n_2})$. In this tuple, i_k counts the number of superchannels occupied with k type 1 flows and no type 2 flows. Let I be defined as $I := \sum_{k=0}^{n_2} i_k$, then $m_2 - I$ is the number of superchannels that are occupied by a type 2 flow. To ensure feasibility, we require that $I \leq m_2$ and that $i_k \geq 0$ for all k in $\{0, \dots, n_2\}$, which yields the detailed state space $\mathcal{X}_{\text{det}} := \{(i_0, \dots, i_{n_2}) \in \mathbb{N}^{n_2+1} : \sum_{k=0}^{n_2} i_k \leq m_2\}$. The total number of states exhibits an $\mathcal{O}(m_2^{n_2}) = \mathcal{O}((m_1/n_2)^{n_2})$ dependency on the total number of channels m_1 and on the number of channels contained in a superchannel n_2 .

2) *Policy-Dependent Transition Rates*: Note that i_0 counts the number of empty superchannels, whereas $R := \sum_{k=0}^{n_2-1} i_k(n_2 - k)$ is the total number of free channels. As

²Note that although the limit probability has a different definition than the blocking probability, it is the customary way to determine the blocking probability as it follows from the well-known PASTA property and the ergodic theorem for (irreducible) Markov chains that these are in fact (almost surely) equal.

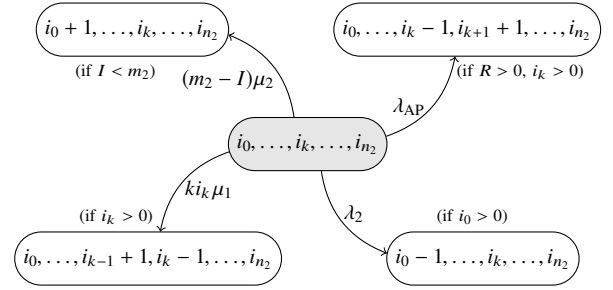


Fig. 2. State transition diagram of the proposed precise and exact Markov chain

depicted in Fig. 2, the following transitions can occur (i.e., have positive rate). If a new type 2 flow arrives and there is a free superchannel (i.e., $i_0 > 0$), the transition $(i_0, i_1, \dots, i_{n_2}) \rightarrow (i_0 - 1, i_1, \dots, i_{n_2})$ takes place. As type 2 flows arrive according to a Poisson process with rate λ_2 , this transition occurs with the same rate. Conversely, when the holding time of a type 2 flow expires, the following transition occurs: $(i_0, i_1, \dots, i_{n_2}) \rightarrow (i_0 + 1, i_1, \dots, i_{n_2})$. There are $(m_2 - I)$ allocated type 2 flows, and the holding time of each of these flows is (independently) exponentially distributed with rate μ_2 . Hence, it follows from the properties of the exponential distribution that this transition occurs with rate $(m_2 - I)\mu_2$. Similarly, the expiration of the holding time of a type 1 flow corresponds to the state transition $(i_0, \dots, i_k, \dots, i_{n_2}) \rightarrow (i_0, \dots, i_{k-1} + 1, i_k - 1, \dots, i_{n_2})$ which, for all k in $[1, n_2]$, has rate $ki_k \mu_1$.

The transition that corresponds to the assignment of a type 1 request, which can only occur if there is a free channel (i.e., $R > 0$), depends on the adopted spectrum allocation policy AP. If we adopt the RA policy, then the state transition $(i_0, \dots, i_k, \dots, i_{n_2}) \rightarrow (i_0, \dots, i_k - 1, i_{k+1} + 1, \dots, i_{n_2})$ occurs for all k in $[0, n_2 - 1]$ with rate $\lambda_{\text{RA}} = (1/R)\lambda_1 i_k (n_2 - k)$. Alternatively, if we adopt the LF or MF policies, the state transition $(i_0, \dots, i_{k_{\text{AP}}}, \dots, i_{n_2}) \rightarrow (i_0, \dots, i_{k_{\text{AP}}} - 1, i_{k_{\text{AP}}+1} + 1, \dots, i_{n_2})$ occurs with rate $\lambda_{\text{LF}} = \lambda_{\text{MF}} = \lambda_1$, where k_{AP} depends on the policy AP. If there is at least one superchannel that is partially occupied by a type 1 flow (i.e., if there is some $k \in [1, n_2 - 1]$ such that $i_k > 0$), then $k_{\text{LF}} = \min\{k \in [1, n_2 - 1] : i_k > 0\}$ and $k_{\text{MF}} = \max\{k \in [1, n_2 - 1] : i_k > 0\}$. Otherwise, that is, if all the superchannels are either completely free or completely occupied (i.e., if $i_k = 0$ for all $k \in [1, n_2 - 1]$), then $k_{\text{LF}} = k_{\text{MF}} = 0$. Note that in case $n_2 = 2$, $k_{\text{LF}} = k_{\text{MF}}$.

3) *Irreducibility and Limit Probabilities*: For each of the three spectrum allocation policies that we consider, the rates that were specified above determine a unique transition rate matrix. One can easily verify that these transition rate matrices \mathcal{Q}_{RA} , \mathcal{Q}_{LF} and \mathcal{Q}_{MF} are irreducible; see Appendix A for a formal proof.

Due to this irreducibility, for every event $A \subseteq \mathcal{X}_{\text{det}}$ they each determine a unique limit probability π_A . The blocking probability BP1 experienced by type 1 requests is (almost surely) equal to the limit probability of the event $A_1 = \{(i_0, \dots, i_{n_2}) \in \mathcal{X}_{\text{d}} : R = 0\}$, whereas the blocking probability BP2 experienced by type 2 requests is (almost surely) equal to the limit probability of the event $A_2 = \{(i_0, \dots, i_{n_2}) \in \mathcal{X}_{\text{d}} : i_0 = 0\}$.

B. Reduced State Space Description

The main reason that using a detailed Markov model—one with the detailed state space \mathcal{X}_{det} —to determine the blocking probabilities is infeasible for large systems, is that determining the blocking probabilities becomes computationally intensive due to the exponential dependency of the number of states on the system dimensions. Hence, one approach to reduce the duration of the computations is to reduce the number of states in the MC models. To that end, we now present alternative policy-dependent precise MC models that adopt the more compact—though less informative—state description that, in case of the random allocation policy, was first introduced by Kim et al. [9].

1) *The Reduced State Space:* A coarser state space description that still allows us to determine whether or not a request is blocked, is the triplet (i, j, e) . In this triplet, $0 \leq i \leq m_1$ (respectively $0 \leq j \leq m_2$) counts the number of type 1 (respectively type 2) flows currently allocated, and $0 \leq e \leq m_2$ counts the number of free superchannels. To ensure feasibility, it must hold that $m_2 \leq i + j + e$ and $i + (j + e)n_2 \leq m_1$. Note that the first inequality is not (explicitly) mentioned by Kim et al. [9], but is nevertheless required to ensure that all superchannels are accounted for. Enforcing these two feasibility constraints yields the reduced state space

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

The number of states in this reduced state space \mathcal{X}_{red} has a $O(m_1 m_2^2) = O(m_1 (m_1/n_2)^2)$ dependency on the total number of channels m_1 and the number of channels that form a superchannel n_2 , which is an improvement over the $O((m_1/n_2)^{n_2})$ dependency of the number of states in the detailed state space \mathcal{X}_{det} . That the reduced state space \mathcal{X}_{red} has fewer states than (or at most the same number of states as) the detailed state space \mathcal{X}_{det} follows from the fact that the function $\Gamma: \mathcal{X}_{\text{det}} \rightarrow \mathcal{X}_{\text{red}}$, defined for all (i_0, \dots, i_{n_2}) in \mathcal{X}_{det} as

$$\Gamma(i_0, \dots, i_{n_2}) := \left(\sum_{k=1}^{n_2} k i_k, m_2 - \sum_{k=0}^{n_2} i_k, i_0 \right), \quad (10)$$

is surjective. Note that in the special case that $n_2 = 2$, the function Γ is invertible: $\Gamma^{-1}(i, j, e) = (e, 2m_2 - i - 2j - 2e, i + j + e - m_2)$. Consequently, in that case, there is a one-to-one correspondence between the detailed and reduced state spaces.

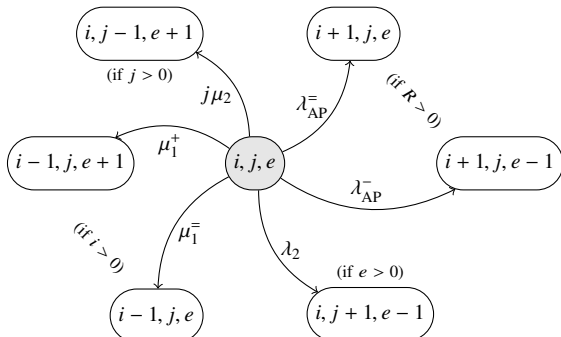


Fig. 3. State transition diagram of the precise but approximate Markov chain and of the proposed imprecise and scalable Markov chain

2) *Policy-Dependent Transition Rates:* All non-zero transition rates are schematically depicted in Fig. 3. If a free superchannel is available (i.e., if $e > 0$), allocating a type 2 request corresponds to the transition $(i, j, e) \rightarrow (i, j + 1, e - 1)$, which therefore occurs with rate λ_2 . Conversely, the expiration of the holding time of an allocated type 2 flow—which can occur if $j > 0$ —corresponds to the transition $(i, j, e) \rightarrow (i, j - 1, e + 1)$, which therefore occurs with rate $j\mu_2$.

