

Towards modular verification of pathways: fairness and assumptions

Peter Drábik

Istituto di Informatica e Telematica
Consiglio Nazionale delle Ricerche
Pisa, Italy

`peter.drabik@iit.cnr.it`

Andrea Maggiolo-Schettini

Dipartimento di Informatica
Università di Pisa
Pisa, Italy

`maggiolo@di.unipi.it`

Paolo Milazzo

Dipartimento di Informatica
Università di Pisa
Pisa, Italy

`milazzo@di.unipi.it`

Modular verification is a technique used to face the state explosion problem often encountered in the verification of properties of complex systems such as concurrent interactive systems. The modular approach is based on the observation that properties of interest often concern a rather small portion of the system. As a consequence, reduced models can be constructed which approximate the overall system behaviour thus allowing more efficient verification.

Biochemical pathways can be seen as complex concurrent interactive systems. Consequently, verification of their properties is often computationally very expensive and could take advantage of the modular approach.

In this paper we report preliminary results on the development of a modular verification framework for biochemical pathways. We view biochemical pathways as concurrent systems of reactions competing for molecular resources. A modular verification technique could be based on reduced models containing only reactions involving molecular resources of interest.

For a proper description of the system behaviour we argue that it is essential to consider a suitable notion of fairness, which is a well-established notion in concurrency theory but novel in the field of pathway modelling. We propose a modelling approach that includes fairness and we identify the assumptions under which verification of properties can be done in a modular way.

We prove the correctness of the approach and demonstrate it on the model of the EGF receptor-induced MAP kinase cascade by Schoeberl et al.

1 Introduction

A big challenge of current biology is understanding the principles and functioning of complex biological systems. Despite the great effort of molecular biologists investigating the functioning of cellular components and networks, we still cannot provide a detailed answer to the question “how a cell works?”.

In the last decades, scientists have gathered an enormous amount of molecular level information. To uncover the principles of functioning of a biological system, just collecting data does not suffice. Actually, it is necessary to understand the functioning of parts and the way these interact in complex systems. The aim of *systems biology* is to build, on top of the data, the science that deals with principles of operation of biological systems. The comprehension of these principles is done by modelling and analysis exploiting mathematical means.

A typical scenario of modelling a biological system is as follows. To build a model that explains the behaviour of a real biological system, first a formalism needs to be chosen. Then a model of the system is created, simulation is performed, and the behaviour is observed. The model is validated by comparing the results with the real experiments. The advantage of simulation is not only validation of laboratory experiments, but also prediction of behaviour under new conditions and automation of the whole process.

Simulation can give either the average system behaviour or a number of possible system behaviours. This may be insufficient when one is interested in analysing all the behaviours of a system.

Model checking may be of help. This technique permits the verification of properties (expressed as logical formulae) by exploring all the possible behaviours of a system. This analysis technique typically relies on a state space representation whose size, unfortunately, makes the analysis often intractable for realistic models. This is true in particular for systems of interest in systems biology (such as metabolic pathways, signalling pathways, and gene regulatory networks), which often consist of a huge number of components interacting in different ways, thus exhibiting very complex behaviours.

Many formalisms originally developed by computer scientists to model systems of interacting components have been applied to biology, also with extensions to allow more precise descriptions of the biological behaviours [2, 4, 6, 9, 18, 19]. Examples of well-established formal frameworks that can be used to model, simulate and model check descriptions of biological systems are [6, 14, 15].

Model checking techniques have traditionally suffered from the state explosion problem. Standard approaches to the solution of this problem are based on abstractions or similar model reduction techniques (e.g. [7]). Moreover, the use of Binary Decision Diagrams (BDDs) [8] to represent the state space (symbolic model checking) often allows significantly larger model to be treated [3].

A method for trying to avoid the state space explosion problem is to consider a decomposition of the system, and to apply a modular verification technique allowing global properties to be inferred from properties of the system components. This approach can be particularly efficient when the modelled systems consist of a high number of components, whereas properties of interest deal only with rather small subset of them. This is often the case for properties of biological systems. Hence, for each property it would be useful to be able to isolate a minimal fragment of the model that is necessary for verifying such a property. If such a fragment can be obtained by working only on the syntax of the model, the application of a standard verification technique on the semantics of the fragment avoids the state explosion.

In previous work we developed a modular verification technique in which the system of interest is described by means of a general automata-based formalism suitable for qualitative description of a large class of biological systems, called sync-programs, which supports modular construction [11, 12]. Sync-programs include a notion of synchronization that enables the modelling of biological systems. The modular verification technique is based on property preservation and allows the verification of properties expressed in the temporal logic $ACTL^-$ to be verified on fragments of models. In order to handle modelling and verification of more realistic biological scenarios, we have proposed a dynamic version of our formalism along with an extension of the modular verification framework [10].

The long-term aim of our research is the development of an efficient modular verification framework specifically designed for biochemical pathways, and of a pathway analysis tool based on such a framework. Presently, we are at the first stages of the development of the modular verification framework. However, we already faced some problems whose solution required the definition of concepts related to the formal modelling of biochemical pathways and that we believe could be interesting not only in the context of modular verification. In particular, we defined a notion of *fairness* for biochemical pathways and a notion of *molecular component* of a pathway. The former is a well-known concept in concurrency theory that could be useful to describe more accurately the dynamics of a pathway (in a qualitative framework). The latter is a notion relating species involved in the same pathway such that two species are considered to be part of the same molecular component if they can be seen as different states of the same molecule. As far as we know, the adoption of a notion of fairness in the context of biology is new. On the other hand, the notion of molecular component has been often implicitly used (for instance in the modelling of biological systems by means of automata), but now we provide new insight on this notion.

In this paper we report preliminary results obtained during the development of the modular verification framework. Modular verification requires either adopting a modular notation for pathway modelling or finding a way to decompose a pathway, simply expressed as a set of biochemical reactions, into a number of modules. The approach that we choose to follow is in between these two alternatives. Actually, we assume the pathway to be expressed as a set of reactions satisfying some modularisation requirements, and then we define a modularisation procedure that allows modules to be inferred from reactions. Modules will be molecular components, hence our modularisation procedure will allow us to consider a pathway not only as a set of reactions, but also as a set of entities interacting with each other (through reactions) and consequently changing state.

Once the molecular components of a pathway are identified, we can use them to decompose the verification of a global pathway property into the verification of a number of sub-properties related with groups of components. To this aim we define a *projection operation* that allows a model fragment describing the behaviour of a group of components to be obtained from a model describing the whole pathway. Such a projection operation is actually an abstraction function, since the behaviour of the group of components will be over-approximated (i.e. the model will include behaviours that are not present in the model of the whole pathway). By considering a suitable temporal logic for the specification of properties (namely $ACTL^-$, a fragment of the CTL logic consisting only of universally quantified formulae) we can prove that properties holding in model fragments obtained by projection also hold in the complete model of the pathway. Nothing can be said, instead, of properties that do not hold in the the model fragment.

In order to verify properties of complete pathway models or of model fragments it is possible to translate them into the input language of an existing model checking tool. Specifically, we use the NuSMV model checker [5], which is a well-established and efficient instrument.

We demonstrate the modular verification approach on the model of the EGF receptor-induced MAP kinase cascade by Schoeberl et al. [20] and we discuss how we plan to continue the development of the approach to improve its efficiency.

2 Modelling Biochemical Pathways with a Notion of Fairness

In biochemistry, metabolic pathways are networks of biochemical reactions occurring within a cell. The reactions are connected by their intermediates: products of one reaction are substrates for subsequent reactions. Reactions are influenced by catalysts and inhibitors, which are molecules (proteins) which can stimulate and block the occurrence of reactions, respectively. For the sake of simplicity we do not consider inhibitors in this paper, although they could be easily dealt with.

