



AVACS – Automatic Verification and Analysis of
Complex Systems

REPORTS
of SFB/TR 14 AVACS

Editors: Board of SFB/TR 14 AVACS

Abstraction-based Computation of Reward
Measures for Markov Automata

by

Bettina Braitleing Luis María Ferrer Fioriti
Hassan Hatefi Bernd Becker
Ralf Wimmer
Holger Hermanns

Publisher: Sonderforschungsbereich/Transregio 14 AVACS
(Automatic Verification and Analysis of Complex Systems)
Editors: Bernd Becker, Werner Damm, Martin Fränzle, Ernst-Rüdiger Olderog,
Andreas Podelski, Reinhard Wilhelm
ATRs (AVACS Technical Reports) are freely downloadable from www.avacs.org

Copyright © October 2014 by the author(s)
Author(s) contact: Bettina Braitling (braitlin@informatik.uni-freiburg.de).

Abstraction-based Computation of Reward Measures for Markov Automata^{*}

Bettina Braitling¹, Luis María Ferrer Fioriti², Hassan Hatefi²,
Ralf Wimmer¹, Bernd Becker¹, and Holger Hermanns²

¹ Albert-Ludwigs-Universität Freiburg, Germany

{braitlin | wimmer | becker}@informatik.uni-freiburg.de

² Saarland University, Saarbrücken, Germany

{ferrer | hhatefi | hermanns}@cs.uni-saarland.de

Abstract. Markov automata allow us to model a wide range of complex real-life systems by combining continuous stochastic timing with probabilistic transitions and nondeterministic choices. By adding a reward function it is possible to model costs like the energy consumption of a system as well.

However, models of real-life systems tend to be large, and the analysis methods for such powerful models like Markov (reward) automata do not scale well, which limits their applicability. To solve this problem we present an abstraction technique for Markov reward automata, based on stochastic games, together with automatic refinement methods for the computation of time-bounded accumulated reward properties. Experiments show a significant speed-up and reduction in system size compared to direct analysis methods.

1 Introduction

During the last few years Markov automata (MA) [1] have become a popular formalism for modelling stochastic systems. Markov automata are compositional, allowing us to model large systems component-wise and to obtain a model for the whole system by combining the models of the components according to fixed composition rules. Markov automata combine nondeterminism with probabilistic behaviour and continuous stochastic timing. Thus they are a generalisation of discrete-time Markov chains (DTMCs), Markov decision processes (MDPs), probabilistic automata (PA), continuous-time Markov chains (CTMCs), and interactive Markov chains (IMCs [2]). Markov automata form the semantic foundation of generalised stochastic Petri nets (GSPNs) [3] and stochastic activity networks (SANs) [4]. For modelling systems as MA, the Markov automata process

^{*} This work was partly supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center AVACS (SFB/TR 14), by the EU 7th Framework Programme under grant agreement no. 295261 (MEALS) and 318490 (SENSATION), and by the CAS/SAFEA International Partnership Program for Creative Research Teams.

algebra (MAPA) [5] has been devised. It is accompanied with tool support: SCOOP [5] transforms descriptions in MAPA into the underlying MA and is able to apply reduction techniques to reduce the MA's size.

Such a powerful modelling formalism is, however, only useful in practice if it is accompanied by efficient analysis algorithms. Model checking of MA against continuous stochastic logic (CSL) has been discussed in [6]. Algorithms for a wide range of properties for MA have been developed: long-run average, expected reachability time, and time-bounded reachability probabilities have been considered in [7, 8]. Markov reward automata (MRA), which is the extension of MA by rewards, have recently been studied in [9]. The analysis algorithms for MA and MRA are implemented in the tool IMCA¹.

Model checking algorithms for many kinds of properties like time-unbounded reachability, expected reachability costs, and long-run averages can be transferred to MA from simpler models like PA and are similarly efficient. In contrast, checking time-bounded properties is much more expensive. The reason is that the methods based on uniformisation, which make checking time-bounded properties on CTMCs efficient, cannot be applied to MA due to the nondeterminism present there. Instead (like for IMCs [10]) one has to resort to discretisation [7–9]: the time until the time bound is split into small intervals such that one can assume that with high probability either none or exactly one step occurs within one interval. For each of these intervals, an unbounded reachability analysis for PA has to be performed. Therefore the analysis of such properties scales badly to large state spaces and is limited to a few thousand states, depending on the structure of the state space and the time bound.

Contributions. To tackle this problem for MRA we present an *abstraction and refinement framework*, the first of its kind. We target *time-bounded accumulated rewards* like “What is the maximal expected cost the system causes within 10 hours of operation?” The MRA at hand is abstracted into a two-player stochastic reward game, which keeps the nondeterminism present in the concrete system separate from the nondeterminism introduced by abstraction. This allows us to compute safe upper and lower bounds on the minimal and maximal reward value of the original system. These bounds are an in-built quality measure for the abstraction, allowing us to refine it. To compute the bounds, we give a fixed point characterisation of time-bounded accumulated rewards on stochastic games, show how to discretise it, and give an estimation of the error caused by the discretisation. Experimental results confirm that our abstraction method yields substantial reductions in system size and reduces the computation times compared to competing tools which work on the concrete state space.

Related Work. This paper continues a series of successful works on abstraction frameworks for simpler probabilistic models, foremost rooted in game-based abstraction for PA [11]. In a preliminary paper [12] we have presented an abstraction framework for time-bounded reachability probabilities for MA, which

¹ The official homepage available at <http://www-i2.informatik.rwth-aachen.de/imca>.

was the first attempt to apply abstraction to MA. Other abstraction methods are restricted to unbounded and step-bounded reachability probabilities in PA: PA-based abstraction [13] abstracts a PA again into a PA whose behaviour is an over-approximation of the behaviour of the original model. It therefore only allows to compute lower bounds on minimal probabilities and upper bounds on maximal probabilities. This is improved by game-based abstraction [11], which yields a probabilistic game. Its advantage is that it yields both upper and lower bounds for minimal and maximal reachability probabilities. Wachter and Zhang [14, 15] have proposed menu-based abstraction for PA, which has the same advantage as game-based abstraction, but yields in many practical cases significantly smaller abstractions.

Structure of the paper. In the next section we briefly review the foundations of MA and stochastic games. In Sec. 3 we present our abstraction and refinement method for MA and show how to compute reward measures for stochastic games. The experimental evaluation follows in Sec. 4. We finally summarise the paper with an outlook to future work in Sec. 5. The proofs of the main propositions are contained in the appendix of this paper.

2 Preliminaries

We first introduce the necessary foundations on stochastic games (SGs), extended by reward functions, which form the basic formalism used in our abstraction framework. We define the properties we consider and give a fixed point characterisation of them for SGs. Finally we define MA as a special case of SGs with a single player.

We denote the set of real numbers by \mathbb{R} , the non-negative real numbers by $\mathbb{R}^{\geq 0}$, and by $\mathbb{R}_{\infty}^{\geq 0}$ the set $\mathbb{R}^{\geq 0} \cup \{\infty\}$. For a finite or countable set S let $\text{Distr}(S)$ denote the set of probability distributions on S , i. e. of all functions $\mu : S \rightarrow [0, 1]$ with $\sum_{s \in S} \mu(s) = 1$. The support of a distribution μ is given by $\text{Supp}(\mu) = \{s \in S \mid \mu(s) > 0\}$; μ is called Dirac if there is $s \in S$ with $\mu(s) = 1$. It is denoted by ξ_s . Given a set $S' \subseteq S$ we write $\mu(S')$ for $\sum_{s \in S'} \mu(s)$.

2.1 Stochastic reward games

Stochastic games are a behavioural model that combines stochastic timing, nondeterminism and probabilistic choices. An SG consists of one or more players who can choose between one or more transitions to change the current state. Each choice may influence the behaviour of the other players. A transition consists of a real-valued or infinite rate $\lambda \in \mathbb{R}_{\infty}^{\geq 0}$ and a probability distribution over the successor states. For our work we need the definition of stochastic two-player games:

Definition 1 (Stochastic game). *A stochastic (continuous-time two-player) game (SG) is a tuple $\mathcal{G} = (V, (V_1, V_2), v_{\text{init}}, \mathbf{T})$ such that $V = V_1 \dot{\cup} V_2$ is a set of states, $v_{\text{init}} \in V$ is the initial state, and $\mathbf{T} \subseteq V \times \mathbb{R}_{\infty}^{\geq 0} \times \text{Distr}(V)$ is a transition relation.*

V_1 and V_2 are the states of player 1 and player 2, respectively; we also denote them as V_1 - and V_2 -states. Transitions $(v, \lambda, \mu) \in \mathbf{T}$ with rate $\lambda < \infty$ are called *Markovian*, transitions with infinite rate *probabilistic*. We denote the set of Markovian and probabilistic transitions by \mathbf{T}_M and \mathbf{T}_P , respectively; it holds that $\mathbf{T} = \mathbf{T}_M \dot{\cup} \mathbf{T}_P$. $\mathbf{T}_M(v)$ and $\mathbf{T}_P(v)$ denote the set of Markovian and probabilistic transitions available at state v , respectively. We use $\mathbf{T}(v) = \mathbf{T}_M(v) \dot{\cup} \mathbf{T}_P(v)$ as the set of all transitions available at state v .

The game starts in state v_{init} . If the current state is $v \in V_1$, then it is player 1's turn, otherwise player 2's. The current player chooses a transition $(v, \lambda, \mu) \in \mathbf{T}(v)$ for leaving state v . The rate $\theta_r((v, \lambda, \mu)) = \lambda \in \mathbb{R}_{\infty}^{\geq 0}$ determines how long we stay at v , whereas $\theta_p((v, \lambda, \mu)) = \mu \in \text{Distr}(V)$ gives us the distribution which leads to the successor states. If $\lambda = \infty$, the transition is taken instantaneously. Otherwise, λ is taken as the parameter of an exponential distribution. In this case, the probability that a transition to state $v' \in V$ happens within $t \geq 0$ time units, is given by $\mu(v') \cdot (1 - e^{-\lambda t})$. For conciseness, we write λ_{tr} instead of $\theta_r(tr)$ and μ_{tr} instead of $\theta_p(tr)$ for $tr \in \mathbf{T}$.