The expiration of the holding time of an allocated type 1 flow—which can occur if $i > 0$ —might free up a superchannel or not, depending on whether or not it was the sole type 1 flow in its superchannel. In order to properly model this behaviour, we need to distinguish three separate cases. If every superchannel that is neither empty nor occupied by a type 2 flow only contains a single type 1 flow (i.e., if $m_2 = i + j + e$), then the departure of a type 1 request always frees a superchannel. This corresponds to the transition $(i, j, e) \rightarrow (i - 1, j, e + 1)$, which in this case has rate $\mu_1^+ = i\mu_1$. Note that in this case $(i - 1, j, e) \notin \mathcal{X}_{\text{red}}$. Conversely, if every superchannel that is neither empty nor occupied by a type 2 flow contains at least two type 1 flows (i.e., if $i \geq n_2(m_2 - j - e - 1) + 2$), then the departure of a type 1 flow will never free a superchannel. This corresponds to the transition $(i, j, e) \rightarrow (i - 1, j, e)$, which in this case has rate $\mu_1^- = i\mu_1$. Note that in this case $(i - 1, j, e + 1) \notin \mathcal{X}_{\text{red}}$. Unfortunately, in the remaining case that $m_2 - j - e < i < n_2(m_2 - j - e - 1) + 2$ —or equivalently, in case $(i - 1, j, e + 1) \in \mathcal{X}_{\text{red}}$ and $(i - 1, j, e) \in \mathcal{X}_{\text{red}}$ —the state representation \mathcal{X}_{red} is in general not sufficiently informative to capture the distribution of the allocated type 1 flows across the superchannels. In these cases it may therefore not be possible to determine the rate μ_1^- of the transition $(i, j, e) \rightarrow (i - 1, j, e)$ and the rate μ_1^+ of the transition $(i, j, e) \rightarrow (i - 1, j, e + 1)$ as a function of (i, j, e) . Consequently, these rates are (possibly unknown) functions of (i, j, e, t) , where t is the time of the transition. All we can say for sure is that they are non-negative, that their sum is equal to $i\mu_1$ and that

$$\mu_1^\pm(i, j, e, t) \in [i_{\min}(i, j, e)\mu_1, i_{\max}(i, j, e)\mu_1], \quad (11)$$

where $i_{\min}(i, j, e) := \max\{0, 2(m_2 - j - e) - i\}$ is the minimum number of allocated type 1 flows that are alone in their superchannel and $i_{\max}(i, j, e) := \left\lfloor \frac{n_2(m_2 - j - e) - i}{n_2 - 1} \right\rfloor$ is the maximum number of such type 1 flows. In order to deal with this indeterminacy, the authors of [9] replace the guaranteed bounds (11) with a precise and constant estimate $\tilde{\mu}_1^+(i, j, e)$ for $\mu_1^+(i, j, e, t)$, which is based on the assumption that all the possible situations—that is, all possible distributions of type 1 flows—that are represented by the state (i, j, e) are equally probable. The approximation error introduced by these estimates will be numerically evaluated in Section VII-D further on, by comparing them with the results obtained using the exact model introduced in Section V-A. The case $n_2 = 2$ again deserves special attention: we immediately verify that $\mu_1^+(i, j, e, t)$ is then a constant function of time t , as $i_{\min} = i_{\max}$.

If there is an available channel (i.e., if $R := m_1 - i - jn_2 > 0$), then a type 1 request can be assigned to a channel in either a non-full superchannel that already contains some type 1 flows or in a completely free superchannel. We again need

to distinguish three cases to properly model this. If there is no empty superchannel (i.e., if $e = 0$) then, regardless of the spectrum allocation policy, the type 1 request is assigned to a free channel in a non-empty superchannel. This corresponds to the transition $(i, j, e) \rightarrow (i + 1, j, e)$, which in this case has rate $\lambda_{\text{AP}}^- = \lambda_1$. Note that in this case $(i + 1, j, e - 1) \notin \mathcal{X}_{\text{red}}$. Conversely, if every superchannel that is neither empty nor occupied by a type 2 flow is completely occupied by type 1 flows (i.e., $i = n_2(m_2 - j - e)$), then a type 1 request is always assigned to an empty superchannel. This corresponds to the transition $(i, j, e) \rightarrow (i + 1, j, e - 1)$, which in this case has rate $\lambda_{\text{AP}}^- = \lambda_1$. Note that in this case $(i + 1, j, e) \notin \mathcal{X}_{\text{red}}$.

In the remaining case, we have that $e > 0$ and $i < n_2(m_2 - j - e)$. The two states $(i + 1, j, e)$ and $(i + 1, j, e - 1)$ are then both feasible, and the rates λ_{AP}^- and λ_{AP}^- of the transitions to these states depend on the used allocation policy. Regardless of the policy, however, both of these rates should be non-negative and their sum should equal λ_1 . Hence, we can focus on determining λ_{AP}^- because $\lambda_{\text{AP}}^- = \lambda_1 - \lambda_{\text{AP}}^-$. For the RA policy, it is shown by Kim et al. [9] that $\lambda_{\text{RA}}^- = \lambda_1 \binom{e n_2}{m_1 - i - j n_2}$. The LF and MF policies only assign a type 1 request to an empty superchannel if all non-empty superchannels containing type 1 requests are completely occupied (i.e., if $i = n_2(m_2 - j - e)$). Consequently, we have that $\lambda_{\text{LF}}^- = \lambda_{\text{MF}}^- = \lambda_1$ if $i = n_2(m_2 - j - e)$ and $\lambda_{\text{LF}}^- = \lambda_{\text{MF}}^- = 0$ if $i < n_2(m_2 - j - e)$.

3) *Irreducibility and Limit Probabilities:* For each of the spectrum allocation policies RA, LF and MF, the approximation $\tilde{\mu}_1^+(i, j, e)$ for $\mu_1^+(i, j, e, t)$ in [9] leads to an approximate transition rate matrix \tilde{Q}_{AP} . Furthermore, since $\lambda_{\text{LF}}^- = \lambda_{\text{MF}}^-$, we know that $\tilde{Q}_{\text{LF}} = \tilde{Q}_{\text{MF}} =: \tilde{Q}_{\text{LM}}$. As shown in Appendix A, \tilde{Q}_{RA} and \tilde{Q}_{LM} are both irreducible. Hence, for every event $A \subseteq \mathcal{X}_{\text{red}}$, each of these policy-dependent approximate transition rate matrices leads to a unique limit probability π_A . An approximation for the blocking probability BP1 (respectively BP2) experienced by type 1 (respectively type 2) traffic requests can therefore be obtained by computing the limit probability of the event $A_1 = \{(i, j, e) \in \mathcal{X}_{\text{red}} : R = 0\}$ (respectively $A_2 = \{(i, j, e) \in \mathcal{X}_{\text{red}} : e = 0\}$).

Note that in case $n_2 = 2$, the limit probabilities derived from \tilde{Q}_{RA} and \tilde{Q}_{LM} provide exact blocking probabilities rather than approximate ones. This is a consequence of the one-to-one correspondence between \mathcal{X}_{det} and \mathcal{X}_{red} and the equality of i_{min} and i_{max} , which were both previously mentioned as special cases.

VI. RELIABLE AND SCALABLE IMPRECISE MC MODELS

We now present our main contributions: policy-dependent imprecise MC models that explicitly account for our indeterminacy about the exact value of $\mu_1^+(i, j, e, t)$ and a common imprecise MC model that is policy-independent. The former will be used to determine guaranteed lower and upper bounds on the performance of each of our three policies, while the latter yields best and worst-case performance bounds for any policy—so not just the ones that we have studied in detail.

A. Policy-Dependent Imprecise MC Models

For each of the three allocation policies AP, we now consider the set of transition rate matrices \mathcal{Q}_{AP} that contains

all the transition rate matrices that are compatible with the relevant (bounds on the) rates specified in Section V-B2. In particular, for every policy AP and any choice of $\mu_1^+(i, j, e) \in [i_{\text{min}}\mu_1, i_{\text{max}}\mu_1]$, we obtain a different matrix $Q_{\text{AP}} \in \mathcal{Q}_{\text{AP}}$. The resulting set \mathcal{Q}_{AP} then characterises a lower transition rate operator $\underline{Q}_{\text{AP}}$ according to (7). In this way, we obtain three sets \mathcal{Q}_{RA} , \mathcal{Q}_{LF} and \mathcal{Q}_{MF} , and their respective associated lower transition rate operators $\underline{Q}_{\text{RA}}$, $\underline{Q}_{\text{LF}}$ and $\underline{Q}_{\text{MF}}$. Note that $\mathcal{Q}_{\text{LF}} = \mathcal{Q}_{\text{MF}} =: \mathcal{Q}_{\text{LM}}$ and $\underline{Q}_{\text{LF}} = \underline{Q}_{\text{MF}} =: \underline{Q}_{\text{LM}}$. For $n_2 = 2$ there is no imprecision—that is $\underline{Q}_{\text{RA}} = \tilde{Q}_{\text{RA}}$ and $\underline{Q}_{\text{LM}} = \tilde{Q}_{\text{LM}}$ —because in this case, as discussed in Section V-B2, the bounds on $\mu_1^+(i, j, e, t)$ are degenerate.