2.1 Syntax and semantics of the modelling notation

Given an infinite set of species S , let us assume biochemical reactions constituting a pathway to have the following form:



where r_j, p_j and c_j , for suitable values of j , are all in S . We have that r_j s are reactants, p_j s are products and c_j s are catalysts of the considered reaction. Given a reaction R we define $re(R) = \{r_1, \dots, r_n\}$, $pro(R) = \{p_1, \dots, p_{n'}\}$, and $cat(R) = \{c_1, \dots, c_m\}$. We denote the set of species involved in reaction R as $species(R) = re(R) \cup pro(R) \cup cat(R)$.

A pathway P is simply a set of reactions, $P = \{R_1, \dots, R_N\}$. Given a pathway P , we can infer the set of species involved in it as $species(P) = \bigcup_{R \in P} species(R)$.

The dynamics of a pathway can be described at several different levels of abstraction. The most precise level consists of a quantitative description in which quantities (or concentrations) of species are taken into account, as well as reaction rates in either a deterministic or a stochastic framework. At a more abstract level reaction rates can be ignored. Ultimately, also quantities of species can be ignored by considering only their presence (or absence) in the considered biochemical solution. The less abstract description level is obviously the most precise, but also the most difficult to treat with formal analysis techniques. The more abstract levels are more suitable for the application of formal analysis techniques and are often precise enough to provide some information on the role of the species and of the reactions involved in the pathway. We choose to adopt the most abstract description level, and hence we define a qualitative formal semantics of pathways in which species can only be either present or absent.

The dynamics of a pathway starts from an initial state representing a biochemical solution and is determined by the reactions. A reaction essentially causes the appearance of some new species in the biochemical solution. Actually, we choose to interpret the effect of a reaction depending on whether it is catalysed or not. In our interpretation a reaction without catalysts creates the products but does not consume the reactants. We choose this interpretation since non-catalysed reactions usually reach a steady-state of dynamic equilibrium in which both reactants and products are present in the biochemical solution. On the other hand, a reaction favoured by catalysts usually tends to be performed as long as there are reactants. Therefore, in our interpretation a reaction with catalysts creates the products and consumes the reactants. This choice implies that a reversible reaction in which both directions are catalysed, which frequently occurs in biological pathways, oscillates between two states. This is realistic in some cases (oscillatory behaviours) but not always. We leave a more detailed treatment of this aspect as future work.

Lastly, we assume that all of the catalysts are required to be present in order for the reaction to occur. Alternative combinations of catalysts that may enable the reaction should be modelled as different reactions having the same reactants and products.

Formally, given a pathway P and a set $\mathbf{s}_0 \subseteq species(P)$ representing species present in the initial state of the system, the *semantics* of P is given by the labelled transition system $(\mathcal{P}(species(P)), \mathbf{s}_0, \rightarrow_R)$, where $\mathcal{P}(species(P))$ is the powerset of the set of species of P , meaning that each state of the LTS is a configuration of the pathway indicating which species are present. We use the boldface notation, e.g. \mathbf{s} to denote states in the semantics, while a simple s denotes a species which means either a reactant, a product or a catalyst. Furthermore, $\rightarrow_R: \mathcal{P}(species(P)) \times P \times \mathcal{P}(species(P))$ is the least transition relation satisfying the following inference rules

$$\frac{re(R) \subseteq \mathbf{s}, pro(R) \not\subseteq \mathbf{s}, \emptyset \neq cat(R) \subseteq \mathbf{s}}{\mathbf{s} \xrightarrow{R} (\mathbf{s} \setminus re(R)) \cup pro(R)} \text{(cat)} \quad \frac{re(R) \subseteq \mathbf{s}, pro(R) \not\subseteq \mathbf{s}, cat(R) = \emptyset}{\mathbf{s} \xrightarrow{R} \mathbf{s} \cup pro(R)} \text{(no-cat)}.$$

Rules (cat) and (no-cat) formalise the dynamics of reactions in the presence and absence of catalysts, respectively. Both rules contain an assumption which states that the reaction does not occur if its products already exist. Note that thanks to this optimisation transitions that do not change the state of the system are excluded, which is convenient for the verification as the size of the transition system is smaller but the set of properties that hold stays the same. We denote the semantic function as *LTS*, i.e. $LTS: P \mapsto (\mathcal{P}(species(P)), \mathbf{s}_0, \rightarrow_R)$.

A path in $LTS(P)$ can be either a finite sequence $\mathbf{s}_0, R_0, \mathbf{s}_1, R_1, \dots, \mathbf{s}_n$ or an infinite sequence $\mathbf{s}_0, R_0, \mathbf{s}_1, R_1, \dots$ where for all i , \mathbf{s}_i is a state and R_i is a reaction and $\mathbf{s}_i \xrightarrow{R_i} \mathbf{s}_{i+1}$ is a transition in $LTS(P)$. The

path consisting only of the initial state \mathbf{s}_0 is denoted ε . In this paper we consider only maximal paths, corresponding to behaviours of the pathway in which as long as some reactions can occur, the pathway activity does not halt. It is worth noting that maximal paths are not necessarily infinite, as a state where no reactions can occur has no successor and a path leading to such a state is finite.

2.2 Fairness

In order to describe the behaviour of a pathway more accurately we consider a notion of fairness. We motivate it by considering a quantitative system consisting of four reactions $A \xrightarrow{k_1} B \{D\}$, $B \xrightarrow{k_2} A \{D\}$, $A \xrightarrow{k_3} C \{D\}$ and $C \xrightarrow{k_4} A \{D\}$, where k_1, k_2, k_3 and k_4 are the reaction rates. By performing the qualitative abstraction, we get a pathway containing reactions $R_1 = A \rightarrow B \{D\}$ and $R_2 = B \rightarrow A \{D\}$, $R_3 = A \rightarrow C \{D\}$ and $R_4 = C \rightarrow A \{D\}$, whose semantics as defined above includes behaviours such as the one where R_3 never occurs. Such a behaviour is a qualitative abstraction which is not correct, since the standard quantitative dynamics ruled by the law of mass action would imply that both R_1 and R_3 occur with a frequency proportional to their kinetic rates. Actually, in a stochastic setting both R_1 and R_3 would infinitely occur with probability 1. A correct qualitative abstraction of our system should therefore only include maximal paths in which both R_1 and R_3 occur infinitely many times.

A concept from concurrency theory that allows to specify the correct behaviour is fairness, which stipulates that reactions should compete in a fair manner. We consider the well-known notion of strong fairness [13], also called compassion, which requires that if a reaction is enabled (ready to occur) infinitely many times, then it will occur infinitely many times.

Technically, fairness is specified by a linear temporal logic (LTL) formula. LTL [17] is built up from formulae over a finite set of atomic propositions S , therefore a $s \in S$ is a LTL formula and if f and g are LTL formulae, then so are $\neg f$, $f \vee g$, $X g$ and $f U g$ where X is read as next and U as until. Additional logical operators can be defined, $true = s \vee \neg s$, $false = \neg true$, $f \wedge g = \neg(\neg f \vee \neg g)$ and $f \rightarrow g = \neg f \vee g$, additional temporal operators eventually $F g = true U g$ and globally $G g = \neg F \neg g$. A LTL formula can be satisfied by a maximal path π in $LTS(P)$ as described by the satisfaction relation \models_{LTL} : $\pi \models_{LTL} s$ if $\pi = \mathbf{s}, R, \pi'$, $\pi \models_{LTL} \neg g$ if not $\pi \models_{LTL} g$, $\pi \models_{LTL} f \vee g$ if $\pi \models_{LTL} f$ or $\pi \models_{LTL} g$, $\pi \models_{LTL} X g$ if $\pi = \mathbf{s}, R, \pi'$ and $\pi' \models_{LTL} g$, and finally $\pi \models_{LTL} f U g$ if there is an $i \geq 0$ such that $\pi = \mathbf{s}_0, R_0, \mathbf{s}_1, R_1, \dots$ and $\mathbf{s}_i, R_i, \pi_i \models_{LTL} f$ and for all $0 \leq k < i$, $s_k, R_k, \pi_k \models_{LTL} g$.

