Analysis of a Post-translational Oscillator Using Process Algebra and Spatio-temporal Logic

Analysis of a Post-translational Oscillator Using Process Algebra and Spatio-temporal Logic

Christopher J. Banks[1], Daniel D. Seaton[2], and Ian Stark[3]

[1] Systems Genomics, The Roslin Institute, University of Edinburgh
[2] Millar Lab, Synthetic and Systems Biology, University of Edinburgh
[3] Laboratory for Foundations of Computer Science, School of Informatics, University of Edinburgh

Abstract. We describe the modelling of a post-translational oscillator using a process algebra and the specification of complex properties of its dynamics using a spatio-temporal logic. We show that specifications in the Logic of Behaviour in Context can be seen as hypotheses about oscillations and other biochemical behaviours, to be tested automatically by model-checking software. By using these techniques we show that the theoretical model behaves in a manner in keeping with known properties of biological circadian oscillators.

1 Introduction

In this paper we describe the encoding of a post-translational oscillator (PTO) model in the Continuous Pi-calculus process algebra (cπ)[12, 11] and the results of computational experiments made on the model including the use of a novel spatio-temporal logic, the Logic of Behaviour in Context ($\mathcal{LBC}$)[4], to specify and check complex properties of the model. The spatio-temporal logic $\mathcal{LBC}$ can be seen as a formal logical language for expressing properties of biochemical systems. There is a long-standing problem of how to express properties of oscillation in temporal logic and one contribution of this paper is to neatly define a temporal logic specification for both general and more specific oscillatory behaviour.

PTOs generate sustained oscillations in the absence of transcription and translation. Such oscillators are of particular interest in the circadian clock field, where PTOs have recently been postulated to generate endogenous 24-hour rhythms in diverse organisms [15]. Here, we investigate a minimal model of a PTO due to Jolley et al. [10]. This model has a simple structure—it consists only of a kinase, a phosphatase, and a substrate—but can exhibit robust oscillatory behaviour similar to that observed in circadian clocks.

The purpose of our study is to further examine the behavioural properties of the PTO when it is coupled with other PTOs, other reaction pathways, and inhibitors. We examine these properties using both simple computational experiments and more complicated, *higher-order* experiments defined by $\mathcal{LBC}$ properties and performed by model checking. The ultimate goal being to evaluate $\mathcal{LBC}$ as a useful logical tool to perform these sorts of analyses and also to draw some

conclusions about the behaviour of the theoretical PTO in relation to real circadian oscillators. In particular circadian clock mechanisms must interact with other systems in an organism; this includes the control of metabolic processes and coupling with the classical transcription-translation feedback loop (TTFL) circadian clocks [1]. This potential to robustly interact with other systems is, to date, unexplored for this model.

A key benefit of using process algebra here is the ability to readily build complex models of interaction through combining simpler components. This composition can be challenging, particularly where components are shared or linked; and this is important as such sharing can be source of significant new behaviours [16]. In this study we compose oscillators through shared enzymes.

The advantage of using a formal logical language to specify and check properties of the model, and its composition with other models, is that it gives us a concise and precise means of expressing the hypothesis we wish to test. Model-checking software then gives us the means to test this. This is especially true where we have a mixture of temporal and spatial behaviour we wish to test; e.g. if we wish to know if the introduction of an inhibitor (*spatial* change) has a given effect regardless of the time at which it is applied (*temporal* change).

Using these techniques we show that the Jolley PTO model does indeed exhibit some of the properties that would be expected of a biological PTO. The oscillator is robust when coupled with other oscillators, using different coupling mechanisms, and crucially when coupling at any point in the oscillation cycle. We also show that the oscillator is robust when perturbed by other simple mechanisms. Finally we demonstrate that $\mathcal{LBC}$ has the potential to specify, at least the qualitative aspects of, even more complex properties of oscillators, such as phase response—that is, how the oscillation is affected by a small perturbation at different points in its cycle.

2 The Jolley PTO model

Jolley, Ode, and Ueda present their model as a set of coupled ODEs. In their paper [10], sets of parameters are identified which give distinct patterns of oscillation in the system. The model aims to provide a framework for analysing and synthesising PTOs and they provide evidence that it is a viable candidate for a minimal circadian clock. However, to date, little further analysis of the properties of the complex behaviour of this oscillator has been done.

The model arose from the observation that PTOs and other oscillatory systems which exist in nature are commonly mediated by multi-site phosphorylation, these include evidence from observations and existing models of the KaiC circadian oscillator [14, 9, 17], the MAP Kinase signalling pathway [6, 13], and others [10]. This motivated the search for the simplest possible phosphorylation-mediated oscillator, to serve as a design principle.

The structure of Jolley's PTO (jPTO), described diagrammatically in Figure 1a, is one molecule with two phosphorylation sites. Therefore the molecule has four states (S00, S01, S10, S11) depending on which of its sites are phos-

phorylated. Two opposing enzymes, a kinase (E) and a phosphatase (F), act to phosphorylate or dephosphorylate a site, respectively.

