Measuring Adaptability and Reliability of Large Scale Systems

Valentina Castiglioni[1], Michele Loreti[2], and Simone Tini[3]

[1] Reykjavik University, Reykjavik, Iceland
[2] University of Camerino, Camerino, Italy
[3] University of Insubria, Como, Italy

Abstract. In this paper we propose a *metric approach* to the analysis and verification of large scale self-organising collective systems. Typically, these systems consist of a large number of agents that have to interact to coordinate their activities and, at the same time, have to adapt their behaviour to the dynamic surrounding environment. It is then natural to apply a probabilistic modelling to these systems and, thus, to use a metric for the comparison of their behaviours. In detail, we introduce the *population metric*, namely a pseudometric measuring the differences in the probabilistic evolution of two systems with respect to some given requirements. We also use this metric to express the properties of *adaptability* and *reliability* of a system, which allow us to identify potential critical issues with respect to perturbations in its initial conditions. Then we show how we can combine our metric with *statistical inference* techniques to obtain a mathematically tractable analysis of large scale systems. Finally, we exploit *mean-field approximations* to measure the adaptability and reliability of large scale systems.

1 Introduction

The ever increasing complexity of the digital world has moved the focus of researchers to new classes of systems that are characterised by a large number of interacting components, or *agents*. These agents, when considered in isolation from the system, usually show a rather simple behaviour. However, the interaction of a massive number of them enables the desired complex behaviour of the system. Most prominent examples of this class of systems are IoT systems [19], wireless sensors networks, and *self-organising collective systems* [2] (SCS). The latter ones are characterised by a large number of interacting agents that coordinate their activities in a decentralised and often implicit way. Each agent may change its behaviour according to the current status of the other agents in order to make the system reach its objectives. However, the dynamic behaviour of a massive number of agents, and the potential interaction of the system with users and physical phenomena, make these changes subject to uncertainties and unpredictable events. For simplicity, given one agent, we call environment the ensemble of all other agents, users and phenomena that can affect its behaviour. We are interested in ensuring that the system, and thus each agent, is able to adjust its behaviour with respect to the current environmental conditions in order to fulfil its tasks.

Due to the unpredictable behaviour of the environment, it is natural to employ a probabilistic model for the formal specification of the behaviour of these systems. In

particular, we can use a *discrete-time Markov chain* (DTMC) to model the semantics of each agent and, thus, of the system (see, e.g., [20]). When quantitative aspects of systems behaviour are considered, we can use metrics for verification purposes [4, 6, 8, 14, 15, 23], as they allow us to quantify how far the current behaviour of a system is from its intended one. In the literature, several formal frameworks have been proposed for modelling and analysing SCS (see among others [7,13,25,31]). However, to the best of our knowledge, so far there have been no proposals of a metric semantics for SCS.

Hence, our first contribution consists in filling this gap as we propose a *metric approach* to the analysis and verification of large scale SCS. For the specification, we consider the probabilistic model from [20]: agents are identical, at any point in time each agent can be in any of finitely many states, and the evolution of the system proceeds in a clock-synchronous fashion. As each agent can change its state probabilistically, at each time step we obtain a probability distribution over the possible configurations of the system. Therefore, we express system's semantics in terms of its *evolution sequence*, i.e., the sequence of probability distributions so obtained. Then, we introduce the *population metric*, a (time-dependent) pseudometric measuring the differences between the evolution sequences of systems. Besides the disparities in the probabilistic behaviour, the distance considers the ability of systems to fulfil their tasks. The population metric consists of two components: a *metric on global states* and the *Wasserstein metric* [29]. The former considers the *global state* of the system, i.e. the identification of the current state of each agent, and is defined in terms of a (time-dependent) *penalty function* comparing two global states only on the base of the objectives of the system (at a given time). The latter lifts this metric to a metric on distributions over global states, and thus on the evolution sequences of systems. We then exploit the population metric to define the notions of *adaptability* and *reliability* of a system, allowing us to analyse its ability to adjust its behaviour to perturbations in the initial environmental conditions.

As systems are constituted by a large number of agents, a direct evaluation of the population metric is generally unfeasible. Hence, as our second contribution, we provide a *randomised algorithm*, based on statistical inference and on results in [26,28,30], to compute the distance between two systems in time $O(TR\log(R))$, where R is the number of samples used to estimate the evolution sequences of the two systems, and T denotes the number of comparisons performed between the estimated evolution sequences, each one at a different time step. We then show an application of our algorithm to evaluate the adaptability of a system in the balancing scenario from [5].

Nevertheless, when the number of agents increases dramatically, the proposed empirical technique falls short of efficiency. For this reason, as our third contribution, we propose a modification of our randomised algorithm based on *mean-field approximation* and the results in [21]. We express system's evolution in terms of the changes in its *occupancy vector*, whose elements correspond to the fraction of agents in a particular state. In [21] it was proven that when the number of agents goes to infinity, the DTMC capturing the evolution of the occupancy vector of the system can be approximated by a deterministic process. This process corresponds to the deterministic solution of a set of difference equations called mean-field, and it can be exploited to obtain a good estimation of the behaviour of the entire system. As an application example, we evaluate the adaptability of the system in the balancing scenario with an infinite number of agents.

2 Background

As general notational conventions, given a set $\mathcal{X}$ we let $|\mathcal{X}|$ denote its cardinality, and given a vector $\mathbf{x} \in \mathcal{X}^n$, we let $\mathbf{x}_{[i]}$ denote the i-th component of $\mathbf{x}$.

Metrics A *metric* on a set $\mathcal{X}$ is a function $m\colon \mathcal{X} \times \mathcal{X} \to \mathbb{R}^{\geq 0}$ with $m(x_1, x_2) = 0$ iff $x_1 = x_2$, $m(x_1, x_2) = m(x_2, x_1)$, and $m(x_1, x_2) \leq m(x_1, x_3) + m(x_3, x_2)$, for all $x_1, x_2, x_3 \in \mathcal{X}$. We obtain a *pseudometric* by relaxing the first property to $m(x_1, x_2) = 0$ if $x_1 = x_2$. As elsewhere in the literature, as the difference in the two notions is not relevant for our purposes, we will not distinguish between metrics and pseudometrics and use the term *metric* as a general term to denote both. A metric m is *l-bounded* if $m(x_1, x_2) \leq l$ for all $x_1, x_2 \in \mathcal{X}$.

Probability distributions Given a countable set $\mathcal{X}$, a *discrete probability distribution*, henceforth simply *distribution*, over $\mathcal{X}$ is a mapping $\mu\colon \mathcal{X} \to [0, 1]$ such that $\sum_{x \in \mathcal{X}} \mu(x) = 1$. The *support* of μ is the set $\mathtt{supp}(\mu) = \{x \in \mathcal{X} \mid \mu(x) > 0\}$. By $\Delta(\mathcal{X})$ we denote the set of all distributions over $\mathcal{X}$, ranged over by $\mu, \pi, \mu', \ldots$ Given an element $x \in \mathcal{X}$, we let $\boldsymbol{\delta}_x$ denote the *Dirac* (or *point*) *distribution on* x, defined by $\boldsymbol{\delta}_x(x) = 1$ and $\boldsymbol{\delta}_x(y) = 0$ for all $y \neq x$. For a finite set of indexes I, weights $p_i \in (0, 1]$ with $\sum_{i \in I} p_i = 1$ and distributions $\mu_i \in \Delta(\mathcal{X})$ with $i \in I$, the distribution $\sum_{i \in I} p_i \mu_i$ is defined by $(\sum_{i \in I} p_i \mu_i)(x) = \sum_{i \in I} p_i \cdot \mu_i(x)$, for all $x \in \mathcal{X}$.

Discrete-time Markov chains A *discrete-time Markov chain* (DTMC) is a pair $\mathcal{M} = \langle \mathcal{X}, \mathbf{P} \rangle$ consisting in a countable set of *states* $\mathcal{X}$ and a $|\mathcal{X}| \times |\mathcal{X}|$ *one step probability matrix* $\mathbf{P}$ such that $\mathbf{P}_{x,y}$ expresses the probability of reaching state y from state x in one computation step. Equivalently, we can define a Markov chain $\mathcal{M}$ as a *stochastic process* $\{\mathbf{X}_t\}_{t \in \mathbb{N}}$ satisfying the Markov property, i.e., the probability of moving to the next state depends only on the current state and not on the previous ones. Formally,

$$\mathtt{Pr}(\mathbf{X}_{t+1} = x \mid \mathbf{X}_0 = x_0, \ldots, \mathbf{X}_t = x_t) = \mathtt{Pr}(\mathbf{X}_{t+1} = x \mid \mathbf{X}_t = x_t) = \mathbf{P}_{x_t, x} \ .$$

3 A calculus of interacting agents

In this section we present a simple formalism that can be used to describe the behaviour of N identical interacting agents. At any point in time, each agent can be in any of its *finitely many* states and the evolution of the system proceeds in a *clock-synchronous* fashion: at each clock tick each member of the population must either execute one of the transitions that are enabled in its current state, or remain in such a state. The presented formalism is a simple adaptation of the one introduced first in [21] and in [20].

Agent specifications An *agent specification* consists of a triple $\langle \mathcal{S}, \mathsf{Act}, \mathcal{D} \rangle$, where:

- $\mathcal{S}$ is a *finite* non-empty set of *state constants*, ranged over by $A, A', A_1, \ldots$;
- Act is a countable non-empty set of *actions*, ranged over by $a, a', a_1, \ldots$;
- $\mathcal{D}$ is a set of *agent definitions* associating each state constant $A_i \in \mathcal{S}$ with a *summation* of enabled actions:

$$A_i := \sum_{j \in J_i} a_{ij}.A_{ij} \ ,$$

with J_i a finite index set, $A_i, A_{ij} \in \mathcal{S}$, and $a_{ij} \in \mathsf{Act}$, for $i \in I$ and $j \in J_i$.

