

# Multi-Scale Verification of Distributed Synchronisation

Paul Gainer · Sven Linker · Clare Dixon ·  
Ullrich Hustadt · Michael Fisher

the date of receipt and acceptance should be inserted later

**Abstract** Algorithms for the synchronisation of clocks across networks are both common and important within distributed systems. We here address not only the formal modelling of these algorithms, but also the formal verification of their behaviour. Of particular importance is the strong link between the very different levels of abstraction at which the algorithms may be verified. Our contribution is primarily the formalisation of this connection between individual models and population-based models, and the subsequent verification that is then possible. While the technique is applicable across a range of synchronisation algorithms, we particularly focus on the synchronisation of (biologically-inspired) pulse-coupled oscillators, a widely used approach in practical distributed systems. For this application domain, different levels of abstraction are crucial: models based on the behaviour of an individual process are able to capture the details of distinguished nodes in possibly heterogenous networks, where each node may exhibit different behaviour. On the other hand, collective models assume homogeneous sets of processes, and allow the behaviour of the network to be analysed at the global level. System-wide parameters may be easily adjusted, for example environmental factors inhibiting the reliability of the shared communication medium. This work provides a formal bridge across the “abstraction gap” separating the individual models and the population-based models for this important class of synchronisation algorithms.

## 1 Introduction

Small computing devices comprising networks, be it commercial wireless sensor networks, or communicating devices in the Internet of Things, become increasingly common. However, to enable these devices to communicate efficiently, they have

---

This work was supported by the Sir Joseph Rotblat Alumni Scholarship at Liverpool and the Engineering and Physical Sciences Research Council, under grants EP/N007565/1 (S4: Science of Sensor Systems Software), EP/L024845/1 (Verifiable Autonomy), and the FAIRSPACE (EP/R026092/1), RAIN (EP/R026084/1). and ORCA (EP/R026173/1) RAI Hubs.

Address(es) of author(s) should be given

to employ methods to use the shared communication medium without too many conflicts, e.g., in the form of collisions. Several protocols to organise shared medium access have been developed and analysed [1, 35]. These protocols typically identify a common time frame and divide this frame into slots associated to each node. Thus every node has an allocated time slot that it may use to communicate its messages onto the shared medium.

Such an approach introduces the need for a common clock between the nodes, i.e., they need to synchronise. A valuable approach to achieve synchrony of nodes is the implementation of biologically-inspired *pulse-coupled oscillators* (PCOs) [25]. A network of PCOs synchronises in the following way: all oscillators have a similar *clock cycle* at the end of which they *fire*. That is, they transmit a broadcast message which is received by all oscillators in their communication range. These oscillators then adjust their own position within their clock cycle according to a *phase response function*. Depending on the concrete implementation, they may move their current position within the clock cycle closer to its end, or closer to its start.

Most analyses of the synchronisation behaviour of PCOs are concerned with continuous clock cycles, i.e., where clocks take real values from the interval  $[0, 1]$ . However, the smaller devices get, the more important it is to save memory and computing time for such a low-level functionality. Even a floating point number may need too much memory, compared to an implementation with, for example, a four-bit vector. Hence, in previous work, we chose to analyse the behaviour of *discrete time PCOs* [16].

In contrast to continuous time PCOs, networks of discrete time PCOs are not always guaranteed to synchronise. Instead, whether they synchronise or not depends on the type of coupling between the oscillators and their common phase-response function. We analysed the behaviour of such networks for different parameters via model-checking, to check both qualitatively for which parameters the networks synchronise, as well as quantitatively for how long they need to achieve a synchronised state and how much energy is used to achieve this [17]. In the context of large numbers of single oscillators, for example in the context of wireless sensor networks, the well-known state-space explosion problem of the model-checking approach is extremely important [9]. We formalised a network of oscillators as *population models* [13] which exploit the behavioural homogeneity of the nodes to encode the global state efficiently. This allows the network size to be increased above what would be feasible when distinguishing each node. But the construction of a population model from a given oscillator specification is not straightforward, and in particular, it is not obvious whether the constructed population model correctly reflects the behaviour of the oscillators. This results in an ‘abstraction gap’: after abstracting into populations, how can we be sure that the abstraction process was correct and that the results of verification of population models actually hold for the concrete models on which they are based?

In this paper, we remedy this lack of certainty, by proving the correspondence of our population model with an explicit formalisation of the oscillators. To that end, we present the concrete oscillator model as well as its formalisation as a discrete-time Markov chain. Subsequently we describe the corresponding population model, and show how we can, in addition to the abstraction created by the populations, reduce the state space even further to facilitate the analysis. Finally, we prove that the behaviour of a network of concrete oscillators can be simulated by the population model. We cannot prove a one-to-one correspondence, since the

concrete model implicitly includes the possibility of identifying individual oscillators, which is exactly what the population model abstracts from. However, by providing a formal notion of abstraction, we prove that population models are a truthful abstraction of concrete models.

The paper is structured as follows. In Sect. 2, we review a selection of related work, both for models of pulse-coupled oscillators, as well as approaches for their verification. After an introduction of preliminary notions in Sect. 3, we present the concrete model of single oscillators, both as an algorithm and as a discrete-time Markov chain derived from this algorithm, in Sect. 4. The abstract model in terms of population models and proofs about their properties are contained in Sect. 5. In Sect. 6, we prove the correspondence between these two types of models, and conclude in Sect. 7.

## 2 Related Work

The canonical model of pulse-coupled oscillators, and their synchronisation, was formulated by Mirollo and Strogatz [25], and based on Peskin's model of a cardiac pacemaker [29]. Here the progression of an oscillator through its oscillation cycle is given by a real value in the interval  $[0, 1]$ . Mirollo and Strogatz proved that with a convex phase response function, a network of mutually coupled oscillators always converges, i.e., their position within the oscillation cycle eventually coincides. Such a model has been shown to be applicable to the clock synchronisation of wireless sensor nodes [31] and swarms of robots [27].

Synchronisation algorithms based on pulse-coupled oscillators are often beneficial in unreliable, decentralised networks, where other synchronisation algorithms are not appropriate. For example, the Flooding Time Synchronisation Protocol (FTSP) [23] requires the use of an arbitrary root node. In situations where the root becomes unavailable due to communication failure or power outage, FTSP will have to assign another root node. When implemented on unreliable, decentralised networks, FTSP may spend considerable resources on repeatedly assigning root nodes, which may slow down or prevent synchronisation [8]. Other algorithms such as the Berkeley algorithm [18] and Cristian's algorithm [11] require the use of centralised time servers, which is problematic for unreliable, decentralised networks.

Several decentralised network algorithms for synchronisation are based on pulse-coupled oscillators [31, 34]. For example, the Gradient Time Synchronisation Protocol (GTSP) by Sommer and Wattenhofer [30] achieves synchronisation by having nodes send their current clock value to their neighbours. Each node then calculates the average of the clock values received and its own clock value. This process is then repeated to maintain synchronisation. Another approach to synchronisation, the Pulse-Coupled Oscillator Protocol [26], makes use of refractory periods after sending messages containing time information. During the refractory period, no more messages are sent, which reduces network bandwidth and energy usage. A similar approach is used in the FiGo protocol [8], which combines biologically inspired synchronisation with information distribution via gossiping. All of these approaches use different phase response functions.

In general, synchronisation algorithms based on PCOs are more robust for unreliable networks, as they do not require centralised nodes and can work with

only partial network connectivity [8]. They are particularly useful for battery-powered nodes in wireless networks, as the node can be placed in a low-power mode during the refractory period, thus reducing energy usage. (The clock keeps ticking even in low-power mode, thanks to the design of microcontrollers such as the ‘Atmel ATmega128L’ [4].)

Synchronisation of clocks for networks of nodes has been investigated from different perspectives. Heidarian et al. [20] analysed the behaviour of a synchronisation protocol based on time allocation slots for up to four nodes and different topologies, from fully connected networks to line topologies. They modelled the protocol as timed automata [2], and used the model-checker UPPAAL [7] to examine its worst-case behaviour. Their model is based on continuous time, and in particular, they did not model pulse-coupled oscillators.

Bartocci et al. [5] described pulse-coupled oscillators as extended timed automata with suitable semantics to model their peculiarities. They defined a dedicated logic to analyse the behaviour of a network of such automata along traces, and used a pacemaker as a case study to verify the eventual synchronisation and the time needed to achieve this.

Our models and methods are slightly different to all of these approaches. This is, of course, evident for all the mentioned work that is not concerned with pulse-coupled oscillators. However, we also define the oscillation cycle to consist of discrete steps. To the best of our knowledge, with the exception the paper by Webster et al. [33] and our previous work [16, 17], there is no other work concerned with PCOs with discrete oscillation cycles. Furthermore, all of these approaches distinguish between single oscillators in the network, while the properties of interest relate to global behaviour. This discrepancy between local modelling and global analysis restricts the size of networks that can be analysed, due to the state-space explosion. To extend the size of analysable networks, we employ *population models*, a counting-abstraction of such networks [12]. Instead of identifying each oscillator on its own, we record how many oscillators are in each step of the oscillation cycle. This reduces the state-space quite tremendously by exploiting the symmetries in the model [13], and we are hence able to extend the size of networks.

The notion of population models should not be confused with *population protocols* [3], a formalism to express distributed algorithms. In contrast to our setting, communication in population protocols is always between two agents, where one agent initiates the communication and the other responds. Furthermore, even though the agents cannot identify the other agents in the network, within the global model each agent is uniquely associated with a state. In our model, we cannot distinguish between two different agents sharing the same state, even at the global level. Finally, our oscillators may change their state without interacting with other oscillators, while the agents in a population protocol must communicate with another agent to change their internal state.

We will present a relation between the concrete models, where each oscillator can be identified, and corresponding population models, and show that these two models are in a *simulation relation* [24]. More precisely, the concrete model *weakly* simulates its abstraction, since the oscillators have to take transitions independently, while in the population model, all oscillators evolve in a single step.

Similarly to typical definitions of counter abstractions [14, 6], we use counters to model concurrent entities that are indistinguishable for our purposes. For example, to analyse the probability of eventually reaching a synchronised state,

we are not interested in an order of oscillators, which would be artificial anyway. However, in contrast to these approaches, we do not include means to introduce new entities into a model. That is, the values within our population models are naturally bounded by the number of oscillators within the network.

### 3 Preliminaries

In this section we define *discrete-time Markov chains* (DTMCs), stochastic processes with discrete state space and discrete time, and introduce Probabilistic Computation Tree Logic (PCTL), a logic that can be used to reason about probabilistic reachability and rewards in these processes.

Throughout this paper, we use the notation  $f \oplus [x \mapsto y]$ , where  $f$  is a function, to express *updating  $f$  at  $x$  by  $y$* . That is, the function that coincides with  $f$ , except for  $x$ , where it takes the value  $y$ .

#### 3.1 Discrete-Time Markov Chains

DTMCs can be used to model systems where the evolution of the system at any moment in time can be represented by a discrete probabilistic choice over several outcomes.

**Definition 1** A discrete-time Markov chain  $D$  is a tuple  $(Q, \sigma_0, \mathbf{P}, L)$  where  $Q$  is a finite set of states.  $\sigma_0$  is the initial state, and  $L : Q \rightarrow \mathbb{P}(\mathcal{V})$  is a labelling function that assigns properties of interest from a set of labels  $\mathcal{V}$  to states.  $\mathbf{P} : Q \times Q \rightarrow [0, 1]$  is the *transition probability matrix* subject to  $\sum_{\sigma' \in Q} \mathbf{P}(\sigma, \sigma') = 1$  for all  $\sigma \in Q$ , where  $\mathbf{P}(\sigma, \sigma')$  gives the probability of transitioning from  $\sigma$  to  $\sigma'$ . We say that there is a transition between two states  $\sigma, \sigma' \in Q$  if  $\mathbf{P}(\sigma, \sigma') > 0$ .

Intuitively, a DTMC is a state transition system where transitions between states are labelled with probabilities greater than 0 and where states are labelled with properties of interest. An execution *path*  $\omega$  of a DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  is a non-empty finite, or infinite, sequence  $\sigma_0 \sigma_1 \sigma_2 \dots$  where  $\sigma_i \in Q$  and  $\mathbf{P}(\sigma_i, \sigma_{i+1}) > 0$  for  $i \geq 0$ . We denote the set of all paths starting in state  $\sigma$  by  $Paths^D(\sigma)$ , and the set of all finite paths starting in  $\sigma$  by  $Paths_f^D(\sigma)$ . For paths where the first state along that path is the initial state  $\sigma_0$  we will simply use  $Paths^D$  and  $Paths_f^D$ , and we will simply use  $Paths$  and  $Paths_f$  if  $D$  is clear from the context. For a finite path  $\omega_f \in Paths_f$  the *cylinder set* of  $\omega_f$  is the set of all infinite paths in  $Paths$  that share prefix  $\omega_f$ . The probability of taking a finite path  $\sigma_0 \sigma_1 \dots \sigma_n \in Paths_f$  is given by  $\prod_{i=1}^n \mathbf{P}(\sigma_{i-1}, \sigma_i)$ . This measure over finite paths can be extended to a probability measure  $Pr$  over the set of infinite paths  $Paths$ , where the smallest  $\sigma$ -algebra over  $Paths$  is the smallest set containing all cylinder sets for paths in  $Paths_f$ . For a detailed description of the construction of the probability measure we refer the reader to [21].

### 3.2 Probabilistic Computation Tree Logic

*Probabilistic Computation Tree Logic* [19] (PCTL) is a probabilistic extension of the temporal logic CTL. Properties for DTMCs can be formulated in PCTL and then checked against the DTMCs using *model checking*.

**Definition 2** The syntax of PCTL is given by:

$$\begin{aligned}\Phi &= p \mid \neg\Phi \mid \Phi \wedge \Phi \mid P_{\bowtie\lambda}[\Psi] \\ \Psi &= X\Phi \mid \Phi U^{\leq k}\Phi\end{aligned}$$

where  $p$  is an atomic proposition,  $\bowtie \in \{<, \leq, \geq, >\}$ ,  $\lambda \in [0, 1]$ , and  $k \in \mathbb{N} \cup \{\infty\}$ .

Formulas denoted by  $\Phi$  are *state formulas* and formulas denoted by  $\Psi$  are *path formulas*. A PCTL formula is always a state formula, and a path formula can only occur inside the P operator. We now give the semantics of PCTL over a DTMC.

**Definition 3** Given a DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$ , we inductively define the satisfaction relation  $\models$  for any state  $\sigma \in Q$  as follows:

$$\begin{aligned}\sigma \models p &\Leftrightarrow p \in L(\sigma) \\ \sigma \models \neg\Phi &\Leftrightarrow \sigma \not\models \Phi \\ \sigma \models \Phi \wedge \Phi' &\Leftrightarrow \sigma \models \Phi \text{ and } \sigma \models \Phi' \\ \sigma \models P_{\bowtie\lambda}[\Psi] &\Leftrightarrow Pr\{\omega \in Paths(\sigma) \mid \omega \models \Psi\} \bowtie \lambda\end{aligned}$$

where  $v \in \mathcal{V}$ , and for any path  $\omega = \sigma_0\sigma_1\sigma_2\cdots$  of  $D$  as follows:

$$\begin{aligned}\omega \models X\Phi &\Leftrightarrow \sigma_1 \models \Phi \\ \omega \models \Phi U^{\leq k}\Phi' &\Leftrightarrow \exists i \in \mathbb{N}(i \leq k \text{ and } \sigma_i \models \Phi' \text{ and } \forall j < i. \sigma_j \models \Phi).\end{aligned}$$

Disjunction, *true*, *false*, and implication are derived as usual, and we define eventuality as  $F^{\leq k}\Phi \equiv true U^{\leq k}\Phi$ . We simply use  $F\Phi$  and  $\Phi U\Phi'$  when  $k = \infty$ .

## 4 Concrete Model of a Network of Pulse-Coupled Oscillators

In this section we give a brief introduction to the formal model of a single pulse-coupled oscillator, as originally presented in previous work [16]. Subsequently, we encode fully-coupled networks of such oscillators as discrete time Markov chains.

### 4.1 Pulse-Coupled Oscillator Model

We consider a fully-coupled network of pulse-coupled oscillators with identical dynamics over discrete time. The *phase* of an oscillator  $u$  at time  $t$  is denoted by  $\phi_u(t)$ . The phase of an oscillator progresses through a sequence of discrete integer values bounded by some  $T \geq 1$ . The phase progression over time of a single uncoupled oscillator is determined by the successor function, where the phase increases over time until it equals  $T$ , at which point the oscillator will fire in the

next moment in time and the phase will reset to one. The phase progression of an uncoupled oscillator is therefore cyclic with period  $T$ , and we refer to one cycle as an *oscillation cycle*.

When an oscillator fires, it may happen that its firing is not perceived by any of the other oscillators coupled to it. We call this a *broadcast failure* and denote its probability by  $\mu \in [0, 1]$ . Note that  $\mu$  is a global parameter, hence the chance of broadcast failure is identical for all oscillators. When an oscillator fires, and a broadcast failure does not occur, it perturbs the phase of all oscillators to which it is coupled; we use  $\alpha_u(t)$  to denote the number of all other oscillators that are coupled to  $u$  and will fire at time  $t$ .

**Definition 4** The *phase response function* is a positive increasing function  $\Delta : \{1, \dots, T\} \times \mathbb{N} \times \mathbb{R}^+ \rightarrow \mathbb{N}$  that maps the phase of an oscillator  $u$ , the number of other oscillators perceived to be firing by  $u$ , and a real value defining the strength of the coupling between oscillators, to an integer value corresponding to the perturbation to phase induced by the firing of oscillators where broadcast failures did not occur. We require  $\Delta(\Phi, 0, \epsilon) = 0$  for all possible phase response functions, that is, oscillators are only perturbed if they perceive at least one firing oscillator.

We can introduce a *refractory period* into the oscillation cycle of each oscillator. A refractory period is an interval of discrete values  $[1, R] \subseteq [1, T]$  where  $R \leq T$  is the size of the refractory period, such that if  $\phi_u(t)$  is inside the interval, for some oscillator  $u$  at time  $t$ , then  $u$  cannot be perturbed by other oscillators to which it is coupled. If  $R = 0$  then we set  $[1, R] = \emptyset$ , and there is no refractory period at all.

**Definition 5** The *refractory function*  $\text{ref} : \{1, \dots, T\} \times \mathbb{N} \rightarrow \mathbb{N}$  is defined as  $\text{ref}(\Phi, \delta) = \Phi$  if  $\Phi \in [1, R]$ , or  $\text{ref}(\Phi, \delta) = \Phi + \delta$  otherwise, and takes as parameters  $\delta$ , the degree of perturbation to the phase of an oscillator, and  $\Phi$ , the phase, and returns  $\Phi$  if it is in the refractory period, or  $\Phi + \delta$  otherwise.