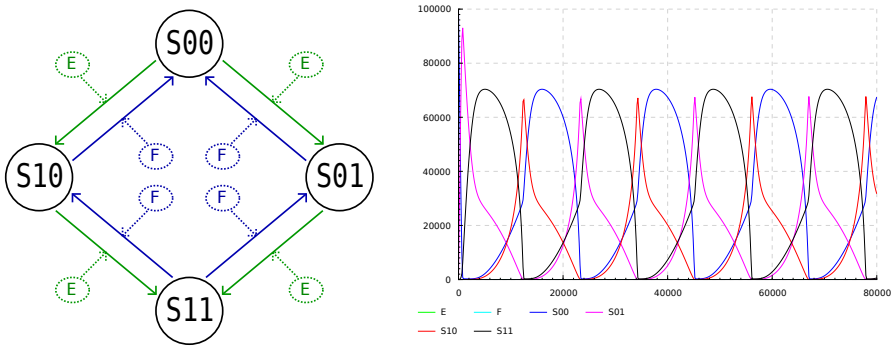

(a) Structure of jPTO, showing the four substrate molecule states, the kinase E, and the phosphatase F.

(b) Oscillatory dynamics of jPTO.

Fig. 1: Structure and dynamics of the Jolley PTO.

The parameters for this model were found by using computational parameter fitting techniques. They then used a clustering algorithm to determine two distinct clusters of parameter sets which produced two different patterns of oscillation. In this study we use one of these; Figure 1b shows the behaviour of the model given these parameters.

3 Process algebra model construction

Model construction in cπ is species-centric. That is the biochemical species, or reagents, are the focus of the modelling process. We first define each species and its binding sites and actions. We then define how different species can interact with each other. Then we define the initial conditions of our mixture, which species are present and in what concentrations. The model can then be executed to determine the behaviour, using numerical simulation. The remainder of this section gives an overview of the construction and execution of the model in cπ. We omit the finer details of cπ syntax and semantics as these are described by Kwiatkowski and Stark [12] and in Kwiatkowski's thesis [11].

3.1 Species

The species in our model are the kinase E, the phosphatase F, and the substrate molecule which has four phosphorylation states $S00$, $S01$, $S10$, and $S11$. The

simplest of these are the two enzymes; they are defined as follows:

$$\begin{aligned} E &\triangleq e(x).x.E \\ F &\triangleq f(x).x.F \end{aligned}$$

The kinase E has a site e and the phosphatase F has a site f. Each can interact on its site with another molecule, perform some other function which depends on the molecule it is bound to, then return to its original state—from which it can perform the same action again. This directly corresponds to the definition of an enzyme.

In our cπ model we represent each of the four states of the substrate as a distinct species. This is simply to break down the syntactic description into smaller parts. In this model a change of state is essentially a change of species, but to the observer these species can be considered as one. The substrate can be defined as follows:

$$\begin{aligned} S00 &\triangleq (\nu M_{00})\ s00a\langle be\rangle.(u.S00 + ra.S01) \\ &\qquad + s00b\langle be\rangle.(u.S00 + rb.S10) \\ S01 &\triangleq (\nu M_{01})\ s01e\langle be\rangle.(u.S01 + r.S11) \\ &\qquad + s01f\langle bf\rangle.(u.S01 + r.S00) \\ S10 &\triangleq (\nu M_{10})\ s10e\langle be\rangle.(u.S10 + r.S11) \\ &\qquad + s10f\langle bf\rangle.(u.S10 + r.S00) \\ S11 &\triangleq (\nu M_{11})\ s11a\langle bf\rangle.(u.S11 + ra.S01) \\ &\qquad + s11b\langle bf\rangle.(u.S11 + rb.S10) \end{aligned}$$

Here each of the states is defined, each containing a definition of the behaviour at each of the two phosphorylation sites. Each of these definitions is similar in structure, reflecting that they in fact represent distinct states of the same molecule. For example, let us examine the definition of $S01$.

One of the two states where one site is phosphorylated, but not the other, is $S01$. The term begins with a ν-term. The ν-term defines a local affinity network M_{01}; this governs the local interactions of unbinding or reacting in the same way as the global affinity network which will be defined below and defines the internal interaction potential of the complexes formed between substrate and enzyme to unbind (u) or react (r).

The structure of $S01$ is then defined as having two sites $s01e$ and $s01f$, each with some behaviour which follows from another molecule binding on that site. Once we have defined which molecules can interact on which sites (below), $s01e$ will accept the kinase E and $s01f$ will accept the phosphatase F. The behaviour which follows binding is defined by the next part of the term; in this case the bound enzyme can either unbind and the substrate returns to state $S01$ or the reaction can occur, changing the substrate either to state $S00$ or to $S11$, depending on whether F or E is bound.

The definition of each of the other states of the substrate follow the same pattern. Full details of definitions, the affinity networks, and their rates can be found in Appendix A.

3.2 Interactions

Now we have the definitions of the molecules and their interaction sites, we need to define which molecules can bind to which sites and at what rate these reactions occur. This is done by means of an affinity network M:

$$\begin{aligned} M = \{ & s00a \leftrightarrow e,\ \ s00b \leftrightarrow e,\ \ s01e \leftrightarrow e,\ \ s10e \leftrightarrow e, \\ & s01f \leftrightarrow f,\ \ s10f \leftrightarrow f,\ \ s11a \leftrightarrow f,\ \ s11b \leftrightarrow f\} \end{aligned}$$

Here we state that each of the substrate sites interacts with either site e of the kinase or site f of the phosphatase. Each of these interactions has a given reaction rate (see Appendix A).

