Systems biology with stochastic process algebras

Jane Hillston.
LFCS, University of Edinburgh

29th June 2006

Joint work with Muffy Calder, Adam Duguid, Stephen Gilmore and Marco Stenico

Outline

Outline

Systems Biology Methodology

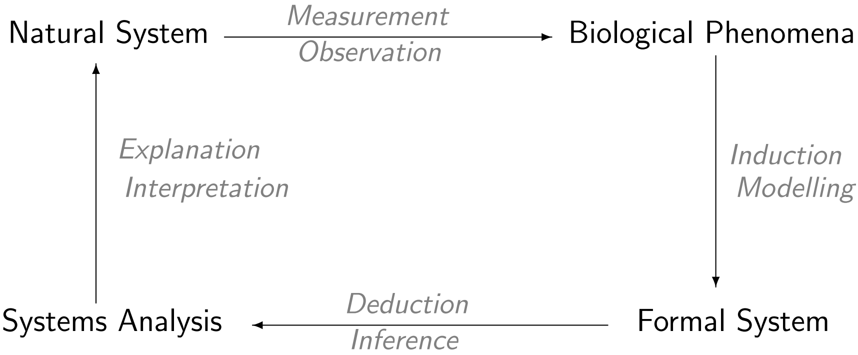

Systems Biology Methodology

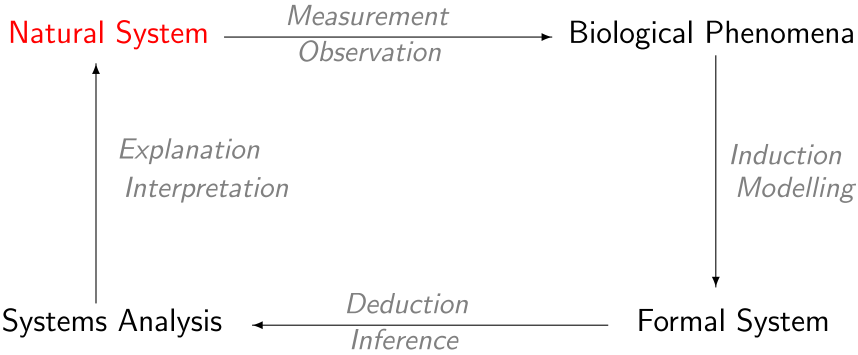

Systems Biology Methodology

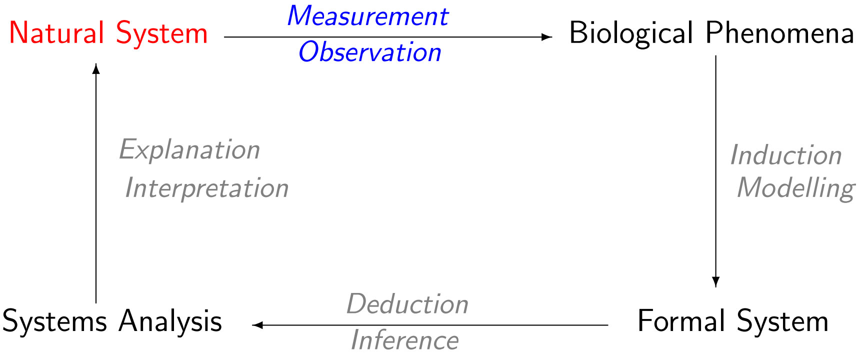

Systems Biology Methodology

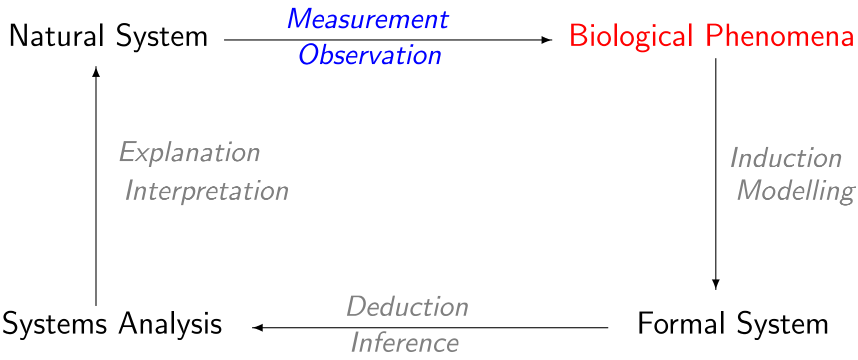

Systems Biology Methodology

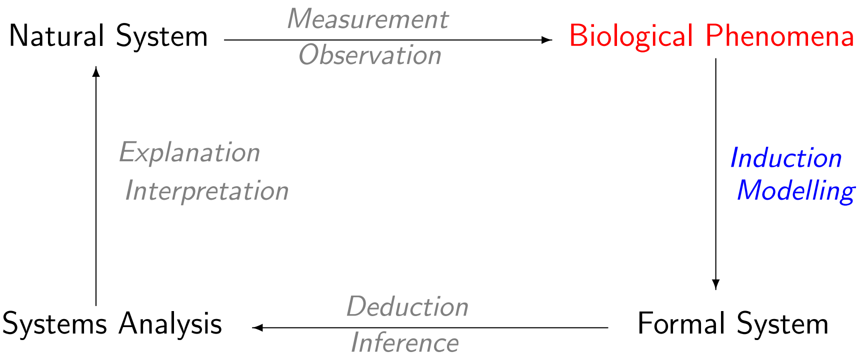

Systems Biology Methodology

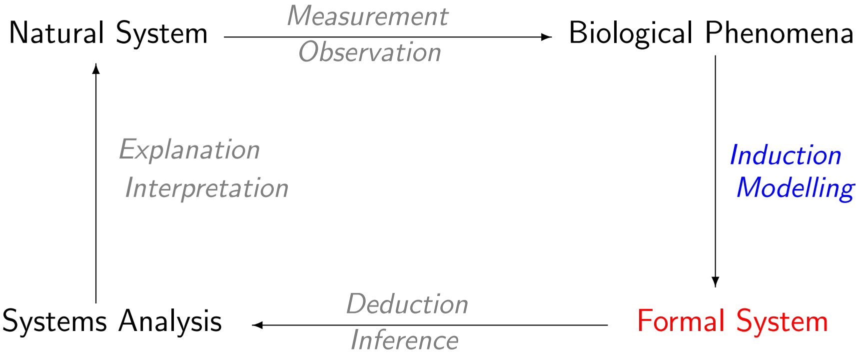

Systems Biology Methodology

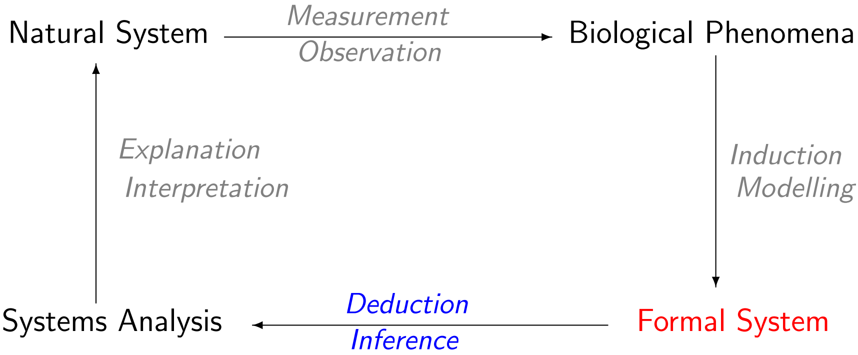

Systems Biology Methodology

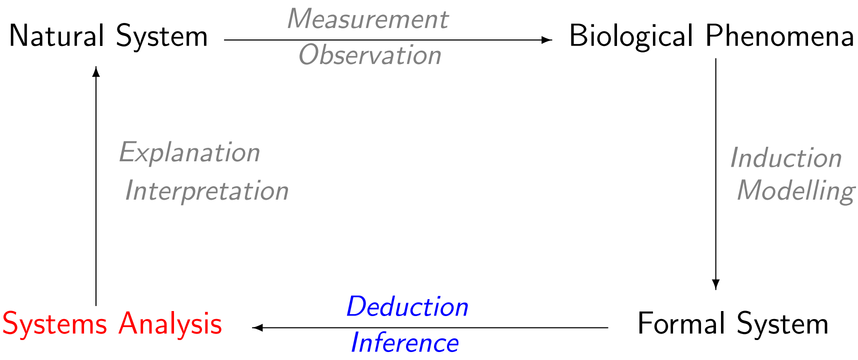

Systems Biology Methodology

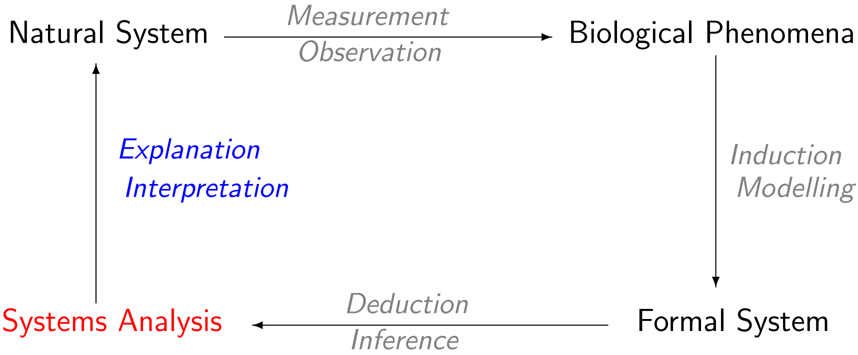

Systems Biology Methodology

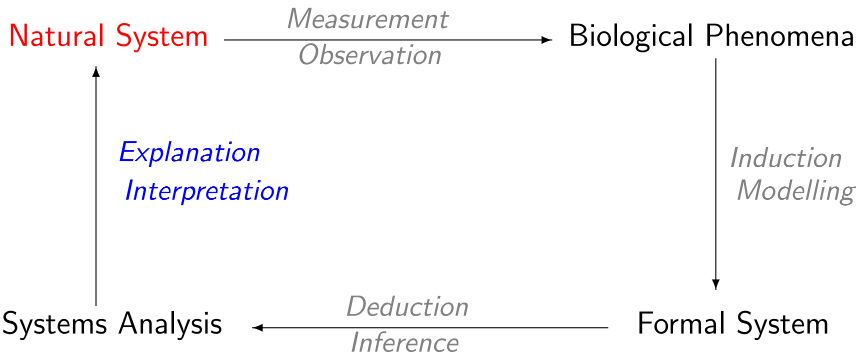

Systems Biology Methodology

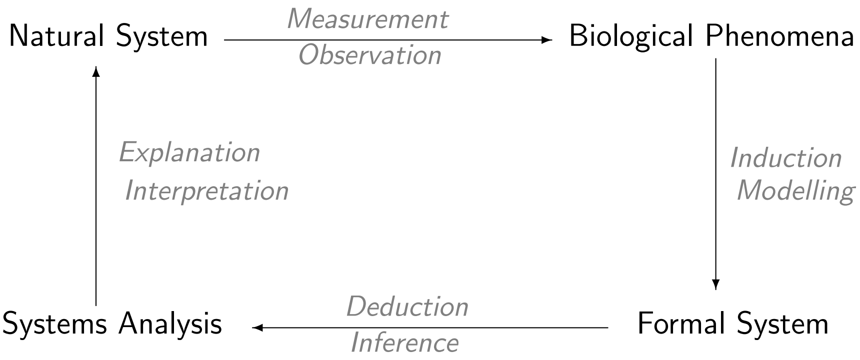

Systems Biology Methodology

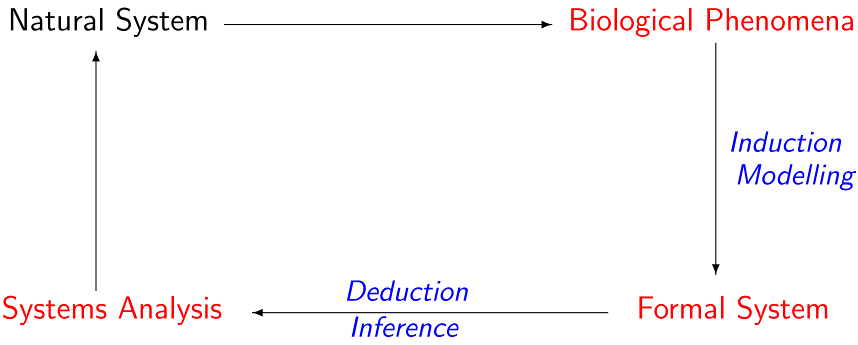

Current Hypothesis

Some of the techniques we have developed over the last thirty years for modelling complex software systems can be beneficially applied to the modelling aspects of systems biology.

Current Hypothesis

Some of the techniques we have developed over the last thirty years for modelling complex software systems can be beneficially applied to the modelling aspects of systems biology.

In particular formalisms which encompass support for

- Abstraction
- Modularity and
- Reasoning

have a key role to play.

Current Hypothesis

Some of the techniques we have developed over the last thirty years for modelling complex software systems can be beneficially applied to the modelling aspects of systems biology.

In particular formalisms which encompass support for

- Abstraction
- Modularity and
- Reasoning

have a key role to play.

Process algebras have mechanisms for each of these, and stochastic extensions which allow dynamic properties to be analysed.

Signal Transduction Pathways

There are many different aspects to systems biology and even the modelling activity can be applied to several different types of biochemical signalling pathways:

Signal Transduction Pathways

There are many different aspects to systems biology and even the modelling activity can be applied to several different types of biochemical signalling pathways:

- gene regulation pathways

Signal Transduction Pathways

There are many different aspects to systems biology and even the modelling activity can be applied to several different types of biochemical signalling pathways:

- gene regulation pathways
- metabolic pathways

Signal Transduction Pathways

There are many different aspects to systems biology and even the modelling activity can be applied to several different types of biochemical signalling pathways:

- gene regulation pathways
- metabolic pathways
- signal transduction pathways.

Signal Transduction Pathways

There are many different aspects to systems biology and even the modelling activity can be applied to several different types of biochemical signalling pathways:

- gene regulation pathways
- metabolic pathways
- signal transduction pathways.

The work we have been doing with PEPA currently is primarily focussed on signal transduction pathways.

Outline

Process Algebras for Systems Biology

- Process algebraic formulations are compositional and make interactions/constraints explicit — not the case with classical ordinary differential equation models.

Process Algebras for Systems Biology

- Process algebraic formulations are compositional and make interactions/constraints explicit — not the case with classical ordinary differential equation models.
- Structure can also be apparent.

Process Algebras for Systems Biology

- Process algebraic formulations are compositional and make interactions/constraints explicit — not the case with classical ordinary differential equation models.
- Structure can also be apparent.
- Equivalence relations allow formal comparison of high-level descriptions.

Process Algebras for Systems Biology

- Process algebraic formulations are compositional and make interactions/constraints explicit — not the case with classical ordinary differential equation models.
- Structure can also be apparent.
- Equivalence relations allow formal comparison of high-level descriptions.