The phase evolution of an oscillator  $u$  over time is then defined as follows, where the *update function* and *firing predicate*, respectively denote the updated phase of oscillator  $u$  at time  $t$  in the next moment in time, and the firing of oscillator  $u$  at time  $t$ ,

$$\begin{aligned} \text{update}_u(t) &= 1 + \text{ref}(\phi_u(t), \Delta(\phi_u(t), \alpha_u(t), \epsilon)), \\ \text{fire}_u(t) &= \text{update}_u(t) > T, \\ \phi_u(t+1) &= \begin{cases} 1 & \text{if } \text{fire}_u(t) \\ \text{update}_u(t) & \text{otherwise.} \end{cases} \end{aligned}$$

## 4.2 Modelling the Network Using a DTMC

We model the whole network of oscillators as a single DTMC  $D = (Q, s_0, \mathbf{P}, L)$ , where each state  $s \in Q$  denotes a global state of the network. More precisely, the labelling function uniquely maps each state  $s$  to an combined encoding of the individual state of each oscillator. For simplicity, we identify the label of a state with the state itself, and hence we omit  $L$  from the DTMC, but describe each member of  $Q$  via its internal state.

We model each transition of an oscillator as a single transition within the DTMC. However, since the oscillators may influence each other within a single time step (that is, when they are firing), we cannot simply allow for arbitrary sequences of transitions. For instance, to model that all the oscillators progress on a similar time-scale, we need to prevent a single oscillator from taking a transition and thus progressing its phase without giving the other oscillators a chance to do the same. We achieve this by the following means:

- we divide the internal computation of each oscillator into two modes: *start* and *update*, and
- we add a counter to the model, containing the number of oscillators that fire.

The counter also possesses both modes, and resets at the start of each “round” of computation. First, in the *start* mode, each oscillator checks whether it would fire, according to its phase response function and the current number of oscillators that already fired, as given by the counter. If it does, it increases the counter and updates its mode to *update*, otherwise it just updates its mode. If all oscillators are in the update mode, they compute their new phases in a single step, according to the phase response function and the current state of the environment counter. Furthermore, we impose an order on the evaluation on the oscillators in the start mode if at least one oscillator fires, starting from the highest phase to the lowest. This ensures that firing oscillators are perceived by the other nodes, and thus may lead to the firing of the latter. This way of modelling the nodes implies the assumption that the time window during which each oscillator listens on the shared medium is long enough to perceive the firing of any other oscillator.

The general idea of the progress of the network of oscillators is visualised in Fig. 1. In the figure, each rounded rectangle shows a state of a network of four oscillators. The circles represents the nodes, where we inscribe its current phase and an abbreviation of its mode. A node that is about to fire is indicated by a starred circle, while a shaded circle indicates a node that is within the refractory period. The rectangle denotes the environment counter, with its corresponding value and mode. The phase response function is arbitrarily chosen, and of minor importance for the example.

In the first state, all outgoing transitions only check whether to increase the counter. Since no oscillator is in the firing phase, all oscillators just update their mode (observe that the single arrow actually denotes four transitions). In the next step, all oscillators increase their phase by one, and reset their mode to *start*. In the next four transitions, oscillator 2 fires and increases the counter, which in turn is sufficient for oscillator 3 to fire as well. Hence they both increased the counter by one, while oscillators 1 and 4 did not. During the last transition of the example, oscillator 2 and 3 reset their phase to one, while oscillator 1 is perturbed and increases its phase by two steps at once. Oscillator 4 is within its refractory period, which means that it is not perturbed, and simply increments its phase. In addition to these transitions, we also need some bookkeeping transitions, to ensure that the counter is reset before the oscillators check their phase response. Furthermore, observe that in the example, it is crucial that oscillator 3 checks its response after oscillator 2 increased the counter, since otherwise 3 would not have been perturbed to fire.

Formally, we conflate the states of the oscillators and the environment into a single state of the DTMC. Each oscillator can be described by a tuple consisting



of the current *phase*  $\Phi$  of the oscillator and the *mode*  $\theta$  within this phase. The phase ranges from 1 to  $T$ , while the mode takes values from  $\{start, update\}$ . Furthermore, we use a single counter to keep track of the number of oscillators that fired successfully within a single phase computation.

For a fixed sequence of  $N$  oscillators, a state of the concrete model consists of a function  $\nu$  that associates a phase and mode with each oscillator,

$$\nu: \{1, \dots, N\} \rightarrow (\{1, \dots, T\} \times \{start, update\}),$$

and the state of the environment  $\eta$  that counts the number of oscillators that fired,

$$\eta \in \{start, update\} \times \{0, \dots, N\}.$$

A state is therefore a tuple  $s = (\eta, \nu)$ , where  $\eta$  is the state of the environment, and  $\nu$  is the state of the network. We denote the set of all concrete system states by  $Q_c$ . For simplicity, we use the notation  $p_\phi$  ( $p_\theta$ , respectively) for the corresponding projection function of the network states, i.e., if  $\nu(u) = (\Phi_u, \theta_u)$ , then  $p_\phi(\nu(u)) = \Phi_u$  and  $p_\theta(\nu(u)) = \theta_u$ . Similarly, for an environment state  $\eta = (\theta, c)$ , we will refer to  $\theta$  by  $p_\theta(\eta)$  and to  $c$  by  $p_c(\eta)$ . We use the notation  $init_\Phi(s) = \{u \mid p_\theta(\nu(u)) = start \wedge p_\phi(\nu(u)) = \Phi\}$  for the set of all oscillators sharing phase  $\Phi$  and mode *start* in the state  $s = (\eta, \nu)$ . Furthermore, we simply use the notation  $init(s) = \{u \mid p_\theta(\nu(u)) = start\}$ .

We now define the transition probabilities between states. To do this we first distinguish the following cases:

1. the environment resets its counter;
2. no oscillator has a clock value of  $T$ ;
3. an oscillator is in the mode *start*, has a clock value lower than  $T$ , is perturbed, but not enough to fire;
4. an oscillator is in the mode *start*, has a clock value lower than  $T$  and is perturbed enough to fire;
5. an oscillator is in the mode *start*, has a clock value of  $T$ , and broadcasts its pulse;
6. an oscillator is in the mode *start*, has a clock value of  $T$ , and fails to broadcast its pulse;
7. all oscillators are in the mode *update*, update their clock and reset their state to *start*.

We will impose an order on certain transitions for two reasons. Firstly, we will restrict transitions that are only used for bookkeeping purposes. For example, we will require that the reset transition of the environment is taken before any of the transitions for the oscillators within a phase are activated. In particular, this means that each computation starts with a transition of the type 1. Secondly, we need to ensure that, if at least one oscillator fires, the phase response of all oscillators is evaluated starting with oscillators in the highest phase, down to the lowest phase, as described above. The cases stated above are reflected in the following definitions for the transition probability between two states  $s = (\eta, \nu)$  and  $s' = (\eta', \nu')$ .

Case 1, where the environment resetting its counter is treated as follows. In the precondition, we require that the mode of the counter is *start*, and the state of the oscillators does not change from  $s$  to  $s'$ . Furthermore, the mode of the counter

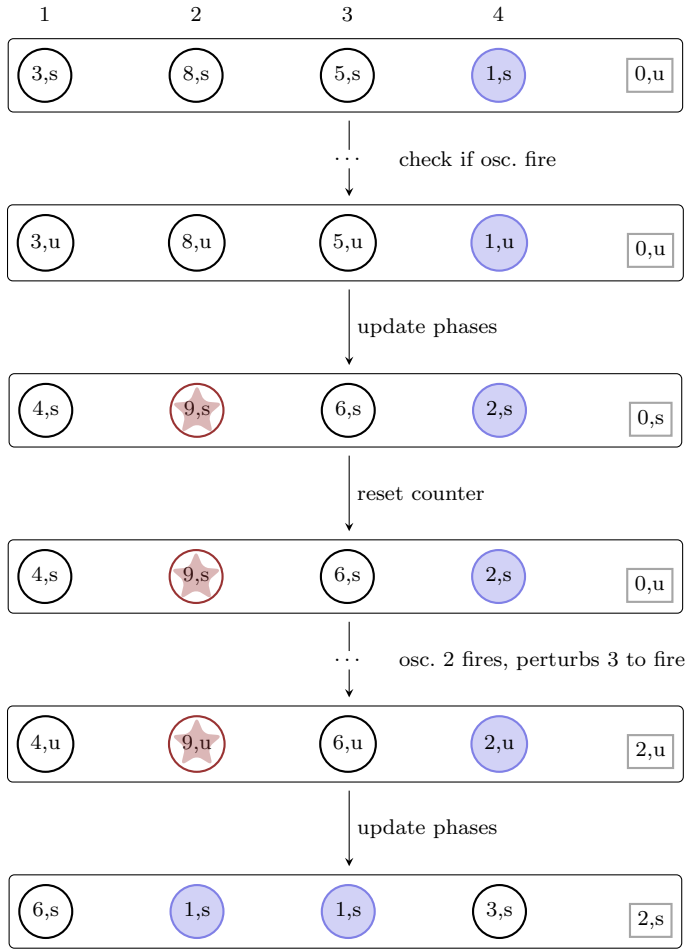


Fig. 1: Transitions in the Concrete Oscillator Model ( $N = 4, T = 9, R = 2$ )

changes to *update* in  $s'$ , and its value is set to 0. Since this transition is mandatory at the beginning of each round, its probability is 1.

$$\text{If } p_\theta(\eta) = \text{start} \wedge p_\theta(\eta') = \text{update} \wedge p_c(\eta') = 0 \wedge \forall u: \nu(u) = \nu'(u), \quad (1)$$

then  $\mathbf{P}(s, s') = 1$ .

Now we turn to the cases 5 and 6 where some oscillator is at the end of its cycle. The preconditions of both cases are similar: the counter is required to be in the *update* mode, and there is an oscillator  $w$ , whose phase is  $T$  and mode is *start*. Furthermore, in  $s'$ , the mode of  $w$  is *update*, and the state of all other oscillators does not change. The difference between the cases is whether the counter is increased, that is, whether the oscillator manages to broadcast its signal. The probability of succeeding is  $\frac{1-\mu}{|\text{init}_T(s)|}$ , since there may be more than one oscillator in phase  $T$  at state  $s$ . Hence we have to normalise the transition probability accordingly.

Similarly, the probability of failing to fire is  $\frac{\mu}{|\text{init}_T(s)|}$ .

If  $p_\theta(\eta) = \text{update}$  and there is a  $w$  s.t. (2)

$$\begin{aligned} p_\theta(\nu(w)) &= \text{start} \wedge p_\phi(\nu(w)) = T \wedge p_\theta(\nu'(w)) = \text{update} \\ &\wedge p_\phi(\nu(w)) = p_\phi(\nu'(w)) \wedge \forall u: u \neq w \rightarrow \nu(u) = \nu'(u) \\ &\wedge p_c(\eta') = p_c(\eta) + 1 \end{aligned}$$

$$\text{then } \mathbf{P}(s, s') = \frac{1 - \mu}{|\text{init}_T(s)|}.$$

If  $p_\theta(\eta) = \text{update}$  and there is a  $w$  s.t. (3)

$$\begin{aligned} p_\theta(\nu(w)) &= \text{start} \wedge p_\phi(\nu(w)) = T \wedge p_\theta(\nu'(w)) = \text{update} \\ &\wedge p_\phi(\nu(w)) = p_\phi(\nu'(w)) \wedge \forall u: u \neq w \rightarrow \nu(u) = \nu'(u) \\ &\wedge p_c(\eta') = p_c(\eta) \end{aligned}$$

$$\text{then } \mathbf{P}(s, s') = \frac{\mu}{|\text{init}_T(s)|}.$$

If no oscillator is at the end of its cycle, that is, in case 2, we define the probability of one oscillator updating its mode as follows. Observe that we have to normalise the transition probability by the number of all oscillators that have not transitioned to their update mode yet. This is correct, since no oscillator fires, which also means that no oscillator can be activated beyond the maximum phase. This implies in particular that the order of oscillator transitions does not matter in this round.

If  $p_\theta(\eta) = \text{update}$  and there is a  $w$  s.t. (4)

$$\begin{aligned} p_\theta(\nu(w)) &= \text{start} \wedge p_\theta(\nu'(w)) = \text{update} \wedge p_\phi(\nu(w)) = p_\phi(\nu'(w)) \\ &\wedge \forall u: p_\phi(\nu(u)) < T \wedge \forall u: u \neq w \rightarrow \nu(u) = \nu'(u) \wedge \eta = \eta' \end{aligned}$$

$$\text{then } \mathbf{P}(s, s') = \frac{1}{|\text{init}(s)|}.$$

Now we will consider the cases 3 and 4, where some oscillator already fired (i.e.,  $p_c(\eta) > 0$ ), and other oscillators are perturbed. We distinguish between two cases: either an oscillator is sufficiently perturbed to also fire or the perturbation does not cause the phase to exceed the firing threshold. One complication arises in these cases: we have to ensure that we only allow the oscillators to update their mode once all oscillators with a higher phase have been considered. Since the perturbation function is increasing, a higher phase may result in a higher perturbation. That is, oscillators with a higher phase need to be perturbed by fewer firing oscillators before their phase is increased beyond the threshold and they in turn fire. Hence, if we did not enforce such an order, oscillators with a lower phases might not be perturbed when oscillators with a higher phase fire. Again, observe that we normalise the transition probabilities according to the number of oscillators satisfying similar conditions. That is, this time we need to normalise on the number of oscillators with the same phase in the *start* mode.

If  $p_\theta(\eta) = \text{update}$  and there is a  $w$  s.t. (5)

$$\begin{aligned}
& p_\theta(\nu(w)) = \text{start} \wedge p_\theta(\nu'(w)) = \text{update} \wedge p_\phi(\nu(w)) = p_\phi(\nu'(w)) \\
& \wedge p_\phi(\nu(w)) < T \wedge \exists u: p_\phi(\nu(u)) = T \\
& \wedge \forall u: u \neq w \rightarrow (p_\theta(\nu(u)) = \text{update} \vee p_\phi(\nu(u)) \leq p_\phi(\nu(w))) \\
& \wedge p_\phi(\nu(w)) + \Delta(p_\phi(\nu(w)), p_c(\eta), \epsilon) + 1 \leq T \\
& \wedge \forall u: u \neq w \rightarrow \nu(u) = \nu'(u) \\
& \wedge \eta' = \eta
\end{aligned}$$

$$\text{then } \mathbf{P}(s, s') = \frac{1}{|\text{init}_{p_\phi(s(w))}(s)|}.$$

The cases where a perturbed oscillator fires are analogous to oscillators with a maximal phase, except for the additional conditions that some other oscillator fired, and that all oscillators with higher phases have already been considered.

If  $p_\theta(\eta) = \text{update}$  and there is a  $w$  s.t. (6)

$$\begin{aligned}
& p_\theta(\nu(w)) = \text{start} \wedge p_\theta(\nu'(w)) = \text{update} \wedge p_\phi(\nu(w)) = p_\phi(\nu'(w)) \\
& \wedge p_\phi(\nu(w)) < T \wedge \exists u: p_\phi(\nu(u)) = T \\
& \wedge \forall u: u \neq w \rightarrow (p_\theta(\nu(u)) = \text{update} \vee p_\phi(\nu(u)) \leq p_\phi(\nu(w))) \\
& \wedge p_\phi(\nu(w)) + \Delta(p_\phi(\nu(w)), p_c(\eta), \epsilon) + 1 > T \\
& \wedge \forall u: u \neq w \rightarrow \nu(u) = \nu'(u) \\
& \wedge \eta' = \eta
\end{aligned}$$

$$\text{then } \mathbf{P}(s, s') = \frac{\mu}{|\text{init}_{p_\phi(s(w))}(s)|}$$

If  $p_\theta(\eta) = \text{update}$  and there is a  $w$  s.t. (7)

$$\begin{aligned}
& p_\theta(\nu(w)) = \text{start} \wedge p_\theta(\nu'(w)) = \text{update} \wedge p_\phi(\nu(w)) = p_\phi(\nu'(w)) \\
& \wedge p_\phi(\nu(w)) < T \wedge \exists u: p_\phi(\nu(u)) = T \\
& \wedge \forall u: u \neq w \rightarrow (p_\theta(\nu(u)) = \text{update} \vee p_\phi(\nu(u)) \leq p_\phi(\nu(w))) \\
& \wedge p_\phi(\nu(w)) + \Delta(p_\phi(\nu(w)), p_c(\eta), \epsilon) + 1 > T \\
& \wedge \forall u: u \neq w \rightarrow \nu(u) = \nu'(u) \\
& \wedge p_c(\eta') = p_c(\eta) + 1 \\
& \wedge p_\theta(\eta') = p_\theta(\eta)
\end{aligned}$$

$$\text{then } \mathbf{P}(s, s') = \frac{1 - \mu}{|\text{init}_{p_\phi(s(w))}(s)|}.$$

The final case 7, where all oscillators update their clock values simultaneously, is given by the following equation. It requires that all oscillators have finished their computation, whether they fire, and both the counter and the oscillators will reset

their mode to *start* after the transition.

$$\begin{aligned}
& \text{If } p_\theta(\eta) = \textit{update} \text{ and } p_\theta(\eta') = \textit{start} \text{ and} & (8) \\
& \text{for all } u \text{ we have } p_\theta(\nu(u)) = \textit{update} \wedge p_\theta(\nu'(u)) = \textit{start} \wedge F_{\textit{update}} \\
& \text{then } \mathbf{P}(s, s') = 1.
\end{aligned}$$

The formula  $F_{\textit{update}}$  is an abbreviation for the conjunction of the following four conditions, which model the update of the phases of the oscillators, according to the phase response function. Observe that the phases of the oscillators had not been updated by the previously defined transitions. Hence, we now update the phases of all oscillators at once.

$$\forall u: p_\phi(\nu(u)) = T \rightarrow \quad (8a)$$

$$p_\phi(\nu'(u)) = 1$$

$$\forall u: p_\phi(\nu(u)) < T \wedge p_\phi(\nu(u)) \leq R \rightarrow \quad (8b)$$

$$p_\phi(\nu'(u)) = p_\phi(\nu(u)) + 1$$

$$\forall u: p_\phi(\nu(u)) < T \wedge p_\phi(\nu(u)) > R \wedge \quad (8c)$$

$$p_\phi(\nu(u)) + \Delta(p_\phi(\nu(u)), p_c(\eta), \epsilon) + 1 \leq T \rightarrow$$

$$p_\phi(\nu'(u)) = p_\phi(\nu(u)) + \Delta(p_\phi(\nu(u)), p_c(\eta), \epsilon) + 1$$

$$\forall u: p_\phi(\nu(u)) < T \wedge p_\phi(\nu(u)) > R \wedge \quad (8d)$$

$$p_\phi(\nu(u)) + \Delta(p_\phi(\nu(u)), p_c(\eta), \epsilon) + 1 > T \rightarrow$$

$$p_\phi(\nu'(u)) = 1.$$

In this formula, (8a) handles the simple case of firing oscillators, while (8b) defines the behaviour of oscillators within their refractory period. The formulas (8c) and (8d) reflect the two cases where oscillators are perturbed, either not exceeding their oscillation cycle, or firing, respectively.

With this model, we could begin to analyse the synchronisation behaviour with respect to different phase response functions or broadcast failure probabilities. However, the state space of the model increases exponentially with the number of oscillators, which makes an analysis beyond small numbers of infeasible. To overcome this restriction, we increase the level of abstraction as presented in the next section.

## 5 Population Model

In this section, we define a *population model* of a network of pulse-coupled oscillators for parameters as defined in Sect. 4.1 as  $\mathcal{S} = (\Delta, N, T, R, \epsilon, \mu)$ . Oscillators in our model have identical dynamics, and two oscillators are indistinguishable if they share the same phase. That is, we can reason about groups of oscillators, instead of individuals. We therefore encode the global state of the model as a tuple  $\langle k_1, \dots, k_T \rangle$  where each  $k_\phi$  is the number of oscillators sharing a phase value of  $\phi$ . The population model does not account for the introduction of additional oscillators to a network, or the loss of existing coupled oscillators. That is, the population  $N$  remains constant.