3.3 Mixture

Having now defined the structure and rate parameters of the model, all that remains to be able to execute the model is a definition of the initial conditions we wish to simulate. Here we define a process $\mathit{\Pi}$ which lists the species present and their initial concentrations.

$$\mathit{\Pi} \ \triangleq \ c_S \cdot S00 \parallel c_E \cdot E \parallel c_F \cdot F$$

Here we have some concentration c_S of substrate in its unphosphorylated state $S00$ and likewise some concentrations c_E and c_F of E and F (see Appendix A).

3.4 Validation

From this description the cπ tool generates a vector-space model of species concentrations over time. This is then compiled into a set of model ODEs and an initial value problem, suitable for numerical simulation. In this case the model description generates precisely the set of ODEs which were defined by Jolley et al. and therefore precisely the same behaviour; as shown in Figure 1b.

4 Basic time series analysis

In this section we describe a number of computational experiments which were performed, aided by the compositional nature of the cπ description of the model.

4.1 Coupled jPTOs

The first experiment determines the behaviour of two identical jPTOs when coupled. The coupling is achieved by the two jPTOs sharing a pool of enzymes E and F.

We achieve the coupling in our cπ model in the following way. First we make a copy of the substrate species, call it T as shown in Appendix B. The process term can then be updated to include our new copy:

$$\Pi \triangleq c_S \cdot S00 \parallel c_T \cdot T00 \parallel c_E \cdot E \parallel c_F \cdot F$$

and the global affinity network M can then allow E and F to interact with the sites of T.

The behaviour of the coupled jPTOs can be seen in Figure 2a. The result of coupling two identical jPTOs is that the two act in synchrony, but the period is doubled. It is clear that the doubling of the period is due to each jPTO only having half the concentration of enzymes available, the other half of the concentration being sequestered by the other jPTO—each is competing equally for the same pool.

4.2 Weaker coupling

It is possible to consider other schemes for coupling. For example, if the coupling was made weaker by only sharing one of the enzymes, does synchronisation still occur?

Here we take two jPTOs in a similar manner to above, however we only share the kinase E. This is achieved in the model simply by having a separate phosphatase for each jPTO, F_S and F_T:

$$\Pi \triangleq c_S \cdot S00 \parallel c_T \cdot T00 \parallel c_E \cdot E \parallel c_{F_S} \cdot F_S \parallel c_{F_T} \cdot F_T$$

Here we set $c_{F_S} = c_{F_T} = c_E$ and we then set the global affinity network accordingly. See Appendix C.

We can see, in Figure 2b, that indeed the jPTOs still synchronise when coupled less strongly. We can also see that each jPTO, given its own pool of phosphatase, spends more time in the less phosphorylated states as it can dephosphorylate at a greater rate than it can phosphorylate. If the concentration of each enzyme was adjusted accordingly, so $c_{F_S} + c_{F_T} = c_E$, then the system behaves as the coupled jPTOs sharing both kinase and phosphatase (as Figure 2a).

4.3 Coupling out of phase

This experiment determines the behaviour of coupling a jPTO with an identical jPTO, but out of phase—that is, when the models are coupled with oscillators beginning at different points in their cycle. To achieve this we take two jPTO models of identical structure, but the second is shifted by a quarter of its cycle, we call this jPTO-90.

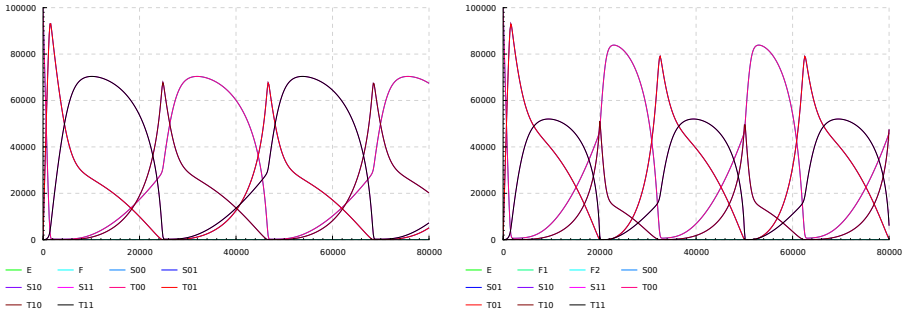

(a) Coupling two identical jPTOs.

(b) Weaker coupling by sharing kinase only.

Fig. 2: jPTO composition dynamics.

When the two models are composed we can see that, after a transient period, the cycles of the two jPTOs begin to synchronise as shown in Figure 3a. For comparison we also coupled jPTOs in various phase states. Synchronisation appears to occur in a number of selected phases. This suggests that the synchronisation of two jPTOs is quite robust. Figure 3b shows synchronisation when jPTOs are coupled in anti-phase: jPTO-180.

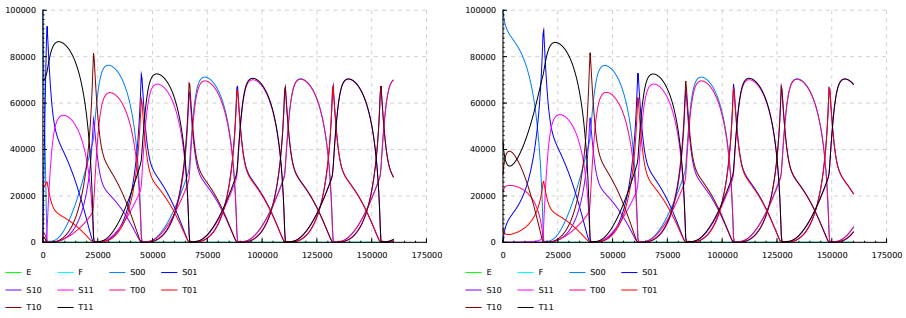

(a) jPTO + jPTO-90

(b) jPTO + jPTO-180

Fig. 3: Coupling of two identical clocks, sharing an enzyme pool, starting out of phase.