Paths. The dynamics of SGs is specified by paths. An infinite path $\pi \in (V \times \mathbb{R}^{\geq 0} \times \mathbf{T})^\omega$ is an infinite sequence of states, sojourn times, and transitions. A finite path is such a sequence which is finite and ending in a state, i. e. $\pi \in (V \times \mathbb{R}^{\geq 0} \times \mathbf{T})^* \times V$. We usually write $v \xrightarrow{t, tr}$ instead of $(v, t, tr) \in (V \times \mathbb{R}^{\geq 0} \times \mathbf{T})$. We use Paths^* and Paths^ω to denote the set of finite and infinite paths, respectively. Given a finite path $\pi = v_0 \xrightarrow{t_0, tr_0} v_1 \xrightarrow{t_1, tr_1} \dots v_n$, v_i is the $(i + 1)$ -th state of π , denoted by $\text{St}(\pi, i)$, t_i is the time of staying at v_i , denoted by $\text{Ti}(\pi, i)$, and $\text{Tr}(\pi, i) = tr_i$ is the executed transition for $i \in \{0, \dots, n - 1\}$. Note that v_i is left instantaneously, i. e. $\text{Ti}(\pi, i) = 0$, if $\text{Tr}(\pi, i)$ has an infinite rate. Moreover, $|\pi|$ refers to n , the length of π , and $\text{last}(\pi)$ to v_n , its last state.

Strategies. The nondeterminism which may occur at a state is resolved by a function, which is called a *strategy* (or policy or scheduler). Each player follows its own strategy in order to accomplish its goal. A strategy of player i ($i = 1, 2$) is a partial function $\sigma_i : \text{Paths}^* \rightarrow \text{Distr}(\mathbf{T})$ such that $\sigma_i(\pi) = \eta$ only if $\text{last}(\pi) \in V_i$ and $\text{Supp}(\eta) \subseteq \mathbf{T}(\text{last}(\pi))$. This strategy class is called *generic*, since it uses the complete path history to resolve the nondeterminism.² We denote the set of all strategies for player i by Strat_i .

Probability measure. Given strategies σ_1, σ_2 for both players and a state $v \in V$, a probability space on the set of infinite paths starting in v can be constructed. The set of measurable events is thereby a σ -algebra that is induced by a standard cylinder set construction [17] together with a unique probability measure $\text{Pr}_{v, \sigma_1, \sigma_2}$

² This class is also known as the class of *early schedulers* [16] because the scheduler makes its choice when entering a state and—in contrast to a late scheduler—may not change its choice while residing in a state. This is the most general scheduler class for MA, since they do not exhibit nondeterminism between Markovian transitions (see Sec. 2.3).

on the events. $\Pr_{v,\sigma_1,\sigma_2}(II)$ is the probability of the set of paths II , starting from state v , given that player 1 and player 2 play with strategies σ_1 and σ_2 , respectively. Both the σ -algebra and the probability measure are constructed by extending the existing techniques used for MA and IMCs. We omit the details here; for more information see, e. g. [6, 18, 19].

Zenoness. It may happen that an SG contains an end component consisting of probabilistic transitions only. Such an end component leads to the existence of infinite paths π with finite sojourn times, i. e. $\lim_{n \rightarrow \infty} \sum_{i=0}^n \text{Ti}(\pi, i) < \infty$. This phenomenon is called *Zenoness*. Since such behaviour has to be considered unrealistic, we assume that the SGs under consideration are non-Zeno, i. e. that they do not contain such end components. Formally, an SG is non-Zeno iff $\forall v \in V : \forall \sigma_1 \in \text{Strat}_1 \wedge \forall \sigma_2 \in \text{Strat}_2 : \Pr_{v,\sigma_1,\sigma_2}(\{\pi \mid \lim_{n \rightarrow \infty} \sum_{i=0}^n \text{Ti}(\pi, i) < \infty\}) = 0$.

For more on strategies and on SGs in general we refer to [20, 21].

Now we extend SGs by rewards (or costs). We consider two kinds of rewards: *transient rewards* for staying in a certain state and *instantaneous rewards* for taking a transition.

Definition 2 (Stochastic reward game). A stochastic reward game (SRG) is a tuple $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$ such that \mathcal{G} is an SG, $\rho_t : \mathbf{T} \rightarrow \mathbb{R}^{\geq 0}$ the transient reward function, and $\rho_i : \mathbf{T} \rightarrow \mathbb{R}^{\geq 0}$ the instantaneous reward function.

The transient reward $\rho_t(tr)$ of a transition $tr = (v, \lambda, \mu)$ is the cost of staying in v for one time unit before taking transition tr , i. e. residing in state v for Δ time units yields a transient reward of $\Delta \cdot \rho_t(tr)$. Since a state is immediately left if a transition with infinite rate is chosen, we can assume that the transient reward of such transitions is zero.

The instantaneous reward $\rho_i(tr)$ is the cost of the state change using transition $tr \in \mathbf{T}$. The accumulated reward along a path π is the sum of the costs for the transitions and the costs for staying in the states of the path. We are interested in time-bounded rewards, i. e. the costs accumulated until a time bound T is reached. It is denoted by Rew_T .

$$\begin{aligned} Rew_T(\pi) = & \sum_{i=0}^{n_T-1} \left(\rho_t(\text{Tr}(\pi, i)) \cdot \text{Ti}(\pi, i) + \rho_i(\text{Tr}(\pi, i)) \right) \\ & + \rho_t(\text{Tr}(\pi, n_T)) \cdot \left(T - \sum_{i=0}^{n_T-1} \text{Ti}(\pi, i) \right), \end{aligned} \quad (1)$$

where n_T is the largest number such that $\sum_{i=0}^{n_T-1} \text{Ti}(\pi, i) \leq T$. Each player can independently of the other try to maximise or minimise the expectation of this reward by choosing an appropriate scheduler.

$$\mathcal{R}_{\text{opt}_2}^{\text{opt}_1}(v, T) = \text{opt}_1_{\sigma_1 \in \text{Strat}_1} \text{opt}_2_{\sigma_2 \in \text{Strat}_2} \int_{\pi \in \text{Paths}^\omega} Rew_T(\pi) \, d\Pr_{v,\sigma_1,\sigma_2}(\pi) \quad (2)$$

is the $\text{opt}_1\text{-opt}_2$ expected time-bounded reward (ETR) when player i tries to optimise according to $\text{opt}_i \in \{\text{inf}, \text{sup}\}$ for $i = 1, 2$ and the game starts in state $v \in V$.

2.2 Time-bounded reward as a fixed point

In this section we provide a fixed point characterisation of the ETR for an SG. In the following we restrict the presentation to the case where player 1 maximises the expected reward, i. e. setting opt_1 to sup in (2). We denote the case by $\mathcal{R}_{\text{opt}}^{\text{sup}}$, where player 2 still has the choice to either minimise or maximise the ETR.

Lemma 1 (Fixed point characterisation). *Given an SRG $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$, a time bound $T \geq 0$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, $\text{opt}_v = \text{sup}$ if $v \in V_1$ and $\text{opt}_v = \text{opt}$ otherwise. $\mathcal{R}_{\text{opt}}^{\text{sup}}(\cdot, T)$ is the least fixed point of the higher order operator $\Omega_{\text{opt}} : (V \times \mathbb{R}^{\geq 0} \mapsto \mathbb{R}^{\geq 0}) \mapsto (V \times \mathbb{R}^{\geq 0} \mapsto \mathbb{R}^{\geq 0})$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, such that*

$$\Omega_{\text{opt}}(F)(v, T) = \underset{tr \in \mathbf{T}(v)}{\text{opt}_v} \begin{cases} \left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr}T}) \\ + \int_0^T \lambda_{tr} e^{-\lambda_{tr}t} \sum_{v' \in V} \mu_{tr}(v') F(v', T-t) dt, & \text{if } tr \in \mathbf{T}_M(v), \\ \rho_i(tr) + \sum_{v' \in V} \mu_{tr}(v') F(v', T), & \text{if } tr \in \mathbf{T}_P(v). \end{cases} \quad (3)$$

This lemma is proven in the appendix. A similar fixed point characterisation can be attained for the case that player 1 minimises the ETR. Both characterisations, however, yield Volterra integral equation systems which are not directly tractable [22]. We demonstrate in Sec. 3.2 how to approximate $\mathcal{R}_{\text{opt}}^{\text{sup}}$ by applying a discretisation technique. Moreover, the characterisation provides a sound theory for the abstraction of MA subject to time-bounded reward analysis (see Sec. 3).

2.3 Markov automata

For consistency reasons we present MA [1, 23, 24] as a special case of SGs. This is possible under two conditions: First, we consider only *closed* MA, i. e. we assume that the model to be analysed is not subject to further composition operations. Then the actions with which transitions are labelled do not carry any information and can be omitted. Second, we make the *maximal progress assumption* [24], which is typically made for closed MA before analysis. It says that probabilistic transitions (which are carried out immediately without progress of time) have precedence over Markovian transitions, which are delayed by an exponentially distributed amount of time. Together with the restriction that there is no nondeterminism between Markovian transitions in MA, we obtain the following definition:

Definition 3 (Markov (reward) automaton). An SG $\mathcal{M} = (V, (V_1, V_2), v_{\text{init}}, \mathbf{T})$ is a Markov automaton (MA) if $V_2 = \emptyset$ and each state $v \in V$ contains either only probabilistic transitions $(v, \infty, \mu) \in \mathbf{T}(v)$ or a single Markovian transition $(v, \lambda, \mu) \in \mathbf{T}(v)$ with $\lambda < \infty$.

A Markov reward automaton (MRA) is an SRG $\mathcal{M}_{\text{rew}} = (\mathcal{M}, \rho_t, \rho_i)$ where \mathcal{M} is an MA.

The fixed point characterisation of the ETR for SRGs is valid for MRA as well. Since the set V_2 is empty for MRA, we can omit the opt subscript from Ω_{opt} as defined in Eq. (3) and simply write it as Ω in this case.

Like SGs can be seen as a generalisation of MA, MA can be seen as the generalisation of some other common models: A (closed) probabilistic automaton (PA) is an MA where all transitions are probabilistic. A PA with $|\{\mu \mid (v, \infty, \mu) \in \mathbf{T}(v)\}| = 1$ for all $v \in V$ is a discrete-time Markov chain (DTMC). An interactive Markov chain (IMC) is an MA where all distributions occurring in probabilistic transitions are Dirac. A continuous-time Markov chain (CTMC) is an MA with only Markovian transitions.