Fig. 2: Evolution of the global state over four discrete time steps.

**Definition 6** A global state of a population model  $\mathcal{S} = (\Delta, N, T, R, \epsilon, \mu)$  is a  $T$ -tuple  $\sigma \in \{0, \dots, N\}^T$ , where  $\sigma = \langle k_1, \dots, k_T \rangle$  and  $\sum_{\phi=1}^T k_\phi = N$ . The set of all global states of  $\mathcal{S}$  is  $\Gamma(\mathcal{S})$ , or simply  $\Gamma$  when  $\mathcal{S}$  is clear from the context.

*Example 1* Figure 2 shows four global states for an instantiated population model of  $N = 8$  oscillators with  $T = 10$  discrete values for their phase and a refractory period of length  $R = 2$ . We assume that the phase response function is linear, that is,  $\Delta(\Phi, \alpha, \epsilon) = [\Phi \cdot \alpha \cdot \epsilon]$ , where  $[\cdot]$  denotes rounding to the closest integer. Furthermore, let  $\epsilon = 0.115$ . For example  $\sigma_0 = \langle 0, 0, 2, 1, 0, 0, 5, 0, 0, 0 \rangle$  is the global state where two oscillators have a phase of three, one oscillator has a phase of four, and five oscillators have a phase of seven. The starred node indicates the number of oscillators with phase ten that will fire in the next moment in time, while the shaded nodes indicate oscillators with phases that lie within the refractory period (one and two). If no oscillators have some phase  $\Phi$  then we omit the 0 in the corresponding node. Observe that, while going from  $\sigma_{i-1}$  to  $\sigma_i$  ( $1 \leq i \leq 3$ ), the oscillator phases increase by one. In the next section, we will explain how transitions between these global states are made. Note that directional arrows indicate cyclic direction, and do not represent transitions.

With every state  $\sigma \in \Gamma$  we associate a non-empty set of *failure vectors*, where each failure vector is a tuple of broadcast failures that could occur in  $\sigma$ .

**Definition 7** A *failure vector* is a  $T$ -tuple  $F = \langle f_1, \dots, f_T \rangle \in (\{0, \dots, N\} \cup \{\star\})^T$ , where  $f_i = \star$  implies  $f_j = \star$  for all  $1 \leq j \leq i$ . We denote the set of all possible failure vectors by  $\mathcal{F}$ .

Given a failure vector  $F = \langle f_1, \dots, f_T \rangle$ ,  $f_\Phi \in \{0, \dots, N\}$  indicates the number of broadcast failures that occur for all oscillators with a phase of  $\Phi$ . If  $f_\Phi = \star$  then no oscillators with a phase of  $\Phi$  fire. Semantically,  $f_\Phi = 0$  and  $f_\Phi = \star$  differ in that the former indicates that all (if any) oscillators with phase  $\Phi$  fire and no broadcast failures occur, while the latter indicates that all (if any) oscillators with a phase of  $\Phi$  do not fire. If no oscillators fire at all in a global state then we have only one possible failure vector, namely  $\{\star\}^T$ .

## 5.1 Transitions

In Section 5.2 we will describe how we can calculate the set of all possible failure vectors for a global state, and thereby identify all of its successor states. However we must first show how we can calculate the single successor state of a global state  $\sigma$ , given some failure vector  $F$ .

*Absorptions.* For real deployments of synchronisation protocols it is often the case that the duration of a single oscillation cycle will be at least several seconds [10, 28]. The perturbation induced by the firing of a group of oscillators may lead to groups of other oscillators to which they are coupled firing in turn. The firing of these other oscillators may then cause further oscillators to fire, and so forth, leading to a “chain reaction”, where each group of oscillators triggered to fire is *absorbed* by the initial group of firing oscillators. Since the whole chain reaction of absorptions may occur within just a few milliseconds, and in our model the oscillation cycle is a sequence of discrete states, when a chain reaction occurs the phases of all perturbed oscillators are updated at one single time step.

Since we are considering a fully connected network of oscillators, two oscillators sharing the same phase will have their phase updated to the same value in the next time step. They will always perceive the same number of other oscillators firing. Therefore, for each phase  $\Phi$  we define the function  $\alpha^\Phi: \Gamma \times \mathcal{F} \rightarrow \{0, \dots, N\}$ , where  $\alpha^\Phi(\sigma, F)$  is the number of oscillators with a phase greater than  $\Phi$  perceived to be firing by oscillators with phase  $\Phi$ , in some global state, incorporating the broadcast failures defined in the failure vector  $F$ . This allows us to encode the aforementioned chain reactions of firing oscillators. Note that our encoding of chain reactions results in a global semantics that differs from typical parallelisation operations, for example, the construction of the cross product of the individual oscillators. Observe that, in the concrete model of Sect. 4.2, we modelled such a behaviour by case 4.

Given a global state  $\sigma = \langle k_1, \dots, k_T \rangle$  and a failure vector  $F = \langle f_1, \dots, f_T \rangle$ , the following mutually recursive definitions show how we calculate the values  $\alpha^1(\sigma, F), \dots, \alpha^T(\sigma, F)$ , and how functions introduced in Sect. 4.1 are modified to indicate the update in phase, and firing, of *all* oscillators sharing the same phase  $\Phi$ . Observe that to calculate any  $\alpha^\Phi(\sigma, F)$  we only refer to definitions for phases greater than  $\Phi$  and the base case is  $\Phi = T$ , that is, values are computed from  $T$  down to 1. The function  $\text{ref}$  is the refractory function as defined in Sect. 4.1.

$$\text{update}^\Phi(\sigma, F) = 1 + \text{ref}(\Phi, \Delta(\Phi, \alpha^\Phi(\sigma, F), \epsilon)) \quad (9)$$

$$\text{fire}^\Phi(\sigma, F) = \text{update}^\Phi(\sigma, F) > T \quad (10)$$

$$\alpha^\Phi(\sigma, F) = \begin{cases} 0 & \text{if } \Phi=T \\ \alpha^{\Phi+1}(\sigma, F) + k_{\Phi+1} - f_{\Phi+1} & \text{if } \Phi < T, f_{\Phi+1} \neq \star \text{ and } \text{fire}^{\Phi+1}(\sigma, F) \\ \alpha^{\Phi+1}(\sigma, F) & \text{otherwise} \end{cases} \quad (11)$$

*Transition Function.* We now define the transition function that maps phase values to their updated values in the next time step. Note that since we no longer distinguish different oscillators with the same phase we only need to calculate a single value for their evolution and perturbation.

**Definition 8** The *phase transition function*  $\tau: \Gamma \times \{1, \dots, T\} \times \mathcal{F} \rightarrow \mathbb{N}$  maps a global state  $\sigma$ , a phase  $\Phi$ , and some possible failure vector  $F$  for  $\sigma$ , to the updated phase in the next discrete time step, with respect to the broadcast failures defined in  $F$ , and is defined as

$$\tau(\sigma, \Phi, F) = \begin{cases} 1 & \text{if } \text{fire}^\Phi(\sigma, F) \\ \text{update}^\Phi(\sigma, F) & \text{otherwise.} \end{cases} \quad (12)$$

Let  $\mathcal{U}_\Phi(\sigma, F)$  be the set of phase values  $\Psi$  where all oscillators with phase  $\Psi$  in  $\sigma$  will have their phase updated to  $\Phi$  in the next time step, with respect to the broadcast failures defined in  $F$ . Formally,

$$\mathcal{U}_\Phi(\sigma, F) = \{\Psi \mid \Psi \in \{1, \dots, T\} \wedge \tau(\sigma, \Psi, F) = \Phi\}. \quad (13)$$

We can now calculate the successor state of a global state  $\sigma$  and define how the model evolves over time.

**Definition 9** The *successor function*  $\vec{\text{succ}} : \Gamma \times \mathcal{F} \rightarrow \Gamma$  maps a global state  $\sigma$  and a failure vector  $F$  to a state  $\sigma'$ , and is defined as  $\vec{\text{succ}}(\langle k_1, \dots, k_T \rangle, F) = \langle k'_1, \dots, k'_T \rangle$ , where  $k'_\Phi = \sum_{\Psi \in \mathcal{U}_\Phi(\sigma, F)} k_\Psi$  for  $1 \leq \Phi \leq T$ .

*Example 2* Recall that the perturbation function of our example was given as  $\Delta(\Phi, \alpha, \epsilon) = [\Phi \cdot \alpha \cdot \epsilon]$ , where  $[\cdot]$  denotes rounding and  $\epsilon = 0.115$ . Consider the global state  $\sigma_2$  of Fig 3 where no oscillators will fire since  $k_{10} = 0$ . We therefore have one possible failure vector for  $\sigma_0$ , namely  $F = \{\star\}^{10}$ . Since no oscillators fire the dynamics of the oscillators are determined solely by their standalone evolution, and all oscillators simply increase their phase by 1 in the next time step. Now consider the global state  $\sigma_3$  and  $F = \langle \star, \star, \star, \star, \star, \star, 1, 0, 0, 0 \rangle$ , a possible failure vector for  $\sigma_3$ , indicating that oscillators with phases of 7 to 10 will fire and one broadcast failure will occur for the single oscillator that will fire with phase 7. Here a chain reaction occurs as the perturbation induced by the firing of the 5 oscillators causes the single oscillator with a phase of 7 to also fire. A broadcast failure occurs when this single oscillator fires, and the perturbation of the 5 firing oscillators is insufficient to cause the 2 oscillators with a phase of 6 to also fire. In the next state the oscillator with phase 7 has been absorbed by the group of the 5 oscillators that had phase 10.

More explicitly, since  $\text{fire}^{10}(\sigma_3, F)$  holds we have that  $\alpha^9(\sigma_3, F) = \alpha^{10}(\sigma_3, F) + k_{10} - f_{10} = 0 + 5 - 0 = 5$ . Now, since  $\Delta(9, 5, 0.14) = [9 \cdot 5 \cdot 0.115] = [5.175] = 5$ , we have  $\text{update}^9(\sigma_3, F) = 15 > 10$ , and thus,  $\text{fire}^9$  holds. Hence, we have that  $\alpha^8(\sigma_3, F) = \alpha^9(\sigma_3, F) + k_9 - f_9 = 0 + 5 - 0 = 5$ , and similarly, due to  $\Delta(8, 5, 0.115) = 5$ ,  $\text{fire}^8$  holds. That is, we have that  $\alpha^7(\sigma_3, F) = \alpha^8(\sigma_3, F) + k_8 - f_8 = 0 + 5 - 0 = 5$ . We then continue calculating  $\alpha^\Phi(\sigma_3, F)$  for  $6 \geq \Phi \geq 1$ , and noting that  $\Delta(6, 5, 0.115) = [3.45] = 3$ . Hence  $\text{fire}^6(\sigma_3, F)$  does not hold, and we obtain  $\alpha^1(\sigma_3, F) = \alpha^2(\sigma_3, F) = \alpha^3(\sigma_3, F) = \alpha^4(\sigma_3, F) = \alpha^5(\sigma_3, F) = \alpha^6(\sigma_3, F) = \alpha^7(\sigma_3, F) = 5$ . We conclude that  $\mathcal{U}_1(\sigma_3, F) = \{10, 9, 8, 7\}$ ,  $\mathcal{U}_{10}(\sigma_3, F) = \{6, 5\}$ ,  $\mathcal{U}_9(\sigma_3, F) = \{4, 3\}$ , and  $\mathcal{U}_\Phi(\sigma_3, F) = \emptyset$  for  $9 > \Phi > 3$ . Since  $R = 2$  we have that  $\mathcal{U}_3(\sigma_3, F) = \{2\}$  and  $\mathcal{U}_2(\sigma_3, F) = \{1\}$ . We calculate the successor of  $\sigma_3$  as  $\sigma_4 = \vec{\text{succ}}(\langle \star, \star, \star, \star, \star, \star, 1, 0, 0, 0 \rangle, F) = \langle k_{10} + k_9 + k_8 + k_7, k_1, k_2, 0, 0, 0, 0, 0, 0, k_4 + k_3, k_6 + k_5 \rangle = \langle 6, 0, 0, 0, 0, 0, 0, 0, 0, 2 \rangle$ .

**Lemma 1** *The number of oscillators is invariant during transitions, i.e., the successor function only creates tuples that are states of the given model. Formally, let  $\sigma = \langle k_1, \dots, k_T \rangle$  and  $\sigma' = \langle k'_1, \dots, k'_T \rangle$  be two states of a model  $\mathcal{S}$  such that  $\sigma' = \vec{\text{succ}}(\sigma, F)$ , where  $F$  is some possible failure vector for  $\sigma$ . Then  $\sum_{\Phi=1}^T k_\Phi = \sum_{\Phi=1}^T k'_\Phi = N$ .*

*Proof* Observe that the range of the function  $\tau$  is bound by  $T$ . By construction we can see that for any  $\sigma$ , for any possible failure vector  $F$  for  $\sigma$ , and for all






Fig. 3: Evolution of the global state over four discrete time steps.

$\Phi \in \{1, \dots, T\}$ , we have that  $1 \leq \tau(\sigma, \Phi, F) \leq T$ . Hence for all  $\Psi$  with  $1 \leq \Psi \leq T$ , there is a  $\Phi$  such that  $\Psi \in \mathcal{U}_\Phi(\sigma, F)$ . This implies  $\bigcup_{\Phi=1}^T \mathcal{U}_\Phi(\sigma, F) = \{1, \dots, T\}$ . Furthermore, there cannot be more than one  $\Phi$  such that  $\Psi \in \mathcal{U}_\Phi(\sigma, F)$ , since  $\tau$  is functional. Now we have  $\sum_{\Phi=1}^T k'_\Phi = \sum_{\Phi=1}^T \sum_{\Psi \in \mathcal{U}_\Phi(\sigma, F)} k_\Psi = \sum_{\Phi=1}^T k_\Phi = N$ .  $\square$

## 5.2 Failure Vector Calculation

We construct all possible failure vectors for a global state by considering every group of oscillators in decreasing order of phase. At each stage we determine if the oscillators would fire. If they fire then we consider each outcome where any, all, or none of the firings result in a broadcast failure. We then add a corresponding value to a partially calculated failure vector and consider the next group of oscillators with a lower phase. If the oscillators do not fire then there is nothing left to do, since by Def. 4 we know that  $\Delta$  is increasing, therefore all oscillators with a lower phase will also not fire. We can then pad the partial failure vector with  $\star$  appropriately to indicate that no failure could happen since no oscillator fired.

Table 1 illustrates how a possible failure vector for global state  $\sigma_3$  in Fig. 3 is iteratively constructed. The first three columns respectively indicate the current iteration  $i$ , the global state  $\sigma_3$  with the currently considered oscillators underlined, and the elements of the failure vector  $F$  computed so far. The fourth column is *true* if the oscillators with phase  $T+1-i$  would fire given the broadcast failures in the partial failure vector. We must consider all outcomes of any or all firings resulting in broadcast failure. The final column therefore indicates whether the value added to the partial failure vector in the current iteration is the only possible value (*false*), or a choice from one of several possible values (*true*).

Initially we have an empty partial failure vector. At the first iteration there are 5 oscillators with a phase of 10. These oscillators will fire so we must consider each

Table 1: Construction of a possible failure vector for a global state  $\sigma_3 = \langle 0, 0, 0, 0, 0, 2, 1, 0, 0, 5 \rangle$ .

iteration ( $i$ )	$\pi_1$	failure vector $B$	fired	branches
0	$\langle 0, 0, 0, 0, 0, \underline{2}, 1, 0, 0, 5 \rangle$	$\langle \rangle$	–	<i>false</i>
1	$\langle 0, 0, 0, 0, 0, \underline{2}, 1, 0, 0, \underline{5} \rangle$	$\langle 0 \rangle$	<i>true</i>	<i>true</i>
2	$\langle 0, 0, 0, 0, 0, \underline{2}, 1, 0, \underline{0}, 5 \rangle$	$\langle 0, 0 \rangle$	<i>true</i>	<i>false</i>
3	$\langle 0, 0, 0, 0, 0, \underline{2}, 1, \underline{0}, 0, 5 \rangle$	$\langle 0, 0, 0 \rangle$	<i>true</i>	<i>false</i>
4	$\langle 0, 0, 0, 0, 0, \underline{2}, 1, 0, 0, \underline{5} \rangle$	$\langle 1, 0, 0, 0 \rangle$	<i>true</i>	<i>true</i>
5	$\langle 0, 0, 0, 0, 0, \underline{2}, 1, 0, 0, 5 \rangle$	$\langle \star, \star, \star, \star, \star, \star, 1, 0, 0, 0 \rangle$	<i>false</i>	–

case where 0, 1, 2, 3, 4 or 5 broadcast failures occur. Here we choose 0 broadcast failures, which is then added to the partial failure vector. At iterations 2 and 3 the oscillators would have fired, but since there are no oscillators with a phase of 9 or 8 we only have one possible value to add to the partial failure vector, namely 0. At iteration 4 a single oscillator with a phase of 7 fires, and we choose the case where the firing resulted in a broadcast failure. In the final iteration oscillators with a phase of 6 do not fire, hence we can conclude that oscillators with phases less than 6 also do not fire, and can fill the partial failure vector appropriately with  $\star$ .

Formally, we define a family of functions  $fail$  indexed by  $\Phi$ , where each  $fail_\Phi$  takes as parameters some global state  $\sigma$ , and  $V$ , a vector of length  $T - \Phi$ .  $V$  represents all broadcast failures for all oscillators with a phase greater than  $\Phi$ . The function  $fail_\Phi$  then computes the set of all possible failure vectors for  $\sigma$  with suffix  $V$ . Here we use the notation  $v \frown v'$  to indicate vector concatenation.

**Definition 10** We define  $fail_\Phi : \Gamma \times \{0, \dots, N\}^{T-\Phi} \rightarrow \mathbb{P}(\{0, \dots, N\} \cup \{\star\}^T)$ , for  $1 \leq \Phi \leq T$ , as the family of functions indexed by  $\Phi$ , where  $\sigma = \langle k_1, \dots, k_T \rangle$  and

$$fail_\Phi(\sigma, V) = \begin{cases} \bigcup_{k=0}^{k_\Phi} fail_{\Phi-1}(\sigma, \langle k \rangle \frown V) & \text{if } 1 < \Phi \leq T \text{ and } fire^\Phi(\sigma, \{\star\}^\Phi \frown V) \\ \bigcup_{k=0}^{k_1} \{\langle k \rangle \frown V\} & \text{if } \Phi = 1 \text{ and } fire^1(\sigma, \{\star\} \frown V) \\ \{\{\star\}^\Phi \frown V\} & \text{otherwise} \end{cases}$$

Observe that the result of  $fail_T$  is always a set of well defined failure vectors, since whenever  $\star$  is introduced into a failure vector at index  $\Phi$ , all preceding indices are also filled with  $\star$ , as required by Definition 7.

**Definition 11** Given a global state  $\sigma \in \Gamma$ , we define  $\mathcal{F}_\sigma$ , the set of all possible failure vectors for that state, as  $\mathcal{F}_\sigma = fail_T(\sigma, \langle \rangle)$ , and define  $next(\sigma)$ , the set of all successor states of  $\sigma$ , as  $next(\sigma) = \{\overrightarrow{\text{succ}}(\sigma, F) \mid F \in \mathcal{F}_\sigma\}$ .