4.4 Perturbation

Another useful property of a circadian oscillator is that it is robust to some perturbations—although others may disrupt it. In this experiment we determine the behaviour of the jPTO when perturbed by a pulse of some inhibitor.

To construct a model for this we first construct an inhibitor molecule which rapidly appears in the system and decays rapidly. The mechanism for inhibitor appearing in the system is to have another molecule which is initially present and autonomously becomes the inhibitor. The inhibitor then decays. We will use the inhibitor to bind and sequester components of the jPTO. See Appendix E.

Figure 4a shows the effect of inhibiting the kinase; there is a transient period—about as long as the pulse—and then the jPTO settles back into its normal oscillation. This shows that the jPTO is somewhat robust to temporary sequestration of its enzymes.

Figure 4b shows the result of the inhibitor sequestering the doubly phosphorylated substrate molecule S11. Here the fact that S11 itself is only present in pulses and the fact that the inhibitor does not decay when it is bound to S11 means that the inhibitor remains in the system for longer. We can see echo pulses as the inhibitor binds and unbinds the fluctuating concentration of S11. However, overall, the inhibitor eventually decays and the system stabilises. This shows that the substrate is also robust to this kind of perturbation.

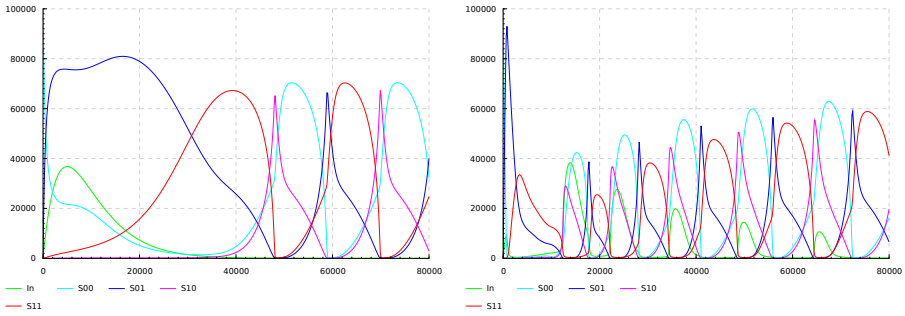

(a) Inhibitor binds E.

(b) Inhibitor binds S11.

Fig. 4: Perturbation of a jPTO with a pulse of inhibitor.

5 Model-checking experiments

The experiments in the previous section show a number of properties which are mostly amenable to analysis by conventional techniques. The compositional nature of $c\pi$ models aids greatly in the model construction for models where we are looking at compositions of two models or composition with an inhibitor; something which is much more difficult to do by working directly with ODEs. However the analysis of these models is little more than the inspection of time series for a relatively small set of models and initial conditions.

We will see that we can use the model checking of $\mathcal{LBC}$ specifications to automate the process of inspecting time series for a given behaviour. Moreover,

and most importantly, we can define higher-order experiments which require many models and many initial conditions. We gain a means to precisely express a set of computational experiments, which in a conventional setting would require case-specific programming, and to have them automatically checked.

5.1 Behaviour under composition

The spatial aspects of $\mathcal{LBC}$ [4] directly take advantage of model compositionality. Specifications about the behaviour of a model when it is composed with another model can be made using the context modality. We can make use of this in analysing the behaviour of coupled oscillators.

A basic assertion in $\mathcal{LBC}$ has the form $P \models \phi$, meaning that system P satisfies property ϕ. A property of the form $Q \triangleright \psi$ is a *guarantee*: if Q is introduced to the system, then the resulting mixture satisfies ψ. For example: $PTO1 \models PTO2 \triangleright \phi$ states that when we couple $PTO1$ and $PTO2$ we have some behaviour ϕ. Likewise we could state $PTO \models Inhib \triangleright \phi$ meaning that our PTO has some behaviour ϕ when we introduce an inhibitor (*Inhib*).

However, the most interesting properties are those which make a statement about introducing something over time. For example: $PTO1 \models \mathbf{G}_t(PTO2 \triangleright \phi)$ which uses the temporal operator $\mathbf{G}$ (for *globally* true) to declare $PTO1$ has the property that for any time up to t, if we add $PTO2$ then the system from that point will satisfy ϕ.

5.2 Complex dynamics

$\mathcal{LBC}$ also has the power to express complex dynamics, such as periodicity and oscillation. Numerous bodies of work have attempted to express oscillation properties in standard temporal logic [3, 2, 5], but all fall short of a general formula for oscillation. It is possible to express oscillation, however, with some prior knowledge of the type of oscillation. Following the idea in Calzone et al. [5] and extending it to a time-bounded logic, we can express oscillation in the temporal fragment of $\mathcal{LBC}$ as follows:

$$PTO \models \mathbf{G}_{[0,t]}(\mathbf{F}_{[0,p]}(([S]' > 0) \wedge \mathbf{F}_{[0,p]}([S]' < 0)))$$

where $[S]'$ is the first derivative of $[S]$ with respect to time. The formula states that at any time up to t the concentration of S will, within a further time p, be rising and then within another additional time p be falling. This describes a repeated rising and falling with period at most p. Whilst this is not a general formula, it does cover a large class of sustained oscillation. However, its weakness is that it does not distinguish from noise—although noise is not a problem when studying ODE models.