- There are well-established techniques for reasoning about the behaviours and properties of models, supported by software. These include qualitative and quantitative analysis, and model checking.

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a Markov Process (CTMC).

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a Markov Process (CTMC).

SPA
MODEL

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a Markov Process (CTMC).

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a Markov Process (CTMC).

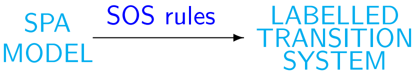

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a Markov Process (CTMC).

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a Markov Process (CTMC).

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a system of ordinary differential equations (ODEs).

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a system of ordinary differential equations (ODEs).

SPA
MODEL

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a system of ordinary differential equations (ODEs).

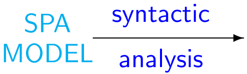

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a system of ordinary differential equations (ODEs).

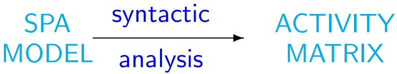

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a system of ordinary differential equations (ODEs).

SPA MODEL —(syntactic analysis)→ ACTIVITY MATRIX —(continuous interpretation)→

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a system of ordinary differential equations (ODEs).

SPA MODEL —(syntactic analysis)→ ACTIVITY MATRIX —(continuous interpretation)→ ODEs

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

SPA
MODEL

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

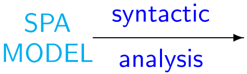

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

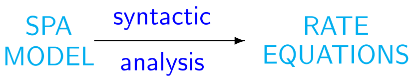

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

SPA MODEL —(syntactic analysis)→ RATE EQUATIONS —(Gillespie's algorithm)→

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

Deriving quantitative data

SPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a stochastic simulation.

Each of these has tool support so that the underlying model is derived automatically according to the predefined rules.

Models as Tools

When designing a modelling formalism it is important to consider two key aspects:

Models as Tools

When designing a modelling formalism it is important to consider two key aspects:

- the ability of the formalism to capture the behaviour of interest and the availability of data to instantiate the model;

Models as Tools

When designing a modelling formalism it is important to consider two key aspects:

- the ability of the formalism to capture the behaviour of interest and the availability of data to instantiate the model;
- the amenability of the formalism to appropriate analysis.

Models as Tools

When designing a modelling formalism it is important to consider two key aspects:

- the ability of the formalism to capture the behaviour of interest and the availability of data to instantiate the model;
- the amenability of the formalism to appropriate analysis.

What is the behaviour of interest and appropriate analysis depends on the question or problem you are seeking to address.

Models as Tools

When designing a modelling formalism it is important to consider two key aspects:

- the ability of the formalism to capture the behaviour of interest and the availability of data to instantiate the model;
- the amenability of the formalism to appropriate analysis.

What is the behaviour of interest and appropriate analysis depends on the question or problem you are seeking to address.

It is not the case the models have to be completely faithful to their subject in order to be useful.

Motivations for Abstraction

Our motivations for seeking more abstraction in process algebra models for systems biology comes from both key aspects of modelling:

Motivations for Abstraction

Our motivations for seeking more abstraction in process algebra models for systems biology comes from both key aspects of modelling:

- The data that we have available to parameterise models is sometimes speculative rather than precise.

Motivations for Abstraction

Our motivations for seeking more abstraction in process algebra models for systems biology comes from both key aspects of modelling:

- The data that we have available to parameterise models is sometimes speculative rather than precise. This suggests that we should use semiquantitative models rather than quantitative ones.