B. A Policy-Independent Imprecise MC Model

Next to elegantly handling the indeterminacy of $\mu_1^+(i, j, e, t)$, we also determine (not necessarily tight) best and worst case blocking probabilities for all possible policies. To this end, we consider the set \mathcal{Q}_{PI} of transition rate matrices that are compatible with all the relevant (bounds on the) rates specified in Section V-B2, with the exception that for $R > 0$ we now only require that

$$\lambda_{\text{AP}}^-(i, j, e, t) \in \begin{cases} \{\lambda_1\} & \text{if } i = n_2(m_2 - j - e), \\ [0, \lambda_1] & \text{otherwise} \end{cases} \quad (12)$$

and that $\lambda_{\text{AP}}^-(i, j, e, t) + \lambda_{\text{AP}}^-(i, j, e, t) = \lambda_1$. In this way, we cover all policies that assign type 1 (type 2) requests if there is at least one empty (super)channel, including time-dependent policies or policies that depend on the order of the allocated flows! The lower transition rate operator that corresponds to \mathcal{Q}_{PI} —again using (7)—is denoted by $\underline{Q}_{\text{PI}}$.

C. Ergodicity and Limit Lower and Upper Probabilities

Each of the aforementioned lower transition operators can be easily evaluated; see Appendix B for details on how to do this. We here only mention that evaluating $\underline{Q}_{\text{RA}}f$, $\underline{Q}_{\text{LM}}f$ or $\underline{Q}_{\text{PI}}f$ is linear in the number of states. Indeed, for any Q in \mathcal{Q}_{RA} , \mathcal{Q}_{LM} or \mathcal{Q}_{PI} and every state $x \in \mathcal{X}_{\text{red}}$, it is evident from the transition diagram of Fig. 3 that $[Qf](x)$ is a linear function with at most six non-zero coefficients, which we then need to minimise with respect to the simple inclusions in (11) and, for $\underline{Q}_{\text{PI}}f$, also (12).

In Appendix C, we show that \mathcal{Q}_{RA} , \mathcal{Q}_{LM} and \mathcal{Q}_{PI} have separately specified rows and that the corresponding lower transition rate operators $\underline{Q}_{\text{RA}}$, $\underline{Q}_{\text{LM}}$ and $\underline{Q}_{\text{PI}}$ are ergodic. Therefore, for every $A \subseteq \mathcal{X}_{\text{red}}$, there are unique lower and upper limit probabilities $\underline{\pi}_A$ and $\bar{\pi}_A$, which can be computed using the method introduced in Section III-B. The resulting bounds $\underline{\pi}_A$ and $\bar{\pi}_A$ do not require us to estimate the unknown transition rates, nor do they require us to specify a spectrum allocation policy in the case of $\underline{Q}_{\text{PI}}$. For all of the above imprecise MCs, the lower/upper bounds on the blocking probability BP1 (respectively BP2) experienced by type 1 (respectively type 2) requests corresponds to choosing $A_1 = \{(i, j, e) \in \mathcal{X}_{\text{red}} : R = 0\}$ (respectively $A_2 = \{(i, j, e) \in \mathcal{X}_{\text{red}} : e = 0\}$).

VII. NUMERICAL ASSESSMENT

We consider the spectrum assignment problem over a link with up to $F = 480$ slices. Type 1 channels are formed by $n_1 = 3$ contiguous slices (including guardbands), whereas type 2 superchannels consist of $n_1 n_2 = 3n_2$ slices, where n_2 varies according to the considered scenario.

A. Scalability of the Detailed Versus the Reduced State Space

We start by evaluating the number of states of the detailed state space \mathcal{X}_{det} as a function of the number of type 1 slots m_1 . As reported in Fig. 4, the number of states in the detailed state

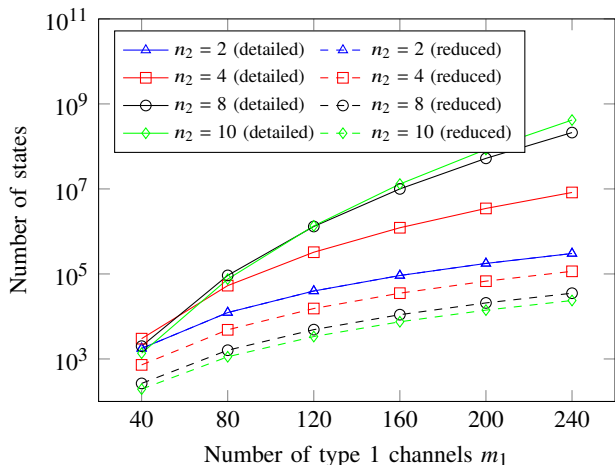


Fig. 4. The number of states in the detailed state space \mathcal{X}_{det} and the reduced state space \mathcal{X}_{red} as a function of m_1 and n_2 .

space rapidly grows with increasing m_1 , whereas the number of states in the reduced state space \mathcal{X}_{red} is at least an order of magnitude smaller if $n_2 > 2$. Moreover, the number of states in the reduced state space grows at a lower rate than the number of states in the detailed state space for $n_2 > 2$. Note also that, as already mentioned before, in case $n_2 = 2$, the number of states in the detailed and reduced state space is equal.

B. Numerically Determining Precise Limit Probabilities

It is well-known that there are multiple methods for computing limit probabilities of precise MCs. We here evaluate the performance of three such methods for the exact and approximate models described in Section V.

a) *Precise version of Alg. 1*: The first method we consider is the precise version of Alg. 1, which is obtained by replacing the lower transition rate operator \underline{Q} with a transition rate matrix Q . This method iteratively determines the limit probability π_A of an event A up to some relative tolerance ϕ . Note that Alg. 1 is similar to the well-known *power method* [29], which determines the whole stationary distribution. We opt to use the precise version of Alg. 1 instead of the power method because this allows for easier comparison in Section VII-E. We implemented Alg. 1 in Python using NumPy and SciPy functions, which are optimised for sparse matrices and dense vectors. For the parameters, we use $\delta = 0.9 \cdot 2 / \|Q_{\text{AP}}\|$, $\phi = 10^{-3}$ and $N = 10^6$.

b) *Generalised Minimal RESidual (GMRES)*: When the limit distribution π of an irreducible MC is interpreted as a row vector whose components are the limit probabilities π_x of the states, then it is well-known—see for instance [29, 30]—that π is the unique probability mass function on \mathcal{X} that satisfies the equilibrium condition $\pi Q = 0$. As explained in [29, Section 10.2.2], one can determine this limit distribution by numerically solving the linear system of equations $\pi R = b$, where the matrix R is obtained from Q by changing the elements of the last column of Q to 1, and b is a row vector where every component is zero except for the last one, which is equal to 1. We solve this system of equations using the LGMRES algorithm implemented in SciPy, which is a slightly adapted version of the GMRES algorithm for the sake of faster convergence [31], and the SciPy ILU function as preconditioner. For all parameters, we simply use the standard SciPy values.

c) *Gillespie*: A third method to determine the limit probability π_A is to simulate a sample path $\omega(t): [0, T] \rightarrow \mathcal{X}$ of the irreducible MC over a time period $T \in \mathbb{R}_{>0}$, as the ergodic theorem for MCs guarantees that $\frac{1}{T} \int_0^T \mathbb{I}_A(\omega(t)) dt$ converges (almost surely) to π_A for $T \rightarrow +\infty$. We here generate a sample path using the Gillespie algorithm, and assess the accuracy of the approximation by means of the batch mean estimation method [32]. By the PASTA property, we only have to observe the system at the arrival epochs of a Poisson process instead of keeping track of all inter-transition times. Therefore, we can simulate the system by generating a sample path of the embedded (discrete-time) Markov chain (sometimes also called the jump chain). If a transition in this sample path corresponds to the arrival of a type 1 or type 2 flow, we observe if one of the two request types would be blocked at that instant. In our C implementation of the combination of the Gillespie algorithm and the batch mean estimation method, we use batches of $4 \cdot 10^7$ arrivals and [32, Rule 1] as a rule to determine the end of the burn in period. After that, we iteratively simulate at least 5 and at most 50 batches. If the number of batches is in between the minimum and maximum number of batches, we pre-emptively stop the simulation if the relative error—taken to be the width of the 95%-confidence interval divided by the mean, as proposed in [32, Eq. (12)]—is smaller than the tolerance $\phi = 10^{-3}$.

C. Determining Blocking Probabilities

We now use each of the aforementioned computational methods—the three methods for precise MCs described in Section VII-B, and Alg. 1 for imprecise MCs—to obtain (lower and upper bounds for) the blocking probabilities for a set of scenarios, which will allow us to make a fair assessment of the performance of each method. The way we go about doing this is the following. We first select four distinct systems by specifying m_1 and setting $n_2 = 4$. Each combination of m_1 and n_2 results in a distinct detailed state space \mathcal{X}_{det} and reduced state space \mathcal{X}_{red} , the size of which is reported in Tab. I. For every system defined this way, we set $\mu_1 = \mu_2 = 1$ and $\lambda_1 = \rho \mu_1 = \rho \mu_2 = \lambda_2$, where $\rho \in \mathbb{R}_{>0}$ is called the *traffic load*.