We partition the state space V of an MA based on the rate of outgoing transitions of the states. A state is called *probabilistic* (*Markovian*) if its outgoing transitions are all probabilistic (Markovian). We assume that there are no *deadlock* states, which have no outgoing transitions. They can be turned into Markovian states by adding a Markovian transition (v, λ, ξ_v) with arbitrary rate $\lambda < \infty$. Then we have $V = MS \dot{\cup} PS$ with MS being the Markovian states and PS the probabilistic states.

3 Abstraction and refinement of MRA

In this section we first describe our abstraction of an MA, then how safe bounds on the maximal (minimal) ETR can be computed on SGs (and therefore also on our abstraction). Finally we will show how an abstraction can be refined in case that it is too coarse, i. e. if it yields bounds that are too far apart.

3.1 Abstraction

The abstraction of an MA $\mathcal{M} = (V, (V, \emptyset), v_{\text{init}}, \mathbf{T})$ is based on a partition of V , which is a set $\mathcal{P} \subseteq 2^V \setminus \{\emptyset\}$ with $\bigcup_{B \in \mathcal{P}} B = V$ and $B \cap B' = \emptyset$ for all $B, B' \in \mathcal{P}$ with $B \neq B'$. For $v \in V$ we denote the unique block B of \mathcal{P} with $v \in B$ by $[v]_{\mathcal{P}}$.

Definition 4 (Lifted distribution). Let $\mu \in \text{Distr}(V)$ be a probability distribution over V and \mathcal{P} a partition of V . The lifted distribution $\bar{\mu} \in \text{Distr}(\mathcal{P})$ is given by $\bar{\mu}(B) = \sum_{v \in B} \mu(v)$ for $B \in \mathcal{P}$.

Definition 5 (Labelling function). Let Lab be a finite set of labels and \mathcal{M} an MA. A labelling function is a function $\text{lab} : \mathbf{T} \rightarrow \text{Lab}$ which is injective at each state, i. e. for all $v \in V$ and all $\text{tr}, \text{tr}' \in \mathbf{T}(v)$ we have either $\text{tr} = \text{tr}'$ or $\text{lab}(\text{tr}) \neq \text{lab}(\text{tr}')$. Additionally we require $\text{lab}(\text{tr}) = \perp$ iff tr is Markovian.

For a set $B \subseteq V$ of states, we define $Lab(B) = \{lab(tr) \mid \exists v \in B : tr \in \mathbf{T}(v)\}$ as the set of actions which are enabled in B . We also write $Lab(v)$ instead of $Lab(\{v\})$.

A simple, basic way to abstract an MA \mathcal{M} is to use the blocks of a partition \mathcal{P} as abstract states and create the abstract transitions by lifting the distributions as in Def. 4. If two transitions coincide after lifting, they are combined into one abstract transition.

Such an abstraction introduces additional nondeterminism into the system: There is now a nondeterministic choice between transitions which belong to different concrete states, but to the same state in the abstraction. In a basic abstraction, this new, abstract nondeterminism cannot be distinguished from the original, concrete nondeterminism. This makes it impossible to obtain both lower and upper bounds on the actual reward value.

To avoid this effect, we introduced a game abstraction for MA in [12], which we now extend to MRA.

Definition 6 (Game abstraction of MA). *Given an MA $\mathcal{M} = (V, (V, \emptyset), v_{\text{init}}, \mathbf{T})$, a labelling function $lab : \mathbf{T} \rightarrow Lab$, and a partition $\mathcal{P} = \{B_1, \dots, B_n\}$ of V such that for all $B \in \mathcal{P}$ either $B \subseteq MS$ or $B \subseteq PS$. We construct the game abstraction $\overline{\mathcal{M}}^{\mathcal{P}, lab} = (\overline{V}, (\overline{V}_1, \overline{V}_2), \overline{v}_{\text{init}}, \overline{\mathbf{T}})$ with:*

- $\overline{V} = \overline{V}_1 \dot{\cup} \overline{V}_2$,
- $\overline{V}_1 = \mathcal{P}$,
- $\overline{V}_2 = \{(B, \alpha) \in \mathcal{P} \times Lab \mid \alpha \in Lab(B)\} \dot{\cup} \{*\}$,
- $\overline{v}_{\text{init}} = [v_{\text{init}}]_{\mathcal{P}}$, and
- $\overline{\mathbf{T}} = \overline{\mathbf{T}}_{\mathcal{P}} \dot{\cup} \overline{\mathbf{T}}_M$ with

$$\begin{aligned} \overline{\mathbf{T}}_{\mathcal{P}} &= \{([v]_{\mathcal{P}}, \infty, \xi_{([v]_{\mathcal{P}}, \alpha)}) \mid v \in V \wedge \alpha \in Lab([v]_{\mathcal{P}})\} \\ &\quad \dot{\cup} \{(([v]_{\mathcal{P}}, \alpha), \infty, \overline{\mu}) \mid v \in PS \wedge \alpha \in Lab(v) \wedge (v, \infty, \mu) \in \mathbf{T}_{\mathcal{P}}\} \\ &\quad \dot{\cup} \{(([v]_{\mathcal{P}}, \alpha), \infty, \xi_*) \mid v \in PS \wedge \alpha \in Lab([v]_{\mathcal{P}}) \setminus Lab(v)\}, \\ \overline{\mathbf{T}}_M &= \{(([v]_{\mathcal{P}}, \perp), \lambda, \overline{\mu}) \mid v \in MS \wedge (v, \lambda, \mu) \in \mathbf{T}_M\} \\ &\quad \dot{\cup} \{(*, 1, \xi_*)\}. \end{aligned}$$

We call \overline{V}_2 -states of the form (B, \perp) with $B \subseteq MS$ *Markovian* and \overline{V}_2 -states of the form (B, α) with $\alpha \neq \perp$ and $B \subseteq PS$ *probabilistic*. The \overline{V}_1 - and \overline{V}_2 -states strictly alternate.

Player 1 resolves the nondeterminism already present in the concrete MA when it selects at state B a label α present in one of the concrete states of B . Player 2 resolves the nondeterminism introduced by the abstraction by selecting at abstract state (B, α) a concrete state $v \in B$ and firing the lifted transition of state v that has the label α . In case there is no transition with label α in state v , the abstraction goes to a special state $*$ that represents the worst outcome for the property under consideration. This is similar to the menu-based abstraction in [14, 15].

In order to abstract an MRA $\mathcal{M}_{rew} = (\mathcal{M}, \rho_t, \rho_i)$ we have to add *abstract reward functions* to the abstraction. For this we need an additional function

$\nu : \overline{\mathbf{T}} \rightarrow 2^{\mathbf{T}}$, which maps the abstract transitions back to the corresponding concrete transitions in \mathcal{M}_{rew} . For a transition $\overline{tr} = (\overline{v}, \lambda, \overline{\mu}) \in \overline{\mathbf{T}}$ we get:

$$\nu(\overline{tr}) = \begin{cases} \{(v', \infty, \mu) \in \mathbf{T} \mid v' \in B\}, & \text{if } \overline{v} = (B, \alpha) \wedge \alpha \neq \perp \wedge B \subseteq PS, \\ \{(v', \lambda, \mu) \in \mathbf{T} \mid v' \in B\}, & \text{if } \overline{v} = (B, \perp) \wedge B \subseteq MS, \\ \emptyset, & \text{otherwise.} \end{cases}$$

The choice of the reward function depends on whether we want to compute a lower or an upper bound on the (minimal or maximal) ETR. In case of a lower bound, player 2 chooses the smallest possible reward among all transitions which were mapped onto the same abstract transition. The case of an upper bound is analogous.

We give the definition of the reward structures for the case that player 1 maximises the ETR. If player 1 minimises the ETR, the only change is that $\rho_t((*, 1, \xi_*))$ is set to ∞ instead of 0.

Definition 7 (opt-Abstraction-induced SRG). *Given an MRA $\mathcal{M}_{rew} = (\mathcal{M}, \rho_t, \rho_i)$, a partition \mathcal{P} of the state space, a labelling function lab for the transitions, and $opt \in \{\inf, \sup\}$. Then the opt-Abstraction-induced SRG (or for short opt-AISRG) with respect to \mathcal{P} and lab is a tuple $opt\text{-}\overline{\mathcal{M}}_{rew}^{\mathcal{P}, lab} = (\overline{\mathcal{M}}^{\mathcal{P}, lab}, \overline{\rho}_t^{opt}, \overline{\rho}_i^{opt})$, where $\overline{\mathcal{M}}^{\mathcal{P}, lab}$ is the game abstraction of \mathcal{M} obtained from Def. 6 and $\overline{\rho}_t^{opt}$ and $\overline{\rho}_i^{opt}$ are abstract transient and instantaneous reward functions defined as:*

$$\overline{\rho}_t^{opt}(\overline{tr}) = \begin{cases} \text{opt}_{tr \in \nu(\overline{tr})} \rho_t(tr), & \text{if } \nu(\overline{tr}) \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

and

$$\overline{\rho}_i^{opt}(\overline{tr}) = \begin{cases} \text{opt}_{tr \in \nu(\overline{tr})} \rho_i(tr), & \text{if } \nu(\overline{tr}) \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

respectively, where $\overline{tr} \in \overline{\mathbf{T}}$ is an abstract transition.

We illustrate the abstraction of an MRA in Example 1.

Example 1. Figure 1(a) shows an MRA \mathcal{M}_{rew} together with a partition $\mathcal{P} = \{B_0, B_1\}$. We assume that all probabilistic transitions are labelled with $lab(tr) = \alpha$ and all Markovian transitions with $lab(tr) = \perp$. It holds $PS = \{v_0, v_1, v_2\} = B_0$ and $MS = \{v_3, v_4, v_5\} = B_1$. For each transition tr the rewards are given in the form “ $(\rho_i(tr) | \rho_t(tr))$ ” next to the transition in red colour.

Figure 1(b) shows the resulting game abstraction. The blocks B_0 and B_1 have become \overline{V}_1 -states, whereas all other states are \overline{V}_2 -states. We show both reward structures: the values in red colour next to each abstract transition are the minimal rewards in the form “ $(\overline{\rho}_i^{\inf}(tr) | \overline{\rho}_t^{\inf}(tr))$ ”; the blue figures below are the maximal rewards, shown as “ $(\overline{\rho}_i^{\sup}(tr) | \overline{\rho}_t^{\sup}(tr))$ ”.