Motivations for Abstraction

Our motivations for seeking more abstraction in process algebra models for systems biology comes from both key aspects of modelling:

- The data that we have available to parameterise models is sometimes speculative rather than precise. This suggests that we should use semiquantitative models rather than quantitative ones.
- Process algebra based analyses such as comparing models (e.g. for equivalence or simulation) and model checking are only possible is the state space is not prohibitively large.

Motivations for Abstraction

Our motivations for seeking more abstraction in process algebra models for systems biology comes from both key aspects of modelling:

- The data that we have available to parameterise models is sometimes speculative rather than precise. This suggests that we should use semiquantitative models rather than quantitative ones.
- Process algebra based analyses such as comparing models (e.g. for equivalence or simulation) and model checking are only possible is the state space is not prohibitively large.

Individuals vs. Populations

Whenever we have a system consisting of a large number of individuals we can consider treating them instead as a population.

Individuals vs. Populations

Whenever we have a system consisting of a large number of individuals we can consider treating them instead as a population.

These alternative styles of model are already available in the systems biology arena:

Individuals vs. Populations

Whenever we have a system consisting of a large number of individuals we can consider treating them instead as a population.

These alternative styles of model are already available in the systems biology arena:

- Stochastic Simulations (Gillespie *et al.*) are individual-based models in which each molecule is treated separately.

Individuals vs. Populations

Whenever we have a system consisting of a large number of individuals we can consider treating them instead as a population.

These alternative styles of model are already available in the systems biology arena:

- Stochastic Simulations (Gillespie *et al.*) are individual-based models in which each molecule is treated separately.
- Ordinary Differential Equations are population-based models in which populations of molecules are represented abstractly as concentrations.

Individuals vs. Populations

Whenever we have a system consisting of a large number of individuals we can consider treating them instead as a population.

These alternative styles of model are already available in the systems biology arena:

- Stochastic Simulations (Gillespie *et al.*) are individual-based models in which each molecule is treated separately.
- Ordinary Differential Equations are population-based models in which populations of molecules are represented abstractly as concentrations.

From our perspectives these are essentially analysis techniques and we would like to have access to both whilst also being able to carry out process algebra-based analyses.

Alternative Representations

ODEs
(Continuous Approximation)

Stochastic Simulation

Alternative Representations

SPA models with discrete levels of population

ODEs
(Continuous Approximation)

SPA models of individuals

Stochastic Simulation

Alternative Representations

CTMCs with discrete levels of population

ODEs (Continuous Approximation)

CTMCs with individuals

Stochastic Simulation

Alternative Representations

CTMCs with discrete levels of population

ODEs
(Continuous Approximation)

Agreement for large numbers of individuals

CTMCs with individuals

Stochastic Simulation

Alternative Representations

CTMCs with discrete levels of population

Asymptotic relationship

as $N \longrightarrow \infty$

CTMCs with individuals

ODEs (Continuous Approximation)

Agreement for large

numbers of individuals

Stochastic Simulation

Alternative Representations

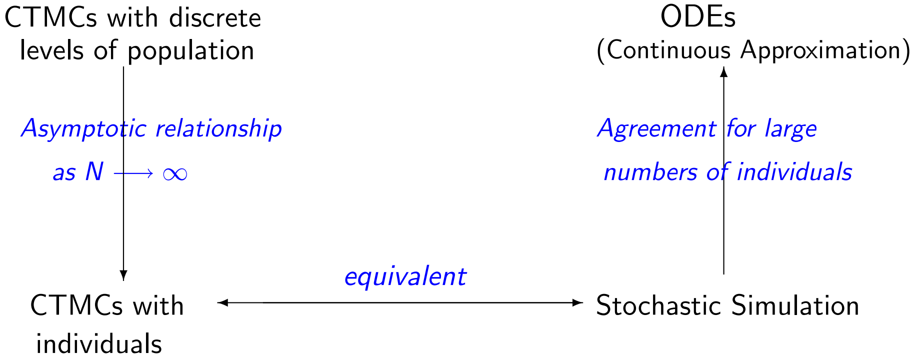

Alternative Representations

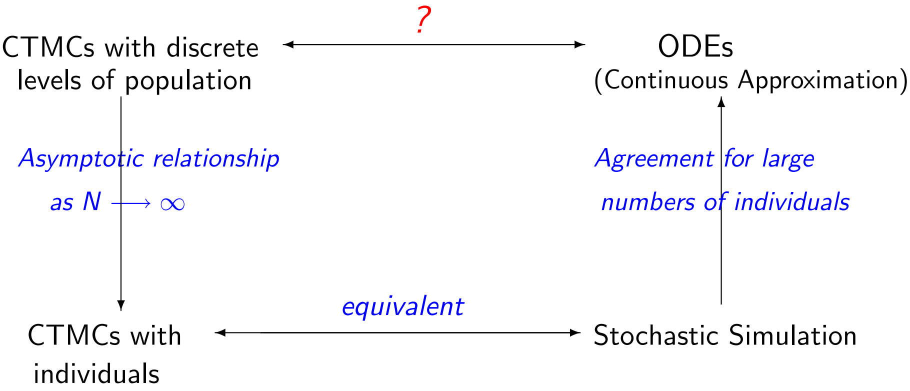

Mapping biological systems to process algebra

The seminal work by Regev *et al.* has stimulated much use of the stochastic π-calculus and related calculi, for modelling biochemical signalling within cells.

Mapping biological systems to process algebra

The seminal work by Regev *et al.* has stimulated much use of the stochastic π-calculus and related calculi, for modelling biochemical signalling within cells.

This work is individual-based: it treats a molecule in a pathway as corresponding to the component in the process algebra description.

Mapping biological systems to process algebra

The seminal work by Regev *et al.* has stimulated much use of the stochastic π-calculus and related calculi, for modelling biochemical signalling within cells.

This work is individual-based: it treats a molecule in a pathway as corresponding to the component in the process algebra description.

Mapping biological systems to process algebra

The seminal work by Regev *et al.* has stimulated much use of the stochastic π-calculus and related calculi, for modelling biochemical signalling within cells.

This work is individual-based: it treats a molecule in a pathway as corresponding to the component in the process algebra description.

In the PEPA modelling we have been doing we are experimented with more population-based mappings between process algebra constructs and elements of signalling pathways.

Mapping biological systems to process algebra

The seminal work by Regev *et al.* has stimulated much use of the stochastic π-calculus and related calculi, for modelling biochemical signalling within cells.

This work is individual-based: it treats a molecule in a pathway as corresponding to the component in the process algebra description.

In the PEPA modelling we have been doing we are experimented with more population-based mappings between process algebra constructs and elements of signalling pathways.

In our first mapping we focus on species (c.f. a type rather than an instance, or a class rather than an object).

Mapping biological systems to process algebra

The seminal work by Regev *et al.* has stimulated much use of the stochastic π-calculus and related calculi, for modelling biochemical signalling within cells.

This work is individual-based: it treats a molecule in a pathway as corresponding to the component in the process algebra description.

In the PEPA modelling we have been doing we are experimented with more population-based mappings between process algebra constructs and elements of signalling pathways.

In our first mapping we focus on species (c.f. a type rather than an instance, or a class rather than an object).

In our second we focus on sub-pathways.

Alternative Mappings: illustration

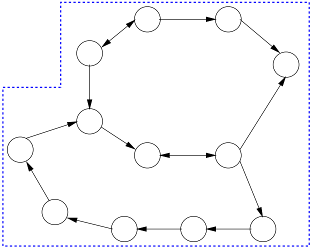

Alternative Mappings: illustration

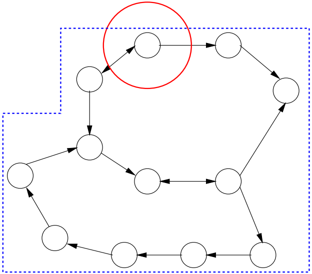

Reagent mapping: Each species is a distinct component in the model with local states to capture differing levels of concentration

Alternative Mappings: illustration

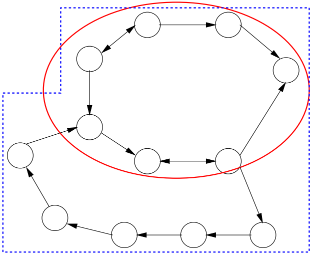

Pathway mapping: Each sub-pathway is a distinct component in the model with local states to capture progress through the pathway

Alternative Mappings: illustration

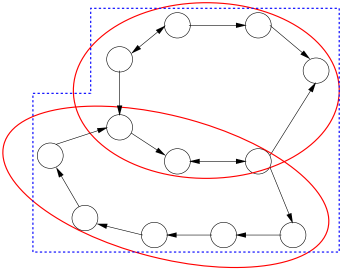

Pathway mapping: Each sub-pathway is a distinct component in the model with local states to capture progress through the pathway

Alternative Mappings: illustration

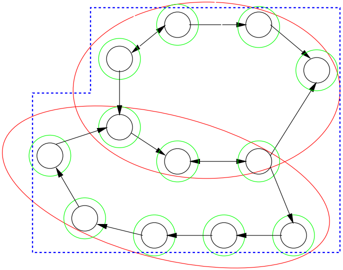

Reasoning based on bisimulation equivalence is able to prove that the two representations are equivalent.