Note that for some global states  $|next(\sigma)| < |\mathcal{F}_\sigma|$ , since we may have that  $\overrightarrow{\text{succ}}(\sigma, F) = \overrightarrow{\text{succ}}(\sigma, F')$  for some  $F, F' \in \mathcal{F}_\sigma$  with  $F \neq F'$ .

Given a global state  $\sigma$  and a failure vector  $F \in \mathcal{F}_\sigma$ , we will now compute the probability of a transition being made to state  $\overrightarrow{\text{succ}}(\sigma, F)$  in the next time step. Recall that  $\mu$  is the probability with which a broadcast failure occurs. Firstly we define the probability mass function PMF :  $\{1, \dots, N\}^2 \rightarrow [0, 1]$ , where  $PMF(k, f)$  gives the probability of  $f$  broadcast failures occurring given that  $k$  oscillators fire, as  $PMF(k, f) = \mu^f (1 - \mu)^{k-f} \binom{k}{f}$ . We then denote by PFV :  $\Gamma \times \mathcal{F}_\sigma \rightarrow [0, 1]$  the function mapping a possible broadcast failure vector  $F$  for  $\sigma$ , to the probability of the failures in  $F$  occurring. That is,

$$PFV(\langle k_1, \dots, k_T \rangle, \langle f_1, \dots, f_T \rangle) = \prod_{\Phi=1}^T \begin{cases} PMF(k_\Phi, f_\Phi) & \text{if } f_\Phi \neq \star \\ 1 & \text{otherwise} \end{cases} \quad (14)$$

**Lemma 2** For any global state  $\sigma$ , PFV is a discrete probability distribution over  $\mathcal{F}_\sigma$ . Formally,  $\sum_{F \in \mathcal{F}_\sigma} PFV(\sigma, F) = 1$ .

*Proof* Given a global state  $\sigma = \langle k_1, \dots, k_T \rangle$  we can construct a tree of depth  $T$  where each leaf node is labelled with a possible failure vector for  $\sigma$ , and each node  $\Lambda$  at depth  $\Phi$  is labelled with a vector of length  $\Phi$  corresponding to the last  $\Phi$  elements of a possible failure vector for  $\sigma$ . We denote the label of a node  $\Lambda$  by  $V(\Lambda)$ . We label each node  $\Lambda_\omega$  with  $\langle \omega \rangle \frown V(\Lambda)$ . We iteratively construct the tree, starting with the root node,  $root$ , at depth 0, which we label with the empty tuple  $\langle \rangle$ . For each node  $\Lambda$  at depth  $0 \leq \Phi < T$  we construct the children of  $\Lambda$  as follows:

1. If oscillators with phase  $\Phi$  fire we define the sample space  $\Omega = \{0, \dots, n_\Phi\}$  to be a set of disjoint events, where each  $\omega \in \Omega$  is the event where  $\omega$  broadcast failures occur, given that  $k_\Phi$  oscillators fired. For each  $\omega \in \Omega$  there is a child  $\Lambda_\omega$  of  $\Lambda$  with label  $\langle \omega \rangle \frown V(\Lambda)$ , and we label the edge from  $\Lambda$  to  $\Lambda_\omega$  with  $\text{PMF}(k_\Phi, \omega)$ .
2. If oscillators with phase  $\Phi$  do not fire then  $\Lambda$  has a single child  $\Lambda_\star$  labelled with  $\langle \star \rangle \frown V(\Lambda)$ , and we label the edge from  $\Lambda$  to  $\Lambda_\star$  with 1.

We denote the label of an edge from a node  $\Lambda$  to its child  $\Lambda'$  by  $L(\Lambda, \Lambda')$ . For case 2 we can observe that if oscillators with phase  $\Phi$  do not fire then we know that oscillators with any phase  $\Psi < \Phi$  will also not fire, since from Def. 4 we know that  $\Delta$  is an increasing function. Hence, all descendants of  $\Lambda$  will also have a single child, with an edge labelled with 1, and each node is labelled with the label of its parent, prefixed with  $\langle \star \rangle$ .

After constructing the tree we have a vector of length  $T$  associated with each leaf node, corresponding to a failure vector for  $\sigma$ . The set  $\mathcal{F}_\sigma$  of all possible failure vectors for  $\sigma$  is therefore the set of all vectors labelling leaf nodes. We denote by  $\mathbf{P}^\downarrow(\Lambda)$  the product of all labels on edges along the path from  $\Lambda$  back to the root. Given a global state  $\sigma = \langle k_1, \dots, k_T \rangle$  and a failure vector  $F = \langle f_1, \dots, f_T \rangle \in \mathcal{F}_\sigma$  labelling some leaf node  $\Lambda$  at depth  $T$ , we can see that

$$\mathbf{P}^\downarrow(\Lambda) = 1 \cdot \prod_{\Phi=1}^T \begin{cases} \text{PMF}(k_\Phi, f_\Phi) & \text{if } f_\Phi \neq \star \\ 1 & \text{otherwise} \end{cases} = \text{PFV}(\sigma, F).$$

Let  $D^\Phi$  denote the set of all nodes at depth  $\Phi$ . We show  $\sum_{d \in D^\Phi} \mathbf{P}^\downarrow(d) = 1$  by induction on  $\Phi$ . For  $\Phi = 0$ , i.e.,  $D^\Phi = \{root\}$ , the property holds by definition. Now assume that  $\sum_{d \in D^\Phi} \mathbf{P}^\downarrow(d) = 1$  holds for some  $0 \leq \Phi < T$ . Let  $\Lambda$  be some node in  $D^\Phi$ , and let  $C^\Lambda$  be the set of all children of  $\Lambda$ . Consider the following two cases: If oscillators with phase  $\Phi$  do not fire then  $|C^\Lambda| = 1$ , and for the only  $c \in C^\Lambda$  we have that  $L(\Lambda, c) = 1$ . If oscillators with phase  $\Phi$  fire observe that  $\text{PMF}$  is a probability mass function for a random variable defined on the sample space  $\Omega = \{0, \dots, k_\Phi\}$ . In either case we can see that  $\sum_{c \in C^\Lambda} L(\Lambda, c) = 1$ . Note that  $D^{\Phi+1} = \bigcup_{d \in D^\Phi} C^d$ , and recall that  $L(d, c) \cdot \mathbf{P}^\downarrow(d) = \mathbf{P}^\downarrow(c)$ . Therefore,

$$\sum_{d \in D^{\Phi+1}} \mathbf{P}^\downarrow(d) = \sum_{d \in D^\Phi} \sum_{c \in C^d} L(d, c) \cdot \mathbf{P}^\downarrow(d) = \sum_{d \in D^\Phi} \left( \mathbf{P}^\downarrow(d) \sum_{c \in C^d} L(d, c) \right).$$

Since  $\sum_{c \in C^d} L(d, c) = 1$  for each  $d \in D^\Phi$ , and from the induction hypothesis, we then have that

$$\sum_{d \in D^\Phi} \left( \mathbf{P}^\downarrow(d) \sum_{c \in C^d} L(d, c) \right) = \sum_{d \in D^\Phi} \mathbf{P}^\downarrow(d) = 1.$$

We have already shown that  $P^\perp(A) = \text{PFV}(\sigma, F)$  for any leaf node  $A$  labelled with a failure vector  $F$ , and since the set of all labels for leaf nodes is  $\mathcal{F}_\sigma$  we can conclude that

$$\sum_{F \in \mathcal{F}_\sigma} \text{PFV}(\sigma, F) = \sum_{d \in D^T} P^\perp(d) = 1.$$

This proves the lemma.  $\square$

*Example 3* We consider again the global states  $\sigma_3 = \langle 0, 0, 0, 0, 0, 2, 1, 0, 0, 5 \rangle$  and  $\sigma_4 = \langle 6, 0, 0, 0, 0, 0, 0, 0, 2 \rangle$ , given in Fig. 3, of the population model instantiated in Example 1, and the failure vector  $F = \langle *, *, *, *, *, *, 1, 0, 0, 0 \rangle$  given in Example 2, noting that  $F \in \mathcal{F}_{\sigma_3}$ ,  $\vec{\text{succ}}(\sigma_3, F) = \sigma_4$ , and  $\mu = 0.1$ . We calculate the probability of a transition being made from  $\sigma_3$  to  $\sigma_4$  as

$$\begin{aligned} & \text{PFV}(\langle 0, 0, 0, 0, 0, 2, 1, 0, 0, 5 \rangle, \langle *, *, *, *, *, *, 1, 0, 0, 0 \rangle) \\ &= 1 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot \text{PMF}(1, 1) \cdot \text{PMF}(0, 0) \cdot \text{PMF}(0, 0) \cdot \text{PMF}(5, 0) \\ &= (0.1^1 \cdot 0.9^0 \cdot 1) \cdot (1) \cdot (1) \cdot (0.1^0 \cdot 0.9^5 \cdot 1) = 0.059049 \end{aligned}$$

We now have everything we need to fully describe the evolution of the global state of a population model over time. An execution path of a population model  $\mathcal{S}$  is an infinite sequence of global states  $\omega = \sigma_0 \sigma_1 \sigma_2 \sigma_3 \dots$ , where  $\sigma_0$  is called the *initial state*, and  $\sigma_{k+1} \in \text{next}(\sigma)$  for all  $k \geq 0$ .

### 5.3 Synchronisation

When all oscillators in a population model have the same phase in a global state we say that the state is *synchronised*. Formally, a global state  $\sigma = \langle k_1, \dots, k_T \rangle$  is *synchronised* if, and only if, there is some  $\Phi \in \{1, \dots, T\}$  such that  $k_\Phi = N$ , and hence  $k_{\Phi'} = 0$  for all  $\Phi' \neq \Phi$ . We will often want to reason about whether some particular run  $\omega$  of a model leads to a global state that is synchronised. We say that a path  $\omega = \sigma_0 \sigma_1 \dots$  synchronises if, and only if, there exists some  $k \geq 0$  such that  $\sigma_k$  is synchronised. Once a synchronised global state is reached any successor states will also be synchronised. Finally we can say that a model synchronises if, and only if, all runs of the model synchronise.

### 5.4 Model Construction

Given a population model  $\mathcal{S} = (\Delta, N, T, R, \epsilon, \mu)$  we construct a DTMC  $D(\mathcal{S}) = (Q, \sigma_0, \mathbf{P}, L)$  where  $L$  ranges over the singleton  $\{\text{synch}\}$ . We define the set of states  $Q$  to be  $\Gamma(\mathcal{S}) \cup \{\sigma_0\}$ , where  $\sigma_0$  is the initial state of the DTMC. For each  $\sigma = \langle k_1, \dots, k_T \rangle \in \Gamma(\mathcal{S})$ , we set  $L(\sigma) = \{\text{synch}\}$  if  $k_T = N$ .

In the initial state all oscillators are *unconfigured*. That is, oscillators have not yet been assigned a value for their phase. For each  $\sigma = \langle k_1, \dots, k_T \rangle \in Q \setminus \{\sigma_0\}$  we define

$$\mathbf{P}(\sigma_0, \sigma) = \frac{1}{T^N} \binom{N}{k_1, \dots, k_T}$$

to be the probability of moving from  $\sigma_0$  to a state where  $k_i$  arbitrary oscillators are configured with the phase value  $i$  for  $1 \leq i \leq T$ . The multinomial coefficient defines the number of possible assignments of phases to distinct oscillators that result in the global state  $\sigma$ . The fractional coefficient normalises the multinomial coefficient with respect to the total number of possible assignments of phases to all oscillators. In general, given an arbitrary set of initial configurations (global states) for the oscillators, the total number of possible phase assignments can be calculated by computing the sum of the multinomial coefficients for each configuration (global state) in that set. Since  $\Gamma$  is the set of all possible global states, we have that

$$\sum_{\langle k_1, \dots, k_T \rangle \in \Gamma} \binom{N}{k_1, \dots, k_T} = T^N.$$

We assign probabilities to the transitions as follows: for every  $\sigma \in Q \setminus \{\sigma_0\}$ , we consider each  $F \in \mathcal{F}_\sigma$ , and set  $\mathbf{P}(\sigma, \vec{\text{succ}}(\sigma, F)) = \text{PFV}(\sigma, F)$ . For every combination of  $\sigma$  and  $\sigma'$  where  $\sigma' \notin \text{next}(\sigma)$  we set  $\mathbf{P}(\sigma, \sigma') = 0$ .

### 5.5 Model Reduction

We now describe a reduction of the population model that results in a significant decrease in the size of the model, but is equivalent to the original model with respect to the reachability of synchronised states. We first distinguish between states where one or more oscillators are about to fire, and states where no oscillators will fire at all. We refer to these states as *firing states* and *non-firing states* respectively.

**Definition 12** Given a population model  $\mathcal{S}$ , a global state  $\langle k_1, \dots, k_T \rangle \in \Gamma$  is a *firing state* if, and only if,  $k_T > 0$ . We denote by  $\Gamma^F$  the set of all firing states of  $\mathcal{S}$ , and denote by  $\Gamma^{\text{NF}} = \Gamma \setminus \Gamma^F$  the set of all non-firing states of  $\mathcal{S}$ . We will again omit  $\mathcal{S}$  if it is clear from the context

Given a DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  let  $|\mathbf{P}| = |\{(t, t') \mid t, t' \in Q^2 \text{ and } \mathbf{P}(t, t') > 0\}|$  be the number of non-zero transitions in  $\mathbf{P}$ , and  $|D| = |Q| + |\mathbf{P}|$  to be the total number of states and non-zero transitions in  $D$ .

**Theorem 1** For every population model  $\mathcal{S}$  and its corresponding DTMC  $D(\mathcal{S}) = (Q, \sigma_0, \mathbf{P}, L)$ , there is a reduced model  $D'(\mathcal{S}) = (Q', \sigma_0, \mathbf{P}', L')$  where  $|D'(\mathcal{S})| < |D(\mathcal{S})|$  and unbounded-time reachability properties with respect to synchronised firing states in  $D(\mathcal{S})$  are preserved in  $D'(\mathcal{S})$ . In particular, the states and transitions in  $D(\mathcal{S})$  are reduced in  $D'(\mathcal{S})$  such that  $Q' = Q \setminus \Gamma^{\text{NF}}$  and

$$\begin{aligned} |Q'| &= 1 + \frac{T^{(N-1)}}{(N-1)!}, \\ |\mathbf{P}'| &\leq |\mathbf{P}| - 2|\Gamma^{\text{NF}}| \end{aligned}$$

where  $x^{(n)}$  is the rising factorial.

We now proceed to prove this theorem. To that end, we need some preliminary properties of non-firing states and their relation to firing states.

**Lemma 3** *Every non-firing state  $\sigma \in \Gamma^{\text{NF}}$  has exactly one successor state, and in that state all oscillator phases have increased by 1.*

*Proof* Given a non-firing state  $\sigma = \langle k_1, \dots, k_T \rangle$  observe that as  $k_T = 0$  there is only one possible failure vector for  $\sigma$ , namely  $\{\star\}^T$ . The set of all successor states of  $\sigma$  is then the singleton  $\{\vec{\text{succ}}(\sigma, \{\star\}^T)\}$ . By construction we can then see that  $\text{update}^\Phi(\sigma, \{\star\}^T) = 1$  and  $\mathcal{U}_\Phi(\sigma, \{\star\}^T) = \{\Phi - 1\}$  for  $1 \leq \Phi \leq T$ . The single successor state is then given by  $\vec{\text{succ}}(\sigma, \{\star\}^T) = \langle 0, k_1, \dots, k_{T-1} \rangle$ .  $\square$

**Corollary 1** *An immediate corollary of Lemma 3 is that a transition from any non-firing state is taken deterministically, since for any  $\sigma \in \Gamma^{\text{NF}}$  we have  $\text{PFV}(\sigma, \{\star\}^T) = 1$ .*

*Reachable State Reduction.* Given a path  $\omega = \sigma_0 \dots \sigma_{n-1} \sigma_n$  where  $\sigma_i \in \Gamma^{\text{NF}}$  for  $0 < i < n$  and  $\sigma_0, \sigma_n \in \Gamma^{\text{F}}$ , we omit transitions  $(\sigma_i, \sigma_{i+1})$  for  $0 \leq i < n$ , and instead introduce a direct transition from  $\sigma_0$ , the first firing state, to  $\sigma_n$ , the next firing state in the sequence. For any  $\sigma = \langle k_1, \dots, k_T \rangle \in \Gamma$  let  $\delta_\sigma = \max\{\Phi \mid k_\Phi > 0 \text{ and } 1 \leq \Phi \leq T\}$  be the highest phase of any oscillator in  $\sigma$ . The successor state of a non-firing state is then the state where all phases have increased by  $T - \delta_\sigma$ . Observe that  $T - \delta_\sigma = 0$  for any  $\sigma \in \Gamma^{\text{F}}$ .

**Definition 13** The *deterministic successor function*  $\vec{\text{succ}} : \Gamma \rightarrow \Gamma^{\text{F}}$ , given by

$$\vec{\text{succ}}(\langle k_1, \dots, k_T \rangle) = \{0\}^{T-\delta_\sigma} \widehat{\langle k_1, \dots, k_{\delta_\sigma} \rangle},$$

maps a state  $\sigma \in \Gamma$  to the next firing state reachable by taking  $T - \delta_\sigma$  deterministic transitions. Observe that for any firing state  $\sigma$  we have  $\delta_\sigma = T$ , and hence that  $\vec{\text{succ}}(\sigma) = \sigma$ .

We now update the definition for the set of all successor states for some global state  $\sigma \in \Gamma$  to incorporate the deterministic successor function.

**Definition 14** Given a global state  $\sigma \in \Gamma$ , we define  $\vec{\text{next}}(\sigma)$  to be the set of all successor states of  $\sigma$ , where

$$\vec{\text{next}}(\sigma) = \{\vec{\text{succ}}(\vec{\text{succ}}(\sigma, F)) \mid F \in \mathcal{F}_\sigma\}.$$

**Definition 15** Given a firing state  $\sigma \in \Gamma^{\text{F}}$  let  $\text{pred}(\sigma)$  be the set of all non-firing predecessors of  $\sigma$ , where  $\sigma$  is reachable from the predecessor by taking some positive number of transitions deterministically. Formally,

$$\text{pred}(\sigma) = \{\sigma' \mid \sigma' \in \Gamma^{\text{NF}} \text{ and } \vec{\text{succ}}(\sigma') = \sigma\}.$$

We refer to all states  $\sigma' \in \text{pred}(\sigma)$  as *deterministic predecessors* of  $\sigma$ .

Then given  $D = (Q, \sigma_0, \mathbf{P}, L)$  with  $Q = \{\sigma_0\} \cup \Gamma$ , we define  $Q' = Q \setminus \bigcup_{\sigma \in \Gamma^{\text{F}}} \text{pred}(\sigma)$  to be the reduction of  $Q$  where all non-firing states from which a firing state can be reached deterministically are removed.