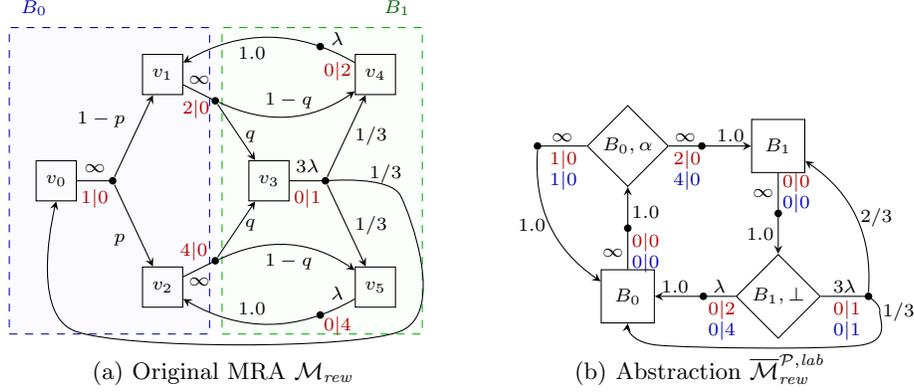


Fig. 1. An example for the game abstraction of an MRA \mathcal{M}_{rew} .

The soundness of our framework follows from the fact that the least fixed points of the abstract semantics are over- and underapproximations of the least fixed point in the concrete semantics. To prove it (see Appendix B) we use the Galois connection framework of abstract interpretation [25, 26], using abstract domains similar to the ones in the work of Wachter [15].

Theorem 1 (Soundness). *Let $\mathcal{M}_{rew} = (\mathcal{M}, \rho_t, \rho_i)$ be an MRA, its high-order operator for the maximal ETR Ω , and $\text{opt-}\overline{\mathcal{M}}_{rew}^{\mathcal{P},lab} = (\overline{\mathcal{M}}^{\mathcal{P},lab}, \overline{\rho}_t^{\text{opt}}, \overline{\rho}_i^{\text{opt}})$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, be a game abstraction with rewards, its high-order operators $\Omega_{\text{inf}}, \Omega_{\text{sup}}$. Then:*

$$\text{lfp}(\Omega_{\text{inf}})([v]_{\mathcal{P}}, T) \leq \text{lfp}(\Omega)(v, T) \leq \text{lfp}(\Omega_{\text{sup}})([v]_{\mathcal{P}}, T)$$

for all $v \in V$, and $T \in \mathbb{R}^{\geq 0}$.

This theorem is proven in the appendix.

3.2 Reward computation

In this section we describe how to compute optimal ETRs for the general class of SRGs. For this purpose a discretisation technique is employed, which is then applied to the fixed point characterisation given in Lemma 1. The technique yields a discretised fixed point characterisation accompanied by a stable numerical algorithm with strict error bound for computing the optimal ETRs in SRGs.

Discretisation. As stated before, it is not generally feasible to directly solve the fixed point characterisation in Lemma 1 due to the complex integrals occurring in Eq. (3). Instead the SRG subject to analysis needs to be *discretised*. For this, the interval $[0, T]$ is first split into k time steps of size $\delta = \frac{T}{k}$. The discretisation then simplifies the computation of $\mathcal{R}_{\text{opt}}^{\text{sup}}$ by assuming that with high probability

at most one Markovian transition fires within each time step. Finally, we provide lower and upper bounds for the error created by the discretisation.

We aim to simplify the reward computation proposed by Eq. (3). For that, we first express $\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T)$ in terms of its behaviour in the first discretisation step $[0, \delta)$ with $\text{opt} \in \{\text{inf}, \text{sup}\}$. As time passes only if a Markovian transition is taken, we assume w. l. o. g. that v only contains Markovian transitions, i. e. $\emptyset \neq \mathbf{T}(v) \subseteq \mathbf{T}_M$. To see why the assumption does not restrict the generality of the discretisation, note that the simplification only applies to the part of Eq. (3) that contains the integral equation, so the case corresponding to the probabilistic transitions remains untouched. We partition the paths from v that take transition $tr \in \mathbf{T}_M(v)$ into the set $\Pi_{v, tr}^{\delta, 0}$ of paths that make no Markovian jump in $[0, \delta)$ and the set $\Pi_{v, tr}^{\delta, >0}$ of paths that do at least one jump in that interval. We therefore write $\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T)$ as the sum of

1. The optimal expected reward attained in $[0, \delta)$ by paths from $\Pi_{v, tr}^{\delta, >0}$
2. The optimal expected reward attained in $[\delta, T]$ by paths from $\Pi_{v, tr}^{\delta, >0}$
3. The optimal expected reward attained in $[0, \delta)$ by paths from $\Pi_{v, tr}^{\delta, 0}$
4. The optimal expected reward attained in $[\delta, T]$ by paths from $\Pi_{v, tr}^{\delta, 0}$

It is not hard to express the last item in terms of $\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T - \delta)$. We further combine the first three items, denoted by $\text{Acc}_{tr}(v, T)$, and finally have:

$$\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) = \text{opt}_v \left(\text{Acc}_{tr}(v, T) + e^{-\lambda_{tr} \cdot \delta} \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T - \delta) \right) \quad (4)$$

We can show (see Appendix C) that $\text{Acc}_{tr}(v, T)$ is obtained by:

$$\begin{aligned} \text{Acc}_{tr}(v, T) &= \left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr} \cdot \delta}) \\ &\quad + \int_0^\delta \lambda_{tr} e^{-\lambda_{tr} \cdot t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \end{aligned} \quad (5)$$

As for the fixed point characterisation in Lemma 1, the exact computation of $\text{Acc}_{tr}(v, T)$, $\text{opt} \in \{\text{inf}, \text{sup}\}$ is in general intractable. However, if the discretisation constant δ is very small, then, with high probability, at most one Markovian jump happens in each discretisation step. Hence the reward gained by paths carrying multiple Markovian jumps within at least one such interval is negligible and can be omitted from the computation. In other words, the reward gained after the first Markovian jump in each discretisation constant is ignored by this approximation. It naturally induces some error and thereby approximates $\text{Acc}_{tr}(v, T)$, denoted by $\widetilde{\text{Acc}}_{\delta, tr}(v, k)$ and $\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T)$, denoted by $\widetilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k)$. As a result we have:

$$\widetilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k) = \text{opt}_v \left(\widetilde{\text{Acc}}_{\delta, tr}(v, k) + e^{-\lambda_{tr} \cdot \delta} \widetilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k - 1) \right) \quad (6)$$

$\widetilde{\text{Acc}}_{\delta, tr}$ and $\widetilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}$ both count the number of discretisation steps instead of real time. This makes their computation tractable.

$$\begin{aligned}
\widetilde{\text{Acc}}_{\delta, \text{opt}}(v, k) &= \left(\rho_i(\text{tr}) + \frac{\rho_t(\text{tr})}{\lambda_{\text{tr}}} \right) (1 - e^{-\lambda_{\text{tr}} \delta}) \\
&\quad + \int_0^{\delta} \lambda_{\text{tr}} e^{-\lambda_{\text{tr}} t} \sum_{v' \in V} \mu_{\text{tr}}(v') \tilde{\mathcal{R}}_{\text{opt}}^{\text{sup}}(v', k-1) dt \\
&= \left(\rho_i(\text{tr}) + \frac{\rho_t(\text{tr})}{\lambda_{\text{tr}}} + \sum_{v' \in V} \mu_{\text{tr}}(v') \tilde{\mathcal{R}}_{\text{opt}}^{\text{sup}}(v', k-1) \right) (1 - e^{-\lambda_{\text{tr}} \delta})
\end{aligned} \tag{7}$$

By using Eq. (6) instead of the Markovian part in Eq. (3) of the fixed point characterisation in Lemma 1, we get a discretised fixed point characterisation which is directly computable.

Definition 8 (Discretised maximum time-bounded reward). *Given an SRG $(\mathcal{G}, \rho_t, \rho_i)$, a time bound $T \geq 0$, and a step size $\delta > 0$ such that $T = k \cdot \delta$ for $k \in \mathbb{N}$. Let $\text{opt} \in \{\inf, \text{sup}\}$, $\text{opt}_v = \text{sup}$ if $v \in V_1$ and $\text{opt}_v = \text{opt}$ otherwise. $\mathcal{R}_{\text{opt}}^{\text{sup}}(\cdot, T)$ is the least fixed point of the higher order operator $\Omega_{\text{opt}}^{\delta} : (V \times \mathbb{N} \mapsto \mathbb{R}^{\geq 0}) \mapsto (V \times \mathbb{N} \mapsto \mathbb{R}^{\geq 0})$ such that*

$$\Omega_{\text{opt}}^{\delta}(F)(v, k) = \underset{\text{tr} \in \mathbf{T}(v)}{\text{opt}_v} \begin{cases} \left(\rho_i(\text{tr}) + \frac{\rho_t(\text{tr})}{\lambda_{\text{tr}}} + \sum_{v' \in V} \mu_{\text{tr}}(v') F(v', k-1) \right) \cdot (1 - e^{-\lambda_{\text{tr}} \delta}) \\ + e^{-\lambda_{\text{tr}} \delta} F(v, k-1), & \text{if } v \in \mathbf{T}_M(v), \\ \rho_i(\text{tr}) + \sum_{v' \in V} \mu_{\text{tr}}(v') F(v', k), & \text{if } v \in \mathbf{T}_P(v). \end{cases} \tag{8}$$

Discretisation error. This section evaluates the precision of the discretisation technique described above. The discretisation technique can be applied to any kind of SRG respecting Def. 2. Its precision can be accordingly assessed for the general class of SRGs. However, opt-AISRGs obtained from MA abstraction have a special structure, namely all their Markovian transitions are controlled by player 2. For this specific structure it is usually possible to find a tighter error bound for time-bounded analysis (see for example [27]). Hence we restrict ourselves to a subclass of SRGs whose Markovian transitions (if there are any) are controlled by one player. In other words, the discretisation of models in this subclass in general introduces a smaller error compared to the general class of SRGs. The subclass is formally defined as:

Definition 9 (Single Markovian Controller SRG). $\mathcal{G}_{\text{rew}} = (\mathcal{G}, \rho_t, \rho_i)$ is called a single Markovian controller SRG (1MC-SRG) iff either $\forall \text{tr} \in \mathbf{T}_M : \text{tr} \in \mathbf{T}_M(v) \Rightarrow v \in V_1$ or $\forall \text{tr} \in \mathbf{T}_M : \text{tr} \in \mathbf{T}_M(v) \Rightarrow v \in V_2$.