Alternative Mappings: illustration

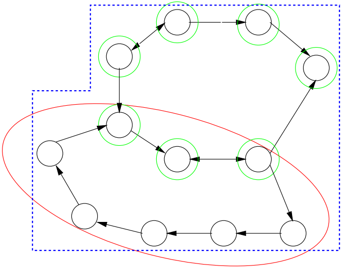

Different parts of the system may use different mappings, reflecting perhaps the level of knowledge (data) available, or the primary interests of the modeller.

Reagent-Centric Modelling

In the reagent-centric modelling we map each species within the pathway to a distinct PEPA component.

Reagent-Centric Modelling

In the reagent-centric modelling we map each species within the pathway to a distinct PEPA component.

The local states of the component capture different levels of concentration. This can be viewed as a discretization of the continuous concentration captured by an ODE.

Reagent-Centric Modelling

In the reagent-centric modelling we map each species within the pathway to a distinct PEPA component.

The local states of the component capture different levels of concentration. This can be viewed as a discretization of the continuous concentration captured by an ODE.

Reactions map to activities in the process algebra and have appropriate impact on the concentration levels of the reagents involved.

Reagent-Centric Modelling

In the reagent-centric modelling we map each species within the pathway to a distinct PEPA component.

The local states of the component capture different levels of concentration. This can be viewed as a discretization of the continuous concentration captured by an ODE.

Reactions map to activities in the process algebra and have appropriate impact on the concentration levels of the reagents involved.

Rates are calculated appropriately (based on mass action).

Reagent Roles in Reactions

Reagent Role	*Reaction impact*	*Concentration impact*
Producer	decrease	has a positive impact proportional to the current concentration level
Product	increase	had no impact on the rate (except at saturation)
Enzyme	unchanged	has a positive impact proportional to current concentration
Inhibitor	unchanged	has a negative impact inversely proportional to current concentration

Reagent Roles in Reactions

Reagent Role	*Reaction impact*	*Concentration impact*
Producer	decrease	has a positive impact proportional to the current concentration level
Product	increase	had no impact on the rate (except at saturation)
Enzyme	unchanged	has a positive impact proportional to current concentration
Inhibitor	unchanged	has a negative impact inversely proportional to current concentration

Reagent Roles in Reactions

Reagent Role	*Reaction impact*	*Concentration impact*
Producer	decrease	has a positive impact proportional to the current concentration level
Product	increase	had no impact on the rate (except at saturation)
Enzyme	unchanged	has a positive impact proportional to current concentration
Inhibitor	unchanged	has a negative impact inversely proportional to current concentration

Reagent Roles in Reactions

Reagent Role	*Reaction impact*	*Concentration impact*
Producer	decrease	has a positive impact proportional to the current concentration level
Product	increase	had no impact on the rate (except at saturation)
Enzyme	unchanged	has a positive impact proportional to current concentration
Inhibitor	unchanged	has a negative impact inversely proportional to current concentration

Small synthetic example network

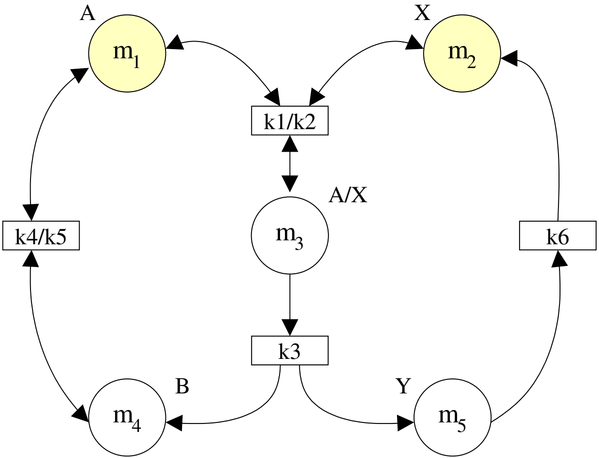

Small synthetic example network in PEPA (1)

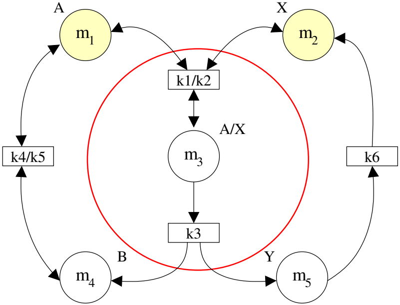

$$A/X_H \stackrel{def}{=} (k2react, k2).A/X_L + (k3react, k3).A/X_L$$

$$A/X_L \stackrel{def}{=} (k1react, k1).A/X_H$$

Small synthetic example network in PEPA (2)

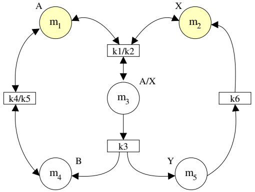

$$
\begin{aligned}
A_H &\stackrel{def}{=} (k1react, k1).A_L + (k5react, k5).A_L \\
A_L &\stackrel{def}{=} (k2react, k2).A_H + (k4react, k4).A_H \\
X_H &\stackrel{def}{=} (k1react, k1).X_L \\
X_L &\stackrel{def}{=} (k2react, k2).X_H + (k6react, k6).X_H \\
A/X_H &\stackrel{def}{=} (k2react, k2).A/X_L + (k3react, k3).A/X_L \\
A/X_L &\stackrel{def}{=} (k1react, k1).A/X_H \\
B_H &\stackrel{def}{=} (k4react, k4).B_L \\
B_L &\stackrel{def}{=} (k5react, k5).B_H + (k3react, k3).B_H \\
Y_H &\stackrel{def}{=} (k6react, k6).Y_L \\
Y_L &\stackrel{def}{=} (k3react, k3).Y_H
\end{aligned}
$$

Small synthetic example network in PEPA (3)

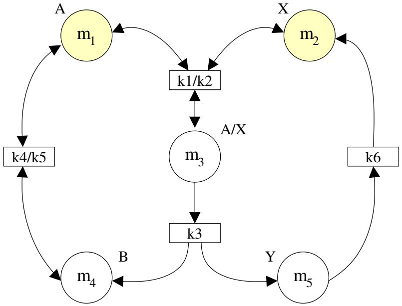

$$(((A_H \underset{\{k1react,k2react\}}{\bowtie} X_H) \underset{\{k1react,k2react\}}{\bowtie} A/X_L) \underset{\{k3react,k4react,k5react\}}{\bowtie} B_L) \underset{\{k3react,k6react\}}{\bowtie} Y_L$$

State space of the reactants model

r_1	$A_H \underset{\{k1react,k2react\}}{\bowtie} X_H \underset{\{k1react,k2react\}}{\bowtie} A/X_L \underset{\{k3react,k4react,k5react\}}{\bowtie} B_L \underset{\{k3react,k6react\}}{\bowtie} Y_L$
r_2	$A_L \underset{\{k1react,k2react\}}{\bowtie} X_H \underset{\{k1react,k2react\}}{\bowtie} A/X_L \underset{\{k3react,k4react,k5react\}}{\bowtie} B_H \underset{\{k3react,k6react\}}{\bowtie} Y_L$
r_3	$A_L \underset{\{k1react,k2react\}}{\bowtie} X_L \underset{\{k1react,k2react\}}{\bowtie} A/X_H \underset{\{k3react,k4react,k5react\}}{\bowtie} B_L \underset{\{k3react,k6react\}}{\bowtie} Y_L$
r_4	$A_L \underset{\{k1react,k2react\}}{\bowtie} X_L \underset{\{k1react,k2react\}}{\bowtie} A/X_L \underset{\{k3react,k4react,k5react\}}{\bowtie} B_H \underset{\{k3react,k6react\}}{\bowtie} Y_H$
r_5	$A_H \underset{\{k1react,k2react\}}{\bowtie} X_L \underset{\{k1react,k2react\}}{\bowtie} A/X_L \underset{\{k3react,k4react,k5react\}}{\bowtie} B_L \underset{\{k3react,k6react\}}{\bowtie} Y_H$

Pathway-Centric Modelling

In the pathway-centric modelling we map distinct sub-pathway to a distinct PEPA component. The granularity at which sub-pathways are defined is in the control of the modeller.

Pathway-Centric Modelling

In the pathway-centric modelling we map distinct sub-pathway to a distinct PEPA component. The granularity at which sub-pathways are defined is in the control of the modeller.

The local states of the component capture progression through the sub-pathway.

Pathway-Centric Modelling

In the pathway-centric modelling we map distinct sub-pathway to a distinct PEPA component. The granularity at which sub-pathways are defined is in the control of the modeller.

The local states of the component capture progression through the sub-pathway.

Concentrations are implicit — there is one copy of the sub-pathway component for each concentration level. Thus the total concentration may be distributed across different states of the sub-pathway.

Pathway-Centric Modelling

In the pathway-centric modelling we map distinct sub-pathway to a distinct PEPA component. The granularity at which sub-pathways are defined is in the control of the modeller.

The local states of the component capture progression through the sub-pathway.

Concentrations are implicit — there is one copy of the sub-pathway component for each concentration level. Thus the total concentration may be distributed across different states of the sub-pathway.

Sub-pathways synchronise on the reactions they have in common at rates dictated by the current concentration in each sub-pathway.