**Lemma 4** *For any  $D(\mathcal{S}) = (Q, \sigma_0, \mathbf{P}, L)$  with  $Q = \Gamma \cup \{\sigma_0\}$ , the reduction  $Q'$  is equal to  $\Gamma^{\text{F}} \cup \{\sigma_0\}$ .*

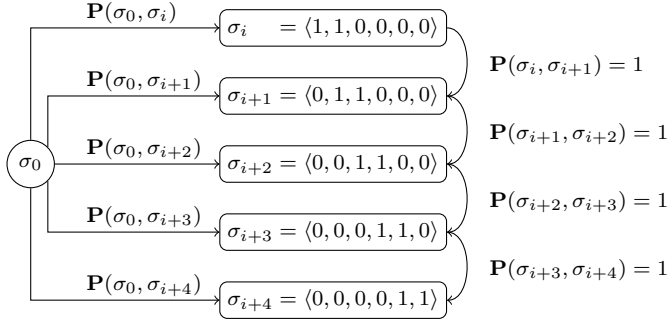


Fig. 4: Five possible initial configurations in  $Q$  for  $N = 2$ ,  $T = 6$ .

*Proof* Let  $P = \bigcup_{\sigma \in \Gamma^F} \text{pred}(\sigma)$  be the set of all predecessors of firing states in  $\Gamma^F$ . Since  $Q = \Gamma \cup \{\sigma_0\}$  and  $Q' = Q \setminus P$  we can see that  $Q' = \Gamma^F \cup \{\sigma_0\}$  if, and only if,  $P = \Gamma^{\text{NF}}$ . From Definition 15 it follows that  $P \subseteq \Gamma^{\text{NF}}$ . In addition, for any  $\sigma \in \Gamma^{\text{NF}}$  there is some state  $\sigma'$  such that  $\sigma \in \text{pred}(\sigma')$  and  $\sigma' = \text{succ}(\sigma) \in \Gamma^F$ , hence  $\Gamma^{\text{NF}} \subseteq P$  and the lemma is proved.  $\square$

**Lemma 5** For a population model  $S = (\Delta, N, T, R, \epsilon, \mu)$  and its corresponding DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  with  $Q = \Gamma \cup \{\sigma_0\}$ , the number of states in the reduction of  $Q$  is given by  $|Q'| = 1 + \frac{T^{(N-1)}}{(N-1)!}$ , where  $x^{(n)}$  is the rising factorial.

*Proof* Observe that there are  $\binom{N+T-1}{N}$  ways to assign  $T$  distinguishable phases to  $N$  indistinguishable oscillators [15]. Since  $Q = \Gamma \cup \{\sigma_0\}$  and  $\Gamma$  is the set of all possible configurations for oscillators we can see that  $|Q| = \binom{N+T-1}{N} + 1$ . For any non-firing state  $\sigma = \langle k_1, \dots, k_T \rangle \in \Gamma^{\text{NF}}$  we know from Definition 6 that  $\sum_{\phi=1}^T k_\phi = N$  and from Definition 12 that  $k_T = 0$ , so it must be the case that  $\sum_{\phi=1}^{T-1} k_\phi = N$ . That is, there must be  $\binom{N+T-2}{N}$  ways to assign  $T-1$  distinguishable phases to  $N$  indistinguishable oscillators, and so  $|\Gamma^{\text{NF}}| = \binom{N+T-2}{N}$ . From Lemma 4 we know that  $Q' = Q \setminus \Gamma^{\text{NF}}$  so it must be the case that  $|Q'| = |Q| - |\Gamma^{\text{NF}}| = 1 + \binom{N+T-1}{N} - \binom{N+T-2}{N} = 1 + \frac{T^{(N-1)}}{(N-1)!}$ .  $\square$

*Transition Matrix Reduction.* Here we describe the reduction in the number of non-zero transitions in the model. We illustrate how initial transitions to non-firing states are removed by using a simple example, and then describe how we remove transitions from firing states to any successor non-firing states..

Figure 4 shows five possible initial configurations  $\sigma_i, \dots, \sigma_{i+4} \in Q$  for  $N = 2$  oscillators with  $T = 6$  values for phase, where a transition is taken from  $\sigma_0$  to each  $\sigma_k$  with probability  $\mathbf{P}(\sigma_0, \sigma_k)$ . Any infinite run of  $D$  where a transition is taken from  $\sigma_0$  to one of the configured states  $\sigma_i, \dots, \sigma_{i+3}$  will pass through  $\sigma_{i+4}$ , since all transitions  $(\sigma_{i+k}, \sigma_{i+k+1})$  for  $0 \leq k \leq 3$  are taken deterministically. Also, observe that states  $\sigma_i, \dots, \sigma_{i+3}$  are not in  $Q'$ , since  $\sigma_{i+4}$  is reachable from each by taking some number of deterministic transitions. We therefore set the probability of moving from  $\sigma_0$  to  $\sigma_{i+4}$  in  $\mathbf{P}'$  to be the sum of the probabilities of moving from

$\sigma_0$  to  $\sigma_{i+4}$  and each of its predecessors in  $\mathbf{P}$ . Generally, given a state  $\sigma \in Q'$  where  $\sigma \neq \sigma_0$ , we set  $\mathbf{P}'(\sigma_0, \sigma) = \mathbf{P}(\sigma_0, \sigma) + \sum_{\sigma' \in \text{pred}(\sigma)} \mathbf{P}(\sigma_0, \sigma')$ .

We now define how we calculate the probability with which a transition is taken from a firing state to each of its possible successors. For each firing state  $\sigma \in Q'$  we consider each possible successor  $\sigma' \in \overrightarrow{\text{next}}(\sigma)$  of  $\sigma$  and define  $\mathcal{F}_{\sigma \rightarrow \sigma'}$  to be the set of all possible failure vectors for  $\sigma$  for which the successor of  $\sigma$  is  $\sigma'$ , given by  $\mathcal{F}_{\sigma \rightarrow \sigma'} = \{F \in \mathcal{F}_\sigma \mid \overrightarrow{\text{succ}}(\overrightarrow{\text{succ}}(\sigma, F)) = \sigma'\}$ . We then set the probability with which a transition from  $\sigma$  to  $\sigma'$  is taken to  $\mathbf{P}'(\sigma, \sigma') = \sum_{F \in \mathcal{F}_{\sigma \rightarrow \sigma'}} \text{PFV}(\sigma, F)$ .

**Lemma 6** *For a population model  $\mathcal{S} = (\Delta, N, T, R, \epsilon, \mu)$ , the corresponding DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  with  $Q = \{\sigma_0\} \cup \Gamma$ , and its reduction  $D'(\mathcal{S}) = (Q', \sigma_0, \mathbf{P}', L')$ , the transitions in  $\mathbf{P}$  are reduced in  $\mathbf{P}'$  such that  $|\mathbf{P}'| \leq |\mathbf{P}| - 2|\Gamma^{\text{NF}}|$*

*Proof* From Lemma 4 we know that  $|Q'| = |Q \setminus \Gamma^{\text{NF}}|$ , and hence that  $|\Gamma^{\text{NF}}|$  transitions from  $\sigma_0$  to non-firing states are not in  $\mathbf{P}'$ , and from Lemma 3 we also know that there is one transition from each non-firing state to its unique successor state that is not in  $\mathbf{P}'$ . Since no additional transitions are introduced in the reduction it is clear that  $|\mathbf{P}'| \leq |\mathbf{P}| - 2|\Gamma^{\text{NF}}|$ .  $\square$

**Lemma 7** *For every population model DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$ , unbounded-time reachability properties with respect to synchronised firing states in  $D$  are preserved in its reduction  $D'$ .*

*Proof* We want to show that for every  $\bowtie \in \{<, \leq, \geq, >\}$  and every  $\lambda \in [0, 1]$ , if  $\sigma_0 \models \text{P}_{\bowtie \lambda}[\text{F synch}]$  holds in  $D$  then it also holds in  $D'$ . From the semantics of PCTL over a DTMC we have

$$\sigma_0 \models \text{P}_{\bowtie \lambda}[\text{F synch}] \quad \Leftrightarrow \quad \text{Pr}\{\omega \in \text{Paths}^D \mid \omega \models \text{F synch}\} \bowtie \lambda.$$

Therefore we need to show that

$$\text{Pr}^D\{\omega \in \text{Paths}^D \mid \omega \models \text{F synch}\} = \text{Pr}^{D'}\{\omega' \in \text{Paths}^{D'} \mid \omega' \models \text{F synch}\},$$

where  $\text{Pr}^D$  and  $\text{Pr}^{D'}$  denote the probability measures with respect to the sets of infinite paths from  $\sigma_0$  in  $D$  and  $D'$  respectively.

Given a firing state  $\sigma^f \in Q$  we denote by  $\text{Paths}_{\sigma^f}^D$  the set of all infinite paths of  $D$  starting in  $\sigma_0$  where the first firing state reached along that path is  $\sigma^f$ . All such sets for all firing states in  $Q$  form a partition, such that  $\bigcup_{\sigma^f \in \Gamma^f} \text{Paths}_{\sigma^f}^D = \text{Paths}^D$ . That is, for all firing states  $\sigma^f, \sigma^{f'} \in Q$  where  $\sigma^f \neq \sigma^{f'}$  we have that  $\text{Paths}_{\sigma^f}^D \cap \text{Paths}_{\sigma^{f'}}^D = \emptyset$ .

Now observe that any infinite path  $\omega$  of  $D$  can be written in the form  $\omega = \sigma_0 \omega_1^{\text{NF}} \sigma_1^f \omega_2^{\text{NF}} \sigma_2^f \dots$  where  $\sigma_i^f$  is the  $i^{\text{th}}$  firing state in the path and each  $\omega_i^{\text{NF}} = \sigma_i^1 \sigma_i^2 \dots \sigma_i^{k_i}$  is a possibly empty sequence of  $k_i$  non-firing states. Then for every such path in  $D$  there is a corresponding path  $\omega'$  of  $D'$  without non-firing states, and of the form  $\omega' = \sigma_0 \sigma_1^f \sigma_2^f \sigma_3^f \dots$ , as for any  $i$  we have  $\sigma_i^j \in \text{pred}(\sigma_i^f)$  for all  $1 \leq j \leq k_i$ . As only deterministic transitions have been removed in  $D'$  we can see that  $\text{Pr}^D\{\sigma_1^f \omega_2^{\text{NF}} \sigma_2^f \dots\} = \text{Pr}^{D'}\{\sigma_1^f \sigma_2^f \sigma_3^f \dots\}$ . Hence, we only have to consider the finite paths from  $\sigma_0$  to  $\sigma_1^f$ . To that end, observe that there are  $|\text{pred}(\sigma_1^f)|$  possible prefixes for each path from  $\sigma_0$  to  $\sigma_1^f$  where the initial transition is taken from  $\sigma_0$  to some non-firing predecessor of  $\sigma_1^f$ , plus the single prefix where the initial transition is taken to  $\sigma_1^f$  itself. Overall there are exactly  $|\text{pred}(\sigma_1^f)| + 1$  distinct finite prefixes



Table 2: Reduction in state space and transitions for different model instances.

$N$	$T$	$D$		$D'$		Reduction (%)	
		States	Transitions	States	Transitions	States	Transitions
3	6	113	188	22	52	80.5	72.3
5	6	505	1030	127	389	74.9	62.2
8	6	2575	7001	793	3154	69.2	54.9
3	8	241	410	37	97	84.6	76.3
5	8	1585	3250	331	1097	79.1	66.2
8	8	12871	34615	3433	14519	73.3	58.1
3	10	441	752	56	156	87.3	79.3
5	10	4005	8114	716	2484	82.1	69.4
8	10	48621	128936	11441	50883	76.5	60.5

that have  $\omega'$  as their corresponding path in  $D'$ . We denote the set of these prefixes for a path  $\omega'$  in  $D'$  by  $Pref(\omega')$ . Since the measure of each finite prefix extends to a measure over the set of infinite paths sharing that prefix, it is sufficient to show that the sum of the probabilities for these finite prefixes is equal to the probability of the unique prefix  $\sigma_0, \sigma_1^f$  of  $\omega'$ , that is  $Pr^D Pref(\omega') = Pr^{D'}\{\sigma_0, \sigma_1^f\}$ . We can then write

$$\begin{aligned} Pr^D Pref(\omega') &= \mathbf{P}(\sigma_0, \sigma_1^f) + \sum_{\sigma' \in pred(\sigma_1^f)} \mathbf{P}(\sigma_0, \sigma') \cdot 1^{k_{\sigma'}} \\ &= \mathbf{P}(\sigma_0, \sigma_1^f) + \sum_{\sigma' \in pred(\sigma_1^f)} \mathbf{P}(\sigma_0, \sigma'), \end{aligned}$$

where  $k_{\sigma'}$  is the number of deterministic transitions that lead from  $\sigma'$  to  $\sigma_1^f$  in  $D$ . Now recall that for any  $\sigma \in Q' \setminus \{\sigma_0\}$  we have

$$\mathbf{P}'(\sigma_0, \sigma) = \mathbf{P}(\sigma_0, \sigma) + \sum_{\sigma' \in pred(\sigma)} \mathbf{P}(\sigma_0, \sigma').$$

So we have shown that  $Pr^D Pref(\omega') = Pr^{D'}\{\sigma_0, \sigma_1^f\}$  and the lemma is proved.  $\square$

*Proof (of Theorem 1)* Follows from Lemmas 5 and 6 for the reduction of states and transitions respectively, and from Lemma 7 for the preservation of unbounded time reachability properties.

## 5.6 Empirical Analysis

Table 2 shows the number of reachable states and transitions of the DTMC, and corresponding reduction, for different population sizes ( $N$ ) and oscillation cycle lengths ( $T$ ), using the Mirolo and Strogatz model of synchronisation [25]. The number of reachable states is stable under changes to the parameters  $R$ ,  $\epsilon$ , and  $\mu$ , since every possible firing state is always reachable from the initial state. For the results shown here the parameters were arbitrarily set to  $R = 1$ ,  $\epsilon = 0.1$ . The underlying graph of the DTMC, and hence the number of transitions, is stable under changes to the parameter  $\mu$ , and is not if interest here.

Table 3: Reduction in transitions for different population model instances.

$N$	$T$	$R$	$\epsilon$	Transitions		Reduction (%)
				$D$	$D'$	
5	10	1	0.1	8114	2484	69.4
5	10	3	0.1	7928	2391	69.8
5	10	5	0.1	7568	2211	70.8
5	10	7	0.1	6976	1915	72.5
5	10	9	0.1	6006	1430	76.2
5	10	1	0.01	6006	1430	76.2
5	10	1	0.05	6426	1640	74.5
5	10	1	0.1	8114	2484	69.4
5	10	1	0.25	8950	2902	67.6
5	10	1	0.5	9382	3118	66.7

Table 3 shows the number of transitions of the DTMC, and corresponding reduction, for various population model instances, and again uses the Mirolo and Strogatz model of synchronisation. Increasing the length of the refractory period ( $R$ ) results in an increase in the reduction of transitions in the model. A longer refractory period leads to more firing states where the firing of a group of oscillators is ignored. This results in successor states having oscillators with lower values for phase, and hence a longer sequence of deterministic transitions (later removed in the reduction) leading to the next firing state. Conversely, increasing the strength of the coupling between oscillators ( $\epsilon$ ) results in a decrease in the reduction of transitions in the model. For the Mirolo and Strogatz model of synchronisation used here, increasing the coupling strength results in a linear increase in the perturbation to phase induced by the firing of an oscillator. This results in successor states of firing states having oscillators with higher values for phase, and hence a shorter sequence of deterministic transitions leading to the next firing state.

### 5.7 Reward Structures for Reductions

While probabilistic reachability properties allow us to quantitatively analyse models with respect to the likelihood of reaching a synchronised state, they do not allow us to reason about other properties of interest, for instance the expected time taken for the network to synchronise [16], or the expected energy consumption of the network [17]. Therefore, we will often want to augment the DTMC corresponding to a population model with rewards. We do this by annotating states and transitions with real-valued rewards (respectively costs, should values be negative) that are awarded when states are visited, or transitions taken.

**Definition 16** Given a DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  a *reward structure* for  $D$  is a pair  $\mathcal{R} = (R_s, R_t)$  where  $R_s : Q \rightarrow \mathbb{R}$  and  $R_t : Q \times Q \rightarrow \mathbb{R}$  are the *state reward* and *state transition* functions that respectively map real valued rewards to states and transitions in  $D$ .

For any finite path  $\omega = \sigma_0 \cdots \sigma_k$  of  $D$  we define the total reward accumulated along that path up to, but not including,  $\sigma_k$  as

$$\text{tot}_{\mathcal{R}}(\sigma_0 \cdots \sigma_k) = \sum_{i=0}^{k-1} (R_s(\sigma_i) + R_t(\sigma_i, \sigma_{i+1})). \quad (15)$$

Given a DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  augmented with a reward structure  $\mathcal{R}$ , and some state  $\sigma \in Q$ , we will often want to reason about the reward that is accumulated along a path  $\omega = \sigma_0 \sigma_1 \sigma_2 \cdots \in \text{Paths}$  that eventually passes through some set of target states  $\Omega \subset Q$ . We first define a random variable over the set of infinite paths  $V_\Omega : \text{Paths} \rightarrow \mathbb{R} \cup \{\infty\}$ . Given the set  $\omega_\Omega = \{j \mid \sigma_j \in \Omega\}$  of indices of states in  $\omega$  that are in  $\Omega$  we define the random variable

$$V_\Omega(\omega) = \begin{cases} \infty & \text{if } \omega_\Omega = \emptyset \\ \text{tot}_{\mathcal{R}}(\sigma_0 \cdots \sigma_k) & \text{otherwise, where } k = \min \omega_\Omega, \end{cases}$$

and define the expectation of  $V_\Omega$  with respect to  $Pr_\sigma$  by

$$E[V_\Omega] = \int_{\omega \in \text{Paths}} V_\Omega(\omega) dPr = \sum_{\omega \in \text{Paths}} V_\Omega(\omega) Pr\{\omega\}.$$

The logic of PCTL can be extended to include reward properties by introducing the state formula  $R_{\bowtie r}[\mathbf{F} \Psi]$ , where  $\bowtie \in \{<, \leq, \geq, >\}$  and  $r \in \mathbb{R}$  [22]. Given a state  $\sigma \in Q$ , a real value  $r$ , and a PCTL path formula  $\Psi$ , the semantics of this formula is given by

$$\sigma \models R_{\bowtie r}[\mathbf{F} \Psi] \Leftrightarrow E[V_{\text{Sat}(\Psi)}] \bowtie r,$$

where  $\text{Sat}(\Phi)$  denotes the set of states in  $Q$  that satisfy  $\Phi$ .