Fairness is expressed by formula Φ , and a maximal path π is fair iff it satisfies Φ , i.e. $\pi \models_{LTL} \Phi$. We have

$$\Phi \iff \bigwedge_{R \in P} (GF \text{ enabled}(R) \rightarrow GF \text{ occurred}(R))$$

where $\text{enabled}(R) \iff ((\bigwedge_{r \in re(R)} r) \wedge (\bigvee_{p \in pro(R)} \neg p) \wedge (\bigwedge_{c \in cat(R)} c))$ and the satisfaction of proposition $\text{occurred}(R)$ is defined as $\pi' \models_{LTL} \text{occurred}(R)$ iff there is a path π such that $\pi = \mathbf{s}, R, \pi'$.

It should be noted that our fairness neither requires all reactions to occur infinitely nor requires fair paths to be infinite.

2.3 Modelling the EGF receptor-induced MAP kinase cascade

We apply our modular verification approach to a well-established computational model of the EGF signalling pathway. We consider the model of the MAP kinase cascade activated by surface and internalised EGF receptors, proposed by Schoeberl et al. in [20]. This model includes a detailed description of the reactions that involve active EGF receptors and several effectors named GAP, ShC, SOS, Grb2, RasGDP/GTP and Raf. Moreover, the model describes the activity of internalised receptors, namely

receptors that are no longer located on the cell membrane, but on a vesicle obtained by endocytosis and floating in the cytoplasm. Such internalised receptors continue to interact with effectors and to contribute to the pathway functioning, but actually the pathway can be seen as composed by two almost identical branches: the first consisting of the reactions stimulated by receptors on the cell membrane and the second consisting of reactions stimulated by internalised receptors.

A diagram representing all of the reactions of the pathway considered in the model is shown in Figure 1. In the figure, species are identified by a short name, but also by a number (in black) in the interval $[1 - 60]$. Arrows represent reactions, which are also associated with an identifier (in grey) in the interval $[v0 - v101]$. Note that the two branches of the pathway are partially combined in the figure. In particular, the representation of most of the species is combined with the representation of its internalised counterpart. In such cases, the number between brackets denotes the number identifying the internalised species. The same holds for reactions: in many cases an arrow denotes both a reaction stimulated by receptors in the cell membrane and the corresponding reaction stimulated by internalised receptors.

The set of reactions constituting the pathway can be trivially reconstructed from the diagram in Figure 1. The only non-trivial aspect is related with the presence in the diagram of some reactions in which one reactant is actually acting as a catalyst. For instance, this happens in the case of the reactions involving Raf^* and MEK, in which Raf^* initially binds MEK and then releases it phosphorylated. We describe these two reactions in the diagram with the following single catalysed reaction:



Other species acting as catalysts are MEK-PP, Phosphatase1, Phosphatase2 and Phosphatase3. By applying the same transformation also to the reactions they are involved in we obtain a pathway constituted by 80 reactions. We call this pathway P_{EGF} .

We recall that fairness requires that a reaction that is infinitely often enabled is also infinitely often performed. This prevents starvation situations to happen among reactions. In the case of P_{EGF} the two branches of the pathway include reactions that could be involved in infinite loops (e.g. the reactions involving MEK and ERK). This means that the semantics of the pathway includes behaviours in which only one branch executes forever even if the other is constantly enabled. Such unrealistic behaviours are excluded by the adoption of fairness.

3 Identification of Molecular Components

In this section we argue, that under conditions often found in practice, a pathway can be decomposed into components, which, as it will be shown in the following sections, can be used for modular verification.

3.1 Assumptions

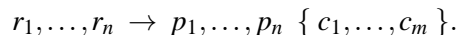
Intuitively, a species can be seen as a part of a “state” or “configuration” of a more general system component, and a reaction can be seen as a synchronised state change of a set of such system components. In order to view a pathway through this optics, it is convenient to assume that the pathway has equal number of reactants and products (which is not the case in general). Moreover, we assume a positional correspondence between the reactants and the products, in particular we assume that product p_j is the result of the transformation of reactant r_j by the reaction. In our experience, it is usually possible to translate a reaction of a pathway into such a “normal form”. Reactions of cellular pathways very often represent bindings (and unbindings) of well-defined macromolecules, such as proteins and genes, to

form (or to break) complexes either with other macromolecules or with small molecules such as ions and nutrients. Also conformational changes are common, in which a protein (or a complex constituted by a few proteins) changes its own “state”. If we consider a complex not as a single entity, but as a combination of macromolecules we have that all of the mentioned kinds of reaction do not change the number of (macro)molecules in the system. Hence, it should be possible to model them with the reactions in the form we assume here. For the moment we leave the translation of reaction into the assumed form to the modeller.

In Section 2.1 we have introduced the syntax of the modelling formalism of biochemical pathways, in which a reaction is allowed to have a different number of reactants and products.

3.2 Components identification

Let us, thus, assume that the pathway P consists of reactions in the following form:



Such a form enables us to identify a set of components I that constitute the pathway. Now we present an algorithm that given a pathway P returns the set of components I along with the partition of the set of species belonging to respective components.

We illustrate the intuitive idea on an example. Each reaction can be seen as a synchronisation of components. For example reaction $r_1, r_2 \rightarrow p_1, p_2 \{ c \}$ can be interpreted as a synchronisation of three components: one that changes its state from a state where r_1 holds into a state where p_1 is present and r_1 is not, another component that changes its state from a state where r_2 holds to a state where p_2 is present and r_2 is not, and a component which participates passively and stays in a state where c is present. Since we suppose that only one reaction takes place at a time in the whole system, the states of all the components do not change other than those involved in the reaction in the way we described. From the example we can see that species r_1 and p_1 belong to the same component. Similarly r_2 belongs to the component that contains p_2 , while c is from a separate one.

The algorithm follows. We start by assuming that each species belongs to a different component and we refine this assumption by iterating over the reactions constituting P . The result of the algorithm is a mapping map assigning each species to its component.

Algorithm 1 Algorithm to partition species into different components

Let $map : S \mapsto I$ be an injective mapping

for all R in P **do**

for all r_j in $re(R)$ **do**

$$map := \begin{cases} p \mapsto map(r_j) & \forall p \in \{s \in S \mid map(s) = map(r_j)\} \\ s \mapsto map(s) & \text{otherwise} \end{cases}$$

end for

end for

return map

The algorithm updates the mapping by unifying the elements assigned to reactants and products in the same position in a reaction, and this is done for all reactions in the pathway.

The set of components $comp(P) = I$ of pathway P is the image of mapping map . Components of a reaction R denoted, using the same notation, as $comp(R)$ are defined as $comp(\{R\})$.

3.3 Initial state

We adopt a semi-automatic heuristic procedure to find an initial state of the pathway. The idea is the following: for each species s in $species(P)$, if there is no reaction creating it (i.e. if $s \notin \bigcup_{R \in P} PRO(R)$) then in the initial state s is present. This means that species that cannot be produced are assumed to be present in the initial state. Otherwise their presence in the model would not be meaningful. Subsequently, we resort again to the partitioning of species according to components to find other species to be inserted. In particular, we find those components containing no species present in the previous phase. These components must contain loops, hence we choose manually some of their species to insert. All other species are assumed absent.

3.4 Visualisation of component interaction

A component interaction graph can be drawn which visualises the components of a pathway and their interactions. It is a directed graph in which vertices are system components (elements of I) and edges connect components that are involved together in a reaction. If two components are both involved as reactants (and consequently products), the edge connecting them will not be oriented (displayed as bidirectional). If one of the two is involved as reactant and the other as catalyst, then the edge will start from the vertex representing the latter to the vertex representing the former. There is no edge between vertices representing components involved in the same reactions only as catalysts.

3.5 The model

The model P_{EGF} is made up of 143 species and 80 reactions. It is in the correct form assumed in Section 3.1 and no preprocessing is needed. After performing the components identification procedure, 14 components are identified. On Figure 2 we can see the component interaction graph of P_{EGF} . Each node of the graph is labelled by the intuitive name of the component that we have chosen.