The accuracy of $\tilde{\mathcal{R}}_{\text{opt}}^{\text{sup}}$ depends on some parameters including the step size δ . The smaller δ is, the better is the quality of the discretisation. In order to assess the accuracy of the discretisation we first need to define some parameters of SRGs.

Definition 10. Given an SRG $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$, we define the maximum (finite) exit rate existing in \mathcal{G} as $\bar{e} = \max_{tr \in \mathbf{T}_M} \lambda_{tr}$, and the maximum transient reward existing in ρ_t as $\bar{r}_t = \max_{tr \in \mathbf{T}} \rho_t(tr)$. Moreover, let \bar{r}_i be the maximum instantaneous reward that can be earned between two consecutive Markovian jumps. This value can be efficiently computed via the Bellman equation given in [9, Theorem 1] after assigning zero value to all transient rewards.

The following theorem quantifies the quality of the discretisation.

Theorem 2. Given an IMC-SRG $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$, time bound $T > 0$ and discretisation step $\delta > 0$ such that $T = k\delta$ for some $k \in \mathbb{N}$. Then for all $v \in V$ we have:

$$\tilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k) \leq \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) \leq \tilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k) + \frac{\bar{e}T}{2}(\bar{r}_t + \bar{e}r_i)(1 + \frac{\bar{e}T}{2})\delta$$

Using Theorem 2 it is possible to find a step size that respects a given predefined accuracy level. The proposed error bound is a linear approximation of the original bound (see Eq. (11) in Appendix D), which is a more complicated function with the same set of parameters. Since the original bound is tighter, in practice it is used for finding an appropriate step size by applying Newton's method.

3.3 Initial abstraction and labelling function

An important part of the abstraction and later the refinement process is starting with a suitable initial partition \mathcal{P} and labelling function lab . On the one hand, if \mathcal{P} is too coarse the resulting game abstraction requires many refinement steps until the desired accuracy is reached. On the other hand, if \mathcal{P} is unnecessarily fine the resulting abstraction will not be able to reduce the size of the state space sufficiently.

A simple way to obtain a partition [12] is by exploiting the actions of the probabilistic transitions in an MA. They are used for synchronisation of different MA. The partition contains one block with all Markovian states and groups the probabilistic states according to the action labels available in a state. The labelling function is given by the action labels of the transitions.

However, experiments have shown that in such a partition transitions in the same block that are labelled with the same action may exhibit very different behaviour. Therefore they trigger refinement steps which can be avoided by a labelling function that takes the similarities of transitions into account. To achieve this, we first define a new initial partition: The Markovian states form one block of \mathcal{P} ; the probabilistic states are grouped according to the number of outgoing transitions:

$$\mathcal{P} = \{MS\} \dot{\cup} \{\{v \in PS \mid |\mathbf{T}(v)| = |\mathbf{T}(v')|\} \mid v' \in PS\}.$$

Based on this partition we compute the labelling function lab , using a greedy strategy as follows: All Markovian transitions are labelled with \perp as required by

Def. 4. For the transitions of a probabilistic block $B \subseteq PS$ we proceed as follows: We take an arbitrary state $v \in B$ (actually the first one in our list of states) and assign each transition $tr \in \mathbf{T}(v)$ a unique label $lab(tr) := \alpha_{tr}$. Running over the transitions $tr \in \mathbf{T}(v)$, we choose from each state $v' \in B$ with $v' \neq v$ a transition $tr' \in \mathbf{T}(v')$ which is not labelled yet and minimises $m_{\mathcal{P}}(\mu_{tr}, \mu_{tr'})$ defined as follows:

$$m_{\mathcal{P}} : \text{Distr}(V) \times \text{Distr}(V) \rightarrow [0, 2] \text{ with } m_{\mathcal{P}}(\mu, \mu') = \sum_{B' \in \mathcal{P}} |\mu(B') - \mu'(B')|.$$

The function $m_{\mathcal{P}}$ is a *pseudo-metric*³ that measures the similarity of probability distributions with respect to a partition \mathcal{P} . Formally we take for each state $v' \in B$ with $v' \neq v$ an arbitrary

$$tr' \in \begin{array}{l} \arg \min \\ tr'' \in \mathbf{T}(v') \\ lab(tr'') \text{ is undefined} \end{array} m_{\mathcal{P}}(\mu_{tr}, \mu_{tr''})$$

and set $lab(tr') := \alpha_{tr}$.

By labelling the transitions in this way we ensure that more similar probabilistic transitions belong to the same probabilistic \bar{V}_2 -state, which prevents unnecessary splitting operations during refinement. Since all states within one probabilistic block B get the same set of labels with this labelling function, we do not need to introduce the *-state.

3.4 Refinement

We approximate the maximal ETR \mathcal{R}^{sup} of an MRA \mathcal{M}_{rew} by computing a lower and an upper bound $\tilde{\mathcal{R}}_{\delta, \text{inf}}^{\text{sup}}$ and $\tilde{\mathcal{R}}_{\delta, \text{sup}}^{\text{sup}}$ with a game abstraction $\bar{\mathcal{M}}^{\mathcal{P}, lab}$ and the abstract reward functions $\bar{\rho}_i^{\text{opt}}$ and $\bar{\rho}_i^{\text{opt}}$. If these bounds are too far apart, i.e. $\tilde{\mathcal{R}}_{\delta, \text{sup}}^{\text{sup}} - \tilde{\mathcal{R}}_{\delta, \text{inf}}^{\text{sup}} > \varepsilon$, with ε being a precision threshold, our abstraction is too coarse and needs to be refined. The result of the refinement is a new partition, which in turn leads to a new game abstraction. This refinement-loop is repeated until the intended precision threshold ε is reached.

The reason behind the difference of the bounds can be two different situations: (1) The difference occurs due to different choices in the player 2 strategies σ_2^{inf} and σ_2^{sup} . (2) The difference is a result of the different reward structures. In case (1) we can use a strategy-based refinement strategy like in [12]. In case (2) we have to use a refinement strategy based on the reward values.

For this *value-based* refinement strategy we first have to search for a \bar{V}_2 -state \bar{v} where the values for $\tilde{\mathcal{R}}_{\delta, \text{inf}}^{\text{sup}}$ and $\tilde{\mathcal{R}}_{\delta, \text{sup}}^{\text{sup}}$ differ and the reward functions give different values. More precisely, this means that we have to search for an abstract transition $\bar{tr} \in \bar{\mathbf{T}}$ which was chosen by σ_2^{inf} and σ_2^{sup} and where $\left[\bar{\rho}_i^{\text{opt}}(\bar{tr}) + \frac{\rho_i^{\text{opt}}(\bar{tr})}{\lambda_{\bar{tr}}} \right] = r_{\text{opt}}$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, and $r_{\text{inf}} \neq r_{\text{sup}}$.

³ The function $m_{\mathcal{P}}$ has the following properties: $m_{\mathcal{P}}(\mu, \mu') = m_{\mathcal{P}}(\mu', \mu)$, $m_{\mathcal{P}}(\mu, \mu'') \leq m_{\mathcal{P}}(\mu, \mu') + m_{\mathcal{P}}(\mu', \mu'')$, and $m_{\mathcal{P}}(\mu, \mu) = 0$ for all distributions μ, μ', μ'' . However, it is not a metric because $m_{\mathcal{P}}(\mu, \mu') = 0$ does not imply that $\mu = \mu'$ holds.

If we have found such an abstract transition \bar{tr} , we split the preceding block B . For this we compute two sets of concrete states, one ($B_{r_{\text{inf}}}$) containing the states with reward r_{inf} , the other ($B_{r_{\text{sup}}}$) containing the states with reward r_{sup} : $B_{r_{\text{opt}}} = \left\{ v \mid v \in B \wedge (v, \lambda, \mu) = tr \in \nu(\bar{tr}) \wedge \left[\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right] = r_{\text{opt}} \right\}$ with $\text{opt} \in \{\text{inf}, \text{sup}\}$.

In case of an abstract Markovian transition \bar{tr} it may occur that $B_{r_{\text{opt}}} = \emptyset$ since it is possible that no concrete transition $tr \in \nu(\bar{tr})$ matches the constraints. Should this happen for both $B_{r_{\text{opt}}}$ we use the concrete transitions tr which optimise $\left[\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right]$ instead.

After we have generated the concrete state sets for the minimal and maximal reward, we replace B with $B_{r_{\text{inf}}}$, $B_{r_{\text{sup}}}$ and $B \setminus (B_{r_{\text{inf}}} \dot{\cup} B_{r_{\text{sup}}})$.

Block B is replaced by at least two and at most three new blocks, which leads to a new, strictly finer partition and thus to a new game abstraction of \mathcal{M}_{rew} , which in turn can be analysed and refined. This refinement-loop is repeated until the precision threshold ε is reached.

Similar to [12] we apply a “precision trick”, i.e. we start with a coarse, temporary precision threshold $\hat{\varepsilon}$ for the refinement-loop. If precision $\hat{\varepsilon}$ is reached, we switch to a higher precision, i.e. we lower $\hat{\varepsilon}$ and continue the refinement-loop. This process is repeated until the final precision ε is reached.

Zenoness. Although we assume the considered MRA \mathcal{M}_{rew} is non-Zeno, i.e. it does not contain probabilistic end components, it may happen that Zenoness is introduced into the abstraction $\overline{\mathcal{M}}_{rew}^{\mathcal{P},lab}$. This occurs, e.g. if a non-cyclic sequence of probabilistic states is partitioned into the same block $B \in \mathcal{P}$. If the instantaneous reward function ρ_i is non-zero within the end component, the value iteration will not terminate since the accumulated reward does not converge.

We avoid this effect with the following method: Before applying value iteration to solve the discretised fixed point characterisation, we employ a standard graph algorithm [28] to search for end components in $\overline{\mathcal{M}}_{rew}^{\mathcal{P},lab}$. If a probabilistic end component is found, we refine the corresponding blocks B into smaller blocks and recompute the abstraction. This process is repeated until all probabilistic end components have been removed. Since there are no probabilistic end components present in \mathcal{M}_{rew} , our method will always terminate.