**Theorem 2** *For every population model  $\mathcal{S}$  with corresponding DTMC  $D = (Q, \sigma_0, \mathbf{P}, L)$  and a reduction  $D' = (Q', \sigma_0, \mathbf{P}', L')$  of  $D$ , and for every reward structure  $\mathcal{R} = (R_s, R_t)$  for  $D$ , there is a reward structure  $\mathcal{R}' = (R'_s, R'_t)$  for  $D'$  such that unbounded-time reachability reward properties with respect to synchronised firing states in  $D$  are preserved in  $D'$ .*

Given a reward structure  $\mathcal{R} = (R_s, R_t)$  for  $D$  we construct the corresponding reward structure  $\mathcal{R}' = (R'_s, R'_t)$  as follows:

- There is no reward for the initial state and we set  $R_s(\sigma_0) = 0$ .
- For every firing state  $\sigma^f$  in  $Q$  with  $R_s(\sigma^f) = r$  we set  $R'_s(\sigma^f) = r$ .
- For every pair of distinct firing states  $\sigma_1^f, \sigma_2^f \in Q'$ , where there is a non-zero transition from  $\sigma_1^f$  to  $\sigma_2^f$  in  $D'$ , there is a (possibly empty) sequence  $\sigma_1^{\text{NF}} \cdots \sigma_k^{\text{NF}}$  of  $k$  deterministic predecessors of  $\sigma_2^f$  in  $Q$  such that  $k > 0$  implies  $\mathbf{P}(\sigma_1^f, \sigma_1^{\text{NF}}) > 0$ ,  $\mathbf{P}(\sigma_k^{\text{NF}}, \sigma_2^f) = 1$ , and  $\mathbf{P}(\sigma_i^{\text{NF}}, \sigma_{i+1}^{\text{NF}}) = 1$  for  $1 \leq i < k$ . We set the reward for taking the transition from  $\sigma_1^f$  to  $\sigma_2^f$  in  $D'$  to be the sum of the rewards that would be accumulated across that sequence by a path in  $D$ , formally

$$R'_t(\sigma_1^f, \sigma_2^f) = \text{tot}_{\mathcal{R}}(\sigma_1^f \sigma_1^{\text{NF}} \cdots \sigma_k^{\text{NF}} \sigma_2^f).$$

- For every firing state  $\sigma^F$  in  $Q'$  there is a non-zero transition from the initial state  $\sigma_0$  to  $\sigma^F$  in  $\mathbf{P}'$ . Therefore, all paths of  $D'$  where  $\sigma^F$  is the first firing state along that path share the same prefix, namely  $\sigma_0, \sigma^F$ . For paths of  $D$  this is not necessarily the case, since  $\sigma^F$  is the first firing state not only along the path where the initial transition is taken to  $\sigma^F$  itself, but also along any path where the initial transition is taken to a non-firing state from which a sequence of deterministic transitions leads to  $\sigma^F$  (that state is a deterministic predecessor of  $\sigma^F$ ). We therefore set the reward along a path  $\omega' = \sigma_0 \sigma_1^F \sigma_2^F \dots$  for taking the initial transition to  $\sigma^F$  in  $D'$  to be the sum of the total rewards accumulated along all distinct path prefixes of the form  $\sigma_0 \omega^{NF} \sigma^F$ , normalised by the total probability of taking any of these paths, where  $\omega^{NF}$  is a possibly empty sequence of deterministic predecessors of  $\sigma^F$ , and where the total reward for each prefix is weighted by the probability of taking the transitions along that sequence,

$$R'_t(\sigma_0, \sigma^F) = \frac{\sum_{\omega_{pre} \in Pref(\omega')} \text{tot}_{\mathcal{R}}(\omega_{pre}) Pr^D\{\omega_{pre}\}}{Pr^{D'}\{\sigma_0 \sigma_1^F\}} \quad (16)$$

*Proof (of Theorem 2)* We want to show that for every reward structure  $\mathcal{R}$  for  $D$  and corresponding reward structure  $\mathcal{R}'$  for  $D'$ , every  $\bowtie \in \{<, \leq, \geq, >\}$  and every  $r \in \mathbb{R}$ , if  $\sigma_0 \models R_{\bowtie r}[\mathbf{F} \text{ synch}]$  holds in  $D$  then it also holds in  $D'$ . Let  $V_{Sat(\mathbf{F} \text{ synch})}$  and  $V'_{Sat(\mathbf{F} \text{ synch})}$  respectively denote the random variables over  $Paths^D(\sigma_0)$  and  $Paths^{D'}(\sigma_0)$  whose expectations correspond to  $\mathcal{R}$  and  $\mathcal{R}'$ . From the semantics of PCTL over a DTMC we have

$$\begin{aligned} \sigma_0 \models R_{\bowtie r}[\mathbf{F} \text{ synch}] &\Leftrightarrow E[V_{Sat(\text{synch})}] \bowtie r \\ &\Leftrightarrow \sum_{\omega \in Paths} V_{Sat(\text{synch})} Pr_{\sigma_0}^D\{\omega\} \bowtie r. \end{aligned}$$

Therefore, we need to show that

$$\sum_{\omega \in Paths^D} V_{Sat(\text{synch})}(\omega) Pr^D\{\omega\} = \sum_{\omega' \in Paths^{D'}} V'_{Sat(\text{synch})}(\omega') Pr^{D'}\{\omega'\}, \quad (17)$$

where  $Pr^D$  and  $Pr^{D'}$  denote the probability measures with respect to the sets of infinite paths from  $\sigma_0$  in  $D$  and  $D'$  respectively. There are two cases:

Firstly, if there exists some path of  $D$  that does not synchronise then by definition  $V_{Sat(\text{synch})} = \infty$ . Also, from Lemma 7 we know that there is a corresponding path of  $D'$  that does not synchronise, and hence that  $V'_{Sat(\text{synch})} = \infty$ . By definition the probability measure of all paths of  $D$  and  $D'$  are strictly positive. Therefore, all summands of Equation 17 are defined, and the expectation of both  $V_{Sat(\text{synch})}$  and  $V'_{Sat(\text{synch})}$  is  $\infty$ .

Secondly, we consider the case where all possible paths of  $D$  and  $D'$  synchronise. First we define the function  $reduce : Paths^D \rightarrow Paths^{D'}$  that maps paths of  $D$  to their corresponding path in the reduction  $D'$ ,

$$reduce(\sigma_0 \omega_1^{NF} \sigma_1^F \omega_2^{NF} \sigma_2^F \dots) = \sigma_0 \sigma_1^F \sigma_2^F \dots,$$

where  $\omega_i^{\text{NF}}$  is the (possibly empty) sequence of deterministic predecessors of the firing state  $\sigma_i^{\text{F}}$ . Let  $\text{reduce}^{-1}(\omega)$  denote the preimage of  $\omega$  under  $\text{reduce}$ . Then, we can rewrite the left side of (17) to

$$\sum_{\omega' \in \text{Paths}^{D'}} \sum_{\omega \in \text{reduce}^{-1}(\omega')} V_{\text{Sat}(\text{synch})}(\omega) Pr^D\{\omega\}.$$

For any path  $\omega$  of  $D$  or  $D'$  let  $\text{pre}_s(\omega)$  be the prefix of that path whose last state is the first firing state along that path that is in the set  $\text{Sat}(\text{synch})$ . So we want to show that the following holds for any path  $\omega'$  of  $D'$ ,

$$\begin{aligned} \sum_{\omega \in \text{reduce}^{-1}(\omega')} V_{\text{Sat}(\text{synch})}(\omega) Pr^D\{\omega\} &= V'_{\text{Sat}(\text{synch})}(\omega') Pr^{D'}\{\omega'\} \\ \sum_{\omega \in \text{reduce}^{-1}(\omega')} \text{tot}_{\mathcal{R}}(\text{pre}_s(\omega)) Pr^D\{\omega\} &= \text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')) Pr^{D'}\{\omega'\}. \end{aligned} \quad (18)$$

Given some path  $\omega$  let  $\omega[i : j]$  denote the sequence of states in  $\omega$  from the  $i^{\text{th}}$  firing state to the  $j^{\text{th}}$  firing state along that path (inclusively). The notation  $\omega[- : j]$  indicates that no states are removed from the start of the path i.e. the first state is  $\sigma_0$ , and the notation  $\omega[i : -]$  indicates that no states are removed from the end of the path. By recalling that  $Pr(\sigma_0\sigma_1 \cdots \sigma_n) = \prod_{i=1}^n \mathbf{P}(\sigma_{i-1}, \sigma_i)$  we can see that  $Pr(\sigma_0\sigma_1 \cdots \sigma_n) = Pr(\sigma_0 \cdots \sigma_i) Pr(\sigma_i \cdots \sigma_n)$  for any  $0 < i < n$ . Also from (15) it is clear that for any reward structure  $\mathcal{R}$ ,  $\text{tot}_{\mathcal{R}}(\sigma_0 \cdots \sigma_n) = \text{tot}_{\mathcal{R}}(\sigma_0 \cdots \sigma_i) + \text{tot}_{\mathcal{R}}(\sigma_i \cdots \sigma_n)$  holds for all  $0 < i < n$ . Now we can rewrite (18) to

$$\begin{aligned} \sum_{\omega \in \text{reduce}^{-1}(\omega')} (\text{tot}_{\mathcal{R}}(\text{pre}_s(\omega)[- : 1]) + \text{tot}_{\mathcal{R}}(\text{pre}_s(\omega)[1 : -])) Pr^D\{\omega[- : 1]\} &= \\ (\text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')[- : 1]) + \text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')[1 : -])) Pr^{D'}\{\omega'[- : 1]\}. \end{aligned} \quad (19)$$

By the definition of  $\mathcal{R}'$  we can write the right hand side of (19) as

$$\begin{aligned} \left( \left( \frac{\sum_{\omega_{\text{pre}} \in \text{Pref}(\omega')} \text{tot}_{\mathcal{R}}(\omega_{\text{pre}}) Pr^D\{\omega_{\text{pre}}\}}{Pr^{D'}\{\omega'[- : 1]\}} \right) + \text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')[1 : -]) \right) Pr^{D'}\{\omega'[- : 1]\} &= \\ \sum_{\omega_{\text{pre}} \in \text{Pref}(\omega')} (\text{tot}_{\mathcal{R}}(\omega_{\text{pre}}) Pr^D\{\omega_{\text{pre}}\}) + \text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')[1 : -]) Pr^{D'}\{\omega'[- : 1]\}. \end{aligned}$$

From Lemma 7 we know that

$$Pr^{D'}\{\omega'[- : 1]\} = Pr^D \text{Pref}(\omega') = \sum_{\omega_{\text{pre}} \in \text{Pref}(\omega')} Pr^D\{\omega_{\text{pre}}\},$$

and hence obtain

$$\begin{aligned} \sum_{\omega_{\text{pre}} \in \text{Pref}(\omega')} (\text{tot}_{\mathcal{R}}(\omega_{\text{pre}}) Pr^D\{\omega_{\text{pre}}\}) + \sum_{\omega_{\text{pre}} \in \text{Pref}(\omega')} \text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')[1 : -]) Pr^D\{\omega_{\text{pre}}\} &= \\ \sum_{\omega_{\text{pre}} \in \text{Pref}(\omega')} (\text{tot}_{\mathcal{R}}(\omega_{\text{pre}}) + \text{tot}_{\mathcal{R}'}(\text{pre}_s(\omega')[1 : -])) Pr^D\{\omega_{\text{pre}}\}. \end{aligned} \quad (20)$$

Since  $Pref(\omega')$  is the set of all possible finite prefixes from the initial state  $\sigma_0$  to the first firing state  $\sigma_1^f$ , and since  $\omega[- : 1] = pre_s(\omega)[- : 1]$  clearly holds, we know that

$$\bigcup_{\omega_{pre} \in Pref(\omega')} \{\omega_{pre}\} = \bigcup_{\omega \in reduce^{-1}(\omega')} \{\omega[- : 1]\} = \bigcup_{\omega \in reduce^{-1}(\omega')} \{pre_s(\omega)[- : 1]\}.$$

Using this fact, and by observing that by definition

$$\text{tot}_{\mathcal{R}'}(pre_s(\omega')[1 : -]) = \text{tot}_{\mathcal{R}}(pre_s(\omega)[1 : -]),$$

we can write (20) as

$$\sum_{\omega \in reduce^{-1}(\omega')} (\text{tot}_{\mathcal{R}}(pre_s(\omega)[- : 1]) + \text{tot}_{\mathcal{R}'}(pre_s(\omega')[1 : -])) Pr^D\{\omega[- : 1]\}.$$

This is the same as the left hand side of (19) and the theorem is proved.  $\square$

## 6 Connecting the Concrete Model and the Population Model

In this section, we define the abstraction function to connect a concrete model with a population model. To that end, let  $D_c = (Q_c, s_0, \mathbf{P}_c)$  be a concrete model of a network of  $N$  PCOs with a clock cycle length  $T$ , a refractory period  $R$ , a phase response function  $\Delta$ , a coupling  $\epsilon$  and broadcast failure probability of  $\mu$ . Furthermore, let  $D_p = (Q_p, \sigma_0, \mathbf{P}_p)$  be the DTMC of a population model for the same parameters. For simpler notation, we introduce some general notation for transitions in DTMCs. If there is a possible transition between two states  $q$  and  $q'$  in a DTMC, that is  $\mathbf{P}(q, q') > 0$ , then we also write  $q \rightarrow q'$ . Observe that for this simplification,  $q$  and  $q'$  are either both in  $Q_c$  or both in  $Q_p$ . We also denote the reflexive, transitive closure of  $\rightarrow$  by  $\Rightarrow$ .

### 6.1 Proving the Correspondence between Concrete and Population Models

We need to associate states in  $D_c$  to states in  $D_p$ . In general, several concrete states will be mapped to a single population state, since we do not distinguish between different orders of oscillators in the latter, while we do in the former.

Furthermore, we want to abstract from different modes of the oscillators. However, it is not sensible to associate all modes within a phase to the same population state, since in the transitions from one mode to the next the system chooses, whether an oscillator fails to broadcast its pulse or not. If we want to be able to define a simulation relation, we need to represent the failures described by the transitions in the population model. To have an exact correspondence, we first collect all the concrete states where the counter and all oscillators are at the start mode into a single set.

$$Q'_c = \{s \in Q_c \mid s = (\eta, \nu) \wedge p_\theta(\eta) = start \wedge \forall u: p_\theta(\nu(u)) = start\}$$




Fig. 5: Weak Simulation Relation of Concrete States and Populations

The abstraction function  $h: Q'_c \rightarrow Q_p$  takes a concrete state  $s$  and counts the number of oscillators sharing the same phase, mapping  $s = (\eta, \nu)$  to the corresponding state of the population model,

$$h(s) = \langle |\{u \mid p_\Phi(\nu(u)) = 1\}|, \dots, |\{u \mid p_\Phi(\nu(u)) = T\}| \rangle.$$

To show that this abstraction is sensibly defined, we need to show that the concrete model can weakly simulate the transitions allowed by the population model, and vice versa. That is, if the abstraction  $\sigma_1$  of a concrete state  $s_1$  allows a transition to another population state  $\sigma_2$ , then there is a sequence of transitions from  $s_1$  leading to  $s_2$ , whose abstraction is  $\sigma_2$ . Furthermore, if there is a transition sequence from one concrete state  $s_1$  to  $s_2$ , where both states can be abstracted to population states  $\sigma_1$  and  $\sigma_2$ , respectively, then there is also a sequence of transitions connecting  $\sigma_1$  with  $\sigma_2$ . This situation is visualised in Fig. 5.

For the first direction, we actually show this condition for a single transition in the population model. However, this result can be straightforwardly extended to transition sequences.

**Lemma 8** *Let  $s_1 \in Q'_c$  and  $\sigma_1, \sigma_2 \in Q_p$  such that  $h(s_1) = \sigma_1$  and  $\sigma_1 \rightarrow \sigma_2$ . Then there is a  $s_2 \in Q'_c$  such that  $s_1 \Rightarrow s_2$  and  $h(s_2) = \sigma_2$ . Furthermore, the sum of the probabilities of transition sequences from  $s_1$  to an instantiation  $s_2$  of  $\sigma_2$  is equal to the probability of the transition from  $\sigma_1$  to  $\sigma_2$ .*

*Proof* Let  $s_1 = (\eta_1, \nu_1) \in Q'_c$ , i.e., a concrete state where  $p_\theta(\eta_1) = \text{start}$  and  $p_\theta(\nu_1(u)) = \text{start}$  for all  $1 \leq u \leq N$ . Then  $h(s_1) = \langle |\{u \mid p_\phi(\nu_1(u)) = 1\}|, \dots, |\{u \mid p_\phi(\nu_1(u)) = T\}| \rangle = \sigma_1$ . Let  $\sigma_1 \rightarrow \sigma_2$ . Note that there is only a single outgoing transition from  $s_1$  according to condition (1). That is, in the successor state of  $s_1$ , we have  $p_c(\eta) = 0$  and  $p_\theta(\eta) = \text{update}$ , while the oscillator states are not changed. To keep the notation tidy, we identify this successor state with  $s_1$  in the following.

Now consider two cases. If  $|\{u \mid p_\phi(\nu_1(u)) = T\}| = 0$ , then  $\sigma_2 = \langle 0, |\{u \mid p_\phi(\nu_1(u)) = 1\}|, \dots, |\{u \mid p_\phi(\nu_1(u)) = T - 1\}| \rangle$ , since no set of oscillators in  $\sigma_1$  is perturbed by a firing oscillator. In particular, there is no  $u$  such that  $p_\phi(\nu_1(u)) = T$ . Hence, for all possible successors  $s'$  of  $s_1$ , we have that only condition (4) is satisfied. Furthermore, this is the case until all oscillators changed their mode to *update*. Let us call this state  $s'_1 = (\eta'_1, \nu'_1)$ . Now, the environment was not changed from  $s_1$  to  $s'_1$ , i.e.,  $p_c(\eta) = p_c(\eta') = 0$ .

Hence, condition (8b) is satisfied for all oscillators. Since  $\Delta(\Phi, 0, \epsilon) = 0$  for all  $\Phi$ , the phase of each oscillator is increased by one. This implies that there is a single successor of  $s'_1$ , which we call  $s_2 = (\eta_2, \nu_2)$  and that  $p_\phi(\nu_2(u)) = p_\phi(\nu'_1(u)) + 1$ . In particular, we have that for all oscillators  $u$ ,  $p_\phi(\nu_2(u)) > 0$ , and  $\{u \mid p_\phi(\nu_2(u)) =$

$\Phi\}$  =  $\{u \mid p_\phi(\nu'_1(u)) = \Phi - 1\}$  for all  $0 < \Phi \leq T$ . Hence  $s_2$  is the required state of the concrete model.

Now let  $|\{u \mid p_\phi(\nu_1(u)) = T\}| > 0$ . Then, each transition in the population model is induced by a failure vector  $F = \langle f_1, \dots, f_T \rangle$ . In particular, there is a maximal number  $k$ , such that for all  $l < k$ , we have  $f_l = \star$ . That is,  $k$  denotes the lowest phase in which oscillators possibly fire.

First, we introduce some notation, where  $\Phi > R$  and  $s = (\eta, \nu)$ .

$$\begin{aligned} Q_c^F(s) &= \{u \mid p_\phi(\nu(u)) = T\} \\ Q_{c,\Phi}^{PF}(s) &= \{u \mid p_\phi(\nu(u)) = \Phi \wedge \Delta(p_\phi(\nu(u)), \alpha^{p_\phi(\nu(u))}(\sigma_1, F), \epsilon) + 1 > T\} \\ Q_{c,\Phi}^P(s) &= \{u \mid p_\phi(\nu(u)) = \Phi \wedge \Delta(p_\phi(\nu(u)), \alpha^{p_\phi(\nu(u))}(\sigma_1, F), \epsilon) + 1 \leq T\} \end{aligned}$$

That is,  $Q_c^F(s)$  denotes the set of oscillators possibly firing in  $s$ . The sets  $Q_c^P(s)$  and  $Q_{c,\Phi}^{PF}(s)$  denote the sets of oscillators being perturbed but not firing (since the perturbation is not sufficient for the oscillators to reach the end of their cycle), and possibly firing, respectively. We can only say that elements of  $Q_c^F(s)$  and  $Q_{c,\Phi}^{PF}(s)$  *possibly* fire, since they may be affected by a broadcast failure.