Visually, we can do some simple observations on the component interaction graph. We can identify enzymes like Phosphatase1, Phosphatase2 and Phosphatase3. We can see the first part of the pathway corresponding to the EGF receptor and its interaction with effectors, and its connection to the MAP kinase cascade through the component RasGDP.

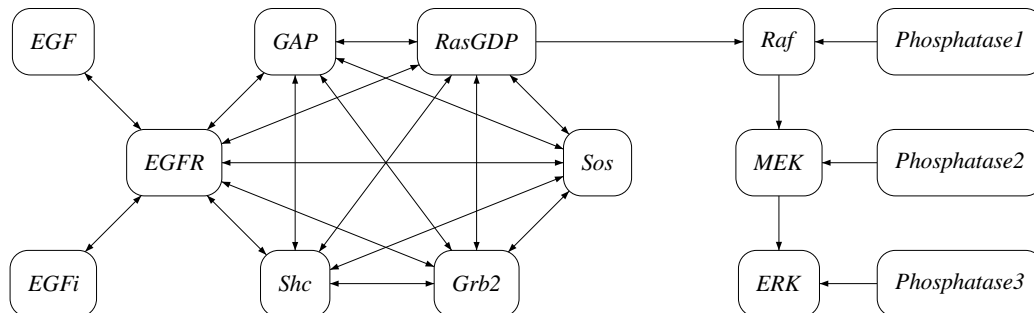


Figure 2: Component interaction graph of P_{EGF}

4 Modular Verification

In this section we define a modular verification technique for pathway models. We proceed by defining the projection of a pathway with help of the identified components. Such a projection can be seen as an abstraction, giving rise to abstract pathways. We prove that a successful verification of a property in the abstraction implies its truth in the original model.

4.1 Abstract pathways: syntax, semantics and fairness

We are interested in analysing only a portion of the entire pathway, in particular a portion induced by only a subset of all components. Let $I = \text{comp}(P)$ and $J \subseteq I$, we define the projection of a pathway P onto J as an abstract pathway $P \upharpoonright J$.

We will need an extension of function *species*, abusing the notation, which operates on a component set: $\text{species}(J) = \{s \in \text{species}(R) \mid R \text{ s.t. } \text{comp}(R) \cap J \neq \emptyset\}$.

Definition 1. An abstract pathway $P \upharpoonright J$ is a pair (PR, AR) , where

- $PR = \{R \in P \mid \text{comp}(R) \subseteq J\}$
- $AR = \bigcup_{\substack{R \in P, \text{comp}(R) \cap J \neq \emptyset \\ \text{comp}(R) \cap (I \setminus J) \neq \emptyset}} \{re(R) \upharpoonright J \rightarrow pro(R) \upharpoonright J \{cat(R) \upharpoonright J\}, re(R) \upharpoonright J \rightarrow re(R) \upharpoonright J \{cat(R) \upharpoonright J\}\}$

where the projection of a set of species $u \subseteq S$ is defined as $u \upharpoonright J = \{s \in u \mid s \in \text{species}(J)\}$.

An abstract pathway consists of two sets of reactions: PR contains reactions which influence only components inside J , and AR contains projections of reactions that influence components both inside and outside J . Reactions in PR are exactly as in P , since all of the species involved in such reactions are considered in the abstract pathway. On the other hand, reactions in AR are obtained from the reactions in P which involve species both from some components in J and from some components not in J . For each of such reactions in P we have two reactions in AP : one describing the situation in which species not in $\text{species}(J)$ are assumed to be configured such that the reaction can occur, and the other describing the opposite situation. In the first case the reaction in AR produces some products; in the second case the reaction in AR performs a self-loop (i.e. it does not change the state).

The abstract pathway semantics is defined as $LTS_\alpha : P \upharpoonright J \mapsto LTS(PR \cup AR)$, that is by using the standard semantics LTS on the projected reactions in both PR and AR .

What changes with respect to the pathway model is the definition of fairness. In fact, in this case fairness constraints can be applied only to reactions in PR since reactions in AR consist of pairs of reactions always applicable at the same time and in which it is reasonable to assume that one of the two is always preferred (describing the situation in which the corresponding reaction in R is always enabled or disabled). We define the notion of abstract fairness as

$$\Phi_\alpha \iff \bigwedge_{R \in PR} (GF \text{ enabled}(R) \rightarrow GF \text{ occurred}(R)).$$

Here the compassion is only required for reactions from PR .

Note that a pathway is a special case of abstract pathway, since the semantics of P is equivalent (isomorphic) to that of $P \upharpoonright I$ and in this case also $\Phi \equiv \Phi_\alpha$ holds.

4.2 Logic for specifying properties

Properties of pathways are specified in temporal logic with species as atomic propositions.

The logic we consider is a fragment of the Computation Tree Logic CTL. Following Attie and Emerson [1], we assume the logic ACTL^- for specification of properties. ACTL is the “universal fragment” of CTL which results from CTL by restricting negation to propositions and eliminating the existential path quantifier and ACTL^- is ACTL without the AX modality.

Definition 2. *The syntax of ACTL^- is defined inductively as follows:*

- *The constants true and false are formulae. s and $\neg s$ are formulae for any atomic proposition s , where the set of atomic propositions AP are the set of all species S .*
- *If f, g are formulae, then so are $f \wedge g$ and $f \vee g$.*
- *If f, g are formulae, then so are $A[f U g]$ and $A[f U_w g]$.*

We define the logic ACTL_J^- to be ACTL^- where the atomic propositions are drawn from $AP_J = \text{species}(J)$. Abbreviations in ACTL^- : $AFf \equiv A[\text{true} U f]$ and $AGf \equiv A[f U_w \text{false}]$.

Properties expressible by ACTL^- formulae represent a significant class of properties investigated in the systems biology literature as identified in [16], such as properties concerning exclusion (“*It is not possible for a state \mathbf{s} to occur*”), necessary consequence (“*If a state \mathbf{s}_1 occurs, then it is necessarily followed by a state \mathbf{s}_2* ”), and necessary persistence (“*A state \mathbf{s} must persist indefinitely*”).

On the other hand, properties as occurrence, possible consequence, sequence and possible persistence are of inherently existential nature, and are not expressible in ACTL^- .

Definition of the semantics of ACTL^- formulae on labelled transition system $LTS(P)$ follows. Note that only fair maximal paths are considered.

Definition 3. *Semantics of ACTL^- . We define $LTS(P), \mathbf{s} \models_{\Phi} f$ (resp. $LTS(P), \pi \models_{\Phi} f$) meaning that f is true in structure $LTS(P)$ at state \mathbf{s} (resp. fair maximal path π). We define \models_{Φ} inductively:*

- *$LTS(P), \mathbf{s} \models_{\Phi} \text{true}$. $LTS(P), \mathbf{s} \not\models_{\Phi} \text{false}$. $LTS(P), \mathbf{s} \models_{\Phi} s$ iff $\mathbf{s}(s) = tt$. $LTS(P), \mathbf{s} \models_{\Phi} \neg s$ iff $\mathbf{s}(s) = ff$.*
- *$LTS(P), \mathbf{s} \models_{\Phi} f \wedge g$ iff $LTS(P), \mathbf{s} \models_{\Phi} f$ and $LTS(P), \mathbf{s} \models_{\Phi} g$.
 $LTS(P), \mathbf{s} \models_{\Phi} f \vee g$ iff $LTS(P), \mathbf{s} \models_{\Phi} f$ or $LTS(P), \mathbf{s} \models_{\Phi} g$.*
- *$LTS(P), \mathbf{s} \models_{\Phi} Af$ iff for every fair maximal path $\pi = (\mathbf{s}, R, \dots)$ in $LTS(P)$: $LTS(P), \pi \models_{\Phi} f$.*
- *$LTS(P), \pi \models_{\Phi} f$ iff $LTS(P), \mathbf{s} \models_{\Phi} f$, where \mathbf{s} is the first state of π*
- *$LTS(P), \pi \models_{\Phi} f \wedge g$ iff $LTS(P), \pi \models_{\Phi} f$ and $LTS(P), \pi \models_{\Phi} g$.
 $LTS(P), \pi \models_{\Phi} f \vee g$ iff $LTS(P), \pi \models_{\Phi} f$ or $LTS(P), \pi \models_{\Phi} g$.*
- *$LTS(P), \pi \models_{\Phi} f U g$ iff $\pi = (\mathbf{s}_0, R_0, \mathbf{s}_1, R_1, \dots)$ and there is $m \in \mathbb{N}$ such that $LTS(P), \mathbf{s}_m \models_{\Phi} g$ and for all $m' < m$: $LTS(P), \mathbf{s}_{m'} \models_{\Phi} f$.*
- *$LTS(P), \pi \models_{\Phi} f U_w g$ iff $\pi = (\mathbf{s}_0, R_0, \mathbf{s}_1, R_1, \dots)$ and for all $m \in \mathbb{N}$, if $LTS(P), \mathbf{s}_{m'} \not\models_{\Phi} g$ for all $m' < m$ then $LTS(P), \mathbf{s}_m \models_{\Phi} f$.*