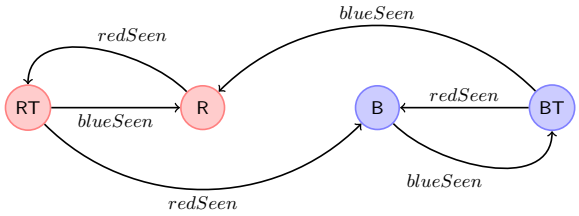

Fig. 1: Behaviour of agents in Example 1.

We let $\mathsf{Act}(A_i) = \{a_{ij} \mid j \in J_i\}$ denote the set of *actions enabled* in state A_i. Notation $\sum_{j \in J_i} a_{ij}.A_{ij}$ can be thought of as the n-ary extension of the standard binary nondeterministic choice operator. We assume that $a_{ij} \neq a_{ij'}$ whenever $j \neq j'$ for $j, j' \in J_i$. Since $\mathcal{S}$ is finite we can assume, without loss of generality, that the state constants are ordered and univocally identified by their index, namely $\mathcal{S} = \{A_1, \ldots, A_{|\mathcal{S}|}\}$.

In order to show how this simple formalism can be used to model the behaviour of a population of agents, and as a running example for an application of our results, we consider the following *balancing scenario* from [5].

Example 1 (Balancing scenario). Le us consider a group of agents that can be either *red* or *blue*. We want to guarantee that the two groups are *balanced* in size, without any centralised control. Each agent can change its colour only by interacting with the other participants in the systems. An agent of a given colour starts a *transitional phase* when it meets another agent of the same colour. In this phase, if another agent with the same colour is met, the agent *changes* its own. The transitional phase is cancelled when an agent with a different colour is found. As typical of SCS, this procedure may never end.

The above behaviour is rendered via state constants B, R, RT and BT. The first two states indicate a *blue* and a *red* agent, respectively; while the last two states describe an agent in a *transition* phase. We let $\mathcal{D}_{RB}$ be the following set of *agent definitions*:

$$\mathsf{B} := blueSeen.\mathsf{BT} \qquad \mathsf{BT} := redSeen.\mathsf{B} + blueSeen.\mathsf{R}$$
$$\mathsf{R} := redSeen.\mathsf{RT} \qquad \mathsf{RT} := redSeen.\mathsf{B} + blueSeen.\mathsf{R}$$

where action $redSeen$ is performed when a *red* agent is met, while $blueSeen$ indicates that the colour of encountered agent is *blue*.

Finally, the agent specification Σ_{RB} is $\Sigma_{RB} = \langle \mathcal{S}_{RB}, \mathsf{Act}_{RB}, \mathcal{D}_{RB} \rangle$, where $\mathcal{S}_{RB} = \{\mathsf{B}, \mathsf{R}, \mathsf{BT}, \mathsf{RT}\}$, $\mathsf{Act}_{RB} = \{blueSeen, redSeen\}$, and $\mathcal{D}_{RB}$ is the one defined above.

An agent specification can be also depicted via a graph as reported in Figure 1. As common in the literature, for $A_i := \sum_{j \in J} a_j.A_j$, we write $A_i \xrightarrow{a_j} A_j$ to denote that if an agent in state A_i performs action a_j then its state changes to A_j, for each $j \in J$.

Global states and occupancy vectors An agent specification $\langle \mathcal{S}, \mathsf{Act}, \mathcal{D} \rangle$ describes the behaviour of a set of agents operating in a system. A *configuration* of a system consists of any representation of the current states of the N agents in it. Two different

levels of abstraction for modelling configurations can be considered. The basic, and more detailed, one is based on *global states*. This consists of a sequence $\mathbf{A}$ in $\mathcal{S}^N$ of the form $(A_{i_1}, \ldots, A_{i_N})$, where A_{i_k} is the current state of agent k, for each $k = 1, \ldots, N$. We refer to N as the *population size*, and we let $\mathbf{A}_{[k]}$ denote the k-th element in $\mathbf{A}$. The state of each agent operating in the system is then univocally identified. Appropriate syntactical shorthands will be used to denote system configurations. For instance, we let $\langle \mathsf{B}[1000], \mathsf{R}[0], \mathsf{BT}[0], \mathsf{RT}[0] \rangle$ denote the configuration with only 1000 blue agents.

However, sometimes we can abstract from the precise state of each single agent while we are interested in considering only the fraction of agents in each state. For this reason, a configuration can be also represented via an *occupancy vector* associating each *state* with the fraction of *agents* in that state. An occupancy vector $\mathbf{o}$ is then an element in $\mathcal{U}^{|\mathcal{S}|} = \{\mathbf{u} \in [0,1]^{|\mathcal{S}|} \mid \sum_{i=1}^{|\mathcal{S}|} \mathbf{u}_{[i]} = 1\}$, namely the *unit simplex* of dimension $|\mathcal{S}|$. We let $\mathrm{OF}_{\mathcal{S},N}$ denote the function mapping each global state in $\mathcal{S}^N$ to the corresponding occupancy vector: $\mathrm{OF}_{\mathcal{S},N} \colon \mathcal{S}^N \to \mathcal{U}^{|\mathcal{S}|}$ is defined for all $\mathbf{A} \in \mathcal{S}^N$ by:

$$\mathrm{OF}_{\mathcal{S},N}(\mathbf{A})_{[i]} = \frac{1}{N} \sum_{k=1}^{N} \mathbf{1}_{[\mathbf{A}_{[k]} = A_i]}$$

where $\mathbf{1}_{[\mathbf{A}_{[k]} = A_i]}$ is 1 if $\mathbf{A}_{[k]} = A_i$, and it is 0 otherwise, $i = 1, \ldots, |\mathcal{S}|$.

We call $\mathrm{OF}_{\mathcal{S},N}$ the *occupancy function* on $\mathcal{S}^N$. We shall drop the subscripts $\mathcal{S}$ and N from $\mathrm{OF}_{\mathcal{S},N}$ when no confusion shall arise.

Probability functions Let us consider an agent specification $\langle \mathcal{S}, \mathsf{Act}, \mathcal{D} \rangle$ and a global state $\mathbf{A} \in \mathcal{S}^N$. The behaviour of $\mathbf{A}$ is modelled via a probabilistic process. Each agent in $\mathbf{A}$ selects probabilistically the next action to perform. The probability of an agent in a given state A_i to perform an action $a \in \mathsf{Act}(A_i)$ in the current time step depends on the distribution of the current states of the other agents, and thus on $\mathrm{OF}(\mathbf{A})$. Clearly, the changes in the distribution of the states of the agents induced by each computation step entail a modification of the (probabilistic) behaviour of each agent at the next step until an equilibrium is eventually reached.

Formally, as a first step in the definition of the dynamic behaviour of agents, we assign a *weight* to each action in a global state $\mathbf{A}$ by means of a *weight function* $\mu_{\mathcal{S}} \colon \mathcal{U}^{|\mathcal{S}|} \times \mathsf{Act} \to \mathbb{R}$. The weight function is built on the *weight expressions* $E \in \mathrm{Exp}$, defined according to the following grammar:

$$E ::= v \mid \mathsf{frc}\,(A) \mid \langle \mathrm{uop} \rangle\, E \mid E\, \langle \mathrm{bop} \rangle\, E \mid (E)$$

Above $v \in [0,1]$ and for each state $A \in \mathcal{S}$, $\mathsf{frc}\,(A)$ denotes the fraction of agents in the system that are currently in state A out of the total number of objects N. Operators $\langle \mathrm{uop} \rangle$ and $\langle \mathrm{bop} \rangle$ are standard arithmetic unary and binary operators.

Each expression is interpreted over a $\mathbf{u} \in \mathcal{U}^{|\mathcal{S}|}$ by means of an expressions interpretation function $[\![\cdot]\!]_{\mathbf{u}} : \mathrm{Exp} \to \mathbb{R}$ defined as follows:

$$[\![v]\!]_{\mathbf{u}} = v \qquad [\![\mathsf{frc}\,(A_i)]\!]_{\mathbf{u}} = \mathbf{u}_{[i]} \qquad [\![\langle \mathrm{uop} \rangle\, E]\!]_{\mathbf{u}} = \langle \mathrm{uop} \rangle\, ([\![E]\!]_{\mathbf{u}})$$
$$[\![E_1\, \langle \mathrm{bop} \rangle\, E_2]\!]_{\mathbf{u}} = ([\![E_1]\!]_{\mathbf{u}})\, \langle \mathrm{bop} \rangle\, ([\![E_2]\!]_{\mathbf{u}}) \qquad [\![(E)]\!]_{\mathbf{u}} = ([\![E]\!]_{\mathbf{u}})$$

Hence, we can associate each action $a \in \mathsf{Act}$ with a weight expression E_a and define the weight of action a with respect to an occupancy vector $\mathbf{o}$ by $\mu_{\mathcal{S}}(\mathbf{o}, a) = [\![E_a]\!]_{\mathbf{o}}$.

We say that a state $A \in \mathcal{S}$ is *probabilistic* in an occupancy vector $\mathbf{o}$ if

$$0 \leq \sum_{a \in \mathsf{Act}\,:\, A \xrightarrow{a} A'} \mu_{\mathcal{S}}(\mathbf{o}, a) \leq 1 \ ,$$

i.e., if the total weight assigned by $\mu_{\mathcal{S}}$ to the actions enabled for A with respect to $\mathbf{o}$ is non-negative and smaller than 1. Then, we say that $\mu_{\mathcal{S}}$ is a *probability function* if all the states in $\mathcal{S}$ are probabilistic in $\mathbf{o}$, for any $\mathbf{o} \in \mathcal{U}^{|\mathcal{S}|}$. In the remainder of the paper we shall consider only functions $\mu_{\mathcal{S}}$ that are probability functions.

We shall drop the subscript $\mathcal{S}$ from $\mu_{\mathcal{S}}$ when no confusion shall arise.