4 Experimental results

We implemented the described abstraction and refinement framework in C++ in a prototype tool called MeGARA. As mentioned earlier, we are currently concentrating on the maximal ETR only, using discretisation (see Sec. 3.2). We compare our experimental results with those of IMCA [10, 7, 9], an analyser for MA and IMCs. For our experiments we used the following case studies:

The **Polling System** (PoS) [7, 29] consists of two stations and one server. Incoming requests are stored within two queues until they are delivered by the

Table 1. Results for the polling system and the queueing system.

Name	IMCA			MeGARA				
	#states	r	time	#states	r_{lb}	r_{ub}	#iter.	time
PoS-3-2	1,547	0.830	0:04	657	0.828	0.830	17	0:02
PoS-3-3	14,322	0.922	1:06	5,011	0.920	0.921	19	1:16
PoS-3-4	79,307	0.985	10:59	11,527	0.982	0.985	20	7:02
PoS-4-2	6,667	0.832	0:20	1,513	0.829	0.831	18	0:12
PoS-4-3	131,529	0.924	13:46	31,992	0.922	0.923	22	8:55
PoS-5-2	27,659	0.833	1:47	2,006	0.830	0.832	18	0:23
QS-2	103	1.768	0:04	76	1.768	1.768	9	0:01
QS-3	163	2.307	0:10	118	2.306	2.306	10	0:02
QS-4	237	2.679	0:19	167	2.678	2.680	13	0:07
QS-8	673	3.351	1:47	455	3.351	3.351	17	0:33
QS-16	2,217	3.530	12:07	1,039	3.530	3.530	24	2:33
QS-32	7,993	TO		3,286	3.532	3.532	40	57:43

server to their station. With a probability of $p = 0.1$ a request erroneously stays in the queue after if it was delivered to a station. In our experiments we varied the queue size Q and the number of different request types J . The rewards in this case study represent costs for processing requests and consuming server memory. The model instances are denoted as “PoS- Q - J ”.

The **Queueing System** (QS) [6] stores requests within two queues of size K , each belonging to a server. Server S_1 handles requests and eliminates them from its queue. Requests processed by server S_2 are either nondeterministically submitted to the queue of S_1 , or with probability $q = 0.3$ re-submitted to the queue of S_2 , or with probability $(1 - q)$ eliminated from the queue. With the help of rewards we explore the average number of jobs in the queues. For our experiments we varied the queue size K . The model instances are denoted as “QS- K ”.

We created the model files with SCOOP [5], a modelling tool for MA. All experiments were done on a Dual Core AMD Opteron processor with 2.4 GHz per core and 64 GB of memory. Computations which took longer than one hour were aborted and are marked with “TO”. Each computation needed less than 4 GB memory, we therefore do not present measurements of the memory consumption.

For all experiments we used time bound $T = 1$ and precision $\varepsilon = 0.01$ in order to compute the maximal ETR.

Table 1 shows the experimental results. The first column contains the name of the respective benchmark instance, the blocks titled “IMCA” and “MeGARA” present the results from IMCA and our abstraction refinement tool, respectively. The columns headed with “#states” give the number of states of the concrete system (in case of IMCA) and the final abstraction (in case of MeGARA). The benchmarks instances are relatively small since the solving of discretised systems is rather expensive [30].

Column “ r ” contains the maximal ETR computed by IMCA, whereas the columns “ r_{lb} ” and “ r_{ub} ” denote the lower and the upper bounds for the abstraction. The columns titled “time” present the total computation time (in format min:s) needed by IMCA and MeGARA, respectively. For our abstraction refinement tool we do not give more detailed time measurements since the

better part of the computation time is needed for the repeated analysis of the abstraction, whereas the time needed for the refinement and re-computation of the abstraction is often negligible. For example, for QS-32 we need 57 min and 43 s, of which 57 min and 41 s are spent on the analysis, whereas the time spent on re-computing and refining the abstraction is less than 2 s.

For the instances of PoS our tool needs about the same amount of time as IMCA or is even faster, for the instances of QS we are always faster. As can be seen from the columns “ r_{lb} ” and “ r_{ub} ” the quality of our abstraction is always very good. In some cases the value “ r ” computed by IMCA is slightly higher than r_{ub} , however this deviation is always well within our precision threshold of $\varepsilon = 0.01$. We always achieve a notable compaction of the state space, even for the smallest instances with only a few hundred states. For the bigger instances we can report on compaction rates up to 92%, e. g. for PoS-5- J we can reduce the system from 27,659 to 2,006 states.

5 Conclusion

We have presented a new abstraction technique for MRA, based on SRGs. We are able to analyse our abstraction regarding ETR properties. Should the quality of the abstraction be too low, we can apply scheduler- and value-based refinement methods. Our experiments show a significant compaction of the state space and a reduction of computation times.

In the future we plan to explore the possibilities of different initial partitions, labelling functions, and refinement techniques. We also plan to work on additional types of properties, e. g. bounded rewards or long-run average.

References

1. Eisentraut, C., Hermanns, H., Zhang, L.: On probabilistic automata in continuous time. In: Proc. of LICS, IEEE CS (2010) 342–351
2. Hermanns, H.: Interactive Markov Chains – The Quest for Quantified Quality. Vol. 2428 of LNCS. Springer (2002)
3. Eisentraut, C., Hermanns, H., Katoen, J.P., Zhang, L.: A semantics for every GSPN. In: Proc. of Petri Nets. Number 7927 in LNCS, Springer (2013) 90–109
4. Meyer, J.F., Movaghar, A., Sanders, W.H.: Stochastic activity networks: Structure, behavior, and application. In: Proc. of PNPM, IEEE CS (1985) 106–115
5. Timmer, M., Katoen, J.P., van de Pol, J., Stoelinga, M.: Efficient modelling and generation of Markov automata. In: Proc. of CONCUR. Vol. 7454 of LNCS, Springer (2012) 364–379
6. Hatefi, H., Hermanns, H.: Model checking algorithms for Markov automata. ECE-ASST **53** (2012)
7. Guck, D., Hatefi, H., Hermanns, H., Katoen, J.P., Timmer, M.: Modelling, reduction and analysis of Markov automata. In: Proc. of QEST. Vol. 8054 of LNCS, Springer (2013) 55–71
8. Guck, D., Hatefi, H., Hermanns, H., Katoen, J., Timmer, M.: Analysis of timed and long-run objectives for markov automata. Logical Methods in Computer Science **10**(3) (2014)

9. Guck, D., Timmer, M., Ruijters, E., Hatefi, H., Stoelinga, M.: Modelling and analysis of Markov reward automata. In: Proc. of ATVA. Vol. 8837 of LNCS, Springer (2014)
10. Guck, D., Han, T., Katoen, J.P., Neuhäüßer, M.R.: Quantitative timed analysis of interactive Markov chains. In: Proc. of NFM. Vol. 7226 of LNCS, Springer (2012) 8–23
11. Kattenbelt, M., Kwiatkowska, M.Z., Norman, G., Parker, D.: A game-based abstraction-refinement framework for Markov decision processes. Formal Methods in System Design **36**(3) (2010) 246–280
12. Braitling, B., Ferrer Fioriti, L.M., Hatefi, H., Wimmer, R., Becker, B., Hermanns, H.: MeGARA: Menu-based game abstraction and abstraction refinement of Markov automata. In: Proc. of QAPL. Vol. 154 of EPTCS (2014) 48–63
13. D’Argenio, P.R., Jeannet, B., Jensen, H.E., Larsen, K.G.: Reduction and refinement strategies for probabilistic analysis. In: Proc. of PAPM-PROBMIV. (2002) 57–76
14. Wachter, B., Zhang, L.: Best probabilistic transformers. In: Proc. of VMCAI. Vol. 5944 of LNCS, Springer (2010) 362–379
15. Wachter, B.: Refined probabilistic abstraction. PhD thesis, Saarland University (2011)
16. Neuhäüßer, M.R., Zhang, L.: Time-bounded reachability probabilities in continuous-time Markov decision processes. In: Proc. of QEST, IEEE CS (2010) 209–218
17. Ash, R.B., Doléans-Dade, C.A.: Probability & Measure Theory. 2nd edn. Academic Press (1999)
18. Neuhäüßer, M.R.: Model checking nondeterministic and randomly timed systems. PhD thesis, RWTH Aachen University and University of Twente (2010)
19. Johr, S.: Model checking compositional Markov systems. PhD thesis, Saarland University, Germany (2008)
20. Shapley, L.S.: Stochastic games. Proceedings of the National Academy of Sciences of the United States of America **39**(10) (1953) 1095
21. Brázdil, T., Forejt, V., Krcál, J., Kretínský, J., Kucera, A.: Continuous-time stochastic games with time-bounded reachability. Information and Computation **224** (2013) 46–70
22. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time markov chains. IEEE Trans. Software Eng. **29**(6) (2003) 524–541
23. Eisentraut, C., Hermanns, H., Zhang, L.: Concurrency and composition in a stochastic world. In: Proc. of CONCUR. Vol. 6269 of LNCS, Springer (2010) 21–39
24. Deng, Y., Hennessy, M.: On the semantics of Markov automata. Information and Computation **222** (2013) 139–168
25. Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Proc. of POPL, ACM (1977) 238–252
26. Cousot, P., Cousot, R.: Abstract interpretation frameworks. J. Log. Comput. **2**(4) (1992) 511–547
27. Fearnley, J., Rabe, M.N., Schewe, S., Zhang, L.: Efficient approximation of optimal control for continuous-time markov games. In: Proc. of FSTTCS. Vol. 13 of LIPIcs, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik (2011) 399–410
28. Chatterjee, K., Henzinger, M.: Faster and dynamic algorithms for maximal end-component decomposition and related graph problems in probabilistic verification. In: Proc. of SODA, SIAM (2011) 1318–1336
29. Timmer, M., van de Pol, J., Stoelinga, M.: Confluence reduction for Markov automata. In: Proc. of FORMATS. Vol. 8053 of LNCS, Springer (2013) 243–257