We assume $\models_{\Phi_{\alpha}}$ to be defined as \models_{Φ} , but with abstract fairness Φ_{α} replacing Φ .

4.3 Modular verification theorems

Now we prove that in order to verify an ACTL_J^- property for a pathway P , it is enough to verify the same property in the abstract semantics of the abstract pathway $P \upharpoonright J$. The principle behind property preservation is that each path in the semantics of the modelled pathway must have a corresponding abstract path in the abstract semantics of a model obtained by projection. This, combined with the fact that ACTL^- properties are universally quantified (namely describe properties that have to be satisfied by all paths) ensure that if an ACTL^- property holds in the abstract semantics of the projection, then it will also hold in the semantics of the original model. In fact, for the components considered in a projection the semantics of the original model will contain essentially a subset of the paths of the projected model.

ACTL^- properties are universally quantified, namely they they deal with all paths starting from a given initial state, and the fact that all original paths (more precisely their projections) are included amongst the paths of the projection. Thus if one proves that the property holds in the projection for all paths it will hold for all paths also in the original system.

First we define the path projection, which from a path in semantics of a pathway with the set of components I removes transitions made by components outside of portion $J \subseteq I$ and restricts the rest of transitions onto J .

Definition 4.

$$\pi \upharpoonright J = \begin{cases} \varepsilon & \text{if } \pi = \varepsilon \\ \pi' \upharpoonright J & \text{if } \pi = \mathbf{s}, R, \pi' \text{ and } \text{comp}(R) \cap J = \emptyset \\ \mathbf{s} \upharpoonright J, R, \pi' \upharpoonright J & \text{if } \pi = \mathbf{s}, R, \pi' \text{ and } \text{comp}(R) \cap J \neq \emptyset \end{cases}$$

Follows the infinite path projection which ensures that the resulting traces are infinite. In case of a finite original trace it adds an infinite looping in the final state.

Definition 5. Given $\pi \upharpoonright J$ with initial state \mathbf{s}_0 we define $\pi^{\infty} \upharpoonright J = \pi \upharpoonright J$ if $\pi \upharpoonright J$ is infinite, otherwise if $\pi \upharpoonright J = \mathbf{s}_0, R_0, \dots, \mathbf{s}_{n-1}, R_{n-1}, \mathbf{s}_n$ we define $\pi^{\infty} \upharpoonright J = \pi \upharpoonright J, (*, \mathbf{s}_n)^{\infty}$, where $(R, \mathbf{s})^{\infty} = R, \mathbf{s}, (R, \mathbf{s})^{\infty}$. We denote $\varepsilon^{\infty} \upharpoonright J = \mathbf{s}_0, (*, \mathbf{s}_0)^{\infty}$ as ε^{∞} .

Now we are in the position to present the crucial result, which states that a fair maximal path in the semantics of a pathway is either projected or infinitely projected into an abstractly fair maximal path in the abstract semantics of an abstract pathway. It is split in two lemmas, where the first one states that at least one of the projections is present as a maximal path in the abstract semantics. The second lemma proves the abstract fairness of the projections.

Lemma 1. $\pi \in \text{LTS}(P)$ with $\pi \models_{\text{LTL}} \Phi$ implies $(\pi \upharpoonright J \in \text{LTS}_{\alpha}(P \upharpoonright J) \text{ or } \pi^{\infty} \upharpoonright J \in \text{LTS}_{\alpha}(P \upharpoonright J))$.

Proof. Let us assume that π is a finite path, then we can prove that either $(\pi \upharpoonright J \in \text{LTS}_{\alpha}(P \upharpoonright J) \text{ or } \pi^{\infty} \upharpoonright J \in \text{LTS}_{\alpha}(P \upharpoonright J))$ by induction on the path length.

Case $\pi = \varepsilon$. We have that $\pi \upharpoonright J = \pi^{\infty} \upharpoonright J = \varepsilon$. Since $\pi = \varepsilon$, the initial state \mathbf{s}_0 is such that no reaction is enabled in \mathbf{s}_0 . By the definition of abstract pathway semantics we know that the initial state of $\text{LTS}_{\alpha}(P \upharpoonright J)$ is $\mathbf{s}_0 \upharpoonright J$.

By definition of abstract pathway we have that $P \upharpoonright J = (PR, AR)$.

- If $AR = \emptyset$, then since $PR \subseteq P$ and there is no reaction in P that is enabled in \mathbf{s}_0 , there is also no reaction in PR that is enabled in $\mathbf{s}_0 \upharpoonright J$, and hence $\varepsilon \in \text{LTS}_{\alpha}(P \upharpoonright J)$.
- If $AR \neq \emptyset$, then for each $R \in P \setminus PR$ we have in AR two reactions R_1, R_2 as follows

$$\begin{aligned} R_1 &= \text{re}(R) \upharpoonright J \rightarrow \text{pro}(R) \upharpoonright J \{ \text{cat}(R) \upharpoonright J \} \\ R_2 &= \text{re}(R) \upharpoonright J \rightarrow \text{re}(R) \upharpoonright J \{ \text{cat}(R) \upharpoonright J \} \end{aligned}$$

Note that if R_1 is enabled, then also R_2 is enabled.

As before, since there is no reaction in P that is enabled in \mathbf{s}_0 , there is also no reaction in PR that is enabled in $\mathbf{s}_0 \upharpoonright J$. If there is some R_2 enabled in $\mathbf{s}_0 \upharpoonright J$, then $(R_2, \mathbf{s}_0 \upharpoonright J)^\infty \in LTS_\alpha(P \upharpoonright J)$, that is $\varepsilon \upharpoonright^\infty J \in LTS_\alpha(P \upharpoonright J)$. On the other hand, if there is no R_2 enabled in $\mathbf{s}_0 \upharpoonright J$, then there is also no R_1 enabled in the same state, hence $\varepsilon \in LTS_\alpha(P \upharpoonright J)$, that is $\varepsilon \upharpoonright J \in LTS_\alpha(P \upharpoonright J)$.