Example 2. For the *red blue* balancing scenario of Example 1, given an occupancy vector $\mathbf{o}$, we can consider the probability function $\mu_{\mathcal{S}_{RB}}$ defined as follows:

$$\begin{aligned} \mu_{\mathcal{S}_{RB}}(\mathbf{o}, blueSeen) &= [\![\alpha \cdot (\mathsf{frc}\,(\mathsf{B}) + \mathsf{frc}\,(\mathsf{BT}))]\!]_{\mathbf{o}} \\ \mu_{\mathcal{S}_{RB}}(\mathbf{o}, redSeen) &= [\![\alpha \cdot (\mathsf{frc}\,(\mathsf{R}) + \mathsf{frc}\,(\mathsf{RT}))]\!]_{\mathbf{o}} \ , \end{aligned}$$

with α a parameter in $[0, 1]$ expressing the probability of an agent *to see another agent*.

Systems semantics A *system specification* Σ is a tuple of the form $\langle \mathcal{S}, \mathsf{Act}, \mathcal{D}, \mu_{\mathcal{S}} \rangle$, where $\langle \mathcal{S}, \mathsf{Act}, \mathcal{D} \rangle$ is an agent specification and $\mu_{\mathcal{S}}$ is a *probability* function. We let Σ^N denote a system Σ composed by N agents.

As outlined above, the behaviour of a *global state* $\mathbf{A} \in \mathcal{S}^N$ can be described in terms of the (probabilistic) evolution of the states of the agents. To this end, let us focus on the behaviour of a single agent in its interaction with the others. Let $\mathbf{o}$ be an occupancy vector in $\mathcal{U}^{|\mathcal{S}|}$. The *agent transition function* $\mathbf{K}$ is used to express the probability under $\mathbf{o}$ of an agent in state A_i to change its state to A_j after one computation step. Formally, we define $\mathbf{K} \colon \mathcal{U}^{|\mathcal{S}|} \times \mathcal{S} \times \mathcal{S} \to [0, 1]$ as follows:

$$\mathbf{K}(\mathbf{u})_{A_i, A_j} = \begin{cases} \displaystyle\sum_{a \in I(A_i)} \mu_{\mathcal{S}}(\mathbf{u}, a) & \text{if } A_i \neq A_j \\ \displaystyle 1 - \sum_{a \in I(A_i)} \mu_{\mathcal{S}}(\mathbf{u}, a) & \text{if } A_i = A_j \end{cases} \ ,$$

where $I(A_i) = \{a \in \mathsf{Act} \mid \exists\, A_j \in \mathcal{S} \colon A_i \xrightarrow{a} A_j \wedge A_i \neq A_j\}$.

As we are assuming that all the states in $\mathcal{S}$ are probabilistic in any occupancy vector $\mathbf{o}$, we can interpret $\mathbf{K}(\mathbf{o})$ as a one step $|\mathcal{S}| \times |\mathcal{S}|$ transition probability matrix, and call it the *agent transition matrix*. Then, the second case in the definition of $\mathbf{K}(\mathbf{o})$ expresses the probability of an agent in state A_i to remain in that state after the clock tick.

Starting from $\mathbf{K}$, we can define the probabilistic behaviour of global states via the *global state transition matrix* $\mathbf{P}^{(N)} \colon \mathcal{S}^N \times \mathcal{S}^N \to [0, 1]$ defined for all $\mathbf{A}, \mathbf{A'} \in \mathcal{S}^N$ as

$$\mathbf{P}^{(N)}_{\mathbf{A}, \mathbf{A}'} = \prod_{k=1}^{N} \mathbf{K}(\mathrm{OF}(\mathbf{A}))_{\mathbf{A}_{[k]}, \mathbf{A}'_{[k]}} \tag{1}$$

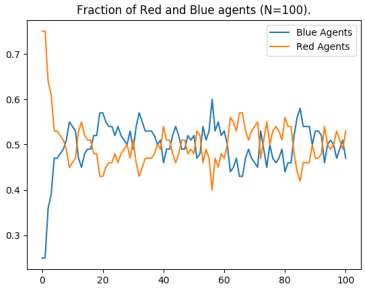

Fig. 2: Single simulation run of red-blue scenario.

Since all states in $\mathcal{S}$ are probabilistic in any occupancy vector $\mathbf{o}$, the $|\mathcal{S}^N| \times |\mathcal{S}^N|$ matrix $\mathbf{P}^{(N)}$ corresponds to the one step transition probability matrix of the (finite state) DTMC modelling a possible single step of the system as result of the parallel execution of a single step of each of the N agents In detail, given a *global state* $\mathbf{A} \in \mathcal{S}^N$, we define the *evolution sequence* of $\mathbf{A}$ as the sequence $\pi_{\mathbf{A},0}, \ldots, \pi_{\mathbf{A},t}, \ldots$ of probability distributions over $\mathcal{S}^N$ such that $\pi_{\mathbf{A},0} = \delta_{\mathbf{A}}$ and $\pi_{\mathbf{A},t+1} = \pi_{\mathbf{A},t}\mathbf{P}^{(N)}$, for each $t > 0$. In this case, we let $\mathbf{X}_{\mathbf{A}}^{(N)}(t)$ denote the Markov chain with transition probability matrix $\mathbf{P}^{(N)}$ as above and $\mathbf{X}_{\mathbf{A}}^{(N)}(0) = \mathbf{A}$, i.e. with initial probability distribution $\boldsymbol{\delta}_{\mathbf{A}}$. From the Markov chain $\{\mathbf{X}_{\mathbf{A}}^{(N)}(t)\}_{t\in\mathbb{N}}$ we can obtain the *occupancy DTMC* $\{\mathbf{O}_{\mathbf{A}}^{(N)}(t)\}_{t\in\mathbb{N}}$, namely the Markov chain modelling the evolution in time of the occupancy vector of the system. Intuitively, and with a slight abuse of notation, $\mathbf{O}_{\mathbf{A}}^{(N)}(t)$ corresponds to $\mathrm{OF}(\mathbf{X}_{\mathbf{A}}^{(N)}(t))$ and its transition probability is defined by:

$$\Pr\{\mathbf{O}_{\mathbf{A}}^{(N)}(t+1) = \mathbf{o}' \mid \mathbf{O}_{\mathbf{A}}^{(N)}(t) = \mathrm{OF}(\mathbf{A})\} = \sum_{\mathbf{A}' :\, \mathrm{OF}(\mathbf{A}')=\mathbf{o}'} \mathbf{P}_{\mathbf{A},\mathbf{A}'}^{(N)} \qquad (2)$$

Notice that $\mathbf{O}^{(N)}$ is well-defined: if $\mathrm{OF}(\mathbf{A}) = \mathrm{OF}(\mathbf{A}'')$, then $\mathbf{A}$ and $\mathbf{A}''$ are two permutations of the same local states. Hence, for all $\mathbf{A}'$ we have $\mathbf{P}_{\mathbf{A},\mathbf{A}'}^{(N)} = \mathbf{P}_{\mathbf{A}'',\mathbf{A}'}^{(N)}$.

Example 3. We can use the presented semantics to generate the stochastic process of our system of Example 1. The result of the simulation of a single computational path is reported in Figure 2. There we consider an initial state composed by $N = 100$ agents out of which 75 are in the state R and 25 are in the state B. We consider $\alpha = 0.5$ (see Example 2). We can observe that, after few steps, the system is able to evolve to balanced configurations. (Data have been obtained by using Spear Tool available at `https://github.com/quasylab/spear`.)

4 Measuring the adaptability and reliability of systems

In this section we provide the tools necessary to study and analyse the differences in the dynamics of the systems described in the previous section. More precisely, we introduce

a *metric* over systems, called the *population metric* $\mathfrak{m}$, that quantifies the dissimilarities over the evolution sequences of the initial configurations of the systems. The definition of the population metric will then make use of a (*time-dependent*) *metric on global states*, measuring the differences of each pair of global states with respect to some parameters of interest, and of the *Wasserstein metric*, allowing for lifting the former metric to probability distributions over global states. We then exploit the population metric to define the notions of *adaptability* and *reliability* of a system, expressing how well the system can adjust its behaviour with respect to changes in the initial conditions.

4.1 A metric over systems: the population metric

A metric over global states We start by proposing a *metric over global states*. This metric expresses the distance between two given configurations by comparing them with respect to some parameters representing the ideal, optimal behaviour of a system. To this end, we introduce a *penalty function* $\rho\colon \mathcal{S}^N \to [0,1]$, namely a function assigning to each global state $\mathbf{A} \in \mathcal{S}^N$ a penalty in $[0,1]$ expressing how far the values of the parameters on interest in $\mathbf{A}$ are from their desired ones. Intuitively, $\rho(\mathbf{A}) = 0$ if $\mathbf{A}$ respects all the parameters. Since some parameters can be time-dependent, so is the penalty function: at any time step t, the t-penalty function ρ_t compares the global states with respect to the values of the parameters expected at time t.

Example 4. In Example 1 one is interested in verifying that *blue* and *red* agents are balanced. It is then natural to define a penalty function on $\mathcal{S}^N_{RB}$ such that: (i) the penalty is zero if the number of *red* agents is equal to the number of the *blue* ones; (ii) the penalty is higher for less balanced systems, (iii) the penalty does not depend on time, (iv) the penalty is a value in $[0,1]$. Given the occupancy function OF of $\mathcal{S}^N_{RB}$, we set

$$\rho_t(\mathbf{A}) = \left|\big(\mathrm{OF}(\mathbf{A})[\mathsf{B}] + \mathrm{OF}(\mathbf{A})[\mathsf{BT}]\big) - \big(\mathrm{OF}(\mathbf{A})[\mathsf{R}] + \mathrm{OF}(\mathbf{A})[\mathsf{RT}]\big)\right| \ ,$$

for all $t \in \mathbb{N}$, where, with a slight abuse of notation, we let $\mathrm{OF}(\mathbf{A})[\mathsf{A}]$ denote the component of the occupancy vector corresponding to state A, for $\mathsf{A} \in \mathcal{S}_{RB}$.

The (*timed*) *metric over global states* is then defined as the difference between the values assigned to them by the penalty function.