30. Zhang, L., Neuhäuser, M.R.: Model checking interactive Markov chains. In: Proc. of TACAS. (2010) 53–68
31. Guck, D., Timmer, M., Ruijters, E., Hatefi, H., Stoelinga, M.: Modelling and analysis of Markov reward automata (extended version). Technical Report TR-CTIT-14-06, Centre for Telematics and Information Technology, University of Twente, Enschede (2014)

A Proof of Lemma 1

Lemma 1 (Fixed point characterisation). *Given an SRG $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$, a time bound $T \geq 0$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, $\text{opt}_v = \text{sup}$ if $v \in V_1$ and $\text{opt}_v = \text{opt}$ otherwise. $\mathcal{R}_{\text{opt}}^{\text{sup}}(\cdot, T)$ is the least fixed point of the higher order operator $\Omega_{\text{opt}} : (V \times \mathbb{R}^{\geq 0} \mapsto \mathbb{R}^{\geq 0}) \mapsto (V \times \mathbb{R}^{\geq 0} \mapsto \mathbb{R}^{\geq 0})$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, such that*

$$\Omega_{\text{opt}}(F)(v, T) = \text{opt}_v \left\{ \begin{array}{l} \left(\rho_i(\text{tr}) + \frac{\rho_t(\text{tr})}{\lambda_{\text{tr}}} \right) (1 - e^{-\lambda_{\text{tr}} T}) \\ + \int_0^T \lambda_{\text{tr}} e^{-\lambda_{\text{tr}} t} \sum_{v' \in V} \mu_{\text{tr}}(v') F(v', T - t) dt, \text{ if } \text{tr} \in \mathbf{T}_M(v), \\ \rho_i(\text{tr}) + \sum_{v' \in V} \mu_{\text{tr}}(v') F(v', T), \text{ if } \text{tr} \in \mathbf{T}_P(v). \end{array} \right.$$

Proof. We first apply a reward preserving transformation to \mathcal{G}_{rew} . For each Markovian transition $\text{tr} \in \mathbf{T}_M(v)$ we introduce an auxiliary state v^{tr} that takes exactly the same transition as tr with the same transient and instantaneous reward values, i. e. we add v^{tr} to the set of states and $(v^{\text{tr}}, \lambda_{\text{tr}}, \mu_{\text{tr}})$ to the set of transitions. Furthermore tr is eliminated from \mathbf{T} and instead a probabilistic transition $(v, \infty, \xi_{v^{\text{tr}}})$, from v to v^{tr} with zero transient and instantaneous reward, is added to \mathbf{T} . Note that v^{tr} must be controlled by the same player as v .

It is not hard to see that the transformation is reward preserving, since the newly added transitions are instantaneous and have zero reward. Moreover, the derived SRG is guaranteed to have at most one Markovian transition for each state by construction, which resembles an MRA (see Sec. 2.3). Consequently, the fixed point characterisation for case $\text{opt} = \text{sup}$ can be directly derived from the one for MRA [31, Lemma 1]. Note that for this case the player roles can be merged since both of the players cooperate. The other case ($\text{opt} = \text{inf}$) can be proven via a slight adaptation of the proof for [31, Lemma 1]. \square

B Proof of Theorem 1

Theorem 1 (Soundness). *Let $\mathcal{M}_{rew} = (\mathcal{M}, \rho_t, \rho_i)$ be an MRA, its high-order operator for maximum reward Ω , and $\text{opt} \cdot \overline{\mathcal{M}}_{rew}^{\mathcal{P}, \text{lab}} = (\overline{\mathcal{M}}^{\mathcal{P}, \text{lab}}, \overline{\rho}_t^{\text{opt}}, \overline{\rho}_i^{\text{opt}})$, $\text{opt} \in \{\text{inf}, \text{sup}\}$, be a game abstraction with rewards, its high-order operators Ω_{inf} , Ω_{sup} . Then:*

$$\text{lfp}(\Omega_{\text{inf}})([v]_{\mathcal{P}}, T) \leq \text{lfp}(\Omega)(v, T) \leq \text{lfp}(\Omega_{\text{sup}})([v]_{\mathcal{P}}, T)$$

for all $v \in V$, and $T \in \mathbb{R}^{\geq 0}$.

Proof. At first we need the definition of Galois connections:

Definition 11 (Galois connection [26]). A Galois connection

$$\langle \mathcal{D}, \preceq \rangle \stackrel{\gamma}{\alpha} \langle \mathcal{D}^{\#}, \preceq^{\#} \rangle$$

consists of a concrete domain \mathcal{D} and an abstract domain \mathcal{D}^\sharp with their respective complete partial orders \preceq and \preceq^\sharp , an abstract function $\alpha : \mathcal{D} \rightarrow \mathcal{D}^\sharp$, and a concretisation function $\gamma : \mathcal{D}^\sharp \rightarrow \mathcal{D}$, such that

$$\alpha(c) \preceq^\sharp a \iff c \preceq \gamma(a)$$

for all elements $a \in \mathcal{D}$ and $c \in \mathcal{D}^\sharp$.

Our soundness argument follows from the following lemma:

Lemma 2 (Semantic Approximation [26]). *Given a Galois connection*

$$\langle \mathcal{D}, \preceq \rangle \xrightarrow[\alpha]{\gamma} \langle \mathcal{D}^\sharp, \preceq^\sharp \rangle,$$

a concrete semantic function $f : \mathcal{D} \rightarrow \mathcal{D}$, and an abstract semantic function $f^\sharp : \mathcal{D}^\sharp \rightarrow \mathcal{D}^\sharp$, such that for all $a \in \mathcal{D}^\sharp$

$$f(\gamma(a)) \preceq \gamma(f^\sharp(a)).$$

Then $\alpha(\text{lfp}_{\preceq}(f)) \preceq^\sharp \text{lfp}_{\preceq^\sharp}(f^\sharp)$ and $\alpha(\text{gfp}_{\preceq}(f)) \preceq^\sharp \text{gfp}_{\preceq^\sharp}(f^\sharp)$.

As we aim to approximate the fixed point of the high order function Ω , our concrete and abstract domains are functions. Namely $\mathcal{D} = S \times \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$, and $\mathcal{D}^\sharp = \mathcal{P} \times \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$. We will use the pointwise order of functions:

$$f \sqsubseteq g \quad :\Leftrightarrow \quad \forall x : f(x) \leq g(x).$$

The Galois connection for the upper bound abstraction is $\langle \mathcal{D}, \sqsubseteq \rangle \xrightarrow[\alpha^u]{\gamma} \langle \mathcal{D}^\sharp, \sqsubseteq \rangle$, where the abstraction function is $\alpha^u(F)(B, t) = \sup_{s \in B} F(s, t)$, and the concretisation function is $\gamma(F^\sharp)(s, t) = F^\sharp([s]_{\mathcal{P}}, t)$. It is easy to see that it fulfils the requirements of Def. 11, but the reader can consult [15] for a detailed proof. Similarly, the Galois connection for the lower bound abstraction is $\langle \mathcal{D}, \sqsupseteq \rangle \xrightarrow[\alpha^l]{\gamma} \langle \mathcal{D}^\sharp, \sqsupseteq \rangle$, where the abstract function is $\alpha^l(F)(B, t) = \inf_{s \in B} F(s, t)$.

We will prove that a lower bound is a sound abstraction, where both players compete against each other. The collaborative case is straightforward, as it is a probabilistic simulation.

Notice that \bar{V}_1 - and \bar{V}_2 - states alternate in the underlying game generated by the abstraction. Also, we are only interested in the values of \bar{V}_1 - states as they are the elements of our abstract domain. Thus, we can rewrite the high order operator Ω_{opt} in Lemma 1 to only use \bar{V}_1 - states. We denote by $tr_{v,a}$ the unique transition in v with label a or the transition (v, ∞, ξ_*) in case $a \notin \text{Lab}(v)$.

For $[v]_{\mathcal{P}} = B \subseteq PS$ we get:

$$\begin{aligned} \Omega(\gamma(F^\sharp))(v, T) &= \sup_{tr \in \mathbf{T}(v)} \rho_i(tr) + \sum_{v' \in V} \mu_{tr}(v') \gamma(F^\sharp)(v', T) \\ &= \sup_{tr \in \mathbf{T}(v)} \rho_i(tr) + \sum_{v' \in V} \mu_{tr}(v') F^\sharp([v']_{\mathcal{P}}, T) \end{aligned}$$

$$\begin{aligned}
&= \sup_{tr \in \mathbf{T}(v)} \rho_i(tr) + \sum_{B' \in \mathcal{P}} \sum_{v' \in B} \mu_{tr}(v') F^\sharp(B', T) \\
&= \sup_{tr \in \mathbf{T}(v)} \rho_i(tr) + \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr}(B') F^\sharp(B', T) \\
&\geq \sup_{tr \in \mathbf{T}(v)} \rho_i^{\inf}(\bar{tr}) + \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr}(B') F^\sharp(B', T) \\
&= \sup_{a \in Lab(v)} \rho_i^{\inf}(\bar{tr}_{v,a}) + \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr_{v,a}}(B') F^\sharp(B', T) \\
&\geq \inf_{w \in B} \sup_{a \in Lab(w)} \rho_i^{\inf}(\bar{tr}_{w,a}) + \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr_{w,a}}(B') F^\sharp(B', T) \\
&\geq \inf_{w \in B} \sup_{a \in Lab(B)} \rho_i^{\inf}(\bar{tr}_{w,a}) + \sum_{B' \in \mathcal{P} \cup \{*\}} \bar{\mu}_{tr_{w,a}}(B') F^\sharp(B', T) \\
&\geq \sup_{a \in Lab(B)} \inf_{w \in B} \rho_i^{\inf}(\bar{tr}_{w,a}) + \sum_{B' \in \mathcal{P} \cup \{*\}} \bar{\mu}_{tr_{w,a}}(B') F^\sharp(B', T) \\
&= \Omega_{\inf}^\sharp(B, T) \\
&= \gamma(\Omega_{\inf}^\sharp)(v, T)
\end{aligned}$$