Case $\pi = \mathbf{s}, R, \pi'$. We distinguish two subcases:

- $comp(R) \cap J = \emptyset$: \mathbf{s}' is the initial state of π' , so by induction hypothesis $\mathbf{s}' \upharpoonright J$ is the initial state of either $\pi' \upharpoonright J$ or $\pi' \upharpoonright^\infty J$. Moreover, state, by definition of path projection, $\mathbf{s} \upharpoonright J = \mathbf{s}' \upharpoonright J$ then $\mathbf{s} \upharpoonright J$ is the initial state of either $\pi' \upharpoonright J$ or $\pi' \upharpoonright^\infty J$, which means that either $\pi' \upharpoonright J = \pi \upharpoonright J \in LTS_\alpha(P \upharpoonright J)$ or $\pi' \upharpoonright^\infty J = \pi \upharpoonright^\infty J \in LTS_\alpha(P \upharpoonright J)$.
- $comp(R) \cap J \neq \emptyset$: \mathbf{s}' is the initial state of π' , so by induction hypothesis $\mathbf{s}' \upharpoonright J$ is the initial state of either $\pi' \upharpoonright J$ or $\pi' \upharpoonright^\infty J$. Moreover, in $P \upharpoonright J$ there are R_1 and R_2 as above. Since R is enabled in \mathbf{s} , R_1 is enabled in $\mathbf{s} \upharpoonright J$. Therefore either $\mathbf{s} \upharpoonright J, R_1, \pi' \upharpoonright J = \pi \upharpoonright J \in LTS_\alpha(P \upharpoonright J)$ or $\mathbf{s} \upharpoonright J, R_1, \pi' \upharpoonright^\infty J = \pi \upharpoonright^\infty J \in LTS_\alpha(P \upharpoonright J)$.

Summarising, if π is a finite path, in all the possible cases we have that either $\pi \upharpoonright J \in LTS_\alpha(P \upharpoonright J)$ or $\pi \upharpoonright^\infty J \in LTS_\alpha(P \upharpoonright J)$.

Let us consider now the case in which π is infinite. We have two subcases:

- $\pi \upharpoonright J$ is finite: By the fairness assumption $\pi \models_{LTL} \Phi$ it follows that if the path $\pi \upharpoonright J$ is finite, then its final state is such that no reaction in PR is enabled. In fact, as a consequence of fairness, the only case in which $\pi \upharpoonright J$ can be finite is when all of the reactions that are infinitely often enabled in π are performed only by components that are not in J . Hence, $\pi \upharpoonright J$ is finite since it contains reactions that are not enabled infinitely often. Let $\pi = \pi_1, \pi_2$, where π_1 is the shortest (finite) prefix such that π_2 contains only moves of reactions enabled infinitely often. It is easy to see that $\pi \upharpoonright J = \pi_1 \upharpoonright J$ and similarly as before we can prove by induction on the length of π_1 that either $\pi_1 \upharpoonright J \in LTS_\alpha(P \upharpoonright J)$ or $\pi_1 \upharpoonright^\infty J \in LTS_\alpha(P \upharpoonright J)$. The fact that no reaction is enabled in the final state of $\pi \upharpoonright J$ ensures that $\pi \upharpoonright J$ is a maximal path of $LTS_\alpha(P \upharpoonright J)$.
- $\pi \upharpoonright J$ is infinite: This case can be proved by showing that the relation between π and $\pi \upharpoonright J$ described in the inductive case of the proof for a finite π is indeed an invariant property in the case of an infinite path π .

□

Now we state and prove the second lemma.

Lemma 2. $\pi \models_{LTL} \Phi$ implies $\pi \upharpoonright J \models_{LTL} \Phi_\alpha$ and $\pi \upharpoonright^\infty J \models_{LTL} \Phi_\alpha$.

Proof. Suppose that $\pi \models_{LTL} \Phi$, i.e. $\pi \models_{LTL} \bigwedge_{R \in P} (GF \text{ enabled}(R) \rightarrow GF \text{ occurred}(R))$. Let $J \subseteq I$ and $P \upharpoonright J = (AR, PR)$. We want to prove that $\pi \upharpoonright J \models_{LTL} \Phi_\alpha$, that is $\pi \upharpoonright J \models_{LTL} \bigwedge_{R \in PR} (GF \text{ enabled}(R) \rightarrow GF \text{ occurred}(R))$. This holds because of two facts (1) and (2) that can be easily checked: for any reaction R from PR

$$\bullet \mathbf{s} \models_{LTL} \text{enabled}(R) \text{ implies } \mathbf{s} \upharpoonright J \models_{LTL} \text{enabled}(R) \quad (1)$$

$$\bullet \mathbf{s} \models_{LTL} \text{occurred}(R) \text{ implies } \mathbf{s} \upharpoonright J \models_{LTL} \text{occurred}(R) \quad (2)$$

Analogously $\pi \upharpoonright^\infty J \models_{LTL} \bigwedge_{R \in PR} (GF \text{ enabled}(R) \rightarrow GF \text{ occurred}(R))$. □

Finally, the property preservation theorem states that a successful verification of a property in the abstraction implies its truth in the original model.

Theorem 3. For a pathway P and a $J \subseteq I$ where I is the component set of P and f an ACTL⁻ formula we have $LTS_\alpha(P \upharpoonright J) \models_{\Phi_\alpha} f$ implies $LTS(P) \models_\Phi f$.

Proof. By induction on the structure of f (for all \mathbf{s}).

$f = s$. By definition of state projection and the fact that AP_R s are pairwise disjoint, for all atomic propositions s from AP_J we get that $LTS_\alpha(P \upharpoonright J), \mathbf{s} \upharpoonright J \models_{\Phi_\alpha} s$ iff $LTS(P), \mathbf{s} \models_\Phi s$. Analogously for $f = \neg s$.

$f = g \wedge h$. From the assumption $LTS_\alpha(P \upharpoonright J), \mathbf{s} \upharpoonright J \models_{\Phi_\alpha} g \wedge h$ by CTL semantics, $LTS_\alpha(P \upharpoonright J), \mathbf{s} \upharpoonright J \models_{\Phi_\alpha} g$ and $LTS_\alpha(P \upharpoonright J), \mathbf{s} \upharpoonright J \models_{\Phi_\alpha} h$. By induction hypothesis $LTS(P), \mathbf{s} \models_\Phi g$ and $LTS(P), \mathbf{s} \models_\Phi h$. Hence, $LTS(P), \mathbf{s} \models_\Phi g \wedge h$. Case $f = g \vee h$ is proved analogously.

$f = A[g U_w h]$. Let π be an arbitrary fair maximal path starting in \mathbf{s} . We establish $LTS(P), \pi \models_\Phi [g U_w h]$. By Lemma 1 at least one of $\pi \upharpoonright J$ or $\pi \upharpoonright^\infty J$ is a path in $LTS_\alpha(P \upharpoonright J)$, and by Lemma 2 both are abstractly fair.

Let us suppose first that $\pi \upharpoonright J$ is the abstractly fair maximal path in $LTS_\alpha(P \upharpoonright J)$. Hence by the assumption $LTS_\alpha(P \upharpoonright J), \pi \upharpoonright J \models_{\Phi_\alpha} [g U_w h]$. There are two cases:

1. $LTS_\alpha(P \upharpoonright J), \pi \upharpoonright J \models_{\Phi_\alpha} G g$. Let t be any state along π . By CTL semantics $LTS_\alpha(P \upharpoonright J), t \upharpoonright J \models_{\Phi_\alpha} g$. By induction hypothesis we have $LTS(P), t \models_\Phi g$. Since t was an arbitrary state of π , we get $LTS(P), \pi \models_\Phi G g$ and thus $LTS(P), \pi \models_\Phi g U_w h$.
2. $LTS_\alpha(P \upharpoonright J), \pi \upharpoonright J \models_{\Phi_\alpha} [g U h]$. Let $\mathbf{s}_j^{m''}$ be the first state along $\pi \upharpoonright J$ that satisfies h . Then there is at least one state $\mathbf{s}_j^{m''}$ along π such that $\mathbf{s}_j^{m''} \upharpoonright J = \mathbf{s}_j^{m''}$. Let $\mathbf{s}_j^{m'}$ be first such state. By induction hypothesis $LTS(P), \mathbf{s}_j^{m'} \models_\Phi h$. From the definition of path projection any \mathbf{s}_j^m with $m < m'$ projects to $\mathbf{s}_j^m \upharpoonright J$ that is before $\mathbf{s}_j^{m'}$ in $\pi \upharpoonright J$. By the assumption $LTS_\alpha(P \upharpoonright J), \mathbf{s}_j^m \upharpoonright J \models_{\Phi_\alpha} g$, hence by induction hypothesis $LTS(P), \mathbf{s}_j^m \models_\Phi g$. By CTL semantics we get $LTS(P), \pi \models_\Phi g U h$.