It has been shown that more expressive logics can express more general formulae for oscillation; for example Dluhoš et al. [7] show that one can use a "freeze operator" to do this. In fact, it is possible to give a general formula for

sustained—and not necessarily regular—limit cycle periodicity using $\mathcal{LBC}$. The formula:

$$PTO \models \mathbf{F}_{[p_{min},p_{max}]}(\widehat{PTO} \triangleright (\mathbf{F}_{[0,s]}\mathbf{G}_{[0,t]}(|[S] - [\widehat{S}]| < \epsilon))) \tag{1}$$

where $\widehat{PTO}$ is a copy of PTO, S is the species being observed, $\widehat{S}$ is the copy of S in $\widehat{PTO}$, and s is a maximum transient period before reaching the limit cycle. The formula states that if we introduce $\widehat{PTO}$ after some period in $[p_{min}, p_{max}]$ then, within s, $[S]$ and $[\widehat{S}]$ will synchronise to within ϵ for at least time t. This essentially takes a copy of the model, shifts it forward in time by $p_{min} \leq t \leq p_{max}$, and determines if it matches up with the original model. If it does, allowing for some initial transient period, then the model is periodic in species S.

In the context of our case study, we can now check if coupled PTOs still oscillate:

$$PTO1 \models \mathbf{G}_{[0,c]}(PTO2 \triangleright \texttt{Osc})$$

where c is the end of the first cycle of *PTO1* and `Osc` is one of our oscillation formulae from above. If coupling the PTOs at any time within the first cycle of *PTO1* gives a system which still oscillates—with some period bounds, as above—then the formula will be true.

5.3 Perturbation response

$\mathcal{LBC}$ can be used to express properties of a system under perturbation. For example, one might wish to determine if some perturbation causes a greater peak concentration in a species S. The formula:

$$PTO \models \mathbf{F}_{[0,t]}(P \triangleright \mathbf{F}_{[0,r]}([S] > pk))$$

states that some peak value pk is exceeded under some perturbation P, within time t, where r is the maximum expected time of the peak after the perturbation. As the perturbation P could be any model, it could simply be a quantity of some species, a constant amount of inhibitor, a pulse of inhibitor, etc.

Of particular interest in the study of oscillators are the *phase response* [8] characteristics of system. That is, given a short perturbation, at any point in the cycle, what is the effect on the phase of the oscillation? Biologists often plot a *phase response curve*, using a large number of experiments, to visualise the phase response. $\mathcal{LBC}$ cannot give such a precise and quantitative account of phase response as this, however it is certainly possible to formulate some more qualitative—or even semi-quantitative—properties of phase response. For example:

$$PTO \models \widehat{PTO} \triangleright \mathbf{F}_{[c_1,c_2]}(P \triangleright (\mathbf{G}_{[t_1,t_2]}([\widehat{S}]' > 0 \implies \mathbf{F}_{[s_1,s_2]}[S]' > 0)))$$

states that some perturbation P applied within $[c_1, c_2]$ will cause a *forward* phase shift in $[s_1, s_2]$. t_1 is a known max transient period after introducing P, t_2 is a sensible maximum time to simulate for, and the formula assumes that we know the perturbed system still oscillates.

5.4 Results

The following results of verifying the above $\mathcal{LBC}$ properties against the cπ models of Jolley's PTO were obtained by using the reference implementation of the $\mathcal{LBC}$ signal-based model checker[4]. First we show a number of formulae which give the same results as the experiments performed above, albeit without the need to manually inspect a simulation trace. These results serve to verify the use of the model checker. Finally we show the results of checking formulae which describe higher-order computational experiments, i.e. those which check properties which would require the inspection of a great number of simulation runs.

Oscillation Our first test was to check for oscillation using Formula 1. We let:

$$\texttt{Osc} = \mathbf{F}_{[p_{min},p_{max}]}(\widehat{jPTO} \triangleright (\mathbf{F}_{[0,s]}\mathbf{G}_{[0,t]}(|[S00] - [\widehat{S00}]| < \epsilon)))$$

where: we know that the period is around 24000 so we set $p_{min} = 23000$ and $p_{min} = 25000$; we know the system will reach limit cycle within $s = 10000$; we must choose an oscillating species and so choose $S = S00$; a reasonable time to simulate for is $t = 80000$—a few cycles; and we choose $\epsilon = 1$ as our concentration accuracy. The copy model $\widehat{jPTO}$ can be constructed in the same manner as the copy model in Section 4.1 or by using the appropriate function in the reference implementation.

Upon checking $jPTO1 \models \texttt{Osc}$ we find that it returns *True*. This confirms what we have been able to determine manually from inspecting the simulation traces in Figure 1b. Moreover, it shows that the $\mathcal{LBC}$ formula is a succinct and precise means of expressing the oscillation property and the model checker provides an automatic means for testing such a hypothesis.

Coupled oscillators The next step is to test coupled oscillators for oscillation, as in Section 4.1. First we take identical PTOs: *jPTO1* and *jPTO2*. Upon checking $jPTO1 \models jPTO2 \triangleright \texttt{Osc}$ using the same formula parameters as above, we find that the result is *False*. This is because, as seen in Figure 2a, the period of the coupled oscillators is doubled. Therefore, upon relaxing the desired period range to $p_{min} = 23000$ and $p_{max} = 49000$ we find the formula is satisfied and the checker returns *True*.