The performance of the methods we want to compare varies with the traffic load ρ , which is why for every system we consider a low, medium and high traffic load, as listed in Tab. I. Note that, as the computational complexity of the iterative approximation methods in general increases with the number of states, using the exact chains—those with transition rate matrices Q_{RA} , Q_{LF} or Q_{MF} —to determine the blocking probabilities is only feasible for relatively low values of m_1 (and/or n_2).

TABLE I
OVERVIEW OF THE PARAMETERS USED IN THE NUMERICAL EXPERIMENTS.

m_1	n_2	$ \mathcal{X}_{\text{det}} $	$ \mathcal{X}_{\text{red}} $	ρ_{low}	ρ_{med}	ρ_{hi}
40	4	3,003	726	2	10	50
80	4	53,130	4,851	8	28	100
120	4	324,632	15,376	16	50	150
160	4	1,221,759	35,301	32	80	200

1) *Exact Blocking Probabilities:* For each of the three methods discussed in Section VII-B and every system in Tab. I, we measure the execution time of the computations that are required to determine all blocking probabilities. We go about this by fixing a numerical method, a value for m_2 and a traffic load ρ , and timing how long it takes to determine the blocking probabilities for all three allocations policies (corresponding to our three distinct MC models). In order to end up with a reliable value for these timings, we determine the average of the execution time over five consequent runs. In order to facilitate comparing the methods and models later on, for every combination of numerical method and values for m_1 and ρ we actually report the average of the execution time divided by 3 (i.e., the number of distinct models). These execution times are reported in Tab. II, where we also report the mean execution time over all three traffic loads. If the (mean) execution time is longer than 120 seconds, then we say that the computations have timed out and denote this with a / in Tab. II. The Python code we used to obtain these numbers is available on-line³ and was run on a workstation with an Intel i5-7600 CPU @ 3.50GHz, running Ubuntu 16.04.

We observe that for Alg. 1, the execution time increases more than linearly with the number of states and times out for the largest system, which is in line with our expectations. We also observe the same for the GMRES method, which actually already times out for the second largest system. From this we conclude that using Alg. 1 or the GMRES method to determine the blocking probabilities is only feasible for small systems. In fact, compared to the Gillespie method, these two methods are considerably faster for small systems. On the other hand, the Gillespie method has an execution time that does not change significantly with the model size, making it more suited for larger systems. A more accurate comparison is hard to make because the execution times are dependent on implementation choices.

2) *Approximate Blocking Probabilities:* We now run the same computations for the precise MCs characterised by \tilde{Q}_{RA} and \tilde{Q}_{LM} . Recall that for every combination of m_1 and ρ , we report the average of the execution time over five consecutive runs divided by the number of MC models, in this case 2. We report our findings in Tab. III. Note that we do not consider the Gillespie method here because its implementation is slightly less straightforward in this case and because the two alternative methods for precise MC models are already sufficiently fast. By comparing the results to those of Tab. II, we observe that the decrease in the number of states—see Tab. I—indeed results in a substantial decrease of the required computational time.

3) *Lower and Upper Blocking Probabilities:* Finally, we determine lower and upper bounds on the blocking probabilities using the imprecise MC models introduced in Section VI. These bounds are determined using Alg. 1, with $\delta = 1/(2\|\underline{Q}_{AP}\|)$, $\phi = 10^{-3}$ and $N = 10^6$. For the two policy-dependent imprecise MC models—characterised by \underline{Q}_{RA} and \underline{Q}_{LM} —we report the average execution time over five consecutive runs divided by 2, whereas for the policy-independent MC model—characterised by \underline{Q}_{PI} —we simply report the average execution time over five consecutive runs. In Tab. III, we report the obtained timings.

Comparing the mean execution times of the policy-dependent and policy-independent MCs, we observe that those of the policy-independent MC are lower for large m_1 . This is rather surprising because—for reasons explained in Appendix B—an iteration for the policy-independent MC requires more computations than an iteration for the policy-dependent MCs. Comparing the average execution times per load, we see that for the low and high load these are shorter for the policy-dependent models than for the policy-independent model, as expected. However, the average execution time for the medium traffic load and the larger systems is longer for the policy-dependent models than for the policy-independent model. Taking a more detailed look at the raw data, we find that this is caused by the extremely slow convergence of the lower and upper bounds for the medium traffic load and the imprecise MC model characterised by \underline{Q}_{LM} (for example, computing all bounds on blocking probabilities for $m_1 = 160$, ρ_{med} and \underline{Q}_{LM} took about 18 minutes). We currently have no intuitive explanation for this behaviour and consider resolving these convergence issues a matter for further study.

For the two smallest systems and the policy-dependent MCs, the mean execution time is approximately equal to 10 times the mean execution time for the approximate MC models in Tab. III, which is as expected. Similarly, for all systems and the policy-independent MCs, the mean execution time is about 10 times the mean execution time for the approximate MC models or approximately equal to the mean execution time for the policy-dependent MCs. This is faster than expected, because—as explained in Appendix B—the policy-independent MCs require more computations per iteration than the policy-dependent MCs.

³<https://github.com/alexander-e/CTMC-flexi-grid-allocation-policies>

TABLE II
AVERAGE EXECUTION TIME (IN SECONDS) FOR DETERMINING THE EXACT BLOCKING PROBABILITIES.

m_1	Alg. 1				GMRES				Gillespie			
	ρ_{low}	ρ_{med}	ρ_{hi}	mean	ρ_{low}	ρ_{med}	ρ_{hi}	mean	ρ_{low}	ρ_{med}	ρ_{hi}	mean
40	0.03	0.04	0.21	0.10	0.16	0.16	0.13	0.15	/	66.11	14.68	71.37
80	1.62	3.23	10.53	5.12	36.97	35.42	25.06	32.48	/	81.27	13.38	73.82
120	15.16	43.92	93.70	50.93	/	/	/	/	/	85.67	16.10	78.43
160	100.18	/	/	/	/	/	/	/	/	86.00	14.79	76.35

TABLE III
AVERAGE EXECUTION TIME (IN SECONDS) FOR DETERMINING APPROXIMATIONS OF AND BOUNDS ON THE BLOCKING PROBABILITIES.

m_1	Approximate blocking probabilities								Bounds on the blocking probabilities							
	Alg. 1				GMRES				Policy-dependent				Policy-independent			
	ρ_{low}	ρ_{med}	ρ_{hi}	mean	ρ_{low}	ρ_{med}	ρ_{hi}	mean	ρ_{low}	ρ_{med}	ρ_{hi}	mean	ρ_{low}	ρ_{med}	ρ_{hi}	mean
40	0.02	0.02	0.11	0.05	0.01	0.01	0.01	0.01	0.16	0.23	0.92	0.44	0.22	0.25	0.97	0.48
80	0.13	0.21	0.87	0.40	0.17	0.16	0.14	0.16	1.24	3.67	6.83	3.91	2.01	2.49	9.53	4.68
120	0.60	1.20	3.97	1.92	1.04	0.87	0.75	0.89	6.95	52.47	37.20	32.21	11.6	16.46	55.16	27.74
160	2.08	4.82	11.69	6.20	3.82	3.52	2.32	3.22	42.58	/	111.43	/	39.93	60.73	/	84.07

D. Accuracy of the Precise But Approximate MC Models

In Figs. 5a and 5b, we depict the blocking probabilities calculated for the same system with $m_1 = 80$ and $n_2 = 4$ using the RA policy or the LF and MF allocation policies, respectively. From Fig. 5a, we conclude that for the RA policy the blocking probabilities obtained using the approximate MC model are in good accordance with those obtained using the exact MC model. This is as expected, as it was previously observed by Kim et al. [9]. However, for the LF and MF allocation policies, we observe something else. On the one hand, these two policies have a different performance, which can be evaluated with their respective exact MCs. On the other hand, the approximate model yields identical approximate blocking probabilities for the two policies.

From Fig. 5b, we conclude that the approximations of BP2 for the LF and MF policies are in good accordance with the actual values of BP2. However, it also becomes evident that there is a range of traffic loads for which the approximation for BP1 is not very good. We conclude that the approximate MC models are affected by approximation errors that may result in incorrect performance evaluation and system dimensioning. Therefore, rigorous bounds are necessary to make correct design choices.

E. Discussion on Probability Bounds Provided by the Imprecise MC Models

In Fig. 6, for each connection type—that is, for type 1 and type 2 flows—and for $m_1 = 80$, we show the blocking probabilities of the RA policy and of the MF/LF policies, all obtained with the approximate MC. We show similar results for $m_1 = 120$ in Fig. 7. In addition, the same figures show lower and upper bounds obtained with \underline{Q}_{RA} or \underline{Q}_{LM} —displayed as dotted lines—as well as policy-independent bounds obtained with \underline{Q}_{PI} —displayed as dashed lines. Note that the graphs in Figs. 6 and 7 all have a double logarithmic scale. Consequently, our discussion of the graphs—and our use

of terms such as tight, wide, close to or in the middle—should be interpreted in a logarithmic sense.