We now have to construct a sequence of transitions, where we draw the firing oscillators from the sets  $Q_c^F(s_1)$  and  $Q_{c,\Phi}^{PF}(s_1)$ , according to the broadcast failure vector  $F$ . Furthermore, all elements of  $Q_c^F(s_1)$  and the sets  $Q_{c,\Phi}^{PF}(s_1)$  have to take transitions such that their phase value in the next iteration is 1.

Let  $\sigma_1 = \langle k_1, k_2, \dots, k_T \rangle$ . Now consider an arbitrary sequence  $u_1, \dots, u_{k_T}$  of all  $k_T$  elements from  $Q_c^F(s_1)$ . Additionally, let  $C_T \subseteq Q_c^F(s_1)$  be the set of oscillators in phase  $T$  with a broadcast failure, i.e.,  $|C_T| = f_T$ . Observe that  $p_\phi(\nu_1(u_j)) = T$  for all  $1 \leq j \leq k_T$ . Furthermore, let  $r_0 = s_1$ . Then we define a sequence of successors of  $r_0 = (\eta_0, \nu_0)$  as follows, where  $1 \leq j \leq k_T$ . If  $k_i \notin C_T$ , then

$$\begin{aligned} \nu_j &= \nu_{j-1} \oplus [u_j \mapsto (T, \text{update})] \\ \eta_j &= (\text{update}, p_c(\eta_{j-1}) + 1) \end{aligned}$$

otherwise

$$\begin{aligned} \nu_j &= \nu_{j-1} \oplus [u_j \mapsto (T, \text{update})] \\ \eta_j &= (\text{update}, p_c(\eta_{j-1})) \end{aligned}$$

Observe that these states define a sequence of transitions from  $r_0$  to  $r_{k_T}$  according to conditions (2) and (3).

Now, for each phase  $\Phi$ , with  $k \leq \Phi < T$ , we proceed similarly. That is, we first choose a sequence  $u_1^\Phi, \dots, u_{k_\Phi}^\Phi$  of oscillators and a set  $C_\Phi \subseteq Q_{c,\Phi}^{PF}(s_1)$  with  $|C_\Phi| = f_\Phi$ .

Subsequently, we define each  $r_j^\Phi$  to be

$$\begin{aligned} \nu_j^\Phi &= \nu_{j-1}^\Phi \oplus [u_j^\Phi \mapsto (\Phi, \text{update})] \\ \eta_j^\Phi &= \begin{cases} (\text{update}, p_c(\eta_{j-1}^\Phi) + 1) & , \text{ if } j \notin C_\Phi \\ (\text{update}, p_c(\eta_{j-1}^\Phi)) & , \text{ otherwise} \end{cases} \end{aligned}$$

where  $r_0^\Phi = r_{k_\Phi}^{\Phi+1}$ . Observe again, that these sequences exhaust  $Q_{c,\Phi}^{PF}(s_1)$  for each phase  $\Phi$ . Furthermore, we claim that the number of firing oscillators that are not inhibited by a broadcast failure in the concrete model coincides with the number of perceived firing oscillators in the population model in this phase.



*Claim* For each  $\Phi$  with  $k \leq \Phi \leq T$ , we have  $p_c(\eta_0^\Phi) = \alpha^\Phi(\sigma_1, F)$ .

*Proof* For  $\Phi = T$ , we have  $p_c(\eta_0^T) = 0 = \alpha^T(\sigma_1, F)$ . Now let  $\Phi < T$  and assume  $p_c(\eta_0^{\Phi+1}) = \alpha^{\Phi+1}(\sigma_1, F)$ . By definition, we have

$$\alpha^\Phi(\sigma_1, F) = \alpha^{\Phi+1}(\sigma, F) + k_{\Phi+1} - f_{\Phi+1},$$

since  $\Phi < T$  and  $f_{\Phi+1} \neq \star$ . Now, in the sequence  $r_0^{\Phi+1}, \dots, r_{k_{\Phi+1}}^{\Phi+1}$ , we increase  $p_c(\eta_0^{\Phi+1})$  exactly  $k_{\Phi+1} - f_{\Phi+1}$  times, i.e.,

$$p_c(\eta_0^\Phi) = p_c(\eta_0^{\Phi+1}) + k_{\Phi+1} - f_{\Phi+1}.$$

By assumption, we then get

$$p_c(\eta_0^\Phi) = \alpha^{\Phi+1}(\sigma_1, F) + k_{\Phi+1} - f_{\Phi+1} = \alpha^\Phi(\sigma_1, F).$$

This proves the claim.  $\square$

This claim in particular states that the perturbation within the population model and the concrete model is the same.

Since  $\Delta$  is a monotonically increasing function in  $\alpha$ , every oscillator in  $Q_{c,\Phi}^{\text{PF}}(s_1)$  is still perturbed to firing after other oscillators in the same phase fired. Hence, for each pair of states  $u_{j-1}^\Phi$  and  $u_j^\Phi$  with  $1 \leq j \leq k_\Phi - f_\Phi$ , a transition according to condition (7) is well-defined. Similarly, for oscillators that should fire, but are affected by a broadcast failure,  $u_{j-1}^\Phi$  and  $u_j^\Phi$  with  $k_\Phi - f_\Phi + 1 \leq j \leq k_\Phi$ , the transition is defined according to condition (6).

Now, for every  $\Phi < k$ , we know that  $\Phi + \Delta(\Phi, \alpha^\Phi(\sigma_1, F), \epsilon) + 1 \leq T$  and  $\alpha^{\Phi-1}(\sigma_1, F) = \alpha^\Phi(\sigma_1, F)$ , according to equation (11). Hence, for every phase  $\Phi < k$ , we arbitrarily enumerate the oscillators of  $Q_{c,\Phi}^{\text{P}}(s_1) = u_1^\Phi, \dots, u_{k_\Phi}^\Phi$  and define the following sequence of states  $r_j^\Phi$ , where  $r_0^\Phi = r_{k_{\Phi+1}}^{\Phi+1}$ .

$$\begin{aligned} \nu_j^\Phi &= \nu_{j-1} \oplus [u_j^\Phi \mapsto (\Phi, \text{update})] \\ \eta_j^\Phi &= (\text{update}, p_c(\eta_{j-1}^\Phi)) \end{aligned}$$

For each  $\Phi$  and pair of states  $r_j^\Phi$  and  $r_{j+1}^\Phi$ , there is a transition according to condition (5). So, all in all, we have a sequence of transitions from  $s_1$  to  $r_{k_0}^0$ .

Now, in  $r_{k_0}^0 = (\eta_{k_0}^0, \nu_{k_0}^0)$ , we have that  $p_\theta(\eta_{k_0}^0) = \text{update}$  and for all  $u$ ,  $p_\theta(\nu_{k_0}^0(u)) = \text{update}$ . Then let  $s_2 = (\eta, \nu)$  be defined by the following formulas.

$$p_\theta(\eta) = \text{start} \text{ and } p_c(\eta) = p_c(\eta_{k_0}^0) \quad (21)$$

$$\forall u: p_\theta(\nu(u)) = \text{start} \quad (22)$$

$$\forall u: p_\phi(\nu_{k_0}^0(u)) = T \rightarrow p_\phi(\nu(u)) = 1 \quad (23)$$

$$\forall u: p_\phi(\nu_{k_0}^0(u)) < T \wedge p_\phi(\nu_{k_0}^0(u)) \leq R \rightarrow p_\phi(\nu(u)) = p_\phi(\nu_{k_0}^0(u)) + 1 \quad (24)$$

$$\begin{aligned} \forall u: p_\phi(\nu_{k_0}^0(u)) < T \wedge p_\phi(\nu_{k_0}^0(u)) > R \wedge \\ p_\phi(\nu_{k_0}^0(u)) + \Delta(p_\phi(\nu_{k_0}^0(u)), p_c(\eta_{k_0}^0), \epsilon) + 1 \leq T \\ \rightarrow p_\phi(\nu(u)) = p_\phi(\nu_{k_0}^0(u)) + \Delta(p_\phi(\nu_{k_0}^0(u)), p_c(\eta_{k_0}^0), \epsilon) + 1 \end{aligned} \quad (25)$$

$$\begin{aligned} \forall u: p_\phi(\nu_{k_0}^0(u)) < T \wedge p_\phi(\nu_{k_0}^0(u)) > R \wedge \\ p_\phi(\nu_{k_0}^0(u)) + \Delta(p_\phi(\nu_{k_0}^0(u)), p_c(\eta_{k_0}^0), \epsilon) + 1 > T \\ \rightarrow p_\phi(\nu(u)) = 1 \end{aligned} \quad (26)$$

Then  $r_{k_0}^0$  and  $s_2$  satisfy all parts of condition (8). Hence, we have a sequence of transitions from  $s_1$  to  $s_2$ . To prove  $h(s_2) = \sigma_2$ , we need to show that the number of oscillators possessing a phase  $\Phi$  in  $s_2$  matches the  $\Phi$ -th entry of  $\sigma_2 = \langle k'_1, \dots, k'_T \rangle$ . To that end, recall that by Def. 9, each  $k'_\Phi = \sum_{\Psi \in \mathcal{U}_\Phi(\sigma, F)} k_\Psi$ , where  $\sigma_1 = \langle k_1, \dots, k_T \rangle$  and  $\mathcal{U}_\Phi(\sigma, F) = \{\Psi \mid \Psi \in \{1, \dots, T\} \wedge \tau(\sigma, \Psi, F) = \Phi\}$ . Observe that both the concrete model and the population model use the same perturbation function  $\Delta$  and that  $\tau$  is defined in terms of  $\Delta$ . In particular, we have

$$\tau(\sigma, \Phi, F) = \begin{cases} 1 & \text{if } \text{fire}^\Phi(\sigma, F) \\ \text{update}^\Phi(\sigma, F) & \text{otherwise.} \end{cases}$$

Now let us distinguish three cases for  $\Phi$ . In the following, let  $s_1 = (\eta_1, \nu_1)$ ,  $s_2 = (\eta_2, \nu_2)$ ,  $\sigma_1 = \langle k_1, \dots, k_T \rangle$  and  $\sigma_2 = \langle k'_1, \dots, k'_T \rangle$ .

1. If  $\Phi \leq R$ , then  $\text{update}^\Phi(\sigma_1, F) = \Phi + 1$ , due to the definition of the refractory function  $\text{ref}$ . Similarly, for all  $u$  such that  $p_\phi(\nu_1(u)) = \Phi$ , we get that  $p_\phi(\nu_2(u)) = \Phi + 1$ . Hence, for all  $\Phi \leq R$ , we have that  $|\{u \mid p_\phi(\nu_2(u)) = \Phi + 1\}| = |\{u \mid p_\phi(\nu_1(u)) = \Phi\}|$ .
2. If  $\Phi > R$  and  $\text{update}^\Phi(\sigma_1, F) = \Psi$ , with  $\Psi \leq T$ . Then  $p_\phi(\nu_2(u)) = \Psi$ , by formula (25). Observe that the number of oscillators in  $s_1$  with a phase of  $\Phi$  is  $k_\Phi$ . So, the number of oscillators that get perturbed to be in  $\Psi$  is the union of the oscillators  $u$  in phases  $\Phi$ , where  $\Delta(\Phi, p_c(\eta_2), \epsilon) + 1 = \Psi$ . That is,  $\{u \mid p_\phi(\nu_2(u)) = \Psi\} = \{u \mid \Delta(p_\phi(\nu_1(u)), p_c(\eta_2), \epsilon) + 1 = \Psi\}$ . By the definition of  $\tau$  and claim 6.1, we get that  $\tau(\sigma_1, \Phi, F) = \Psi$ . That is, for a specific  $\Psi$ , we have that the phases  $\Phi$  of oscillators perturbed to  $\Psi$  are in  $\mathcal{U}_\Psi(\sigma_1, F)$ . Hence, since the sets of oscillators in each phase are disjoint,  $|\{u \mid p_\phi(\nu_2(u)) = \Psi\}| = \sum_{\Phi \in \mathcal{U}_\Psi(\sigma_1, F)} k_\Phi$ .
3. Finally, let  $\text{update}^\Phi(\sigma_1, F) = \Psi$  and  $\Psi > T$ . Then  $\tau(\sigma_1, \Phi, F) = 1$ . Furthermore, by formulas (23) and (26), we have  $p_\phi(\nu_2(u)) = 1$  for all  $u$  with phase  $\Phi$ . With similar reasoning as above, we get that  $|\{u \mid p_\phi(\nu_2(u)) = 1\}| = \sum_{\Phi \in \mathcal{U}_1(\sigma_1, F)} k_\Phi$ .

Hence, we get  $h(s_2) = \sigma_2$ , and we are done.  $\square$

Now we turn our attention to the other direction. That is, if we have a sequence of transitions in the concrete model, we can find a corresponding transition sequence in the population model.

**Lemma 9** *Let  $D_c = (Q_c, s_0, \mathbf{P}_c)$  be a concrete network of oscillators and  $D_p = (Q_p, \sigma_0, \mathbf{P}_p)$  be its abstraction as a population model. Furthermore, let  $s_1, s_2 \in Q'_c$  and  $\sigma_1 \in Q_p$  such that  $h(s_1) = \sigma_1$  and  $s_1 \Rightarrow s_2$ . Then there is a  $\sigma_2 \in Q_p$  such that  $h(s_2) = \sigma_2$  and  $\sigma_1 \Rightarrow \sigma_2$ .*

*Proof* If  $s_1 = s_2$ , then the lemma holds trivially. Otherwise, assume that for all  $s' = (\eta, \nu)$  different from  $s_1$  and  $s_2$ , such that  $s_1 \Rightarrow s' \Rightarrow s_2$ , we have  $p_\theta(\eta) = \text{update}$ . Furthermore, let  $s_1 = (\eta_1, \nu_1)$  and  $s_2 = (\eta_2, \nu_2)$ . By definition of the abstraction function, we have

$$\begin{aligned} \sigma_1 &= \langle |\{u \mid p_\phi(\nu_1(u)) = 1\}|, \dots, |\{u \mid p_\phi(\nu_1(u)) = T\}| \rangle \\ \sigma_2 &= \langle |\{u \mid p_\phi(\nu_2(u)) = 1\}|, \dots, |\{u \mid p_\phi(\nu_2(u)) = T\}| \rangle \end{aligned}$$

We now distinguish two cases. First, assume that  $\{u \mid p_\phi(\nu_1(u) = T)\} = \emptyset$ , and let  $s = (\eta_s, \nu_s)$  be such that  $s_1 \Rightarrow s$  and  $p_\theta(\nu_s(u)) = \text{update}$  for all  $u$ . Then there is exactly one transition  $s \rightarrow s_2$ , which is defined according to equation (8). Furthermore, due to the assumption that no oscillator fires, we have  $p_c(\eta_s) = 0$ , which implies  $\Delta(\Phi, p_c(\eta_s), \epsilon) = 0$  for all  $\Phi$  by Definition 4. Hence, for all  $u$ , we have  $p_\phi(\nu_2(u)) = p_\phi(\nu_s(u)) + 1 = p_\phi(\nu_1(u)) + 1$ . That is,

$$\begin{aligned} \sigma_2 &= \langle 0, |\{u \mid p_\phi(\nu_1(u) + 1 = 2)\}|, \dots, |\{u \mid p_\phi(\nu_1(u) + 1 = T)\}| \rangle \\ &= \langle 0, |\{u \mid p_\phi(\nu_1(u) = 1)\}|, \dots, |\{u \mid p_\phi(\nu_1(u) = T - 1)\}| \rangle . \end{aligned}$$

That is, we have  $\sigma_1 \rightarrow \sigma_2$  due to a deterministic transition, which, in particular, implies  $\sigma_1 \Rightarrow \sigma_2$ .

The second case is more involved. Let us assume  $\{u \mid p_\phi(\nu_1(u) = T)\} \neq \emptyset$ , that is, at least one oscillator fires. Hence, due to the preconditions of the transitions, we can divide the transition sequence from  $s_1$  to  $s_2$  as follows:

$$s_1 \Rightarrow r_T \Rightarrow r_{T-1} \Rightarrow \dots \Rightarrow r_1 \Rightarrow s_2 ,$$

where  $r_\Phi = (\eta_{r_\Phi}, \nu_{r_\Phi})$  denotes the state where all oscillators with phase  $\Phi$  changed their mode to *update*. Our goal now is to find a broadcast failure vector  $F$ , such that  $\text{succ}(\sigma_1, F) = \sigma_2$ . To that end, let

$$f_\Phi = |\{u \mid p_\phi(\nu_1(u)) \geq \Phi\}| - \left( p_c(\eta_{r_\Phi}) + \sum_{\Psi=\Phi+1}^T f_\Psi \right)$$

for all  $\Phi$  where  $\Phi + \Delta(\Phi, p_c(\eta_{r_\Phi}), \epsilon) + 1 \geq T$ . For the remaining  $\Phi$ , set  $f_\Phi = \star$ . Then  $F = \langle f_1, \dots, f_T \rangle$ . With this broadcast failure vector at hand, we now have to show that

$$\sum_{\Psi \in \mathcal{U}_\Phi(\sigma_1, F)} |\{u \mid p_\phi(\nu_1(u)) = \Psi\}| = |\{u \mid p_\phi(\nu_2(u)) = \Phi\}| .$$

Recall that  $\mathcal{U}_\Phi(\sigma_1, F) = \{\Psi \mid \tau(\sigma_1, \Psi, F) = \Phi\}$ . Together with the condition we want to prove, this implies, that we need to show  $p_\phi(\nu_2(u)) = \tau(\sigma_1, p_\phi(\nu_1(u)), F)$  for all oscillators  $u$ . We now need again to distinguish several cases, according to the different cases of the transition defined by condition (8).

First, let  $u$  be such that  $p_\phi(\nu_1(u)) \leq R$ , i.e., oscillator  $u$  is within its refractory period. If  $p_\phi(\nu_1(u)) = T$ , then we have

$$\begin{aligned} p_\phi(\nu_2(u)) &= 1 && \{\text{Cond. (8a)}\} \\ &= \tau(\sigma_1, p_\phi(\nu_1(u)), F) && \{\text{Def. 8}\} \end{aligned}$$

Otherwise, if  $p_\phi(\nu_1(u)) < T$ , we have

$$\begin{aligned} p_\phi(\nu_2(u)) &= p_\phi(\nu_1(u)) + 1 && \{\text{Cond. (8b)}\} \\ &= \tau(\sigma_1, p_\phi(\nu_1(u)), F) && \{\text{Def. 8}\} \end{aligned}$$

Now assume that  $p_\phi(\nu_1(u)) \geq R$ , i.e., oscillator  $u$  is outside of its refractory period and thus will be perturbed by firing oscillators. If  $p_\phi(\nu_1(u)) = T$ , then we proceed as in the previous case. So, let us assume  $p_\phi(\nu_1(u)) < T$ . To show that the transition function of the population model coincides with the result within the concrete model, we need to ensure that the perceived firing oscillators are equal in both models for each oscillator.

*Claim* Within each phase, the perceived oscillators in the population model coincide with the oscillators that fired up to the next higher phase in the concrete model. Formally, for each  $1 \leq \Phi < T$ , we have  $p_c(\eta_{r\Phi+1}) = \alpha^\Phi(\sigma_1, F)$ .