Out-of-phase coupling In Section 4.3 we showed that for a limited number of out-of-phase couplings of *jPTO1* and *jPTO2* the two systems did indeed synchronise and oscillate together after an initial transient period. This however does not confirm whether this is the case for *all* phase shifts.

Using the test $jPTO1 \models \mathbf{G}_{[0,c]}(jPTO2 \triangleright \texttt{Osc})$ we can use the model checker to give a greater guarantee that coupling the oscillators in any phase shift, up to c times the length of one cycle. Here we know that the length of one cycle is no more than, say, $c = 26000$.

[4] Part of the CPiWorkBench: http://banks.ac/software/

Upon checking, again with the above formula parameters and the relaxed period range, we find that the result is *False*. This is because we have not accounted for the lengthened transient period when coupling out of phase. If we increase the parameter s to 120000 we find the formula is now satisfied, the result is *True*. This higher-order property gives a much stronger guarantee that all out-of-phase couplings oscillate than a limited number of manually inspected simulation traces would give.

Phase response Another higher-order property is the phase response characteristic. Using the inhibitor pulse model in Section 4.4—except using a pulse which lasts for fewer than 2000 time units to match the kind of pulse which would be used to plot a phase response curve—and the formula in Section 5.3 we can place some bounds on the phase response characteristics of the model:

$$jPTO \models \widehat{jPTO} \triangleright \mathbf{F}_{[c_1,c_2]}(Pulse \triangleright (\mathbf{G}_{[t_1,t_2]}([\widehat{S00}]' > 0 \implies \mathbf{F}_{[s_1,s_2]}[S00]' > 0)))$$

where: $[c_1, c_2] = [10000, 34000]$ which is roughly one cycle, this limits the computation; the maximum expected transient period is $t_1 = 10000$; the maximum time to compare oscillations is $t_2 = 80000$; and $[s_1, s_2] = [0, 1000]$ ensures that the whole formula states that: "there is always a forward phase response of no more than 1000 time units".

The model checker confirms that this statement is true for this model. So our small pulse may delay the cycle, but only by a relatively small time; it does not speed up the cycle.

6 Conclusions

We have shown that, using a combination of cπ and $\mathcal{LBC}$, we can express a variety of complex properties of biochemical models. We have shown that precise and succinct statements of complex properties can be built up in a modular fashion. One can even think of higher-order $\mathcal{LBC}$ properties as precise statements of an experimental hypothesis, to be tested by the model checker.

We have also shown that the Jolley PTO model does indeed interact robustly with other oscillators and inhibitors. This includes showing that the oscillator can be coupled at any point in its cycle and that it shows a robust inhibitor response at any point in its cycle. These latter properties are shown using $\mathcal{LBC}$ statements describing higher-order experiments; these are automatically tested by the model checker without any human intervention nor the necessity to write explicit programs for the necessary inspection of large numbers of simulation runs. Extensions to this work could readily include investigating the results of coupling with oscillators of a different type and of coupling with downstream networks.

References

1. Abraham, U., Granada, A. E., Westermark, P. O., Heine, M., Kramer, A., and Herzel, H. Coupling governs entrainment range of circadian clocks. *Molecular systems biology 6*, 1 (Nov. 2010).
2. Ballarini, P., and Guerriero, M. L. Query-based verification of qualitative trends and oscillations in biochemical systems. *Theoretical Computer Science 411*, 20 (Apr. 2010), 2019–2036.
3. Ballarini, P., Mardare, R., and Mura, I. Analysing Biochemical Oscillation through Probabilistic Model Checking. *Electronic Notes in Theoretical Computer Science 229*, 1 (Feb. 2009), 3–19.
4. Banks, C. J., and Stark, I. A Logic of Behaviour in Context. *Information and Computation 236* (2014), 3–18.
5. Calzone, L., Chabrier-Rivier, N., Fages, F., and Soliman, S. Machine learning biochemical networks from temporal logic properties. *Transactions on Computational Systems Biology VI* (2006), 68–94.
6. Chickarmane, V., Kholodenko, B. N., and Sauro, H. M. Oscillatory dynamics arising from competitive inhibition and multisite phosphorylation. *Journal of theoretical biology 244*, 1 (Jan. 2007), 68–76.
7. Dluhoš, P., Brim, L., and Šafránek, D. On Expressing and Monitoring Oscillatory Dynamics. *Electronic Proceedings in Theoretical Computer Science 92* (Aug. 2012), 73–87.
8. Granada, a., Hennig, R. M., Ronacher, B., Kramer, A., and Herzel, H. Phase response curves elucidating the dynamics of coupled oscillators. *Methods in Enzymology 454* (Jan. 2009), 1–27.
9. Johnson, C. H., Mori, T., and Xu, Y. A cyanobacterial circadian clockwork. *Current Biology 18*, 17 (Sept. 2008), R816–R825.
10. Jolley, C. C., Ode, K. L., and Ueda, H. R. A Design Principle for a Post-translational Biochemical Oscillator. *Cell reports 2*, 4 (Oct. 2012), 938–950.
11. Kwiatkowski, M. *A formal computational framework for the study of molecular evolution*. PhD thesis, University of Edinburgh, 2010.
12. Kwiatkowski, M., and Stark, I. The continuous π-calculus: A process algebra for biochemical modelling. In *Computational Methods in Systems Biology, LNCS 5307* (2008), Springer, pp. 103–122.
13. Liu, P., Kevrekidis, I. G., and Shvartsman, S. Y. Substrate-dependent control of ERK phosphorylation can lead to oscillations. *Biophysical journal 101*, 11 (Dec. 2011), 2572–81.
14. Nakajima, M., Imai, K., Ito, H., Nishiwaki, T., Murayama, Y., Iwasaki, H., Oyama, T., and Kondo, T. Reconstitution of circadian oscillation of cyanobacterial KaiC phosphorylation in vitro. *Science (New York, N.Y.) 308*, 5720 (Apr. 2005), 414–5.
15. O'Neill, J. S., van Ooijen, G., Dixon, L. E., Troein, C., Corellou, F., Bouget, F.-Y., Reddy, A. B., and Millar, A. J. Circadian rhythms persist without transcription in a eukaryote. *Nature 469*, 7331 (Jan. 2011), 554–8.
16. Seaton, D. D., and Krishnan, J. The coupling of pathways and processes through shared components. *BMC systems biology 5*, 1 (2011), 103.
17. van Zon, J. S., Lubensky, D. K., Altena, P. R. H., and ten Wolde, P. R. An allosteric model of circadian KaiC phosphorylation. *Proceedings of the National Academy of Sciences of the United States of America 104*, 18 (May 2007), 7420–5.