Definition 1 (Metric over global spaces). *For any time step t, let $\rho_t\colon \mathcal{S}^N \to [0,1]$ be the t-penalty function on $\mathcal{S}^N$. The* timed metric over global states *in $\mathcal{S}^N$, $m_{\rho,t}\colon \mathcal{S}^N \times \mathcal{S}^N \to [0,1]$, is defined, for all global states $\mathbf{A}_1, \mathbf{A}_2 \in \mathcal{S}^N$, by*

$$m_{\rho,t}(\mathbf{A}_1, \mathbf{A}_2) = |\rho_t(\mathbf{A}_1) - \rho_t(\mathbf{A}_2)| \ .$$

When no confusion shall arise, we shall drop the ρ, t subscript. It is not hard to see that for all $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3 \in \mathcal{S}^N$ we have that (i) $0 \leq m(\mathbf{A}_1, \mathbf{A}_2) \leq 1$, (ii) $m(\mathbf{A}_1, \mathbf{A}_1) = 0$, (iii) $m(\mathbf{A}_1, \mathbf{A}_2) = m(\mathbf{A}_2, \mathbf{A}_1)$, and (iv) $m(\mathbf{A}_1, \mathbf{A}_2) \leq m(\mathbf{A}_1, \mathbf{A}_3) + m(\mathbf{A}_3, \mathbf{A}_2)$, thus ensuring that m is well defined.

Proposition 1. *Function m is a 1-bounded pseudometric over $\mathcal{S}^N$.*

We remark that the use of the penalty functions allows us to define the distance between two global states, which are elements in $\mathcal{S}^N$, in terms of a distance on $\mathbb{R}$. As we will discuss in Section 5, this feature significantly lowers the complexity of the evaluation of the population metric. Moreover, thanks to the penalty function, the metric on global states m could be directly generalised to a metric over $\mathcal{S}^{N_1} \times \mathcal{S}^{N_2}$ with $N_1 \neq N_2$, i.e., a metric over global states of different dimensions. However, to simplify the presentation in the upcoming sections and ease of notation, we preferred to consider only the case of global states of the same dimension.

Lifting m to distributions The second step in the definition of the population metric consists in lifting m to a metric over probability distributions over global states. In the literature, we can find a wealth of notions of distances over probability measures (see [22] for a survey). For our purposes, the most suitable one is the *Wasserstein metric* [29]. This metric has been applied in several different contexts, from image processing to economics, and it is known under different names, accordingly. Among the most prominent ones, we recall its use in optimal transport problems [30], where it is called the *Earth mover's distance*, and in the definition of bisimulation metrics (see among others [6,10,14,15]), where it is usually referred to as the *Kantorovich metric* [18]. More recently, the Wasserstein metric has been successfully implemented in privacy problems [9, 11] and has found a wealth of applications in machine learning for improving the stability of generative adversarial networks training [3, 17, 27].

Definition 2 (Wasserstein metric). *For any two probability distributions μ, ν on $\mathcal{S}^N$, the* Wasserstein lifting *of m to a distance between μ and ν is defined by*

$$\mathbf{W}(m)(\mu,\nu) = \min_{\mathfrak{w} \in \mathfrak{W}(\mu,\nu)} \sum_{\mathbf{A},\mathbf{A}' \in \mathcal{S}^N} \mathfrak{w}(\mathbf{A},\mathbf{A}') \cdot m(\mathbf{A},\mathbf{A}')$$

where $\mathfrak{W}(\mu,\nu)$ is the set of the couplings for μ and ν*, i.e., the set of joint distributions $\mathfrak{w}$ over the product space $\mathcal{S}^N \times \mathcal{S}^N$ having μ and ν as left and right marginal respectively, i.e., $\sum_{\mathbf{A}' \in \mathcal{S}^N} \mathfrak{w}(\mathbf{A},\mathbf{A}') = \mu(\mathbf{A})$ and $\sum_{\mathbf{A}' \in \mathcal{S}^N} \mathfrak{w}(\mathbf{A}',\mathbf{A}) = \nu(\mathbf{A})$, for all $\mathbf{A} \in \mathcal{S}^N$.*

Thus, the inifimum in Definition 2 is always achieved and it is, in fact, a minimum.

Due to the convexity of the Wasserstein lifting and in light of Proposition 1, we are guaranteed that $\mathbf{W}(m)$ is a well-defined pseudometric.

Proposition 2. *Function $\mathbf{W}(m)$ is a 1-bounded pseudometric over $\Delta(\mathcal{S}^N)$.*

The population metric We are now ready to lift the distance $\mathbf{W}(m)$ to a distance over systems, which we call the *population metric*. This is obtained from the comparison of the evolution sequences of the initial configurations of two systems. To favour computational tractability, we will *not* compare *all* the probability distributions in the evolution sequences, but *only* those that are reached at certain time steps, called the *observation times* (OT). To perform such a comparison, we propose a sort of *weighted infinity norm* of the tuple of the Wasserstein distances between the distributions in the evolution

sequences. More precisely, we consider a non-increasing function $\lambda\colon \mathrm{OT} \to (0,1]$ allowing us to express how much the distance at time t, namely $\mathbf{W}(m_{\rho,t})(\pi_{\mathbf{A}_1,t}, \pi_{\mathbf{A}_2,t})$, affects the overall distance between global states $\mathbf{A}_1$ and $\mathbf{A}_2$. Following the terminology used for behavioural metrics [1, 16], we refer to λ as to the *discount function*, and to $\lambda(t)$ as to the *discount factor at time* t.

Definition 3 (Population metric). *Assume a finite set* OT *of observation times, a penalty function* $\rho_t\colon \mathcal{S}^N \to [0,1]$ *for each* $t \in \mathrm{OT}$ *and a discount function* $\lambda\colon \mathrm{OT} \to (0,1]$. *The* λ-population metric *over* OT, $\mathfrak{m}^{\lambda}_{\mathrm{OT}}\colon \mathcal{S}^N \times \mathcal{S}^N \to [0,1]$ *is defined, for all* $\mathbf{A}_1, \mathbf{A}_2 \in \mathcal{S}^N$, *by*

$$\mathfrak{m}^{\lambda}_{\mathrm{OT}}(\mathbf{A}_1, \mathbf{A}_2) = \sup_{t \in \mathrm{OT}} \lambda(t) \cdot \mathbf{W}(m_{\rho,t})\Big(\pi_{\mathbf{A}_1,t}, \pi_{\mathbf{A}_2,t}\Big).$$

We remark that although both $m_{\rho,t}$ and $\mathfrak{m}^{\lambda}_{\mathrm{OT}}$ are formally defined as metrics over global states in $\mathcal{S}^N$, their expressive power is totally different. On one hand, $m_{\rho,t}$ compares the global states $\mathbf{A}_1$ and $\mathbf{A_2}$ seen as *static* objects. In fact, $m_{\rho,t}(\mathbf{A}_1, \mathbf{A}_2)$ is based on the evaluation of ρ_t on the current states of $\mathbf{A}_1$ and $\mathbf{A}_2$. On the other hand, $\mathfrak{m}^{\lambda}_{\mathrm{OT}}$ compares $\mathbf{A}_1$ and $\mathbf{A}_2$ as *dynamic* objects. The distance $\mathfrak{m}^{\lambda}_{\mathrm{OT}}(\mathbf{A}_1, \mathbf{A}_2)$ is in fact evaluated by considering the evolution sequences of the two global states.

The following proposition is a direct consequence of Proposition 2.

Proposition 3. *Function* $\mathfrak{m}^{\lambda}_{\mathrm{OT}}$ *is a 1-bounded pseudometric over* $\mathcal{S}^N$.

Notice that if λ is a *strictly* decreasing function, then it specifies how much the distance of *future events* is mitigated and, moreover, it guarantees that to obtain upper bounds on the population metric only a *finite* number of observations is needed.

Furthermore, as for the metric m, the population metric could be easily generalised to a metric over systems composed by a different number of agents. For consistency with the choice made for Definition 1, we considered only the simple case of systems having the same number of agents. We leave as future work an in-depth analysis of the population metric over $\mathcal{S}^{N_1} \times \mathcal{S}^{N_2}$, for $N_1 \neq N_2$. In particular, our metric could be used to measure the differences between the systems Σ^N and Σ^{N+1}, namely to analyse the impact of the addition (or removal) of a single agent from the system.

4.2 System Adaptability and Reliability

We now apply the population metric to verify whether a system is able to *adjust* its behaviour to changes in the initial conditions. For instance, we are interested in verifying whether a small perturbation to the initial distribution of states produces a controlled perturbation to the dynamics of the system. We express this kind of properties in terms of the notions of *adaptability* and *reliability* of a system. The main difference between these notions is in how time is taken into account.

The notion of adaptability imposes some constraints on the *long term* behaviour of systems, disregarding their possible initial dissimilarities. Given the thresholds $\eta_1, \eta_2 \in [0,1)$ and an observable time $\tilde{t}$, we say that a system Σ^N is adaptable around a *global state* $\mathbf{A}_0$ if whenever the computation is started from a global state $\mathbf{A}'$ that differs from $\mathbf{A}_0$ for at most η_1, then we are guaranteed that the distance that we can observe between the evolution sequences of the two systems after time $\tilde{t}$ is bounded by η_2.

1: **function** ESTIMATE($\mathbf{A}, T, R$)
2: $\quad\forall t : (0 \leq t \leq T) : \mathcal{O}_t \leftarrow \emptyset$
3: $\quad counter \leftarrow 0$
4: $\quad$**while** $counter \leq R$ **do**
5: $\quad\quad(\mathbf{A}_0, \ldots, \mathbf{A}_T) \leftarrow$ SIMULATE($\mathbf{A}, T$)
6: $\quad\quad\forall t : \mathcal{O}_t \leftarrow \mathcal{O}_t, \mathbf{A}_t$
7: $\quad\quad counter \leftarrow counter + 1$
8: $\quad$**end while**
9: $\quad$**return** $\mathcal{O}_0, \ldots, \mathcal{O}_T$
10: **end function**

Fig. 3: Function giving R samples of the evolution sequence of $\mathbf{A}$ with time horizon T.