Small synthetic example network in PEPA (4)

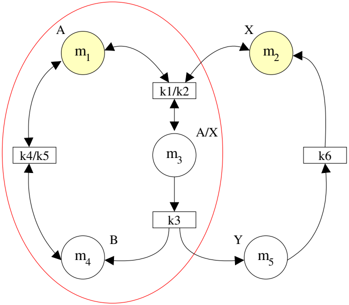

$$
\begin{aligned}
PathwayA_1 &\stackrel{def}{=} (k1react, k1).PathwayA_2 \\
&\quad + (k5react, k5).PathwayA_3 \\
PathwayA_2 &\stackrel{def}{=} (k2react, k2).PathwayA_1 \\
&\quad + (k3react, k3).PathwayA_3 \\
PathwayA_3 &\stackrel{def}{=} (k4react, k4).PathwayA_1
\end{aligned}
$$

Small synthetic example network in PEPA (5)

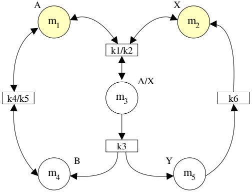

$$
\begin{aligned}
PathwayA_1 &\stackrel{def}{=} (k1react, k1).PathwayA_2 \\
&\quad + (k5react, k5).PathwayA_3 \\
PathwayA_2 &\stackrel{def}{=} (k2react, k2).PathwayA_1 \\
&\quad + (k3react, k3).PathwayA_3 \\
PathwayA_3 &\stackrel{def}{=} (k4react, k4).PathwayA_1 \\
PathwayX_1 &\stackrel{def}{=} (k1react, k1).PathwayX_2 \\
PathwayX_2 &\stackrel{def}{=} (k2react, k2).PathwayX_1 \\
&\quad + (k3react, k3).PathwayX_3 \\
PathwayX_3 &\stackrel{def}{=} (k6react, k6).PathwayX_1
\end{aligned}
$$

Small synthetic example network in PEPA (6)

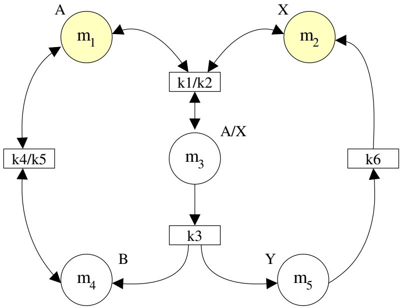

$$PathwayA_1 \underset{\{k1react,k2react,k3react\}}{\bowtie} PathwayX_1$$

State space of the pathway model

p_1	$PathwayA_1 \underset{\{k1react,k2react,k3react\}}{\bowtie} PathwayX_1$
p_2	$PathwayA_2 \underset{\{k1react,k2react,k3react\}}{\bowtie} PathwayX_2$
p_3	$PathwayA_3 \underset{\{k1react,k2react,k3react\}}{\bowtie} PathwayX_1$
p_4	$PathwayA_3 \underset{\{k1react,k2react,k3react\}}{\bowtie} PathwayX_3$
p_5	$PathwayA_1 \underset{\{k1react,k2react,k3react\}}{\bowtie} PathwayX_3$

Relationship

These are easily shown to be bisimulation equivalent (in fact they are isomorphic).

Relationship

These are easily shown to be bisimulation equivalent (in fact they are isomorphic).

Moreover this remains true when we increase the discretisation levels of the concentrations eg with three levels instead of two:

Relationship

These are easily shown to be bisimulation equivalent (in fact they are isomorphic).

Moreover this remains true when we increase the discretisation levels of the concentrations eg with three levels instead of two: the reactant-based model of A becomes:

$$
\begin{aligned}
A_2 &\stackrel{def}{=} (k1react, 2 \times k1).A_1 + (k5react, 2 \times k5).A_1 \\
A_1 &\stackrel{def}{=} (k1react, k1).A_0 + (k5react, k5).A_0 \\
&\quad + (k2react, k2).A_2 + (k4react, k4).A_2 \\
A_0 &\stackrel{def}{=} (k2react, 2 \times k2).A_1 + (k4react, 2 \times k4).A_1
\end{aligned}
$$

Relationship

These are easily shown to be bisimulation equivalent (in fact they are isomorphic).

Moreover this remains true when we increase the discretisation levels of the concentrations eg with three levels instead of two: the configuration of the pathway model becomes:

$$(PathwayA_1 \parallel PathwayA_1) \underset{\{k1react, k2react, k3react\}}{\bowtie} (PathwayX_1 \parallel PathwayX_1)$$

Example: The Ras/Raf-1/MEK/ERK pathway

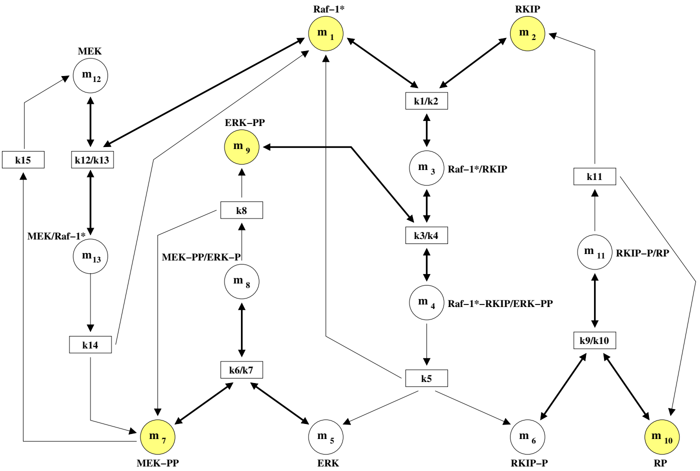

PEPA components of the reagent-centric model

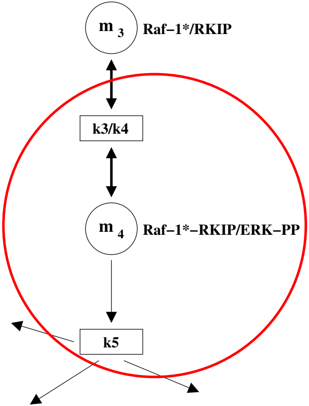

$$\begin{aligned}
\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{H}} &\stackrel{\mathit{def}}{=} \\
&(\mathit{k5product}, k_5).\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{L}} \\
&\quad + (\mathit{k4react}, k_4).\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{L}}
\end{aligned}$$

$$\begin{aligned}
\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{L}} &\stackrel{\mathit{def}}{=} \\
&(\mathit{k3react}, k_3).\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{H}}
\end{aligned}$$

PEPA components of the reagent-centric model

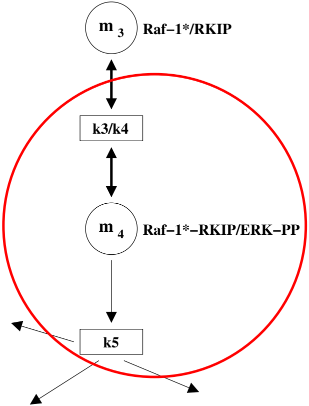

$$\begin{aligned} \text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{H}} &\stackrel{\mathit{def}}{=} (\mathit{k5product}, k_5).\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{L}} \\ &\quad + (\mathit{k4react}, k_4).\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{L}} \end{aligned}$$

$$\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{L}} \stackrel{\mathit{def}}{=} (\mathit{k3react}, k_3).\text{Raf-1}^*/\text{RKIP}/\text{ERK-PP}_{\mathrm{H}}$$

Each reagent gives rise to a pair of PEPA definitions, one for high concentration and one for low concentration.

PEPA components of the pathway-centric model

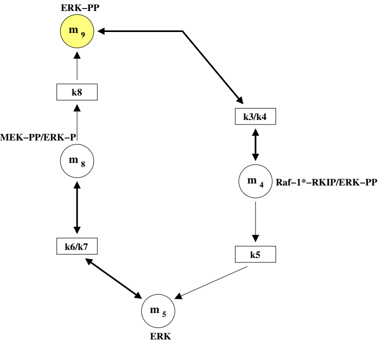

$$\begin{aligned}
Pathway_{30} &\stackrel{def}{=} (k3react, k_3).Pathway_{31} \\
Pathway_{31} &\stackrel{def}{=} (k5product, k_5).Pathway_{32} \\
&\quad +(k4react, k_4).Pathway_{30} \\
Pathway_{32} &\stackrel{def}{=} (k6react, k_6).Pathway_{33} \\
Pathway_{33} &\stackrel{def}{=} (k8product, k_8).Pathway_{30} \\
&\quad +(k7react, k_7).Pathway_{32}
\end{aligned}$$

PEPA components of the pathway-centric model

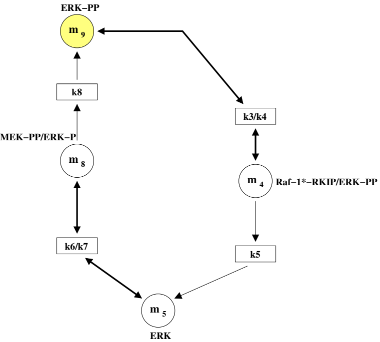

$$
\begin{aligned}
Pathway_{30} &\stackrel{def}{=} (k3react, k_3).Pathway_{31} \\
Pathway_{31} &\stackrel{def}{=} (k5product, k_5).Pathway_{32} \\
&\quad +(k4react, k_4).Pathway_{30} \\
Pathway_{32} &\stackrel{def}{=} (k6react, k_6).Pathway_{33} \\
Pathway_{33} &\stackrel{def}{=} (k8product, k_8).Pathway_{30} \\
&\quad +(k7react, k_7).Pathway_{32}
\end{aligned}
$$

For each reagent that has an initial concentration we define the sub-pathway generated by the progression of that reagent.