A Basic Jolley model

The basic Jolley PTO model is constructed in cπ as follows:

$$\begin{aligned} E &\triangleq e(x).x.E \\ F &\triangleq f(x).x.F \end{aligned}$$

$$\begin{aligned} S00 &\triangleq (\nu M_{00})\ s00a\langle be\rangle.(u.S00 + ra.S01) \\ &\qquad + s00b\langle be\rangle.(u.S00 + rb.S10) \\ S01 &\triangleq (\nu M_{01})\ s01e\langle be\rangle.(u.S01 + r.S11) \\ &\qquad + s01f\langle bf\rangle.(u.S01 + r.S00) \\ S10 &\triangleq (\nu M_{10})\ s10e\langle be\rangle.(u.S10 + r.S11) \\ &\qquad + s10f\langle bf\rangle.(u.S10 + r.S00) \\ S11 &\triangleq (\nu M_{11})\ s11a\langle bf\rangle.(u.S11 + ra.S01) \\ &\qquad + s11b\langle bf\rangle.(u.S11 + rb.S10) \end{aligned}$$

$$\Pi \triangleq c_S \cdot S00 \parallel c_E \cdot E \parallel c_F \cdot F$$

where

$$c_S = 10^5, c_E = 1, c_F = 1.$$

$$\begin{aligned} M_{00} = \{&be \leftrightarrow u : 10.02, \\ &be \leftrightarrow ra : 163.31, \\ &be \leftrightarrow rb : 0\} \end{aligned}$$

$$\begin{aligned} M_{01} = \{&be \leftrightarrow u : 10.02, \\ &be \leftrightarrow r : 40.83, \\ &bf \leftrightarrow u : 10.02, \\ &bf \leftrightarrow r : 8.17\} \end{aligned}$$

$$\begin{aligned} M_{10} = \{&be \leftrightarrow u : 10.02, \\ &be \leftrightarrow r : 8.17, \\ &bf \leftrightarrow u : 10.02, \\ &bf \leftrightarrow r : 40.83\} \end{aligned}$$

$$\begin{aligned} M_{11} = \{&bf \leftrightarrow u : 10.02, \\ &bf \leftrightarrow ra : 0, \\ &bf \leftrightarrow rb : 163.31\} \end{aligned}$$

$$\begin{aligned} M = \{&s00a \leftrightarrow e : 818.18, \\ &s00b \leftrightarrow e : 0, \\ &s01e \leftrightarrow e : 13.64, \\ &s10e \leftrightarrow e : 4903.17, \\ &s01f \leftrightarrow f : 4903.17, \\ &s10f \leftrightarrow f : 13.64, \\ &s11a \leftrightarrow f : 0, \\ &s11b \leftrightarrow f : 818.18\} \end{aligned}$$

B Coupled jPTOs model

The coupled model is constructed from the same substrate and enzyme species as the basic model in Appendix A. The second jPTO is a copy of the original

substrate, renamed so it forms a distinct species:

$$\begin{aligned}
T00 &\triangleq (\nu M_{00})\ t00a\langle be\rangle.(u.T00 + ra.T01) \\
&\qquad + t00b\langle be\rangle.(u.T00 + rb.T10) \\
T01 &\triangleq (\nu M_{01})\ t01e\langle be\rangle.(u.T01 + r.T11) \\
&\qquad + t01f\langle bf\rangle.(u.T01 + r.T00) \\
T10 &\triangleq (\nu M_{10})\ t10e\langle be\rangle.(u.T10 + r.T11) \\
&\qquad + t10f\langle bf\rangle.(u.T10 + r.T00) \\
T11 &\triangleq (\nu M_{11})\ t11a\langle bf\rangle.(u.T11 + ra.T01) \\
&\qquad + t11b\langle bf\rangle.(u.T11 + rb.T10)
\end{aligned}$$