Definition 4 (Adaptability). *Consider a system specification Σ over N agents. Let $\tilde{t} \in \mathrm{OT}$ and $\eta_1, \eta_2 \in [0, 1)$. We say that Σ^N is* $(\tilde{t}, \eta_1, \eta_2)$-adaptable *around $\mathbf{A}_0$ if*

$$\forall\, \mathbf{A}' \in \mathcal{S}^N \text{ with } m_{\rho,0}(\mathbf{A}_0, \mathbf{A}') \leq \eta_1 \text{ it holds } \mathfrak{m}^{\lambda}_{\{t \in \mathrm{OT} \mid t \geq \tilde{t}\}}(\mathbf{A}_0, \mathbf{A}') \leq \eta_2 \ .$$

Roughly, Σ^N is adaptable if whenever the starting conditions are changed, then Σ^N is able to return *close* to the original behaviour within the time threshold $\tilde{t}$. The notion of reliability strengthens that of adaptability by bounding the distance on the evolution sequences of systems from the beginning. A system is reliable if it ensures that small variations in the initial conditions cause only bounded variations in its evolution.

Definition 5 (Reliability). *Consider a system specification Σ over N agents. Let $\eta_1, \eta_2 \in [0, 1)$. We say that Σ^N is* (η_1, η_2)-reliable *around state $\mathbf{A}_0$ if*

$$\forall\, \mathbf{A}' \in \mathcal{S}^N \text{ with } m_{\rho,0}(\mathbf{A}_0, \mathbf{A}') \leq \eta_1 \text{ it holds } \mathfrak{m}^{\lambda}_{\mathrm{OT}}(\mathbf{A}_0, \mathbf{A}') \leq \eta_2 \ .$$

5 Statistical estimation of adaptability and reliability

Given two evolution sequences one could explicitly compute the distance among them. However, this approach is unfeasible when the number of agents in the involved states increases. For this reason, in this section, we discuss an empirical technique that given two global states $\mathbf{A}_1$ and $\mathbf{A}_2$ allows us to generate their evolution sequences and then evaluate the distance between them. The same technique will be used to verify the *adaptability* and *reliability* of a system *around* a given global state $\mathbf{A}$.

5.1 Computing empirical evolution sequences

To compute the *empirical evolution sequence* of a global state $\mathbf{A}$ the function ESTIMATE in Figure 3 can be used. This function invokes R times function SIMULATE, i.e., any simulation algorithm sampling a sequence of global states $\mathbf{A}_0, \ldots, \mathbf{A}_T$, modelling T steps of a computation from $\mathbf{A} = \mathbf{A}_0$. Then, a sequence of observations $\mathcal{O}_0, \ldots, \mathcal{O}_T$ is computed, where each $\mathcal{O}_t$ is the tuple $\mathbf{A}^1_t, \ldots, \mathbf{A}^R_t$ of global states observed at time t in each of the R sampled computations. Each $\mathcal{O}_t$ can be used to estimate the real

function DISTANCE($\mathbf{A}_1, \mathbf{A}_2, \rho, \lambda, \mathrm{OT}, R, \ell$)
  $T \leftarrow \max_{\mathrm{OT}}$
  $\mathcal{O}^1_1, \ldots, \mathcal{O}^1_T \leftarrow$ ESTIMATE($\mathbf{A}_1, T, R$)
  $\mathcal{O}^2_1, \ldots, \mathcal{O}^2_T \leftarrow$ ESTIMATE($\mathbf{A}_2, T, \ell R$)
  $m \leftarrow \infty$
  for all $t \in \mathrm{OT}$ **do**
    $m_t \leftarrow$ COMPUTEW($\mathcal{O}^1_t, \mathcal{O}^2_t, \rho_t$)
    $m \leftarrow \min\{m, \lambda(t) \cdot m_t\}$
  end for
  return m
end function

function COMPUTEW($\mathcal{O}_1, \mathcal{O}_2, \rho$)
  $(\mathbf{A}^1_1, \ldots, \mathbf{A}^R_1) \leftarrow \mathcal{O}_1$
  $(\mathbf{A}^1_2, \ldots, \mathbf{A}^{\ell R}_2) \leftarrow \mathcal{O}_2$
  $\forall l : (1 \leq l \leq R) : \omega_l \leftarrow \rho(\mathbf{A}^l_1)$
  $\forall h : (1 \leq h \leq \ell R) : \nu_h \leftarrow \rho(\mathbf{A}^h_2)$
  re index $\{\omega_l\}$ such that $\omega_l \leq \omega_{l+1}$
  re index $\{\nu_h\}$ such that $\nu_h \leq \nu_{h+1}$
  return $\frac{1}{\ell R} \sum_{h=1}^{\ell R} |\omega_{\lceil \frac{h}{\ell} \rceil} - \nu_h|$
end function

Fig. 4: Functions used to estimate the population metric on systems.

probability distribution $\pi_{\mathbf{A},t}$. For any t, with $0 \leq t \leq T$, we let $\hat{\pi}^R_{\mathbf{A},t}$ be the probability distribution such that, for any $\mathbf{A}' \subseteq \mathcal{S}^N$, $\hat{\pi}^R_{\mathbf{A},t}(\mathbf{A}') = \frac{|\mathbf{A}'|_{\mathcal{O}_t}}{R}$ where $|\mathbf{A}'|_{\mathcal{O}_t}$ denotes the number of occurrences of $\mathbf{A}'$ in $\mathcal{O}_t$. Since the R samples are i.i.d, we can apply to them the weak law of large numbers, obtaining that $\hat{\pi}^R_{\mathbf{A},t}$ converges weakly to $\pi_{\mathbf{A},t}$:

$$\lim_{R \to \infty} \hat{\pi}^R_{\mathbf{A},t} = \pi_{\mathbf{A},t} \ . \tag{3}$$

5.2 Computing distance between two configurations

We now evaluate the distance between the evolution sequences of two global states $\mathbf{A}_1$ and $\mathbf{A}_2$ by exploiting the independent samples collected via function ESTIMATE. To this end, we apply the approach of [26] to estimate the Wasserstein distance between two (unknown) distributions. In order to approximate the distance $\mathbf{W}(m_{\rho,t})(\pi_{\mathbf{A}_1,t}, \pi_{\mathbf{A}_2,t})$, for any $0 \leq t \leq T$, we consider R independent samples $\mathcal{O}_{1,t} = \{\mathbf{A}^1_1, \ldots, \mathbf{A}^R_1\}$ taken from $\pi_{\mathbf{A}_1,t}$ and ℓR samples $\mathcal{O}_{2,t} = \{\mathbf{A}^1_2, \ldots, \mathbf{A}^{\ell R}_2\}$ taken from $\pi_{\mathbf{A}_2,t}$. We then apply the t-penalty function ρ_t to them, obtaining the two sequences of values $\{\omega_l = \rho_t(\mathbf{A}^l_1) \mid 1 \leq l \leq R\}$ and $\{\nu_h = \rho_t(\mathbf{A}^h_2) \mid 1 \leq h \leq \ell R\}$. Without loss of generality, we can assume that $\omega_l \leq \omega_{l+1}$ and $\nu_h \leq \nu_{h+1}$, i.e., the two sequences are ordered. In light of the next theorem, which is based on results from [26, 28, 30], we have that $\mathbf{W}(m_{\rho,t})(\pi_{\mathbf{A}_1,t}, \pi_{\mathbf{A}_2,t})$ can be approximated by $\frac{1}{\ell R} \sum_{h=1}^{\ell R} |\omega_{\lceil \frac{h}{\ell} \rceil} - \nu_h|$, and that the latter value converges to the real distance when $R \to \infty$.

Theorem 1. *Let $\pi_{\mathbf{A}_1,t}, \pi_{\mathbf{A}_2,t} \in \Delta(\mathcal{S}^N)$ be unkonwn. Let $\{\mathbf{A}^1_1, \ldots, \mathbf{A}^R_1\}$ be independent samples taken from $\pi_{\mathbf{A}_1,t}$, and $\{\mathbf{A}^1_2, \ldots, \mathbf{A}^{\ell R}_2\}$ independent samples taken from $\pi_{\mathbf{A}_2,t}$. Let $\{\omega_l = \rho_t(\mathbf{A}^l_1) \mid 1 \leq l \leq R\}$ and $\{\nu_h = \rho_t(\mathbf{A}^h_2) \mid 1 \leq h \leq \ell R\}$ be the ordered sequences obtained from the samples and the t-penalty function. Then*

$$\mathbf{W}(m_{\rho,t})(\pi_{\mathbf{A}_1,t}, \pi_{\mathbf{A}_2,t}) \overset{a.s.}{=} \lim_{R \to \infty} \frac{1}{\ell R} \sum_{h=1}^{\ell R} |\omega_{\lceil \frac{h}{\ell} \rceil} - \nu_h| \ .$$

Functions DISTANCE and COMPUTEW in Figure 4 realise the procedure outlined above. The former takes as input the two global states to compare, the penalty function (seen as the sequence of the t-penalty functions), the discount function λ, the

bounded set OT of observation times, and the parameters R and ℓ used to obtain the samplings of the computations. It calls function ESTIMATE to collect the samples $\mathcal{O}_t$ of possible computations during the observation period $[0, \max_{\mathrm{OT}}]$. Then, for each observation time $t \in \mathrm{OT}$, the distance at time t is computed via the function COMPUTEW$(\mathcal{O}_{1,t}, \mathcal{O}_{2,t}, \rho_t)$. Since the penalty function allows us to evaluate the Wasserstein distance on $\mathbb{R}$, the complexity of function COMPUTEW is $O(\ell R \log(\ell R))$ due to the sorting of $\{\nu_h \mid h \in [1, \ldots, \ell R]\}$ (cf. [26]).