The case $[v]_{\mathcal{P}} = B \subseteq MS$ is straightforward:

$$\begin{aligned}
\Omega(\gamma(F^\sharp))(v, T) &= (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}})(1 - e^{-\lambda_{tr}T}) \\
&\quad + \int_0^T \lambda_{tr} e^{-\lambda_{tr}t} \sum_{v' \in V} \mu_{tr}(v') \gamma(F^\sharp)(v', T-t) dt \\
&= (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}})(1 - e^{-\lambda_{tr}T}) \\
&\quad + \int_0^T \lambda_{tr} e^{-\lambda_{tr}t} \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr}(B') F^\sharp(B', T-t) dt \\
&= (\rho_i^{\inf}(\bar{tr}_{v,\perp}) + \frac{\rho_t^{\inf}(\bar{tr}_{v,\perp})}{\lambda_{tr_{v,\perp}}})(1 - e^{-\lambda_{tr_{v,\perp}}T}) \\
&\quad + \int_0^T \lambda_{tr_{v,\perp}} e^{-\lambda_{tr_{v,\perp}}t} \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr_{v,\perp}}(B') F^\sharp(B', T-t) dt \\
&\geq \inf_{w \in B} (\rho_i^{\inf}(\bar{tr}_{w,\perp}) + \frac{\rho_t^{\inf}(\bar{tr}_{w,\perp})}{\lambda_{tr_{w,\perp}}})(1 - e^{-\lambda_{tr_{w,\perp}}T}) \\
&\quad + \int_0^T \lambda_{tr_{w,\perp}} e^{-\lambda_{tr_{w,\perp}}t} \sum_{B' \in \mathcal{P}} \bar{\mu}_{tr_{w,\perp}}(B') F^\sharp(B', T-t) dt \\
&= \Omega^\sharp(F^\sharp)(B, T) \\
&= \gamma(\Omega^\sharp(F^\sharp))(v, T)
\end{aligned}$$

□

C Correctness of Equation (4) and (5)

We first recall the problem; Let $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$ be an SRG, $T \geq 0$ be a time bound and $0 \leq \delta \leq T$ be a time step. We want to express sup-opt expected time-bounded reward $\mathcal{R}_{\text{opt}}^{\text{sup}}$ ($\text{opt} = \{\text{inf}, \text{sup}\}$) in terms of its behaviour after taking time step δ . As explained before, w.l.o.g. we consider a state v with only Markovian outgoing transitions, that is $\emptyset \neq \mathbf{T}(v) \subseteq \mathbf{T}_M$. We want to show that the following equations hold:

$$\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) = \text{opt}_v \left(\text{Acc}_{tr}(v, T) + e^{-\lambda_{tr} \cdot \delta} \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T - \delta) \right) \quad (4)$$

$$\begin{aligned} \text{Acc}_{tr}(v, T) &= \left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr} \cdot \delta}) \\ &+ \int_0^\delta \lambda_{tr} e^{-\lambda_{tr} t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \end{aligned} \quad (5)$$

Proof. By Lemma 1, $\mathcal{R}_{\text{opt}}^{\text{sup}}$ is the least fixed point of operator Ω_{opt} given in Eq. (3), we therefore have:

$$\begin{aligned} \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) &= \text{opt}_v \left(\left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr} T}) \right. \\ &\quad \left. + \int_0^T \lambda_{tr} e^{-\lambda_{tr} t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \right) \end{aligned} \quad (9)$$

We use notation $A_{\text{opt}}^{\text{sup}}$ to denote the argument of opt_v in Eq. (9), then it holds that:

$$\mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) = \text{opt}_v A_{\text{opt}}^{\text{sup}}(v, tr, T) \quad (10)$$

Now we split the time horizon into two intervals: the one that spans up to time point δ and the other that spans from time point δ up to time bound T . Then we have:

$$\begin{aligned} A_{\text{opt}}^{\text{sup}}(v, tr, T) &= \left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr} T}) + \int_0^T \lambda_{tr} e^{-\lambda_{tr} t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \\ &= \left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr} T}) + \int_0^\delta \lambda_{tr} e^{-\lambda_{tr} t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \\ &\quad + \int_\delta^T \lambda_{tr} e^{-\lambda_{tr} t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \\ &= \left(\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}} \right) (1 - e^{-\lambda_{tr} T}) + \int_0^\delta \lambda_{tr} e^{-\lambda_{tr} t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) dt \end{aligned}$$

$$\begin{aligned}
& + \int_0^{T-\delta} \lambda_{tr} e^{-\lambda_{tr}(\tau+\delta)} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - \delta - \tau) \, d\tau \\
= & (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}}) (1 - e^{-\lambda_{tr}\delta} + e^{-\lambda_{tr}\delta} - e^{-\lambda_{tr}T}) \\
& + \int_0^{\delta} \lambda_{tr} e^{-\lambda_{tr}t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) \, dt \\
& + e^{-\lambda_{tr}\delta} \int_0^{T-\delta} \lambda_{tr} e^{-\lambda_{tr}\tau} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - \delta - \tau) \, d\tau \\
= & (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}}) (1 - e^{-\lambda_{tr}\delta}) + e^{-\lambda_{tr}\delta} (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}}) (1 - e^{-\lambda_{tr}(T-\delta)}) \\
& + \int_0^{\delta} \lambda_{tr} e^{-\lambda_{tr}t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) \, dt \\
& + e^{-\lambda_{tr}\delta} \int_0^{T-\delta} \lambda_{tr} e^{-\lambda_{tr}\tau} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - \delta - \tau) \, d\tau \\
= & (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}}) (1 - e^{-\lambda_{tr}\delta}) + \int_0^{\delta} \lambda_{tr} e^{-\lambda_{tr}t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) \, dt \\
& + e^{-\lambda_{tr}\delta} \left((\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}}) (1 - e^{-\lambda_{tr}(T-\delta)}) \right. \\
& \left. + \int_0^{T-\delta} \lambda_{tr} e^{-\lambda_{tr}\tau} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - \delta - \tau) \, d\tau \right) \\
= & (\rho_i(tr) + \frac{\rho_t(tr)}{\lambda_{tr}}) (1 - e^{-\lambda_{tr}\delta}) + \int_0^{\delta} \lambda_{tr} e^{-\lambda_{tr}t} \sum_{v' \in V} \mu_{tr}(v') \mathcal{R}_{\text{opt}}^{\text{sup}}(v', T - t) \, dt \\
& + e^{-\lambda_{tr}\delta} \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T - \delta)
\end{aligned}$$

Finally, putting the result into Eq. (10) demonstrates the correctness of Eq. (4) and (5). \square

D Proof of Theorem 2

Theorem 2. *Given an 1MC-SRG $\mathcal{G}_{rew} = (\mathcal{G}, \rho_t, \rho_i)$, time bound $T > 0$ and discretisation step $\delta > 0$ such that $T = k\delta$ for some $k \in \mathbb{N}$. Then for all $v \in V$ we have:*

$$\tilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k) \leq \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) \leq \tilde{\mathcal{R}}_{\delta, \text{opt}}^{\text{sup}}(v, k) + \frac{\bar{\epsilon}T}{2} (\bar{r}_t + \bar{\epsilon}r_i) (1 + \frac{\bar{\epsilon}T}{2}) \delta$$

Proof. Similar to the proof of Lemma 1 we first apply a reward preserving transformation to \mathcal{G}_{rew} . For each Markovian transition $tr \in \mathbf{T}_M(v)$ we introduce an auxiliary state v^{tr} that takes exactly the same transition as tr with the same transient and instantaneous reward values, i. e. we add v^{tr} to the set of states and

$(v^{tr}, \lambda_{tr}, \mu_{tr})$ to the set of transitions. Furthermore tr is eliminated from \mathbf{T} and instead a probabilistic transition $(v, \infty, \xi_{v^{tr}})$, from v to v^{tr} with zero transient and instantaneous reward, is added to \mathbf{T} . Note that v^{tr} must be controlled by the same player as v .

It is not hard to see that the transformation is reward preserving, since the newly added transitions are instantaneous and have zero reward. Moreover, the derived SRG is guaranteed to have at most one Markovian transition for each state by construction, which resembles an MRA (see Sec. 2.3). Consequently, the error bound for case $\text{opt} = \text{sup}$ can be directly derived from the one for MRA [31, Theorem 3]. Note that for this case the player roles can be merged since both of the players cooperate.

Here we elaborate more on the error bound. Following from [31, Lemma C2, C5 and C6], for all $v \in V$ it holds that

$$\tilde{\mathcal{R}}_{\text{opt}}^{\text{sup}}(v, k) \leq \mathcal{R}_{\text{opt}}^{\text{sup}}(v, T) \leq \tilde{\mathcal{R}}_{\text{opt}}^{\text{sup}}(v, k) + \text{Err}(\delta, T, \bar{\mathbf{e}}, \bar{\mathbf{r}}_t, \bar{\mathbf{r}}_i) \quad (11)$$

where $\text{Err}(\delta, T, \bar{\mathbf{e}}, \bar{\mathbf{r}}_t, \bar{\mathbf{r}}_i)$ is the error function given by

$$\text{Err}(\delta, T, \bar{\mathbf{e}}, \bar{\mathbf{r}}_t, \bar{\mathbf{r}}_i) = (\bar{\mathbf{r}}_t + \bar{\mathbf{e}}\bar{\mathbf{r}}_i)T - \left(\frac{\bar{\mathbf{r}}_t}{\bar{\mathbf{e}}} + \bar{\mathbf{r}}_i\right)(1 - e^{-\bar{\mathbf{e}}\delta}) \sum_{i=0}^{k-1} e^{-\bar{\mathbf{e}}i\delta} (1 + \bar{\mathbf{e}}\delta)^i. \quad (12)$$

By straightforward mathematics it can be shown that

$$\lim_{\delta \rightarrow 0} \text{Err}(\delta, T, \bar{\mathbf{e}}, \bar{\mathbf{r}}_t, \bar{\mathbf{r}}_i) = 0.$$

Moreover, the bound given in Theorem 2 is a linear approximation and a safe upper bound of the error function, i. e.

$$\text{Err}(\delta, T, \bar{\mathbf{e}}, \bar{\mathbf{r}}_t, \bar{\mathbf{r}}_i) \leq \frac{\bar{\mathbf{e}}T}{2} (\bar{\mathbf{r}}_t + \bar{\mathbf{e}}\bar{\mathbf{r}}_i) \left(1 + \frac{\bar{\mathbf{e}}T}{2}\right) \delta.$$

It remains to show the error bound for the other case ($\text{opt} = \text{inf}$). The proof for this case is fairly similar to the one for the previous case. It can be done via a slight adaptation of the proof for [31, Theorem 3]. \square