The process term is the same as above, but with the addition of the new (copy) substrate:

$$\Pi \triangleq c_S \cdot S00 \parallel c_T \cdot T00 \parallel c_E \cdot E \parallel c_F \cdot F$$

where

$$c_S = 10^5, c_T = 10^5, c_E = 1, c_F = 1,$$

and the global affinity net is then extended to allow the new substrate to interact with the enzymes:

$$\begin{aligned}
M = \{&s00a \leftrightarrow e : 818.18, && t00a \leftrightarrow e : 181.18, \\
&s00b \leftrightarrow e : 0, && t00b \leftrightarrow e : 0, \\
&s01e \leftrightarrow e : 13.64, && t01e \leftrightarrow e : 13.64, \\
&s10e \leftrightarrow e : 4093.17, && t10e \leftrightarrow e : 4093.17, \\
&s01f \leftrightarrow f : 4093.17, && t01f \leftrightarrow f : 4093.17, \\
&s10f \leftrightarrow f : 13.64, && t10f \leftrightarrow f : 13.64, \\
&s11a \leftrightarrow f : 0, && t11a \leftrightarrow f : 0, \\
&s11b \leftrightarrow f : 818.18, && t11b \leftrightarrow f : 818.18\}.
\end{aligned}$$

C Weaker coupled jPTOs

For the weaker coupled model we have a separate phosphatase for each substrate. The model in Section B is extended by replacing species F with the following:

$$\begin{aligned}
F_S &\triangleq fs(x).x.F_S \\
F_T &\triangleq ft(x).x.F_S
\end{aligned}$$

and the process term is extended:

$$\Pi \triangleq c_S \cdot S00 \parallel c_T \cdot T00 \parallel c_E \cdot E \parallel c_{F_S} \cdot F_S \parallel c_{F_T} \cdot F_T$$

where

$$c_S = 10^5, c_T = 10^5, c_E = 1, c_{F_S} = 1, c_{F_T} = 1,$$

and the affinity net is altered so each substrate only has affinity for one of the phosphatases:

$$\begin{aligned}
M = \{ & s00a \leftrightarrow e : 818.18, & & t00a \leftrightarrow e : 181.18, \\
& s00b \leftrightarrow e : 0, & & t00b \leftrightarrow e : 0, \\
& s01e \leftrightarrow e : 13.64, & & t01e \leftrightarrow e : 13.64, \\
& s10e \leftrightarrow e : 4093.17, & & t10e \leftrightarrow e : 4093.17, \\
& s01f \leftrightarrow fs : 4093.17, & & t01f \leftrightarrow ft : 4093.17, \\
& s10f \leftrightarrow fs : 13.64, & & t10f \leftrightarrow ft : 13.64, \\
& s11a \leftrightarrow fs : 0, & & t11a \leftrightarrow ft : 0, \\
& s11b \leftrightarrow fs : 818.18, & & t11b \leftrightarrow ft : 818.18\}.
\end{aligned}$$

D Driving other reactions

To construct the model which drives another phosphorylation reaction, we first construct P which is the molecule to be phosphorylated:

$$\begin{aligned}
P &\triangleq (\nu M_P)\ p\langle x\rangle.(u.P + r.P') \\
P' &\triangleq \tau_d.P
\end{aligned}$$

where $d = 10^{-4}$ and $M_P = \{x \leftrightarrow u : 1, x \leftrightarrow r : 1\}$.

The model is then the same as the basic model in Appendix A, but with a new site, which interacts with the P molecule, added to the *S11* state of the substrate:

$$\begin{aligned}
S11 \triangleq (\nu M_{11})\ & s11a\langle bf\rangle.(u.S11 + ra.S01) \\
& + s11b\langle bf\rangle.(u.S11 + rb.S10) \\
& + s11p(x).x.S11
\end{aligned}$$

the new molecule added to the process:

$$\Pi \triangleq c_S \cdot S00 \parallel c_E \cdot E \parallel c_F \cdot F \parallel c_P \cdot P$$

where

$$c_S = 10^5, c_E = 1, c_F = 1, c_P = 10^5$$

and the affinity net is extended with

$$\begin{aligned} M = \{&s00a \leftrightarrow e : 818.18, \\ &s00b \leftrightarrow e : 0, \\ &s01e \leftrightarrow e : 13.64, \\ &s10e \leftrightarrow e : 4903.17, \\ &s01f \leftrightarrow f : 4903.17, \\ &s10f \leftrightarrow f : 13.64, \\ &s11a \leftrightarrow f : 0, \\ &s11b \leftrightarrow f : 818.18, \\ &s11p \leftrightarrow p : 3 \times 10^{-4}\}. \end{aligned}$$

E Perturbation

To construct the model with a pulse of inhibitor, we take the model in Appendix D and replace the driven species P with an inhibitor In which decays and a species $ProdIn$ which autonomously produces the inhibitor:

$$\begin{aligned} In &\triangleq (\nu M_{In})\ p\langle x\rangle u.In + \tau_d.0 \\ ProdIn &\triangleq \tau_d.P \end{aligned}$$

where $M_{In} = \{x \leftrightarrow u : 0.1\}$ and $d = 5 \times 10^{-3}$ and the inhibitor producer added to the process:

$$\Pi \triangleq c_S \cdot S00 \parallel c_E \cdot E \parallel c_F \cdot F \parallel c_P \cdot ProdIn$$

where

$$c_S = 10^5, c_E = 1, c_F = 1, c_P = 10^5$$

In this model the inhibitor binds to the substrate in its $S11$ state. The models where the inhibitor binds to one or the other of the enzymes is constructed in a similar way, with a corresponding new site on the enzyme instead of the substrate. When binding to the enzyme, however the rate should be adjusted from 3×10^{-4} to 5.