Example 5. We use function COMPUTEW to evaluate the impact of a perturbation in the initial configuration of the system in Example 3. There, we have considered an initial state of the form $\mathbf{A}_1 = \langle \mathsf{B}[25], \mathsf{R}[75], \mathsf{BT}[0], \mathsf{RT}[0] \rangle$. Consider the new initial configuration $\mathbf{A}_2 = \langle \mathsf{B}[0], \mathsf{R}[100], \mathsf{BT}[0], \mathsf{RT}[0] \rangle$. Figure 5a shows the variation in time of the distance between $\mathbf{A}_1$ and $\mathbf{A}_2$, for $R = 100$ and $\ell = 10$: after around 10 steps, the two systems cannot be distinguished. Therefore, we can infer that, after around 10 steps, the actual distance between the two systems will be bounded by the approximation error $e_{\mathbf{W}} = |\, \mathbf{W}(m_{\rho,10})(\hat{\pi}^{100}_{\mathbf{A}_1,10}, \hat{\pi}^{1000}_{\mathbf{A}_2,10}) - \mathbf{W}(m_{\rho,10})(\pi_{\mathbf{A}_1,10}, \pi_{\mathbf{A}_2,10})|$. We refer the interested reader to [24, Corollary 3.5, Equation (3.10)] for an estimation of $e_{\mathbf{W}}$.

5.3 Estimating adaptability and reliability

We can use a randomised algorithm to verify the adaptability and reliability of a system. Given a global state $\mathbf{A}$, a set OT of observation times and a given threshold $\eta_1 \geq 0$, we can sample M variations $\{\mathbf{A}_1, \ldots, \mathbf{A}_M\}$ of $\mathbf{A}$ such that, for any $i = 1, \ldots, M$, $m_{\rho,0}(\mathbf{A}, \mathbf{A}_i) \leq \eta_1$. Then, for each sample we can estimate the distance between $\mathbf{A}$ and $\mathbf{A}_i$ at the different time steps in OT, namely $\mathfrak{m}^{\lambda}_{\{t \in \mathrm{OT} \mid t \geq \tilde{t}\}}(\mathbf{A}, \mathbf{A}_i)$ for any $\tilde{t} \in \mathrm{OT}$. Finally, for each $\tilde{t} \in OT$, we let $L_{\tilde{t}} = \max_i \{\mathfrak{m}^{\lambda}_{\{t \in \mathrm{OT} \mid t \geq \tilde{t}\}}(\mathbf{A}, \mathbf{A}_i)\}$. We can observe that, for the chosen η_1, each $L_{\tilde{t}}$ gives us a lower bound to the $\tilde{t}$-adaptability of our system. Similarly, for $t_{\min} = \min_{\mathrm{OT}} t$, $L_{t_{\min}}$ gives a lower bound for its reliability.

Example 6. Figure 5b shows the evaluation of $L_{\tilde{t}}$ for configuration $\mathbf{A}_1$ from Example 3 with parameters $M = 20$ and $\eta_1 = 0.25$. Observe that the initial perturbation is not amplified and after 15 steps it is less than 0.05. Hence, our system of Example 3 is $(15, 0.25, \eta_2)$-adaptable around $\mathbf{A}_1$ for any $\eta_2 \geq e_{\mathbf{W}}$.

6 Mean-field approximation of adaptability and reliability

When the number of agents increases dramatically, sampling the behaviour of each global state may become unfeasible. Hence, in this section we strengthen our randomised algorithm by means of *mean-field approximation* [21]. We have seen in Section 3 that we can express the behaviour of a system via the Markov chain $\{\mathbf{X}^{(N)}_{\mathbf{A}}(t)\}_{t \in \mathbb{N}}$ taking values in $\mathcal{S}^N$. However, we can also abstract from the identity of each single agent while focusing only on the fraction of agents in a given state, and model systems behaviour via the occupancy DTMC, namely the Markov chain $\{\mathbf{O}^{(N)}_{\mathbf{A}}(t)\}_{t \in \mathbb{N}}$ taking values in $\mathcal{U}^{|\mathcal{S}|}$. We recall that, for each N, the occupancy DTMC $\{\mathbf{O}^{(N)}_{\mathbf{A}}(t)\}_{t \in \mathbb{N}}$ is given by $\mathbf{O}^{(N)}_{\mathbf{A}}(t) = \mathrm{OF}_{\mathcal{S},N}(\mathbf{X}^{(N)}_{\mathbf{A}}(t))$, with initial distribution $\boldsymbol{\delta}_{\mathrm{OF}_{\mathcal{S},N}(\mathbf{A})}$.

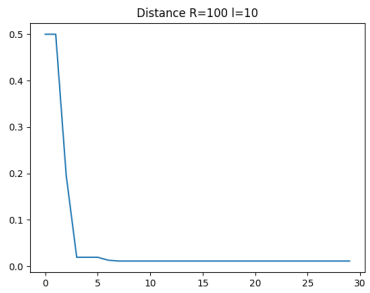

(a) Estimated distance between configurations $\mathbf{A}_1$ and $\mathbf{A}_2$ in Example 5.

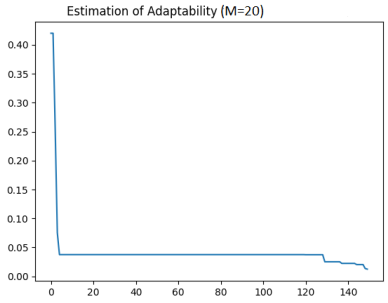

(b) Adaptability of $\mathbf{A}_1$ (Example 3) for $M = 20$, $\eta_1 = 0.25$.

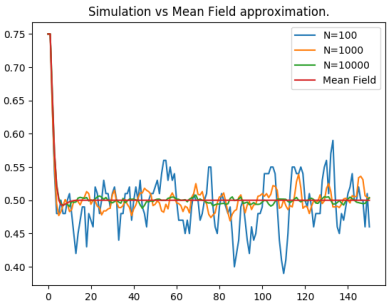

(c) Simulation vs Mean-Field Approximation

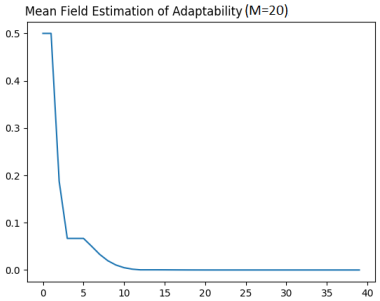

(d) Mean-field adaptability of $\mathbf{A}_1$ (Example 3) for $M = 20$, $\eta_1 = 0.25$.

Fig. 5: Some experiments carried out with the Spear Tool.

In the following we use the fundamental result due to Le Boudec et al. [21] that guarantees that when N goes to infinite, the Markov chain $\mathbf{O}^{(N)}_{\mathbf{A}}(t)$ converges to a *deterministic behaviour*. Given a global state $\mathbf{A} \in \mathcal{S}^N$, we let $c\mathbf{A}$ denote the *global state* in $\mathcal{S}^{cN}$ such that $(c\mathbf{A})_{[\gamma \cdot N + k]} = \mathbf{A}_{[k]}$ for any $1 \leq k \leq N$ and $0 \leq \gamma < c$. Intuitively, $c\mathbf{A}$ consists of c copies of $\mathbf{A}$ and it is called a *c-scale* of $\mathbf{A}$. We can observe that for any $\mathbf{A} \in \mathcal{S}^N$ and for any $c \geq 1$, $\mathrm{OF}_{\mathcal{S},N}(\mathbf{A}) = \mathrm{OF}_{\mathcal{S},cN}(c\mathbf{A})$. Hence, for all $A_i, A_j \in \mathcal{S}$, $\mathbf{K}(\mathrm{OF}_{\mathcal{S},N}(\mathbf{A}))_{A_i,A_j} = \mathbf{K}(\mathrm{OF}_{\mathcal{S},cN}(c\mathbf{A}))_{A_i,A_j}$. Consequently, *scaling* a state by a factor c has no effect on the behaviour of an agent since the probability to select an action does not change. We can consider the sequence of occupancy DTMCs $\{\mathbf{O}^{(cN)}_{c\mathbf{A}}(t)\}$ obtained by increasing the scale c of our system. By Theorem 4.1 of [21], this sequence converges *almost surely* to the following deterministic process in $\mathcal{U}^{|\mathcal{S}|}$:

- $\boldsymbol{D}_{\mathbf{A}}(0) = \mathbf{O}^{(N)}_{\mathbf{A}}(0)$;
- $\boldsymbol{D}_{\mathbf{A}}(t+1) = \boldsymbol{D}_{\mathbf{A}}(t) \cdot \mathbf{K}(\boldsymbol{D}_{\mathbf{A}}(t))$.

In other words,

$$\lim_{c \to \infty} \mathbf{O}^{(cN)}_{c\mathbf{A}}(t) = \boldsymbol{D}_{\mathbf{A}}(t) \tag{4}$$

Example 7. In Figure 5c we can observe how when the scale of the system of Example 3 is increased (we consider $N = 100$, $N = 1000$ and $N = 10000$), the single sampled simulation run gets close to the mean-field approximation.

We say that a penalty function $\rho : \mathcal{S}^N \to [0,1]$ is *scale invariant* if there exists a function $\rho_{\mathbf{o}} : \mathcal{U}^{|\mathcal{S}|} \to [0,1]$ such that $\rho(\mathbf{A}) = \rho_{\mathbf{o}}(\mathrm{OF}_{\mathcal{S},N}(\mathbf{A}))$. We can use mean-field approximation to study the *adaptability* and *reliability* around a given state $\mathbf{A}$ with respect to a scale invariant penalty function.

Proposition 4. *Assume a penalty function ρ_t that is scale invariant for each $t \in \mathbb{N}$. Let $\mathfrak{m}^{\lambda}_{\mathrm{OT}}$ be the population metric defined on $m_{\rho,t}$. Then,*

$$\lim_{c\to\infty} \mathfrak{m}^{\lambda}_{\mathrm{OT}}(c\mathbf{A}_1, c\mathbf{A}_2) = \sup_{t\in\mathrm{OT}} \lambda(t) \cdot |\rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_1}(t)) - \rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_2}(t))| \ . \tag{5}$$

We let $\mathfrak{mf}^{\lambda}_{\mathrm{OT}}(\mathbf{A}_1, \mathbf{A}_2) = \sup_{t\in\mathrm{OT}} \lambda(t) \cdot |\rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_1}(t)) - \rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_2}(t))|$.