In both cases we showed $LTS(P), \pi \models_\Phi g U_w h$. Since π was arbitrary fair maximal path starting in \mathbf{s} , we conclude $LTS(P), \mathbf{s} \models_\Phi A[g U_w h]$.

The reasoning for the case in which the abstractly fair maximal path in $LTS_\alpha(P \upharpoonright J)$ is $\pi \upharpoonright^\infty J$ is analogous to the considered case.

$f = A[g U h]$. Let π be an arbitrary fair maximal path starting in \mathbf{s} . By Lemmas 1 and 2 we have that $\pi \upharpoonright J$ or $\pi \upharpoonright^\infty J$ is a fair maximal path in $LTS_\alpha(P \upharpoonright J)$ and by the assumption $LTS_\alpha(P \upharpoonright J), \pi \upharpoonright J \models_{\Phi_\alpha} [g U h]$ or $LTS_\alpha(P \upharpoonright J), \pi \upharpoonright^\infty J \models_{\Phi_\alpha} [g U h]$. By the above case we get $LTS(P), \mathbf{s} \models_\Phi A[g U h]$. \square

5 Experiments

In this section we exploit the NuSMV model checker to perform some experiments on the model of the EGF pathway. NuSMV includes model checking algorithms that allow fairness constraints to be taken into account. We rely on such algorithms to manage fairness constraints introduced in this paper. Moreover, in order to carry out the projection and encode the resulting abstract pathway in the NuSMV format we have developed a tool (available upon request).

The first experiment is aimed at showing how modular verification could be applied to verify a global property of the pathway, namely that the final product of the pathway is always produced. This can be done in a modular way by proving sub-properties in three different model fragments obtained by projection.

Subsequently, a number of experiments are performed with the aim of showing how the molecular components we identified in the pathway can be used to better understand the pathway dynamics. In particular, we check whether there are some molecular components that are not really necessary to obtain the final product of the pathway. This will be done by applying model checking on models in which

molecular components are selectively disabled by setting their initial states to false. Also in this case the modular verification approach is adopted.

In this case study modular verification allows properties to be verified faster than on the complete model. However, modular verification is still not significantly more efficient than verification on the complete model. This is due to the projection operation we are considering at the moment, which is rather rough. In Section 6 we discuss why this modular verification is a promising approach for the analysis of pathways, and how we plan to improve the approach to make it substantially more efficient.

To run the experiment we used NuSMV 2.5.4 on a workstation equipped with an Intel i5 CPU 2.80 Ghz, with 8GB RAM and running Ubuntu GNU/Linux. In order to make verification faster NuSMV was executed in batch mode by enabling dynamic reordering of BDD variables and by disabling the generation of counterexamples.

5.1 Modular verification of a global property

The final product of the MAP kinase cascade activated by surface and internalised EGF receptors is species ERK-PP. Since surface and internalised receptors activate two different branches of the pathway, we denote by ERK-PP the product of the branch activated by the surface receptors and by ERK-PPi the product of the branch activated by the internalised receptors.

The property to be verified is

$$AF(ERK-PP \vee ERK-PPi) \quad (1)$$

The property holds in the complete model and its verification required 260 seconds. By looking at the diagram in Figure 1 we noticed that the pathway could be partitioned in three parts, with two species acting as “gates”. These two species are (EGF-EGFR*)2-GAP and Raf*. Hence, we decided to try to apply modular verification by splitting property 1 into the following three sub-properties:

$$AF((EGF-EGFR^*)2-GAP) \quad (2)$$

$$AG((EGF-EGFR^*)2-GAP \rightarrow (AF \text{ Raf}^*)) \quad (3)$$

$$AG(\text{Raf}^* \rightarrow AF(ERK-PP \vee ERK-PPi)) \quad (4)$$

Property (2) states that in all paths of the system a state in which species (EGF-EGFR*)2-GAP is present is eventually reached. Property (3) states that whenever a state is reached in which species (EGF-EGFR*)2-GAP is present, then a state in which Raf* is present is eventually reached. Finally, property (4) states that whenever a state is reached in which in which species Raf* is present, then a state in which either ERK-PP or ERK-PPi is present is eventually reached. It is easy to see that the conjunction of (2), (3) and (4) implies (1).

We considered three projections of the complete model to be used to verify properties (2), (3) and (4), respectively. In particular, from the component interaction graph of the model (shown in Figure 2) we extracted the following subsets to be used for projections:

- in order to verify (2) we considered the subset J_1 consisting of components *EGF*, *EGFi*, *EGFR* and *GAP*;
- in order to verify (3) we considered the subset J_2 consisting of components *EGFR*, *GAP*, *Shc*, *RasGDP*, *Grb2* and *Sos*;
- in order to verify (4) we considered the subset J_3 consisting of components *RasGDP*, *Raf*, *MEK*, *ERK*, *Phosphatase1*, *Phosphatase2* and *Phosphatase3*.

| Disabled component | Verification complete model | | | Modular Verification | | | |
|--------------------|-----------------------------|--------|------|----------------------|----------------------|------------------|------------|
| | Property | Result | Time | Property | Result | Time | Total time |
| none | (1) | true | 260s | (2) (3) (4) | true true true | 3s 213s 1s | 217s |
| <i>GAP</i> | (1) | false | 252s | (2),(5) | false,true | 2s | 2s |
| <i>Sos</i> | (1) | false | 253s | (2) (3),(6) | true false,true | 3s 207s | 210s |
| <i>Shc</i> | (1) | true | 252s | (2) (3) (4) | true true true | 3s 208s 1s | 212s |
| <i>Grb2</i> | (1) | false | 253s | (2) (3),(6) | true false,true | 3s 208s | 211s |

Table 1: Model checking results and comparison of verification times

We obtained that (2), (3) and (4) hold in the abstract semantics of the abstract pathways $P|J_1$, $P|J_2$ and $P|J_3$, respectively. Moreover, model checking required less than three seconds for (2), 213 seconds for (3) and less than one second for (4). Overall, modular verification required 217 seconds, that is 43 seconds less than verification on the complete model.

5.2 Reasoning on molecular components

As it can be seen in the component interaction graph and in the diagram in Figure 1, some molecular components are involved in complex interactions. This is true in particular for components *EGFR*, *GAP*, *RasGDP*, *Sos*, *Shc* and *Grb2* which form a clique in the component interaction graph. We are interested in understanding whether all of these components are really necessary in order to obtain the final products of the pathway. The idea is to test whether the final species are produced when the components of interest are assumed one by one as disabled. Molecular components *EGFR* and *RasGDP* are for sure necessary since they connect the clique with the other molecular components of the pathway. Consequently, we focus our analysis on *GAP*, *Sos*, *Shc* and *Grb2*.

In order to disable a molecular component we consider as absent all of its species in the initial state of the systems. Hence, we consider a set of four (complete) models, each with one of the four components under study disabled. On each model we try to verify property (1): if the property does not hold, then the component that is disabled in such a model is necessary for the pathway; on the other hand, if the property holds, then the component turns out to be not necessary since the products of the pathway can be obtained even without it. The same tests can be also done in a modular way by decomposing the pathway and the property as in Section 5.1.

In Table 1 we summarise the property verification results and compare verification times obtained by model checking the complete models and by following the modular approach. The first row of data in the table reports verification results in which no component is disabled (as in Section 5.1). The other results show that *Shc* is not a necessary component, whereas all of the other three are. As previously, the time required by modular verification is smaller than the one required by model checking the complete model. This is true in particular in the case in which *GAP* is disabled since property (2), the verification of which is very fast, turns out to be false.