Along with the approximate blocking probability, the calculated bounds make it possible to evaluate the performance of the system without solving the exact model. For type 2 flows, the bounds calculated with the policy-dependent model are very tight and show that evaluating the performance with the reduced-state model yields accurate results for all the considered parameters and traffic loads.

Different considerations hold for type 1 flows. We know, from our comparison with the exact solution in Fig. 5a, that for $m_1 = 80$ the reduced-state model for the RA policy is accurate. Nevertheless, the calculated policy-dependent lower and upper bounds are relatively wide, especially in the intermediate loads. In contrast, for the LF/MF model, we know already that the reduced-state model approximates multiple policies with different performance, and we therefore expect that the bounds cannot be very tight. However, the bounds calculated with \underline{Q}_{LM} are narrower than those calculated with \underline{Q}_{RA} .

By observing the bounds of the policy-independent model in Figs. 6b and 7b, we can see that the MF and LF policies yield a blocking probability very close to the policy-independent lower bound for type 2 flows and very close to the policy-independent upper bound for type 1 flows. This demonstrates that these policies favour type 2 flows. On the other hand, Figures 6a and 7a show that the RA policy is situated more or less in the middle of the two bounds, for both connection types. This allows us to infer that the RA policy shows no preference for one type of flow over the other.

VIII. CONCLUSION

We have studied the problem of spectrum assignment in a two-service flexi-grid optical link. First, we provided exact Markov chain models for the cases of random, most-filled, and least-filled allocation policies. We observed that the number of states in these models scales exponentially with the dimensions of the system under study. An unfortunate consequence of this

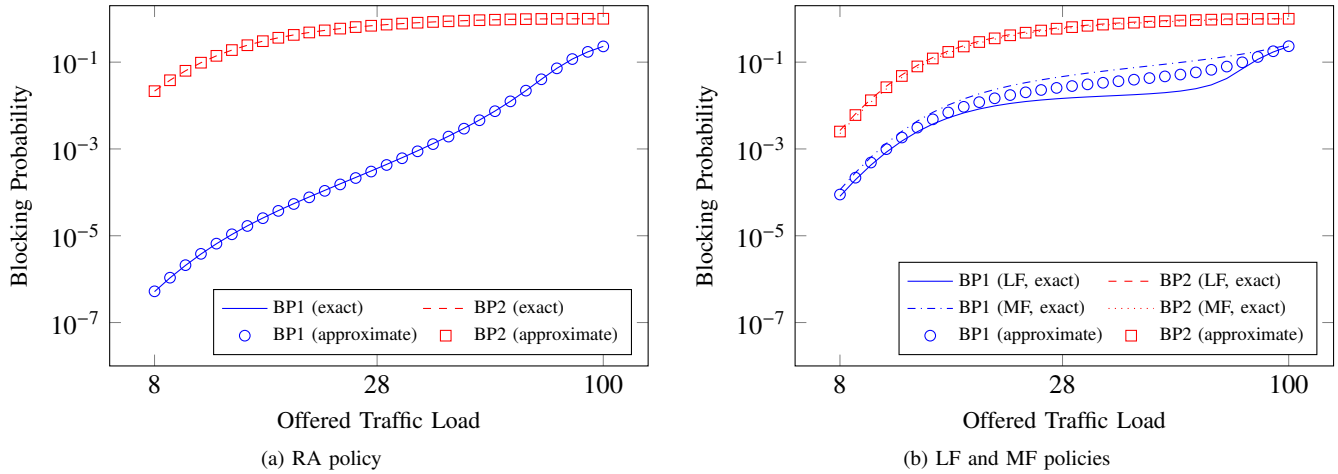


Fig. 5. Blocking probabilities for a system with $m_1 = 80$ and $n_2 = 4$, for the Random Allocation (RA), Least Filled (LF) and Most Filled (MF) allocation policies.

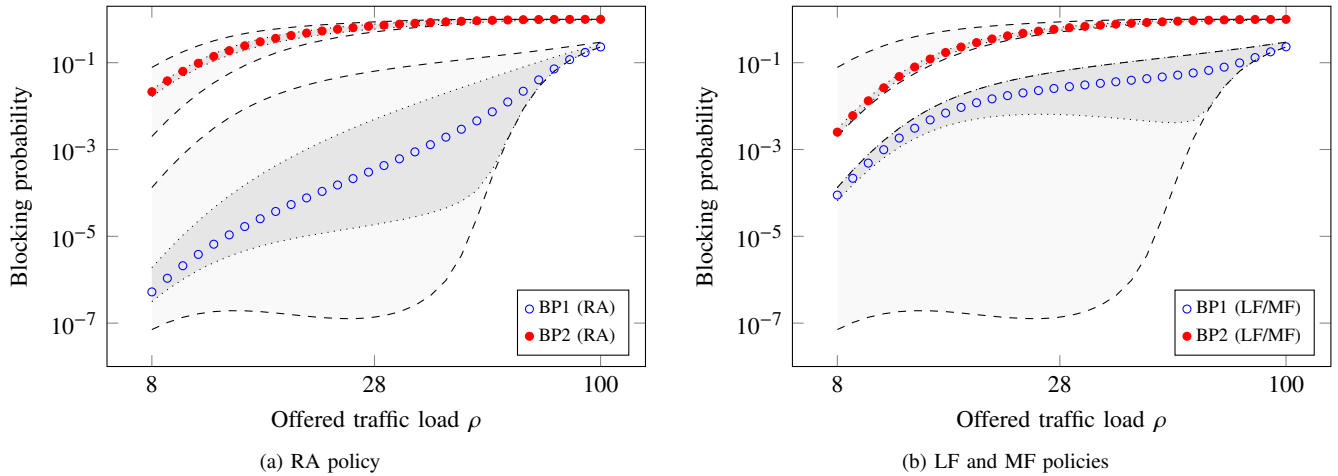


Fig. 6. Blocking probabilities for a system with $m_1 = 80$ and $n_2 = 4$. The lower and upper bounds obtained with \underline{Q}_{RA} (respectively \underline{Q}_{LM}) are displayed as dotted lines, those obtained with \underline{Q}_{PI} are displayed as dashed lines.

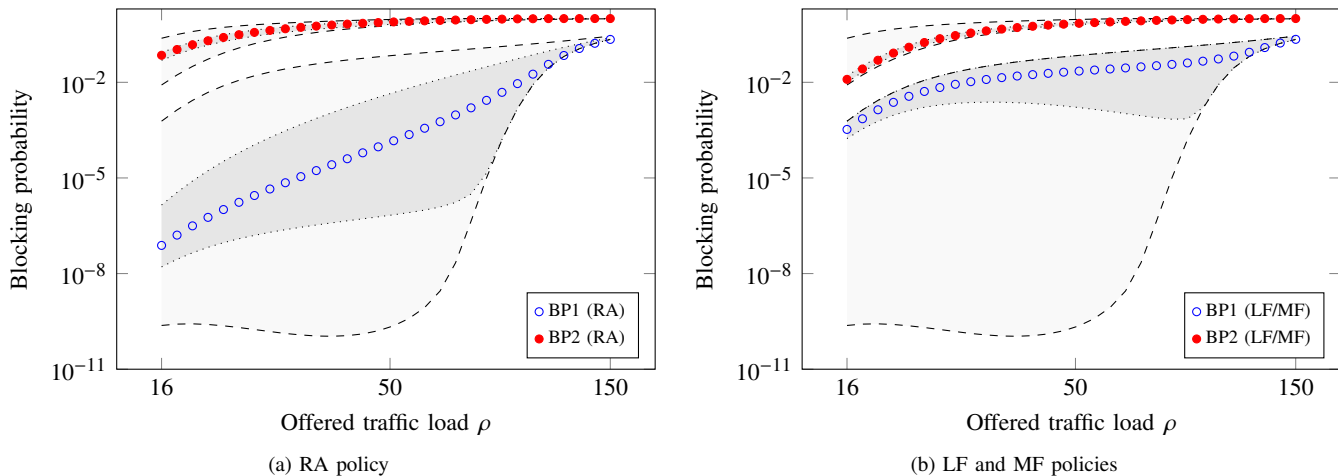


Fig. 7. Blocking probabilities for a system with $m_1 = 120$ and $n_2 = 4$. The lower and upper bounds obtained with \underline{Q}_{RA} (respectively \underline{Q}_{LM}) are displayed as dotted lines, those obtained with \underline{Q}_{PI} are displayed as dashed lines.

exponential scaling is that the computations required to determine the blocking probabilities—the performance measures we are interested in—become intractable for large systems. Then, we evaluated an approximate, reduced-state Markov chain model that is available in the literature for a random allocation policy, and extended it to the case of the most-filled and least-filled policies. For a small scale case, we observed that the blocking probabilities obtained using these approximate models were in good accordance with the exact results for the random allocation policy, but that they were not that accurate for the least-filled and most-filled policies. Furthermore, it is difficult to obtain an evaluation of the precision of the approximated blocking probabilities for large scale problems. For this reason, finally, we introduced four imprecise Markov models: one model for each of the three considered allocation policies and one policy-independent model. These models scale to large systems and provide guaranteed lower and upper bounds on the blocking probabilities instead of a single precise estimate without any measure of accuracy. Hence, the policy-dependent imprecise Markov models can be used to evaluate the precision of their approximate counterparts for large systems. Furthermore, the policy-independent model, which does not make assumptions about the allocation policy, provides guaranteed bounds on the achievable blocking probabilities. As such, for a given system, this policy-independent model can be used to assess the performance of a specific policy with respect to the set of all possible policies.