The same randomised algorithm outlined in Section 5.3 can be adapted to estimate adaptability and reliability via mean-field approximation. Given a global state $\mathbf{A}$, a set OT of observation times and a threshold $\eta_1 \geq 0$, we can sample M variations $\{\mathbf{A}_1, \ldots, \mathbf{A}_M\}$ of $\mathbf{A}$, such that for any i, $m_{\rho,0}(\mathbf{A}, \mathbf{A}_i) \leq \eta_1$. Then, for each sampled global state we can compute the distance between $\mathbf{A}$ and $\mathbf{A}_i$ by using their *mean field* approximants $\boldsymbol{D}_{\mathbf{A}_1}(t)$ and $\boldsymbol{D}_{\mathbf{A}_2}(t)$. Finally, a *lower bound* to the $\tilde{t}$-adaptability of our system can be computed as $L_{\tilde{t}} = \max_i\{\mathfrak{mf}^{\lambda}_{\{t\in\mathrm{OT}|t\geq\tilde{t}\}}(\mathbf{A}, \mathbf{A}_i)\}$. Similarly, for $t_{\min} = \min_{\mathrm{OT}} t$, $L_{t_{\min}}$ gives a lower bound for its reliability.

Example 8. Figure 5d shows the evaluation of $L_{\tilde{t}}$ for configuration $\mathbf{A}_1$ from Example 3 with parameters $M = 20$ and $\eta_1 = 0.25$. Observe that the initial perturbation is not amplified and after 15 steps it is absorbed. Hence, while we increase the scale of our system, we can guarantee that it is $(15, 0.25, \eta_2)$-adaptable around $\mathbf{A}_1$ for any $\eta_2 \geq e_{\mathbf{W}}$.

7 Concluding remarks

We have proposed the population metric, a pseudometric allowing us to compare the behaviour of self-organising collective systems. This metric quantifies the differences in the evolution sequences of two systems, i.e., the sequences of probability distributions over global states describing the (probabilistic) behaviour of each system. Then we have provided a randomised algorithm for the evaluation of the metric over large scale systems. Moreover, we have shown how we can use the population metric to verify the properties of adaptability and reliability of system, expressing its ability of adjusting its behaviour to perturbations in its initial configuration. We have then modified our algorithm to obtain an estimation of the adaptability and reliability of a system via mean-field approximations.

In this work we have considered a discrete-time approach to systems modelling. Hence, as future work, it would be interesting to provide an adaptation of our framework to system with a continuous model of time. In particular, we could exploit the *fluid-flow* approximation based on [12], in place of the mean-field one, to deal with system with a dramatically large population. Another interesting direction for future research, would

be to use our metric to analyse systems with a different number of agents. From the technical point of view, our definitions and results can be directly lifted to cover this case. However, it would allow us to analyse the impact of a new agent (or a new set of agents) on the dynamics of the system. Finally, we could extend our metric approach to study other properties than adaptability and reliability, and thus obtain some measures on systems performance.

References

1. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the Future in Systems Theory. In: Proceedings of ICALP'03, Lecture Notes in Computer Science, vol. 2719, pp. 1022–1037. Springer (2003). https://doi.org/10.1007/3-540-45061-0_79
2. Anderson, S., Bredeche, N., Eiben, A., Kampis, G., van Steen, M.: Adaptive collective systems: Herding black sheep. Bookprints (2013)
3. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: Proceedings of ICML 2017. pp. 214–223 (2017), http://proceedings.mlr.press/v70/arjovsky17a.html
4. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: On the total variation distance of semi-Markov chains. In: Proceedings of FoSSaCS 2015. Lecture Notes in Computer Science, vol. 9034, pp. 185–199 (2015). https://doi.org/10.1007/978-3-662-46678-0_12
5. Bortolussi, L., Hillston, J., Loreti, M.: Fluid approximation of broadcasting systems. Theor. Comput. Sci. **816**, 221–248 (2020). https://doi.org/10.1016/j.tcs.2020.02.020
6. van Breugel, F.: A behavioural pseudometric for metric labelled transition systems. In: Proceedings of CONCUR 2005. pp. 141–155 (2005). https://doi.org/10.1007/11539452_14
7. Bures, T., Plasil, F., Kit, M., Tuma, P., Hoch, N.: Software abstractions for component interaction in the internet of things. Computer **49**(12), 50–59 (2016)
8. Castiglioni, V.: Trace and testing metrics on nondeterministic probabilistic processes. In: Proceedings of EXPRESS/SOS 2018. EPTCS, vol. 276, pp. 19–36 (2018). https://doi.org/10.4204/EPTCS.276.4, https://doi.org/10.4204/EPTCS.276.4
9. Castiglioni, V., Chatzikokolakis, K., Palamidessi, C.: A logical characterization of differential privacy via behavioral metrics. In: Proceedings of FACS 2018. Lecture Notes in Computer Science, vol. 11222, pp. 75–96 (2018). https://doi.org/10.1007/978-3-030-02146-7_4
10. Castiglioni, V., Loreti, M., Tini, S.: The metric linear-time branching-time spectrum on nondeterministic probabilistic processes. Theor. Comput. Sci. **813**, 20–69 (2020). https://doi.org/10.1016/j.tcs.2019.09.019
11. Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L.: Generalized bisimulation metrics. In: Proceedings of CONCUR 2014. pp. 32–46 (2014). https://doi.org/10.1007/978-3-662-44584-6_4
12. Darling, R., Norris, J.: Differential equation approximations for Markov chains. Probab. Surveys **5**, 37–79 (2008). https://doi.org/10.1214/07-PS121
13. De Nicola, R., Loreti, M., Pugliese, R., Tiezzi, F.: A formal approach to autonomic systems programming: The SCEL language. ACM Trans. Auton. Adapt. Syst. **9**(2), 7:1–7:29 (2014). https://doi.org/10.1145/2619998
14. Deng, Y., Chothia, T., Palamidessi, C., Pang, J.: Metrics for action-labelled quantitative transition systems. Electr. Notes Theor. Comput. Sci. **153**(2), 79–96 (2006). https://doi.org/10.1016/j.entcs.2005.10.033
15. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labeled Markov systems. In: Proceedings of CONCUR '99. pp. 258–273 (1999). https://doi.org/10.1007/3-540-48320-9_19

16. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. Theor. Comput. Sci. **318**(3), 323–354 (2004). https://doi.org/10.1016/j.tcs.2003.09.013
17. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: Proceedings of Advances in Neural Information Processing Systems. pp. 5767–5777 (2017), http://papers.nips.cc/paper/7159-improved-training-of-wasserstein-gans
18. Kantorovich, L.V.: On the transfer of masses. Doklady Akademii Nauk **37**(2), 227–229 (1942)
19. Kopetz, H.: Internet of Things, pp. 307–323. Springer US, Boston, MA (2011). https://doi.org/10.1007/978-1-4419-8237-7_13
20. Latella, D., Loreti, M., Massink, M.: On-the-fly PCTL fast mean-field approximated model-checking for self-organising coordination. Sci. Comput. Program. **110**, 23–50 (2015). https://doi.org/10.1016/j.scico.2015.06.009
21. Le Boudec, J.Y., McDonald, D., Mundinger, J.: A generic mean field convergence result for systems of interacting objects. In: Proceedings of QEST'07. pp. 3–18. IEEE Computer Society (2007). https://doi.org/10.1109/QEST.2007.8
22. Rachev, S.T., Klebanov, L.B., Stoyanov, S.V., Fabozzi, F.J.: The Methods of Distances in the Theory of Probability and Statistics. Springer (2013)
23. Song, L., Deng, Y., Cai, X.: Towards automatic measurement of probabilistic processes. In: Proceedings of QSIC 2007. pp. 50–59 (2007). https://doi.org/10.1109/QSIC.2007.65
24. Sriperumbudur, B.K., Fukumizu, K., Gretton, A., Schölkopf, B., Lanckriet, G.R.G.: On the empirical estimation of integral probability metrics. Electronic Journal of Statistics **6**, 1550–1599 (2021). https://doi.org/10.1214/12-EJS722
25. Talcott, C., Nigam, V., Arbab, F., Kappé, T.: Formal Specification and Analysis of Robust Adaptive Distributed Cyber-Physical Systems, pp. 1–35. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-34096-8_1
26. Thorsley, D., Klavins, E.: Approximating stochastic biochemical processes with Wasserstein pseudometrics. IET Syst. Biol. **4**(3), 193–211 (2010). https://doi.org/10.1049/iet-syb.2009.0039
27. Tolstikhin, I.O., Bousquet, O., Gelly, S., Schölkopf, B.: Wasserstein auto-encoders. In: Proceedings of ICLR 2018 (2018), https://openreview.net/forum?id=HkL7n1-0b
28. Vallender, S.S.: Calculation of the Wasserstein distance between probability distributions on the line. Theory Probab. Appl. **18**(4), 784–786 (1974)
29. Vaserstein, L.N.: Markovian processes on countable space product describing large systems of automata. Probl. Peredachi Inf. **5**(3), 64–72 (1969)
30. Villani, C.: Optimal transport: old and new, vol. 338. Springer (2008)
31. Wirsing, M., Hölzl, M.M., Koch, N., Mayer, P. (eds.): Software Engineering for Collective Autonomic Systems - The ASCENS Approach, Lecture Notes in Computer Science, vol. 8998. Springer (2015). https://doi.org/10.1007/978-3-319-16310-9

newpage

A Proof of Proposition 4

Proposition 4. *Assume a penalty function ρ_t that is scale invariant for each $t \in \mathbb{N}$. Let $\mathfrak{m}^{\lambda}_{\mathrm{OT}}$ be the population metric defined on $m_{\rho,t}$. Then,*

$$\lim_{c\to\infty} \mathfrak{m}^{\lambda}_{\mathrm{OT}}(c\mathbf{A}_1, c\mathbf{A}_2) = \sup_{t\in\mathrm{OT}} \lambda(t)\cdot\left|\rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_1}(t)) - \rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_2}(t))\right| \ . \tag{5}$$