Commentary on the models

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels:

Commentary on the models

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.

Commentary on the models

- Neither model currently “correctly” captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.

Commentary on the models

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.
- The pathway-centric model can be regarded as a more structural, coarse-grained view of the system.

Commentary on the models

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.
- The pathway-centric model can be regarded as a more structural, coarse-grained view of the system.
- Applying the structured operational semantics reveals that they are strongly bisimilar

Commentary on the models

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.
- The pathway-centric model can be regarded as a more structural, coarse-grained view of the system.
- Applying the structured operational semantics reveals that they are strongly bisimilar (in fact, in this case, isomorphic).

Quantitative Analysis

Using 8 discrete levels of concentration we have been able to obtain good agreement with the previously published ODE model.

Quantitative Analysis

Using 8 discrete levels of concentration we have been able to obtain good agreement with the previously published ODE model. For example, using CTMC analysis we can assess the impact of the varying the rate at which RKIP binds to Raf-1*.

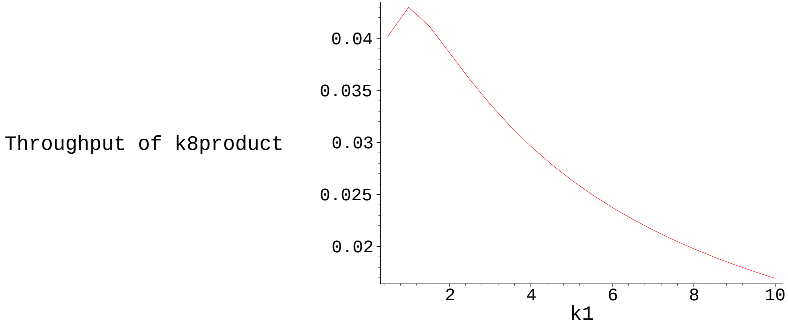

Transformation

- We have defined a systematic transformation between the two representations.

Transformation

- We have defined a systematic transformation between the two representations.
- It is based on an intermediate form, the activity matrix which is also integral to our translation to ODEs.

Transformation

- We have defined a systematic transformation between the two representations.
- It is based on an intermediate form, the activity matrix which is also integral to our translation to ODEs.
- The matrix captures the impact of each reaction on each local state of each component.

Transformation

- We have defined a systematic transformation between the two representations.
- It is based on an intermediate form, the activity matrix which is also integral to our translation to ODEs.
- The matrix captures the impact of each reaction on each local state of each component.
- Each reagent in the system corresponds to a single row in the matrix: as a producer only at the highest level of concentration and consumer only at the lowest. It exhibits both behaviours at intermediate levels.

Transformation

- We have defined a systematic transformation between the two representations.
- It is based on an intermediate form, the activity matrix which is also integral to our translation to ODEs.
- The matrix captures the impact of each reaction on each local state of each component.
- Each reagent in the system corresponds to a single row in the matrix: as a producer only at the highest level of concentration and consumer only at the lowest. It exhibits both behaviours at intermediate levels.
- Identifying the subpathways is more involved but is based on a colouring algorithm which tracks cycles through the causality graph of reactions.