ACKNOWLEDGEMENTS

Much of Alexander Erreygers’s research was conducted while he was a member of the SMACS Research Group of the Department of Telecommunications and Information Processing at Ghent University, and he would therefore like to acknowledge their support. Jasper De Bock’s contribution was partially supported by the H2020-MSCA-ITN-2016 UTOPIAE, grant agreement 722734.

REFERENCES

- [1] O. Gerstel, M. Jinno, A. Lord, and S. B. Yoo, “Elastic optical networking: A new dawn for the optical layer?” vol. 50, no. 2, pp. s12–s20, 2012.
- [2] *Spectral grids for WDM applications: DWDM frequency grid*, ITU Telecommunication Standardization Sector Recommendation ITU-T G.694.1, Feb. 2012. [Online]. Available: <http://www.itu.int/>
- [3] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, “Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies,” vol. 47, no. 11, pp. 66–73, 2009.
- [4] M. D. Feuer and S. L. Woodward, “Advanced ROADM networks,” in *National Fiber Optic Engineers Conference*. Optical Society of America, 2012, paper NW3F.3.
- [5] Y. Yin, Z. Zhu, S. B. Yoo *et al.*, “Fragmentation-aware routing, modulation and spectrum assignment algorithms in elastic optical networks,” in *Optical Fiber Communication Conference*. Optical Society of America, 2013, paper OW3A.5.
- [6] Y. Yu, J. Zhang, Y. Zhao, X. Cao, X. Lin, and W. Gu, “The first single-link exact model for performance analysis of flexible grid WDM networks,” in *National Fiber Optic Engineers Conference*. Optical Society of America, 2013, paper JW2A.68.
- [7] Z. Shen, H. Hasegawa, K. Sato, T. Tanaka, and A. Hirano, “A novel elastic optical path network that utilizes bitrate-specific anchored frequency slot arrangement,” *Optics express*, vol. 22, no. 3, pp. 3169–3179, 2014.
- [8] J. Comellas and G. Junyent, “On the worthiness of flexible grid in elastic optical networks,” in *41st European Conference on Optical Communication (ECOC’15)*. IEEE, Sep. 2015, pp. 1–3.
- [9] J. Kim, S. Yan, A. Fumagalli, E. Oki, and N. Yamanaka, “An analytical model of spectrum fragmentation in a two-service elastic optical link,” in *2015 IEEE Global Communications Conference (GLOBECOM’15)*. IEEE, Dec. 2015, pp. 1–6.
- [10] —, “Two-service analytical model for partially-shared elastic optical link spectrum,” in *2016 IEEE 17th International Conference on High Performance Switching and Routing (HPSR’16)*. IEEE, Jun. 2016, pp. 75–80.
- [11] C. Rottondi, A. Erreygers, G. Verticale, and J. De Bock, “Modelling spectrum assignment in a two-service flexi-grid optical link with imprecise continuous-time Markov chains,” in *13th International Conference on the Design of Reliable Communication Networks (DRCN 2017)*. VDE Verlag, 2017, pp. 39–46.
- [12] A. Erreygers, C. Rottondi, G. Verticale, and J. De Bock, “Imprecise markov models for scalable and robust performance evaluation of flexi-grid spectrum allocation policies,” 2018, submitted.
- [13] G. Zhang, M. De Leenheer, A. Morea, and B. Mukherjee, “A survey on OFDM-based elastic core optical networking,” vol. 15, no. 1, pp. 65–87, 2013.
- [14] I. Tomkos, E. Palkopoulou, and M. Angelou, “A survey of recent developments on flexible/elastic optical networking,” in *14th International Conference on Transparent Optical Networks (ICTON’12)*. IEEE, Jul. 2012, pp. 1–6.
- [15] S. Azodolmolky, M. Klinkowski, E. Marin, D. Careglio, J. S. Pareta, and I. Tomkos, “A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks,” *Computer Networks*, vol. 53, no. 7, pp. 926–944, 2009.
- [16] Z. Shen, H. Hasegawa, and K. Sato, “A novel flexible grid/semi-flexible grid optical path network design algorithm that reserves exclusive frequency slots for high bitrate signals,” in *18th International Conference on Optical Network Design and Modeling (ONDM’14)*, May 2014, pp. 287–292.
- [17] J. Kim, X. Wang, S. WangYan, M. Razo, M. Tacca, and A. Fumagalli, “Blocking fairness in two-service EONs with uneven arrival rates,” in *17th International Conference on Transparent Optical Networks (ICTON’15)*, Jul. 2015, pp. 1–4.
- [18] X. Wang, J. Kim, S. Yan, M. Razo, M. Tacca, and A. Fumagalli, “Blocking probability and fairness in two-

- rate elastic optical networks,” in *16th International Conference on Transparent Optical Networks (ICTON'14)*, Jul. 2014, pp. 1–4.
- [19] T. Augustin, F. P. Coolen, G. de Cooman, and M. C. Troffaes, Eds., *Introduction to Imprecise Probabilities*, ser. Wiley Series in Probability and Statistics. Wiley, 2014.
- [20] D. Škulj, “Efficient computation of the bounds of continuous time imprecise Markov chains,” *Applied Mathematics and Computation*, vol. 250, pp. 165–180, Jan. 2015.
- [21] T. Krak, J. D. Bock, and A. Siebes, “Imprecise continuous-time Markov chains,” *International Journal of Approximate Reasoning*, vol. 88, pp. 452 – 528, 2017.
- [22] M. C. Troffaes, J. Gledhill, D. Škulj, and S. Blake, “Using imprecise continuous time Markov chains for assessing the reliability of power networks with common cause failure and non-immediate repair,” in *Proc. of the 9th International Symposium on Imprecise Probability: Theories and Applications (ISIPTA'15)*, Jul. 2015, pp. 287–294.
- [23] A. Erreygers and J. De Bock, “Imprecise continuous-time Markov chains: Efficient computational methods with guaranteed error bounds,” in *Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications*, ser. Proceedings of Machine Learning Research, vol. 62. PMLR, Jul. 2017, pp. 145–156.
- [24] J. De Bock, “The limit behaviour of imprecise continuous-time Markov chains,” *Journal of Nonlinear Science*, pp. 1–38, 2016.
- [25] D. Škulj and R. Hable, “Coefficients of ergodicity for Markov chains with uncertain parameters,” *Metrika*, vol. 76, no. 1, pp. 107–133, 2013.
- [26] S. Lopatzidis, J. De Bock, G. de Cooman, S. De Vuyst, and J. Walraevens, “Robust queueing theory: an initial study using imprecise probabilities,” *Queueing Systems*, vol. 82, no. 1, pp. 75–101, Feb. 2016.
- [27] P. Walley, *Statistical reasoning with imprecise probabilities*. Chapman and Hall, 1991.
- [28] T. Krak, J. De Bock, and A. Siebes, “Efficient computation of updated lower expectations for imprecise continuous-time hidden Markov chains,” in *Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications*, ser. Proceedings of Machine Learning Research, vol. 62. PMLR, Jul. 2017, pp. 193–204.
- [29] W. J. Stewart, *Probability, Markov chains, queues, and simulation: the mathematical basis of performance modelling*. Princeton University Press, 2009.
- [30] J. R. Norris, *Markov Chains*. Cambridge University Press, 1997.
- [31] A. H. Baker, E. R. Jessup, and T. Manteuffel, “A technique for accelerating the convergence of restarted gmres,” *SIAM Journal on Matrix Analysis and Applications*, vol. 26, no. 4, pp. 962–984, 2005.
- [32] K. Pawlikowski, “Steady-state simulation of queueing processes: Survey of problems and solutions,” *ACM*

Computing Surveys, vol. 22, no. 2, pp. 123–170, Jun. 1990.

APPENDIX A

ERGODICITY OF TRANSITION RATE MATRICES

Let Q be a transition rate matrix with state space \mathcal{X} . One well-known necessary and sufficient condition for the ergodicity of Q is based on the reflexive and transitive binary accessibility relation $\cdot \rightsquigarrow \cdot$. For any $x, y \in \mathcal{X}$, we say that x is accessible from y , denoted by $y \rightsquigarrow x$, if (i) $x = y$, or (ii) $x \neq y$ and there is a sequence $y = x_0, x_1, \dots, x_n = x$ in \mathcal{X} such that $Q(x_{i-1}, x_i) > 0$ for all $i \in \{1, \dots, n\}$. Then Q is irreducible if and only if it is top class

$$\mathcal{R}_Q := \{x \in \mathcal{X} : (\forall y \in \mathcal{X}) y \rightsquigarrow x\}$$

is equal to the state space \mathcal{X} , see [30, Chapter 3].

A. Irreducibility of the Exact Transition Rate Matrices

In Section V-A, we defined a transition rate matrix Q_{AP} with state space \mathcal{X}_{\det} for three assignment policies AP: RA, LF and MF. We posit that for all three assignment policies, the associated transition rate matrix Q_{AP} is irreducible. This is true if for every two states $x, y \in \mathcal{X}_{\det}$, $y \rightsquigarrow x$.