Proof. First we prove that

$$\lim_{c\to\infty} \mathfrak{m}^{\lambda}_{\mathrm{OT}}(c\mathbf{A}_1, c\mathbf{A}_2) \geq \sup_{t\in \mathrm{OT}} \lambda(t)\cdot |\rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_1}(t)) - \rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_2}(t))| \ .$$

We have

$$\begin{aligned}
&\lim_{c\to\infty} \mathfrak{m}^{\lambda}_{\mathrm{OT}}(c\mathbf{A}_1, c\mathbf{A}_2)\\
&= \lim_{c\to\infty} \sup_{t\in \mathrm{OT}} \lambda(t) \min_{\mathfrak{w}\in\mathfrak{W}(\pi_{c\mathbf{A}_1,t},\pi_{c\mathbf{A}_2,t})} \sum_{\substack{\mathbf{A}\in \mathsf{supp}(\pi_{c\mathbf{A}_1,t})\\ \mathbf{A}'\in \mathsf{supp}(\pi_{c\mathbf{A}_2,t})}} \mathfrak{w}(\mathbf{A},\mathbf{A}')|\rho_t(\mathbf{A}) - \rho_t(\mathbf{A}')|\\
&= \lim_{c\to\infty} \sup_{t\in \mathrm{OT}} \lambda(t) \min_{\mathfrak{w}\in\mathfrak{W}(\pi_{c\mathbf{A}_1,t},\pi_{c\mathbf{A}_2,t})} \sum_{\substack{\mathbf{A}\in \mathsf{supp}(\pi_{c\mathbf{A}_1,t})\\ \mathbf{A}'\in \mathsf{supp}(\pi_{c\mathbf{A}_2,t})}} \mathfrak{w}(\mathbf{A},\mathbf{A}')|\rho_{t,\mathbf{o}}(\mathrm{OF}(\mathbf{A})) - \rho_{t,\mathbf{o}}(\mathrm{OF}(\mathbf{A}))|\\
&= \lim_{c\to\infty} \sup_{t\in \mathrm{OT}} \lambda(t) \min_{\mathfrak{w}\in\mathfrak{W}(\pi_{c\mathbf{A}_1,t},\pi_{c\mathbf{A}_2,t})} \sum_{\mathbf{u},\mathbf{u}'\in\mathcal{U}^{|\mathcal{S}|}} \left(\sum_{\substack{\mathbf{A}\in \mathsf{supp}(\pi_{c\mathbf{A}_1,t}):\ \mathrm{OF}(\mathbf{A})=\mathbf{u}\\ \mathbf{A}'\in \mathsf{supp}(\pi_{c\mathbf{A}_2,t}):\ \mathrm{OF}(\mathbf{A}')=\mathbf{u}'}} \mathfrak{w}(\mathbf{A},\mathbf{A}') \right) |\rho_{t,\mathbf{o}}(\mathbf{u}) - \rho_{t,\mathbf{o}}(\mathbf{u}')|\\
&= \lim_{c\to\infty} \sup_{t\in \mathrm{OT}} \lambda(t) \min_{\mathfrak{w}\in\mathfrak{W}(\pi_{c\mathbf{A}_1,t},\pi_{c\mathbf{A}_2,t})} \sum_{\substack{\mathbf{u}\in \mathsf{supp}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t})\\ \mathbf{u}'\in \mathsf{supp}(\pi^{\mathcal{U}}_{c\mathbf{A}_2,t})}} \tilde{\mathfrak{w}}(\mathbf{u},\mathbf{u}')|\rho_{t,\mathbf{o}}(\mathbf{u}) - \rho_{t,\mathbf{o}}(\mathbf{u}')|\\
&\geq \lim_{c\to\infty} \sup_{t\in \mathrm{OT}} \lambda(t) \min_{\tilde{\mathfrak{w}}\in\mathfrak{W}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t},\pi^{\mathcal{U}}_{c\mathbf{A}_2,t})} \sum_{\substack{\mathbf{u}\in \mathsf{supp}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t})\\ \mathbf{u}'\in \mathsf{supp}(\pi^{\mathcal{U}}_{c\mathbf{A}_2,t})}} \tilde{\mathfrak{w}}(\mathbf{u},\mathbf{u}')|\rho_{t,\mathbf{o}}(\mathbf{u}) - \rho_{t,\mathbf{o}}(\mathbf{u}')|\\
&= \sup_{t\in \mathrm{OT}} \lambda(t) \lim_{c\to\infty} \min_{\tilde{\mathfrak{w}}\in\mathfrak{W}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t},\pi^{\mathcal{U}}_{c\mathbf{A}_2,t})} \sum_{\substack{\mathbf{u}\in \mathsf{supp}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t})\\ \mathbf{u}'\in \mathsf{supp}(\pi^{\mathcal{U}}_{c\mathbf{A}_2,t})}} \tilde{\mathfrak{w}}(\mathbf{u},\mathbf{u}')|\rho_{t,\mathbf{o}}(\mathbf{u}) - \rho_{t,\mathbf{o}}(\mathbf{u}')|\\
&= \sup_{t\in \mathrm{OT}} \lambda(t)\, \mathbf{W}(m_{\rho_{\mathbf{o}},t})(\lim_{c\to\infty} \pi^{\mathcal{U}}_{c\mathbf{A}_1,t}, \lim_{c\to\infty} \pi^{\mathcal{U}}_{c\mathbf{A}_2,t})\\
&= \sup_{t\in \mathrm{OT}} \lambda(t)\, \mathbf{W}(m_{\rho_{\mathbf{o}},t})(\boldsymbol{\delta}_{\mathbf{D}_{\mathbf{A}_1}(t)}, \boldsymbol{\delta}_{\mathbf{D}_{\mathbf{A}_2}(t)})\\
&= \sup_{t\in \mathrm{OT}} \lambda(t)\cdot |\rho_{t,\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_1}(t)) - \rho_{t,\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_2}(t))|
\end{aligned}$$

where:

- The second step follows by ρ_t being scale invariant.
- The fourth step follows by defining, for all $\mathbf{u},\mathbf{u}'\in\mathcal{U}^{|\mathcal{S}|}$

$$\tilde{\mathfrak{w}}(\mathbf{u},\mathbf{u}') = \sum_{\substack{\mathbf{A}\in \mathsf{supp}(\pi_{c\mathbf{A}_1,t}):\ \mathrm{OF}(\mathbf{A})=\mathbf{u}\\ \mathbf{A}'\in \mathsf{supp}(\pi_{c\mathbf{A}_2,t}):\ \mathrm{OF}(\mathbf{A}')=\mathbf{u}'}} \mathfrak{w}(\mathbf{A},\mathbf{A}') \ .$$

- The fifth step follow by noticing that $\tilde{\mathfrak{w}}$ is not necessarily the optimal coupling in $\mathfrak{W}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t}, \pi^{\mathcal{U}}_{c\mathbf{A}_2,t})$, where the two distributions are defined, for each $\mathbf{u}\in\mathcal{U}^{|\mathcal{S}|}$, $i\in\{1,2\}$, as

$$\pi^{\mathcal{U}}_{c\mathbf{A}_i,t}(\mathbf{u}) = \sum_{\mathbf{A}\in \mathsf{supp}(\pi_{c\mathbf{A}_i,t}):\ \mathrm{OF}(\mathbf{A})=\mathbf{u}} \pi_{c\mathbf{A}_i,t}(\mathbf{A}) \ .$$

– The sixth step follows by noticing that time and scaling are independent.
– The seventh step follows by noticing that:
  • $\mathbf{O}^{cN}_{\mathbf{A}_i}(t)$ converges in distribution to $\mathbf{D}_{\mathbf{A}_i}(t)$ (Equation (4));
  • since occupancy vectors are scale invariant we can take the limit inside the minimum;
  • the linearity of the limit.

We can then apply the same steps backwards to obtain

$$\sup_{t\in \mathrm{OT}} \lambda(t)\cdot|\rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_1}(t))-\rho_{\mathbf{o}}(\boldsymbol{D}_{\mathbf{A}_2}(t))| \geq \lim_{c\to\infty} \mathfrak{m}^{\lambda}_{\mathrm{OT}}(c\mathbf{A}_1, c\mathbf{A}_2) \ ,$$

and thus that the equality holds. In particular, starting from the optimal coupling $\tilde{\mathfrak{w}} \in \mathfrak{W}(\pi^{\mathcal{U}}_{c\mathbf{A}_1,t}, \pi^{\mathcal{U}}_{c\mathbf{A}_2,t})$ we can obtain a generic coupling $\mathfrak{w} \in \mathfrak{W}(\pi_{c\mathbf{A}_1,t}, \pi_{c\mathbf{A}_2,t})$ as follows:

$$\mathfrak{w}(\mathbf{A},\mathbf{A}') = \frac{\pi_{c\mathbf{A}_1,t}(\mathbf{A})}{\sum\limits_{\mathbf{B}_1\in \mathrm{OF}^{-1}(\mathrm{OF}(\mathbf{A}))} \pi_{c\mathbf{A}_1,t}(\mathbf{B}_1)} \cdot \tilde{\mathfrak{w}}(\mathrm{OF}(\mathbf{A}),\mathrm{OF}(\mathbf{A}'))\cdot \frac{\pi_{c\mathbf{A}_2,t}(\mathbf{A}')}{\sum\limits_{\mathbf{B}_2\in \mathrm{OF}^{-1}(\mathrm{OF}(\mathbf{A}'))} \pi_{c\mathbf{A}_2,t}(\mathbf{B}_2)} \ .$$