Outline

A simple circadian clock

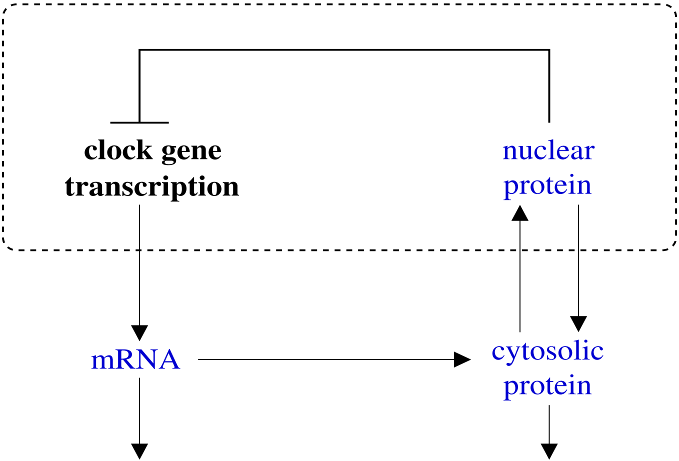

A simple circadian clock

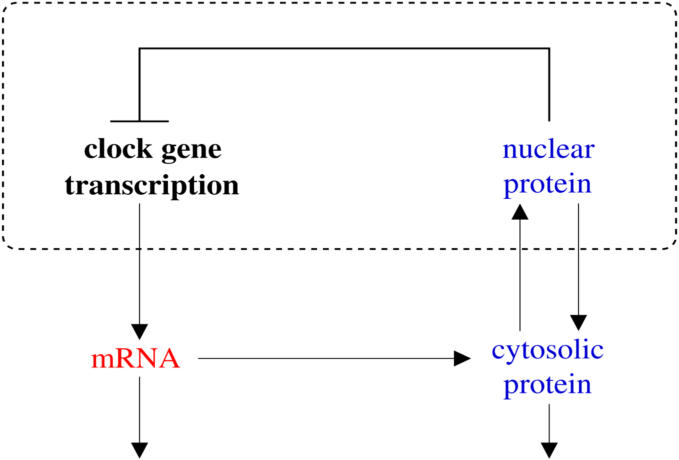

A simple circadian clock

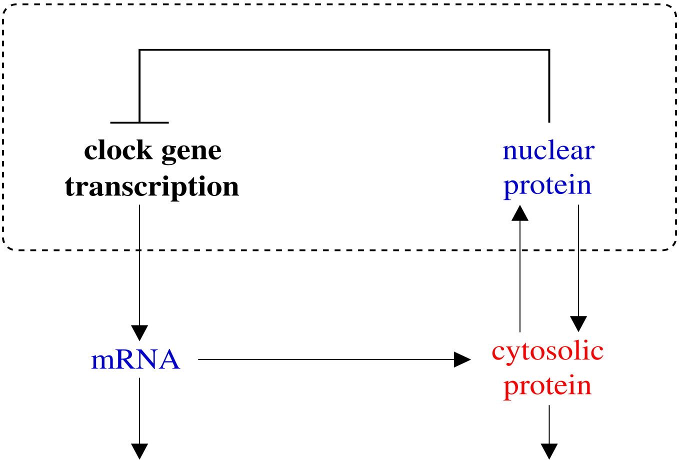

A simple circadian clock

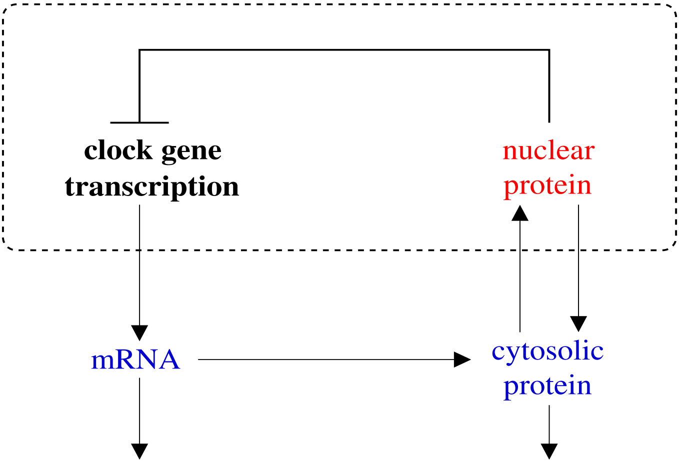

A simple circadian clock

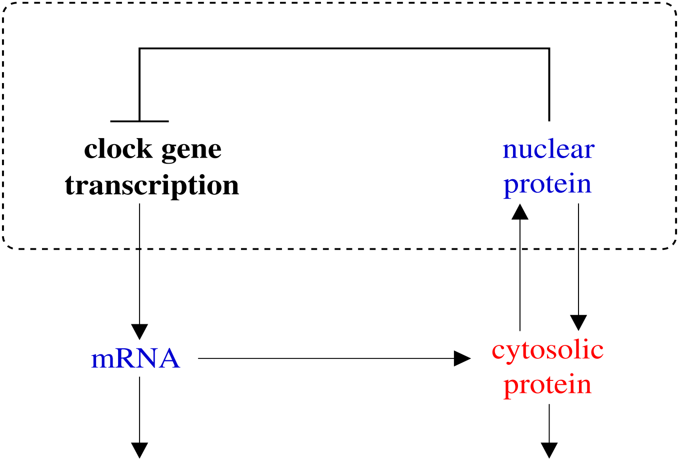

A simple circadian clock

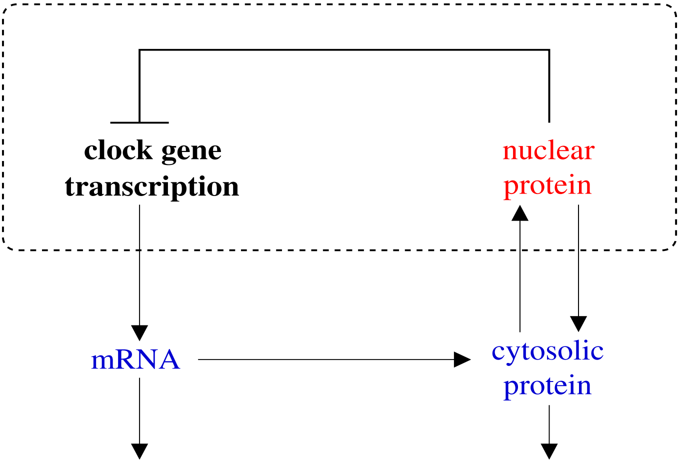

A simple circadian clock

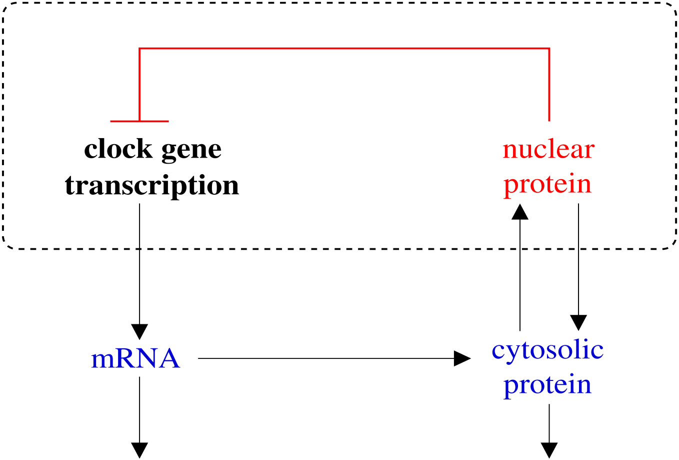

A simple circadian clock

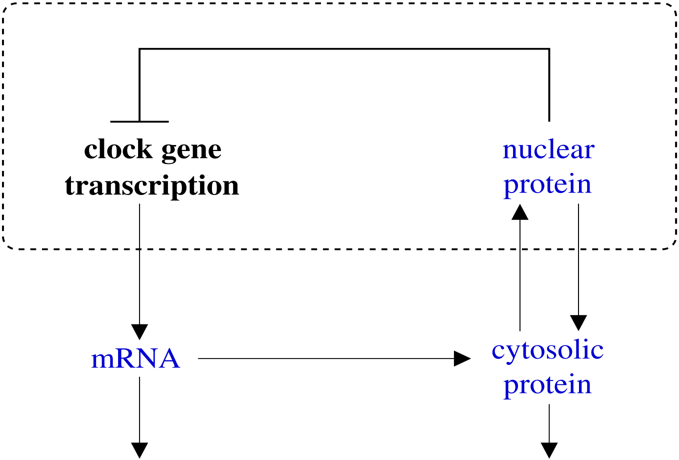

A simple circadian clock

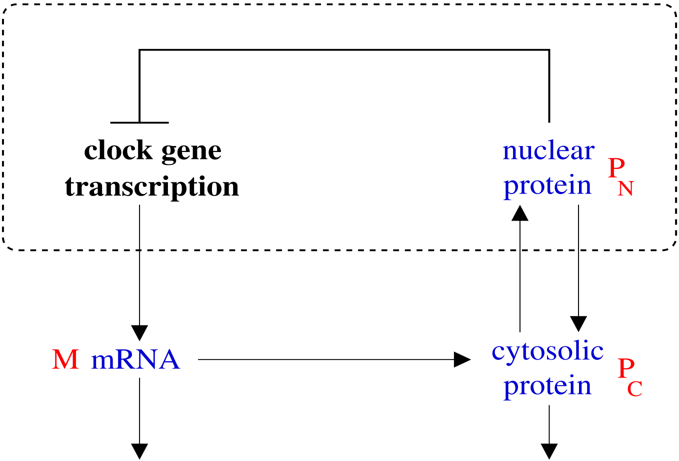

A simple circadian clock

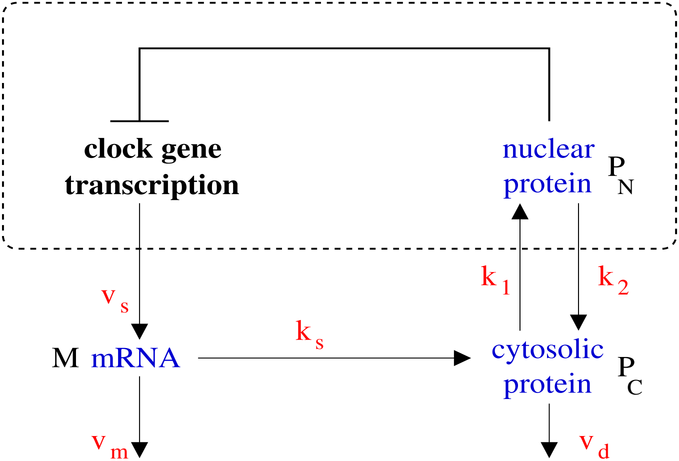

A simple circadian clock

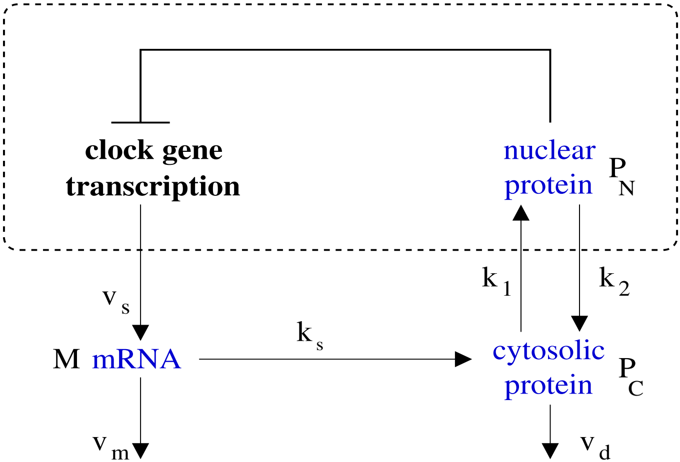

Handcrafted ODEs

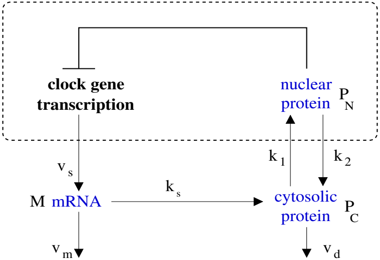

$$\frac{d[M]}{dt} = v_s \frac{k_I^n}{k_I^n + [P_N]^n} - v_m \frac{[M]}{k_m + [M]}$$

$$\frac{d[P_C]}{dt} = k_s[M] - v_d \frac{[P_C]}{k_d + [P_C]} - k_1[P_C] + k_2[P_N]$$

$$\frac{d[P_N]}{dt} = k_1[P_C] - k_2[P_N]$$

Handcrafted ODEs

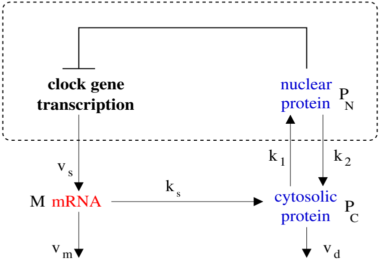

$$\frac{d[M]}{dt} = v_s \frac{k_I^n}{k_I^n + [P_N]^n} - v_m \frac{[M]}{k_m + [M]}$$

$$\frac{d[P_C]}{dt} = k_s[M] - v_d \frac{[P_C]}{k_d + [P_C]} - k_1[P_C] + k_2[P_N]$$

$$\frac{d[P_N]}{dt} = k_1[P_C] - k_2[P_N]$$

Handcrafted ODEs

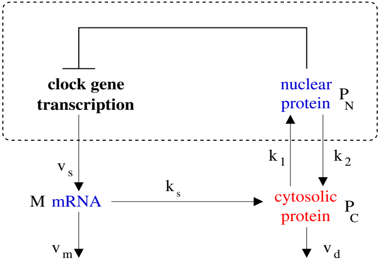

$$\frac{d[M]}{dt} = v_s \frac{k_I^n}{k_I^n + [P_N]^n} - v_m \frac{[M]}{k_m + [M]}$$

$$\frac{d[P_C]}{dt} = k_s[M] - v_d \frac{[P_C]}{k_d + [P_C]} - k_1[P_C] + k_2[P_N]$$

$$\frac{d[P_N]}{dt} = k_1[P_C] - k_2[P_N]$$

Handcrafted ODEs

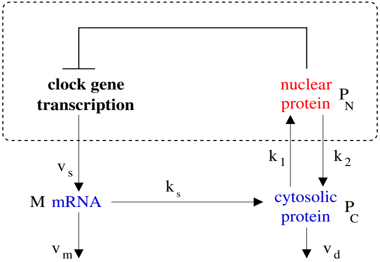

$$\frac{d[M]}{dt} = v_s \frac{k_I^n}{k_I^n + [P_N]^n} - v_m \frac{[M]}{k_m + [M]}$$

$$\frac{d[P_C]}{dt} = k_s[M] - v_d \frac{[P_C]}{k_d + [P_C]} - k_1[P_C] + k_2[P_N]$$

$$\frac{d[P_N]}{dt} = k_1[P_C] - k_2[P_N]$$

Representing the circadian clock in PEPA

- Some of the “steps” in the biological representation (diagram) and the corresponding ODEs do not correspond to elementary reaction steps:

Representing the circadian clock in PEPA

- Some of the “steps” in the biological representation (diagram) and the corresponding ODEs do not correspond to elementary reaction steps:
  - To use our current mappings we need to decompose the enzyme-substrate and gene-repressor reactions into elementary steps.