Instead of proving this, we will prove that for any state $x^* = (i_0^*, \dots, i_{n_2}^*) \in \mathcal{X}_{\det}$, any state of the form $x_k^+ := (i_0^*, \dots, i_k^* + 1, \dots, i_{n_2}^*) \in \mathcal{X}_{\det}$ can be accessed, and similarly for any $x_k^- := (i_0^*, \dots, i_k^* - 1, \dots, i_{n_2}^*) \in \mathcal{X}_{\det}$.

If this is the case, then using the transitivity of the accessibility relation we also have that from any state $y := (j_0, \dots, j_{n_2}) \in \mathcal{X}_{\det}$ we can access any other state $x := (i_0, \dots, i_{n_2}) \in \mathcal{X}_{\det}$ as follows. Let $K_-(y, x) := \{k \in [0, n_2] : j_k > i_k\}$ and $K_+(y, x) := \{k \in [0, n_2] : j_k < i_k\}$, then first for all $k \in K_-(y, x)$ we decrement the k -th component of the state by $j_k - i_k$, and second for all $k \in K_+(y, x)$ we increment the k -th component of the state by $i_k - j_k$. This detour ensures that any tuple along the way is a state, as any such tuple is dominated component wise by either x or y , and any tuple of non-negative natural numbers that is dominated component wise by a state is also a state. Hence, we have that $y \rightsquigarrow x$; and as x and y were arbitrary distinct states, we can conclude that the transition rate operator Q_{AP} is indeed irreducible.

As the rate λ_{AP} in the transition diagram of Fig. 2 is policy dependent, we need to treat every policy separately. Due to length constraints, we here only formally prove the irreducibility of Q_{RA} . The irreducibility of Q_{LF} and Q_{MF} can be proven in similar fashion. The part of the proof where we argue that $x^* \rightsquigarrow x_0^+$ and $x^* \rightsquigarrow x_k^-$ for $k \in [0, n_2]$ remains the same, it is only the part of the proof where we show that $x^* \rightsquigarrow x_k^+$ for $k \in [1, n_2]$ that needs to be modified.

B. Irreducibility of Q_{RA}

Fix some state $x^* \in \mathcal{X}_{\det}$ and some $k \in [0, n_2]$ such that $x_k^+ \in \mathcal{X}_{\det}$ and/or $x_k^- \in \mathcal{X}_{\det}$. We now prove that $x^* \rightsquigarrow x_k^+$ and/or $x^* \rightsquigarrow x_k^-$. Observe that for x^* the number of super-channels not occupied by a type 2 flow is $I^* := \sum_{k=0}^{n_2} i_k^* \leq m_2$.

First, we consider the case $k = 0$. From the feasibility conditions it follows that $x_0^+ \in \mathcal{X}_{\text{det}}$ if and only if $I^* < m_2$. If this is the case, then

$$Q_{\text{RA}}(x^*, x_0^+) = (m_2 - I^*)\mu_2 > 0,$$

where the strict positivity follows from $I^* < m_2$ and $\mu_2 > 0$. We infer from this that in this case indeed $x^* \rightsquigarrow x_0^+$. Also, $x_0^- \in \mathcal{X}_{\text{det}}$ if and only if $i_0^* > 0$. In this case

$$Q_{\text{RA}}(x^*, x_0^-) = \lambda_2 > 0,$$

such that indeed $x^* \rightsquigarrow x_0^-$.

Next, we consider the case $k > 0$. The state x_k^+ is then feasible if and only if $I^* < m_2$. In this case, we can also immediately verify that x_ℓ^+ is a proper state for all $\ell \in \{0, \dots, k-1\}$. Recall from before that as $I^* < m_2$,

$$Q_{\text{RA}}(x^*, x_0^+) = (m_2 - I^*)\mu_2 > 0.$$

Next, let $R^* := \sum_{k=0}^{m_2} i_k^*(n_2 - k)$ and observe that for all $\ell \in \{0, \dots, k-1\}$,

$$Q_{\text{RA}}(x_\ell^+, x_{\ell+1}^+) = \lambda_1 \frac{(i_\ell^* + 1)(n_2 - \ell)}{R^* + n_2 - \ell} > 0,$$

where the inequality follows from $\lambda_1 > 0$, $i_\ell^* \geq 0$, $\ell < k \leq n_2$ and $R^* \geq 0$. We infer from this that in this case $x^* \rightsquigarrow x_k^+$.

Also, x_k^- is feasible if and only if $i_k^* > 0$, and we now assume that this is indeed the case. For the sake of brevity, for all $\ell \in \{0, \dots, k-1\}$ we denote by x'_ℓ the $(n_2 + 1)$ -tuple of non-negative natural numbers that is obtained from x^* by decrementing the i_k^* component by 1 and incrementing the i_ℓ^* component by 1. As we have assumed that $i_k^* > 0$, all states x'_{k-1}, \dots, x'_0 are feasible. We now prove that in this case $x^* \rightsquigarrow x'_{k-1} \rightsquigarrow \dots \rightsquigarrow x'_0 \rightsquigarrow x_k^-$, which implies $x^* \rightsquigarrow x_k^-$ by transitivity. First, observe that

$$Q_{\text{RA}}(x^*, x'_{k-1}) = ki_k^*\mu_1 > 0,$$

where the inequality follows from $k \geq 1$, $i_k^* \geq 1$ and $\mu_1 > 0$. Second, observe that for all $\ell \in \{1, \dots, k-1\}$,

$$Q_{\text{RA}}(x'_\ell, x'_{\ell-1}) = \ell(i_\ell^* + 1)\mu_1 \geq \mu_1 > 0,$$

where the inequality follows from $\ell \geq 1$ and $i_\ell^* \geq 0$. Third, observe that

$$Q_{\text{RA}}(x'_0, x_k^-) = \lambda_2 > 0.$$

From these three observations, we infer that indeed $x^* \rightsquigarrow x_k^-$.

C. Irreducibility of the Approximate Transition Rate Matrices

For each of the considered assignment policies AP, the irreducibility of the approximate transition rate matrix \tilde{Q}_{AP} follows (almost immediately) from the irreducibility of Q_{AP} . In order to prove this, we first argue that a strictly positive transition rate in the detailed state space description induces a strictly positive transition rate in the reduced state space description. That is, for any $x, y \in \mathcal{X}_{\text{det}}$, $Q_{\text{AP}}(x, y) > 0$ implies that $\tilde{Q}_{\text{AP}}(\Gamma(x), \Gamma(y)) > 0$. The reason for this is that any strictly positive transition rate in the detailed state space corresponds to the arrival or departure of a type 1 or type 2 request, and that in the reduced state space the arrival or departure of a type 1 or

type 2 request also corresponds to a transition with non-zero rate. While this argument is rather intuitive, it can easily be verified more formally. Take for instance the RA policy. One can immediately verify that for any generic state $x \in \mathcal{X}_{\text{det}}$ and all $y \in \mathcal{X}_{\text{det}}$ such that $Q_{\text{RA}}(x, y) > 0$, $\tilde{Q}_{\text{RA}}(\Gamma(x), \Gamma(y)) > 0$.

Now that we have established this implication, proving the irreducibility of \tilde{Q}_{AP} becomes almost trivial. We consider the RA policy, but the reasoning for the LF and MF policies is entirely similar. Fix any two distinct states x' and y' in the reduced state space \mathcal{X}_{red} . Because of the surjectivity of the map Γ between both state spaces, there are some $x, y \in \mathcal{X}_{\text{det}}$ such that $\Gamma(x) = x'$, $\Gamma(y) = y'$ and $x \neq y$. As Q_{RA} is irreducible, there exists a sequence $x = x_0, \dots, x_n = y$ in \mathcal{X}_{RA} such that for any $i \in \{1, \dots, n\}$, $Q_{\text{RA}}(x_{i-1}, x_i) > 0$. By the argument made above, this implies that $\tilde{Q}_{\text{RA}}(\Gamma(x_{i-1}), \Gamma(x_i)) > 0$ for all $i \in \{1, \dots, n\}$. Recall that by construction $\Gamma(x_0) = x'$ and $\Gamma(x_n) = y'$, such that we can infer from this that y' is accessible from x' . As x' and y' were arbitrary distinct states in \mathcal{X}_{red} , this proves that \tilde{Q}_{RA} is indeed irreducible.