Note that in the case of modular verification of the models in which *GAP*, *Sos* and *Grb2* were disabled we needed to verify some additional properties. In particular, in the case of *GAP* we have that property (2) does not hold in the abstract semantics of $P \downarrow J_1$, and in the cases of *Sos* and *Grb2* property (3) does not hold in the abstract semantics of $P \downarrow J_2$. We remark that our modular verification approach guarantees only that properties proved to hold in a model fragment also hold in the complete model. Nothing can be said, instead, of properties that does not hold in the model fragments. In order to avoid applying model checking on the complete model to check whether these properties hold there, we consider some new properties whose satisfaction in suitable model fragments implies that properties (2) and (3) actually do not hold. In order to prove that (2) is actually false when *GAP* is disabled we consider the following property:

$$AG(\neg(\text{EGF-EGFR}^*)2\text{-GAP}) \quad (5)$$

In order to prove that (3) is actually false when either *Sos* or *Grb2* is disabled we consider the following property:

$$AG(\neg\text{Raf}^*) \quad (6)$$

Note that it is convenient to verify properties (5) and (6) together with (2) and (3), respectively. This avoids spending twice the time needed by the model checker to construct the data structure necessary to perform the verification. In the case of our experiments the construction of such data structures takes usually the 98%-99% of the verification time. Times reported in Table 1 are based on this optimisation.

6 Discussion and conclusions

In this paper we presented preliminary results in the development of a modular verification framework for biochemical pathways. We defined a modelling notation for pathways associated with a formal semantics and a notion of fairness that allows the dynamics to be accurately described by avoiding starvation situations among reactions. Moreover, we investigated a notion of molecular component of a pathway and we provided a methodology to infer molecular components from pathways the reactions of which satisfy some assumptions. Molecular components were then used by a projection operation that allows abstract pathways modelling an over-approximation of the behaviour of a group of components to be obtained from a pathway model. The fact that a property expressed by means of the ACTL⁻ logic holds in an abstract pathway was shown to imply that they hold also in the complete pathway model. This preservation is at the basis of the modular verification approach which was demonstrated on a well-established model of the EGF pathway.

The results of experiments given in Section 5 show that our modular verification approach allows properties to be verified in a shorter time than in the case of verification of the complete pathway model. However, in most of the cases the time saved was relatively small ($\sim 15\%$). We believe that the cause of this limited gain in efficiency is due to the projection operation we are considering at the moment, which is still somewhat rough. Our plan to improve efficiency is to define a projection operation that combines the current one (that essentially removes some molecular components from the model) with another that somehow minimises the description of components not removed by the model, but whose role in the property to be verified is marginal. In the case of the considered case study this would allow, for example, to reduce the size of the model of the components constituting the clique in the component interaction graph in Figure 2 by focusing on components *EGFR* and *RasGDP*, and by minimising the description of components *GAP*, *Shc*, *Sos* and *Grb2*. This would allow for a significant improvement in modular verification efficiency.

References

- [1] Paul C. Attie & E. Allen Emerson (1998): *Synthesis of concurrent systems with many similar processes*. *ACM Transactions on Programming Languages and Systems* 20(1), pp. 51–115, doi:10.1145/271510.271519.
- [2] Roberto Barbuti, Andrea Maggiolo-Schettini, Paolo Milazzo & Angelo Troina (2006): *A Calculus of Looping Sequences for Modelling Microbiological Systems*. *Fundamenta Informaticae* 72(1-3), pp. 21–35.
- [3] Jerry R. Burch, Edmund M. Clarke, Kenneth L. McMillan, David L. Dill & L. J. Hwang (1992): *Symbolic Model Checking: 10²⁰ States and Beyond*. *Information and Computation* 98(2), pp. 142–170, doi:10.1016/0890-5401(92)90017-A.
- [4] Luca Cardelli (2005): *Brane Calculi*. *Computational Methods in Systems Biology*, pp. 257–278, doi:10.1007/978-3-540-25974-9_24.
- [5] Alessandro Cimatti, Edmund Clarke, Enrico Giunchiglia, Fausto Giunchiglia, Marco Pistore, Marco Roveri, Roberto Sebastiani & Armando Tacchella (2002): *NuSMV Version 2: An OpenSource Tool for Symbolic Model Checking*. In: *Proc. International Conference on Computer-Aided Verification (CAV 2002)*, LNCS 2404, Springer, Copenhagen, Denmark, pp. 241–268, doi:10.1007/3-540-45657-0_29.
- [6] Federica Ciocchetta & Jane Hillston (2009): *Bio-PEPA: A framework for the modelling and analysis of biological systems*. *Theoretical Computer Science* 410(33-34), pp. 3065–3084, doi:10.1016/j.tcs.2009.02.037.
- [7] Edmund M. Clarke, Orna Grumberg & David E. Long (1994): *Model checking and abstraction*. *ACM Transactions on Programming Languages and Systems* 16(5), pp. 1512–1542, doi:10.1145/186025.186051.
- [8] Edmund M. Clarke, Orna Grumberg & Doron Peled (1999): *Model Checking*. MIT Press.
- [9] Vincent Danos & Cosimo Laneve (2004): *Formal molecular biology*. *Theoretical Computer Science* 325(1), pp. 69–110, doi:10.1016/j.tcs.2004.03.065.
- [10] Peter Drábik, Andrea Maggiolo-Schettini & Paolo Milazzo (2010): *Dynamic Sync-programs for Modular Verification of Biological Systems*. In: *2nd Int. Workshop on Non-Classical Models of Automata and applications (NCMA'10)*, 263, Austrian Computer Society, Jena, Germany, pp. 71–83.
- [11] Peter Drábik, Andrea Maggiolo-Schettini & Paolo Milazzo (2010): *Modular Verification of Interactive Systems with an Application to Biology*. *Electronic Notes in Theoretical Computer Science* 268, pp. 61–75, doi:10.1016/j.entcs.2010.12.006.
- [12] Peter Drábik, Andrea Maggiolo-Schettini & Paolo Milazzo (2011): *Modular Verification of Interactive Systems with an Application to Biology*. *Scientific Annals of Computer Science* 21, pp. 39–72.
- [13] E. Allen Emerson & Chin-Laung Lei (1987): *Modalities for model checking: branching time logic strikes back*. *Science of Computer Programming* 8, pp. 275–306, doi:10.1016/0167-6423(87)90036-0.
- [14] François Fages, Sylvain Soliman & Nathalie Chabrier-Rivier (2004): *Modelling and querying interaction networks in the biochemical abstract machine biocham*. *Journal of Biological Physics and Chemistry* 4, pp. 64–73.
- [15] John Heath, Marta Kwiatkowska, Gethin Norman, David Parker & Oksana Tymchyshyn (2008): *Probabilistic model checking of complex biological pathways*. *Theoretical Computer Science* 391(3), pp. 239–257, doi:10.1016/j.tcs.2007.11.013.
- [16] Pedro T. Monteiro, Delphine Ropers, Radu Mateescu, Ana T. Freitas & Hidde de Jong (2008): *Temporal logic patterns for querying dynamic models of cellular interaction networks*. *Bioinformatics* 24(16), pp. 227–233, doi:10.1093/bioinformatics/btn275.
- [17] Amir Pnueli (1981): *The temporal semantics of concurrent programs*. *Theoretical Computer Science* 13(1), pp. 45 – 60, doi:10.1016/0304-3975(81)90110-9.
- [18] Corrado Priami, Aviv Regev, Ehud Shapiro & William Silverman (2001): *Application of a stochastic name-passing calculus to representation and simulation of molecular processes*. *Information Processing Letters* 80(1), pp. 25–31, doi:10.1016/S0020-0190(01)00214-9.

- [19] Aviv Regev, Ekaterina M. Panina, William Silverman, Luca Cardelli & Ehud Shapiro (2004): *BioAmbients: an abstraction for biological compartments*. *Theoretical Computer Science* 325(1), pp. 141–167, doi:10.1016/j.tcs.2004.03.061.
- [20] Birgit Schoeberl, Claudia Eichler-Jonsson, Ernst Dieter Gilles & Gertraud Muller (2002): *Computational modeling of the dynamics of the MAP kinase cascade activated by surface and internalized EGF receptors*. *Nature Biotechnology* 20(4), pp. 370–375, doi:10.1038/nbt0402-370.