*Proof* In the following, we use the notation  $\sigma_1 = \langle k_1, \dots, k_T \rangle$ . Let  $f_\Phi \neq \star$ . Then for  $\Phi = T - 1$ , we have

$$\begin{aligned} p_c(\eta_{rT}) &= |\{u \mid p_\phi(\nu_1(u)) = T\}| - f_T && \{\text{Def. of } f_\Phi\} \\ &= 0 + k_T - f_T && \{\text{Def. of } \sigma_1\} \\ &= \alpha^T(\sigma_1, F) + k_T - f_T && \{\text{Def. of } \alpha^\Phi, \text{ Eq. (11)}\} \\ &= \alpha^{T-1}(\sigma_1, F) \end{aligned}$$

Now assume that  $p_c(\eta_{r\Phi+1}) = \alpha^\Phi(\sigma_1, F)$  for some  $1 < \Phi < T$  with  $f_\Phi \neq \star$ . Then

$$\begin{aligned} p_c(\eta_{r\Phi}) &= |\{u \mid p_\phi(\nu_1(u)) \geq \Phi\}| - \left( \sum_{\Psi=\Phi}^T f_\Psi \right) && \{\text{Def of } f_\Phi\} \\ &= \left( \sum_{\Psi=\Phi}^T k_\Psi \right) - \left( \sum_{\Psi=\Phi}^T f_\Psi \right) && \{\text{Def. of } \sigma_1\} \\ &= \left( \sum_{\Psi=\Phi+1}^T k_\Psi \right) - \left( \sum_{\Psi=\Phi+1}^T f_\Psi \right) + k_\Phi - f_\Phi \\ &= p_c(\eta_{r\Phi+1}) + k_\Phi - f_\Phi && \{\text{Def. of } f_\Phi\} \\ &= \alpha^\Phi(\sigma_1, F) + k_\Phi - f_\Phi && \{\text{Ass.}\} \\ &= \alpha^{\Phi-1}(\sigma_1, F) && \{\text{Def. of } \alpha^\Phi, \text{ Eq. (11)}\} \end{aligned}$$

If  $f_\Phi = \star$ , then the claim immediately holds.  $\square$

Furthermore, observe that for all  $\Phi$ , if  $f_\Phi = \star$ , then we have  $p_c(\eta_{r\Phi+1}) = p_c(\eta_{r1})$ . We can now proceed to prove the final two cases.

First, let  $p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), p_c(\eta_{r1}), \epsilon) + 1 \leq T$ . This implies  $f_{p_\phi(\nu_1(u))} = \star$ . Then, by the observation and the claim above, we also get

$$\begin{aligned} &\tau(\sigma_1, p_\phi(\nu_1(u)), F) \\ &= 1 + \text{ref}(p_\phi(\nu_1(u)), \Delta(p_\phi(\nu_1(u)), \alpha^{p_\phi(\nu_1(u))}(\sigma_1, F), \epsilon)) && \{p_\phi(\nu_1(u)) > R\} \\ &= 1 + p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), \alpha^{p_\phi(\nu_1(u))}(\sigma_1, F), \epsilon) && \{\text{Claim}\} \\ &= 1 + p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), p_c(\eta_{rp_\phi(\nu_1(u))+1}), \epsilon) && \{\text{Obs.}\} \\ &= p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), p_c(\eta_{r1}), \epsilon) + 1 && \{\text{Eq. (8c)}\} \\ &= p_\phi(\nu_2(u)) . \end{aligned}$$

Finally, let  $p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), p_c(\eta_{r1}), \epsilon) + 1 > T$ , i.e.  $p_\phi(\nu_2(u)) = 1$ . Then it has also to be the case that  $p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), p_c(\eta_{rp_\phi(\nu_1(u))+1}), \epsilon) + 1 > T$ . By the claim above, this means  $p_\phi(\nu_1(u)) + \Delta(p_\phi(\nu_1(u)), \alpha^{p_\phi(\nu_1(u))}(\sigma_1, F), \epsilon) + 1 > T$ . Hence,  $\tau(\sigma_1, p_\phi(\nu_1(u)), F) = 1$  as well.

Now recall that we assumed initially that for all intermediate states  $s = (\eta, \nu)$  of the transition sequence, we have  $p_\theta(\eta) = \text{update}$ . If this is not the case, we can partition the sequence into distinct subsequences, where this assumption holds for each subsequence, and apply the arguments above. This proves the lemma.  $\square$

Now we compare the probabilities of transition sequences in the different models.

**Lemma 10** *Let  $D_c = (Q_c, s_0, \mathbf{P}_c)$  be a concrete network of oscillators and  $D_p = (Q_p, \sigma_0, \mathbf{P}_p)$  be its abstraction as a population model, as well as  $s_1, Q'_c$  and  $\sigma_1, \sigma_2 \in Q_p$ , with  $h(s_1) = \sigma_1$ . Then, the sum of the probabilities of transition sequences from  $s_1$  to all instantiations  $s_2$  with  $h(s_2) = \sigma_2$  is equal to the probability of the transition from  $\sigma_1$  to  $\sigma_2$ .*

*Proof* Let  $\sigma = \langle k_1, \dots, k_T \rangle$ . Furthermore, let  $N = \sum_{i=1}^T k_i$ . Now, let  $s$  be an arbitrary state corresponding to  $\sigma$ . If no oscillator fires, we have  $N!$  possibilities to create a transition sequence, each of which has a probability of  $\frac{1}{N!}$  to happen. Hence, we get that the probability that one of these transitions happen is  $N! \cdot \frac{1}{N!} = 1$ , which coincides with the definition in the population model.

For the case that at least one oscillator fires and thus perturbs the other oscillators, we consider the construction in the proof of Lemma 8 with respect to a failure vector  $F = \langle f_1, \dots, f_T \rangle$  for  $\sigma$ . During each phase  $\Phi$ , we have to choose the particular order of the  $k_\Phi$  and in addition, we have to choose the set  $C_\Phi$ . That is, we have  $k_\Phi!$  possible orders, and  $\binom{k_\Phi}{f_\Phi}$  possibilities for the choice of  $C_\Phi$ . Furthermore, the combined probability for the transitions of the oscillators that should fire but are inhibited by a broadcast failure is

$$\frac{1}{|\text{init}_\Phi(s)|!} \cdot (1 - \mu)^{k_\Phi - f_\Phi} \cdot \mu^{f_\Phi} .$$

Observe that at the start of the construction of each phase,  $|\text{init}_\Phi(s)| = k_\Phi$ . Hence the probability above simplifies to

$$\frac{1}{k_\Phi!} \cdot (1 - \mu)^{k_\Phi - f_\Phi} \cdot \mu^{f_\Phi} .$$

Due to the possible choices during the construction of the transition sequence, we have that the probability of *one* of these sequences to happen is

$$\binom{k_\Phi}{f_\Phi} \cdot k_\Phi! \cdot \frac{1}{k_\Phi!} \cdot (1 - \mu)^{k_\Phi - f_\Phi} \cdot \mu^{f_\Phi} = \binom{k_\Phi}{f_\Phi} \cdot (1 - \mu)^{k_\Phi - f_\Phi} \cdot \mu^{f_\Phi} ,$$

which is exactly the function  $\text{PMF}(k_\Phi, f_\Phi)$  as in the population model. Furthermore, with similar reasoning as above, the transition probability for the sequences, where no oscillator is perturbed anymore, is 1. Hence, we get that the combined probability of the set of paths from one instantiation  $s_1$  of a population model state  $\sigma_1$  to an instantiation of a successor  $\sigma_2$  of  $\sigma_1$  is equal to the probability of the the transition from  $\sigma_1$  to  $\sigma_2$ .  $\square$

From Lemmas 8, 9 and 10, we immediately get that the same weak simulation relation holds between a population model and the concrete network of oscillators it represents.

**Theorem 3** *Let  $D_c = (Q_c, s_0, \mathbf{P}_c)$  be a concrete network of oscillators and  $D_p = (Q_p, \sigma_0, \mathbf{P}_p)$  be its abstraction as a population model. If we have  $s_1, s_2 \in Q'_c$  and  $\sigma_1, \sigma_2 \in Q_p$ , with  $h(s_i) = \sigma_i$ , then  $\sigma_1 \rightarrow^* \sigma_2$  if, and only if,  $s_1 \Rightarrow s_2$ . Furthermore, the probabilities over all paths in both models coincide. In particular, we have  $h(s_0) = \sigma_0$ , and that  $s_0$  weakly bisimulates  $\sigma_0$ .*

Table 4: Model Construction Times and Model Checking Times for Both the Concrete and Population Model with  $N = 4$ ,  $T = 10$ ,  $\mu = 0.2$  and  $\epsilon = 0.1$  (in seconds).

$R$	Concrete Model		Population Model	
	Constr.	Check.	Constr.	Check
1	0.014	0.89	0.388	0.017
5	0.069	10.38	0.420	0.047
8	0.056	1.46	0.356	0.013

Hence, we can use population models to analyse the global properties of a network of pulse-coupled oscillators following the concrete model as defined in Sect. 4 without loss of precision. In particular, this allows us to increase the size of the network to check such properties, while still giving us the opportunity to analyse the internal behaviour of nodes, if we restrict the network size.

## 6.2 Experimental Validation

As Theorem 3 implies, the synchronisation probabilities for a concrete model and its corresponding population model coincide. However, we have to keep in mind that the PCTL formulas describing synchronisation are of course different. For a concrete model with four nodes and a cycle length  $T = 10$ , the synchronisation probability can be queried with the following formula.

$$\text{sync}_c \equiv P_{=?}[\text{F} p_\phi(1) = p_\phi(2) \wedge p_\phi(1) = p_\phi(3) \wedge p_\phi(1) = p_\phi(4)]$$

For a population model, the corresponding property is

$$\text{sync}_p \equiv P_{=?}[\text{F} \bigvee_{1 \leq i \leq 10} k_i = 4]$$

For both types of models, we defined a suitable input for the model checker PRISM, and compared the results for different values of  $R$ ,  $\epsilon$  and  $\mu$ . As expected, the model checking results were matching exactly. Table 4 shows the model construction and checking times for some exemplary parameter combinations of the models as reported by PRISM<sup>1</sup>. In the concrete model, the bulk of time is spent in the model checking phase, while the construction is much faster. For the analysis of the population model, however, the situation is reversed. The model construction phase is an order of magnitude longer than the model checking phase. As expected, the model checker needs less time for the analysis of the population model, if we add the time needed for model construction and checking.

## 7 Conclusion

In this paper we have introduced a formal concrete model for a network of nodes synchronising their clocks over a set of discrete values. Furthermore, we developed a population model that can alleviate state-space explosion when reasoning about

<sup>1</sup> The experiments were run on a computer equipped with an Intel Core i7-7700 CPU at 3.6 GHz and with 16GB of RAM. The version of PRISM used was 4.4 beta.

significantly larger networks. We encoded both models as discrete-time Markov chains, and formally connected them by showing that a concrete model of a network weakly simulates a population model of that same network. We then showed that these two models are equivalent with respect to the reachability of distinguished states, namely those where all nodes in the network have synchronised their clocks.

Formalising the individual nodes of a network allows for the analysis of their internal properties. However, this internal structure also inhibits the verification of global network properties. Modelling the whole network as the product of the models for the individual nodes quickly, and unsurprisingly, results in a model that is too large to analyse with existing tools and techniques. While the use of appropriate collective abstractions, such as population models, allow for the analysis of larger networks, they often impose restrictions on the topologies of the network that can be considered. We could, of course, simply take the product of individual population models to represent network structures more specialised than the fully-connected graphs considered here, but again we face the consequences of this approach when trying to analyse the resulting model. In addition, when using population models we lose the possibility to distinguish between nodes having the same internal state. However, this does not restrict our analysis when considering networks of homogeneous nodes where the properties of interest relate to global behaviours of the network itself.

Our current definition of pulse-coupled oscillators only allows for non-negative results of the phase response function. However, there are also oscillator definitions with phase response functions with possibly negative values [32]. That is, instead of shifting the state of an oscillator towards the end of the cycle, the perturbation may reduce the value of the oscillator's state. It would be interesting to study the impact of negative-valued phase response functions in the setting of discrete clock values.

While a concrete model can be instantiated to incorporate different topologies by explicit encoding of possible perturbances in the nodes' transitions, it is by no means obvious how to incorporate topologies into a population model. By design, the nodes in the latter are indistinguishable, hence the differences in the connections between nodes are lost. We could alleviate this restriction slightly, by modelling the connection between networks of strongly connected components. That is, each component can be modelled by a different population model, and the firings within one model can perturb different models. However, this would mean to compute the cross-product of the population model, and hence we are back at the state-space explosion problem. Furthermore, our abstract relation would need to take the mapping of single nodes into different components into account.

Deductive approaches might serve as an additional way to verify larger systems. In particular, due to the regularity of population models, we conjecture the existence of an inductive invariant that holds from a certain size of models onwards. That is, as soon as the population grows to a size to be treated as a single entity, we can increase this size by one node, and guarantee that synchronisation still occurs. For the population model sizes below this threshold, we could still use our proposed model-checking technique as the induction base. However, it is not clear what such an invariant should be, and how it can be verified.

## References

1. Akyildiz IF, Su W, Sankarasubramaniam Y, Cayirci E (2002) Wireless sensor networks: a survey. *Computer Networks* 38(4):393–422
2. Alur R, Dill DL (1994) A theory of timed automata. *Theoretical Computer Science* 126(2):183–235
3. Angluin D, Aspnes J, Diamadi Z, Fischer MJ, Peralta R (2006) Computation in networks of passively mobile finite-state sensors. *Distributed Computing* 18(4):235–253
4. Atmel Corporation (2018) ATmega128L: 8-bit Atmel microcontroller with 128 kBytes in-system programmable flash. <http://www.atmel.com/images/doc2467.pdf>, last accessed 6th April 2018
5. Bartocci E, Corradini F, Merelli E, Tesi L (2010) Detecting synchronisation of biological oscillators by model checking. *Theor Comput Sci* 411(20):1999–2018
6. Basler G, Mazzucchi M, Wahl T, Kroening D (2009) Symbolic Counter Abstraction for Concurrent Software. In: CAV 2009, Springer, Heidelberg, LNCS, vol 5643, pp 64–78, DOI 10.1007/978-3-642-02658-4\_9
7. Behrmann G, David A, Larsen KG, Hakansson J, Petterson P, Yi W, Hendriks M (2006) Uppaal 4.0. In: QEST 2006, IEEE Computer Society, pp 125–126
8. Breza M (2013) Bio-inspired tools for a distributed wireless sensor network operating system. PhD thesis, Imperial College, London
9. Chen Z, Zhang D, Zhu R, Ma Y, Yin P, Xie F (2013) A review of automated formal verification of ad hoc routing protocols for wireless sensor networks. *Sensor Letters* 11(5):752–764
10. Christensen AL, Grady RO, Dorigo M (2009) From fireflies to fault-tolerant swarms of robots. *IEEE Trans Evolut Comput* 13(4):754–766
11. Cristian F (1989) Probabilistic clock synchronization. *Distributed Computing* 3(3):146–158
12. Delzanno G (2003) Constraint-based Verification of Parametrized Cache Coherence Protocols. *Formal Methods in System Design* 23(3):257–301
13. Donaldson AF, Miller A (2006) Symmetry reduction for probabilistic model checking using generic representatives. In: ATVA 2006, Springer, Berlin, Heidelberg, LNCS, vol 4218, pp 9–23, DOI 10.1007/11901914\_4
14. Emerson EA, Trefler RJ (1999) From asymmetry to full symmetry: New techniques for symmetry reduction in model checking. In: CHARME 1999, Springer, Berlin, Heidelberg, LNCS, vol 1703, pp 142–156, DOI 10.1007/3-540-48153-2\_12
15. Feller W (1968) An introduction to probability theory and its applications: volume I, vol 3. John Wiley & Sons New York
16. Gainer P, Linker S, Dixon C, Hustadt U, Fisher M (2017) Investigating parametric influence on discrete synchronisation protocols using quantitative model checking. In: QEST 2017, Springer, Cham, LNCS, vol 10503, pp 224–239, DOI 10.1007/978-3-319-66335-7\_14
17. Gainer P, Linker S, Dixon C, Hustadt U, Fisher M (2018) The power of synchronisation: Formal analysis of power consumption in networks of pulse-coupled oscillators. In: ICFEM 2018, Springer, accepted for publication
18. Gusella R, Zatti S (1989) The accuracy of the clock synchronization achieved by TEMPO in Berkeley UNIX 4.3BSD. *IEEE Trans Soft Eng* 15(7):847–853



19. Hansson H, Jonsson B (1994) A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6(5):512–535
20. Heidarian F, Schmaltz J, Vaandrager F (2012) Analysis of a clock synchronization protocol for wireless sensor networks. *Theor Comput Sci* 413(1):87–105
21. Kemeny JG, Snell JL, Knapp AW (2012) Denumerable Markov chains: with a chapter of Markov random fields by David Griffeath, Graduate Texts in Mathematics, vol 40. Springer, New York
22. Kwiatkowska M, Norman G, Parker D (2007) Stochastic model checking. In: SFM 2007, Springer, Berlin, Heidelberg, LNCS, vol 4486, pp 220–270, DOI 10.1007/978-3-540-72522-0\_6
23. Maróti M, Kusy B, Simon G, Lédeczi A (2004) The flooding time synchronization protocol. In: Proceedings of SenSys 2004, ACM, pp 39–49
24. Milner R (1971) An Algebraic Definition of Simulation Between Programs. In: Proceedings of IJCAI 1971, Morgan Kaufmann, pp 481–489
25. Mirollo RE, Strogatz SH (1990) Synchronization of pulse-coupled biological oscillators. *SIAM J App Math* 50(6):1645–1662
26. Pagliari R, Scaglione A (2007) Design and implementation of a PCO-based protocol for sensor networks. In: Proceedings of SenSys 2007, ACM, pp 387–388
27. Perez-Diaz F, Zillmer R, Groß R (2015) Firefly-inspired synchronization in swarms of mobile agents. In: Proceedings of AAMAS 2015, International Foundation for Autonomous Agents and Multiagent Systems, pp 279–286
28. Perez-Diaz F, Trenkwalder SM, Zillmer R, Groß R (2018) Emergence and inhibition of synchronization in robot swarms. In: DARS 2016, Springer, Cham, pp 475–486, DOI 10.1007/978-3-319-73008-0\_33
29. Peskin C (1975) Mathematical aspects of heart physiology. Courant Lecture Notes, Courant Institute of Mathematical Sciences, New York University
30. Sommer P, Wattenhofer R (2009) Gradient clock synchronization in wireless sensor networks. In: Proceedings of IPSN 2009, IEEE, pp 37–48
31. Tyrrell A, Auer G, Bettstetter C (2006) Fireflies as role models for synchronization in ad hoc networks. In: Proceedings of BIONETICS 2006, ACM, pp 1–7
32. Wang Y, Nuñez F, Doyle FJ (2012) Energy-efficient pulse-coupled synchronization strategy design for wireless sensor networks through reduced idle listening. *IEEE Trans Sig Proc* 60(10):5293–5306
33. Webster M, Breza M, Dixon C, Fisher M, McCann J (2018) Formal verification of synchronisation, gossip and environmental effects for critical IoT systems. In: Proceedings of AVoCS 2018, EasyChair, EasyChair Preprint no. 377
34. Werner-Allen G, Tewari G, Patel A, Welsh M, Nagpal R (2005) Firefly-inspired sensor network synchronicity with realistic radio effects. In: Proceedings of SenSys 2005, ACM, pp 142–153
35. Yick J, Mukherjee B, Ghosal D (2008) Wireless sensor network survey. *Computer Networks* 52(12):2292–2330