APPENDIX B

A CLOSER LOOK AT THE LOWER TRANSITION RATE OPERATORS OF INTEREST

From a practical point of view, an important question is how to efficiently execute the minimisation in (7). To discuss this, we fix some $x = (i, j, e) \in \mathcal{X}_{\text{red}}$ and define

$$\begin{aligned} x_{1,+}^- &:= (i+1, j, e), & x_{1,+}^- &:= (i+1, j, e-1), \\ x_{1,-}^+ &:= (i-1, j, e), & x_{1,-}^+ &:= (i-1, j, e+1), \\ x_{2,+} &:= (i, j+1, e-1), & x_{2,-} &:= (i, j-1, e+1). \end{aligned}$$

For convenience' sake, we first assume that all these states are feasible, or equivalently that all these states are elements of the reduced state space \mathcal{X}_{red} . If we use AP as a place holder for RA, LF or MF, it holds that

$$\begin{aligned} [\underline{Q}_{\text{AP}}f](x) &= \min_{Q \in \mathcal{Q}_{\text{AP}}} [Qf](x) \\ &= \min \left\{ -(\lambda_1 + \lambda_2 + i\mu_1 + j\mu_2)f(x) \right. \\ &\quad + \lambda_2 f(x_{2,+}) + j\mu_2 f(x_{2,-}) \\ &\quad + (\lambda_1 - \lambda_{\text{AP}}^-)f(x_{1,+}^-) + \lambda_{\text{AP}}^- f(x_{1,-}^+) \\ &\quad \left. + (i - i^+)\mu_1 f(x_{1,-}^-) + i^+ \mu_1 f(x_{1,-}^+) \right\}, \\ &\quad : i^+ \in [i_{\min}, i_{\max}], \end{aligned}$$

where λ_{AP}^- , i_{\min} and i_{\max} are functions of (i, j, e) as defined in Section V-B2. The solution to the resulting minimisation is trivial: the minimising value of i^+ is either i_{\min} or i_{\max} . The edge cases where one or more of the six states above are not feasible can be treated similarly.

Computing such minima in practice can be done in multiple ways. In our numerical computations, we a priori construct two sparse—because every row has at most 7 non-zero elements—transition rate matrices Q_{\min} and Q_{\max} : for Q_{\min} we use i_{\min} for all rows and for Q_{\max} we use i_{\max} for all rows. It then trivially holds that $[\underline{Q}_{\text{AP}}f](x) = \min\{[Q_{\min}f](x), [Q_{\max}f](x)\}$. If one is only interested in the x -component of $\underline{Q}_{\text{AP}}f$, then one simply computes the dot product of the x -th row of Q_{\min} and Q_{\max} with f and selects the minimum of the two obtained

values. If one is interested in all components of $\underline{Q}_{\text{AP}}f$, then one computes the product of each of the two matrices $\underline{Q}_{\text{min}}$ and $\underline{Q}_{\text{max}}$ with the column vector f , and computes the component-wise minimum of the two obtained column vectors.

The policy-independent case is very similar to the policy-dependent case. Under similar assumptions on $x = (i, j, e)$ as for the policy-dependent case, the minimisation in (7) reduces to

$$\begin{aligned} [\underline{Q}_{\text{PI}}f](x) &= \min_{Q \in \mathcal{Q}_{\text{PI}}} [Qf](x) \\ &= \min \left\{ -(\lambda_1 + \lambda_2 + i\mu_1 + j\mu_2)f(x) \right. \\ &\quad + \lambda_2 f(x_{2,+}) + j\mu_2 f(x_{2,-}) \\ &\quad + (1 - \alpha)\lambda_1 f(x_{1,+}^-) + \alpha\lambda_1 f(x_{1,+}^+) \\ &\quad + (i - i^+)\mu_1 f(x_{1,-}^-) + i^+\mu_1 f(x_{1,-}^+) \\ &\quad \left. : i^+ \in [i_{\text{min}}, i_{\text{max}}], \alpha \in [0, 1] \right\}. \end{aligned}$$

Note that the cases $e = 0$ and $i = n_2(m_2 - j - e)$ cannot occur under the assumptions we made. Solving the above minimisation is again trivial, as we can independently choose the minimising values of i^+ and α . Here too, the edge cases where one or more of the six adjacent states are not feasible can be treated similarly. In our numerical computations, we execute the required minimisations by determining the component-wise minimum of four column vectors. These four vectors are obtained by multiplying four (sparse) transition rate matrices—one for every combination of extremal values of i^+ and α —with the column vector f .

APPENDIX C LOWER TRANSITION RATE OPERATORS

Let \mathcal{R} denote the set of all transition rate matrices, and let $\mathcal{Q} \subset \mathcal{R}$ be some set of transition rate matrices. For all $x \in \mathcal{X}$, we let \mathcal{Q}_x denote the set of x -rows in \mathcal{Q} :

$$\mathcal{Q}_x := \{Q(x, \cdot) : Q \in \mathcal{Q}\},$$

where $Q(x, \cdot)$ denotes the x -row of Q . Then Krak et al. [21, Definition 7.3] say that \mathcal{Q} has *separately specified rows* if

$$\mathcal{Q} = \{Q \in \mathcal{R} : (\forall x \in \mathcal{X}) Q(x, \cdot) \in \mathcal{Q}_x\}.$$

In words, we require that if we construct a transition rate matrix Q by selecting an arbitrary x -row $Q(x, \cdot)$ in the set of x -rows \mathcal{Q}_x independently for each row, then Q is contained in \mathcal{Q} .

The sets \mathcal{Q}_{RA} , \mathcal{Q}_{LM} and \mathcal{Q}_{PI} of transition rate matrices that are compatible with the relevant (bounds on the) transition rates have separately specified rows. Indeed, the only transitions with linked rates—in this case the departure (or arrival) of a type 1 message—start in the same state, such that the rates of a compatible transition rate matrix can be chosen on a per-row basis. Therefore, the sets \mathcal{Q}_{RA} , \mathcal{Q}_{LM} and \mathcal{Q}_{PI} can be used to define lower transition rate operators $\underline{Q}_{\text{RA}}$, $\underline{Q}_{\text{LM}}$ and $\underline{Q}_{\text{PI}}$ according to (7), and these are guaranteed to satisfy (6).

A. Ergodicity

As explained in [24], a lower transition rate operator \underline{Q} is ergodic if $\lim_{t \rightarrow \infty} [T_s^t f](x_s)$ is independent of the initial state $x_s \in \mathcal{X}$. De Bock provides a necessary and sufficient condition for the ergodicity of a lower transition rate operator in [24, Theorem 19]. In order to state this result, we first need to introduce some notation and terminology.

Similarly to the precise case, we define a binary accessibility relation $\cdot \rightarrow \cdot$. Let \underline{Q} be a lower transition rate operator with state space \mathcal{X} . A state $x \in \mathcal{X}$ is said to be upper reachable from the state $y \in \mathcal{X}$, denoted by $y \rightarrow x$, if either (i) $x = y$, or (ii) $x \neq y$ and there is a sequence $y = x_0, \dots, x_n = x \in \mathcal{X}$ such that $[\underline{Q}\mathbb{I}_{x_k}](x_{k-1}) > 0$ for all $k \in \{1, \dots, n\}$. In the second condition, \underline{Q} denotes the upper transition rate operator, defined for all real-valued functions f on \mathcal{X} by the conjugacy relation $\underline{Q}f := -\underline{Q}(-f)$. Note that the upper accessibility relation \rightarrow is just the generalisation to upper transition rate operators of the accessibility relation \rightsquigarrow for transition rate matrices.

A transition rate operator \underline{Q} is called *top class regular* if the top class \mathcal{X}_{top} is non-empty, where

$$\mathcal{X}_{\text{top}} := \{x \in \mathcal{X} : (\forall y \in \mathcal{X}) y \rightarrow x\}.$$

In the imprecise case, however, top class regularity is not a sufficient condition for the ergodicity of a lower transition rate operator [24]. A second condition called *top class absorption* is also necessary. However, if the top class \mathcal{X}_{top} is equal to the state space \mathcal{X} , then top class absorption is trivially satisfied. Hence, irreducibility— $\mathcal{X}_{\text{top}} = \mathcal{X}$ —is a sufficient condition for ergodicity.

B. Irreducibility of the Lower Transition Rate Operators of Interest

The irreducibility of the policy-dependent lower transition rate operators $\underline{Q}_{\text{RA}}$ and $\underline{Q}_{\text{LM}}$, as well as the irreducibility of the policy-independent lower transition rate operator $\underline{Q}_{\text{PI}}$, follows from the irreducibility of the transition rate matrices \tilde{Q}_{RA} and \tilde{Q}_{LM} .

Fix any two distinct states x and y in \mathcal{X}_{red} . By the irreducibility of \tilde{Q}_{RA} , there exists a sequence $y = x_0, \dots, x_n = x$ in \mathcal{X}_{red} such that $\tilde{Q}_{\text{RA}}(x_{i-1}, x_i) > 0$ for all $i \in \{1, \dots, n\}$. Note that as $\tilde{Q}_{\text{RA}} \in \mathcal{Q}_{\text{RA}}$, it follows from (7) and the conjugacy relation between $\underline{Q}_{\text{RA}}$ and \tilde{Q}_{RA} that, for any $i \in \{1, \dots, n\}$,

$$\begin{aligned} [\underline{Q}_{\text{RA}}\mathbb{I}_{x_i}](x_{i-1}) &= -[\underline{Q}_{\text{RA}}(-\mathbb{I}_{x_i})](x_{i-1}) \\ &\geq -[\tilde{Q}_{\text{RA}}(-\mathbb{I}_{x_i})](x_{i-1}) = \tilde{Q}_{\text{RA}}(x_{i-1}, x_i). \end{aligned}$$

This inequality implies that $[\underline{Q}_{\text{RA}}\mathbb{I}_{x_i}](x_{i-1}) > 0$ for all $i \in \{1, \dots, n\}$, or equivalently that x is upper accessible from y . As x and y were arbitrary distinct states, we conclude that $\underline{Q}_{\text{RA}}$ is irreducible. Entirely similar reasoning can be used to prove the irreducibility of $\underline{Q}_{\text{LM}}$ and that of $\underline{Q}_{\text{PI}}$.