Representing the circadian clock in PEPA

- Some of the "steps" in the biological representation (diagram) and the corresponding ODEs do not correspond to elementary reaction steps:
  - To use our current mappings we need to decompose the enzyme-substrate and gene-repressor reactions into elementary steps.
- PEPA does not have combinators to express repression or catalysis:

Representing the circadian clock in PEPA

- Some of the “steps” in the biological representation (diagram) and the corresponding ODEs do not correspond to elementary reaction steps:
  - To use our current mappings we need to decompose the enzyme-substrate and gene-repressor reactions into elementary steps.
- PEPA does not have combinators to express repression or catalysis:
  - We introduce additional abstract components to the PEPA model which do not correspond to species but to transcription and repression.

Enzymatic degradation

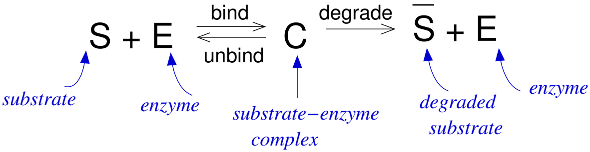

Enzymatic degradation

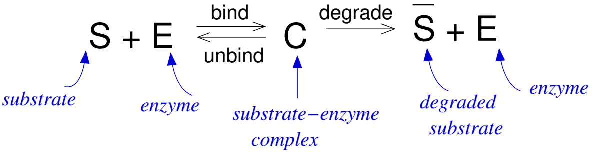

$$S^h \stackrel{def}{=} (b, r_b).S^l \qquad S^l \stackrel{def}{=} (u, r_u).S^h$$

$$E_S^h \stackrel{def}{=} (b, r_b).E_S^l \qquad E_S^l \stackrel{def}{=} (u, r_u).E_S^h$$

$$C_S \stackrel{def}{=} (u, r_u).C_S^l + (c, r_c).C_S^l \qquad C_S^l \stackrel{def}{=} (b, r_b).C_S^h$$

$$\overline{S}^l \stackrel{def}{=} (c, r_c).\overline{S}^h$$

Enzymatic degradation

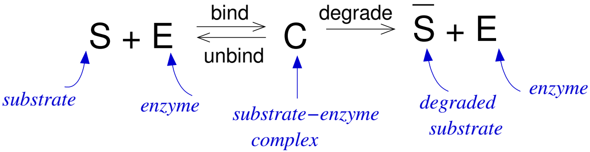

$$
\begin{aligned}
S^h &\stackrel{def}{=} (b, r_b).S^l & \qquad S^l &\stackrel{def}{=} (u, r_u).S^h \\
E_S^h &\stackrel{def}{=} (b, r_b).E_S^l & \qquad E_S^l &\stackrel{def}{=} (u, r_u).E_S^h \\
C_S &\stackrel{def}{=} (u, r_u).C_S^l + (c, r_c).C_S^l & \qquad C_S^l &\stackrel{def}{=} (b, r_b).C_S^h \\
\overline{S}^l &\stackrel{def}{=} (c, r_c).\overline{S}^h & &
\end{aligned}
$$

$$
\left(\left(S^h \underset{\{b,u\}}{\bowtie} E_S^h\right) \underset{\{b,u\}}{\bowtie} C_S^l \right) \underset{\{c\}}{\bowtie} \overline{S}^l
$$

Repression mechanism

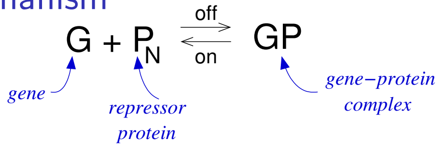

Repression mechanism

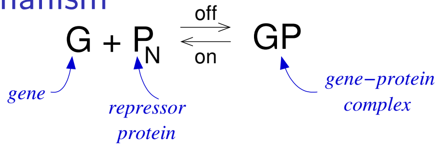

Transcription
$$\begin{cases} T^h \stackrel{def}{=} (transcribe, v_s).T^h + (off, \top).T^l \\ T^l \stackrel{def}{=} (on, \top).T^h \end{cases}$$

Repression
$$\begin{cases} R^h \stackrel{def}{=} (on, v_{on}).R^l \\ R^l \stackrel{def}{=} (off, \top).R^h \end{cases}$$

$$P_N^h \stackrel{def}{=} (off, v_{off}).P_N^l + \dots$$

Repression mechanism

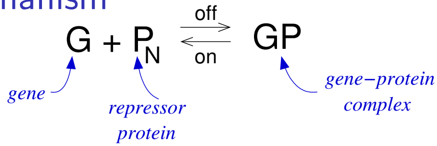

Transcription

$$\begin{cases} T^h \stackrel{def}{=} (transcribe, v_s).T^h + (off, \top).T^l \\ T^l \stackrel{def}{=} (on, \top).T^h \end{cases}$$

Repression

$$\begin{cases} R^h \stackrel{def}{=} (on, v_{on}).R^l \\ R^l \stackrel{def}{=} (off, \top).R^h \end{cases}$$

$$P_N^h \stackrel{def}{=} (off, v_{off}).P_N^l + \ldots$$

Only P_N is explicitly modelled; T and R are abstract entities.

PEPA model of the circadian clock

$$T^h \underset{\{transcribe, on, off\}}{\bowtie} \left(R^l \underset{\{off\}}{\bowtie} \left(\left(M \underset{\{translate\}}{\bowtie} (P_C \underset{\{trans_1, trans_2\}}{\bowtie} P_N)\right)\right)\right)$$

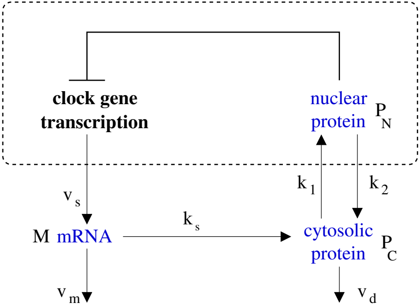

Results of quantitative analysis

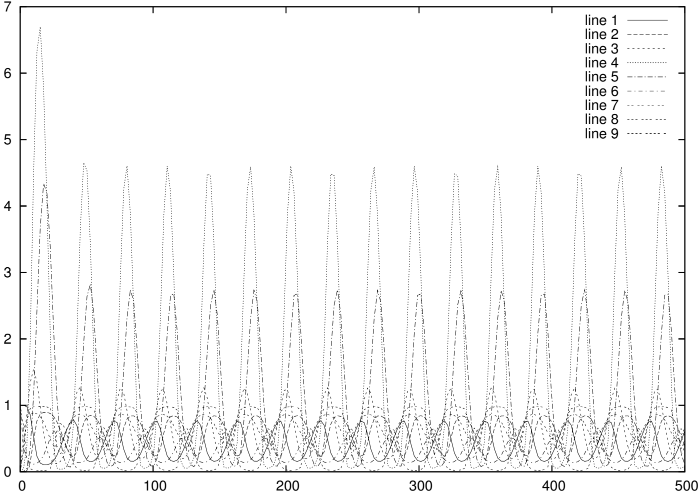

Outline

- Introduction
- Process Algebras for Systems Biology
  - Synthetic Pathway
  - MAPK Pathway
- Other Roles for Abstraction
  - Circadian Clock
- Future Perspectives

Conclusions

- Ultimately we want to understand the functioning of cells at useful levels of abstraction, and to predict unknown behaviour.

Conclusions

- Ultimately we want to understand the functioning of cells at useful levels of abstraction, and to predict unknown behaviour.
- It remains an open and challenging problem to define a set of basic and general primitives for modelling biological systems, inspired by biological processes.

Conclusions

- Ultimately we want to understand the functioning of cells at useful levels of abstraction, and to predict unknown behaviour.
- It remains an open and challenging problem to define a set of basic and general primitives for modelling biological systems, inspired by biological processes.
- Achieving this goal is anticipated to have two broad benefits:

Conclusions

- Ultimately we want to understand the functioning of cells at useful levels of abstraction, and to predict unknown behaviour.
- It remains an open and challenging problem to define a set of basic and general primitives for modelling biological systems, inspired by biological processes.
- Achieving this goal is anticipated to have two broad benefits:
  - Better models and simulations of living phenomena

Conclusions

- Ultimately we want to understand the functioning of cells at useful levels of abstraction, and to predict unknown behaviour.
- It remains an open and challenging problem to define a set of basic and general primitives for modelling biological systems, inspired by biological processes.
- Achieving this goal is anticipated to have two broad benefits:
  - Better models and simulations of living phenomena
  - New models of computations that are biologically inspired.

Conclusions

- Ultimately we want to understand the functioning of cells at useful levels of abstraction, and to predict unknown behaviour.
- It remains an open and challenging problem to define a set of basic and general primitives for modelling biological systems, inspired by biological processes.
- Achieving this goal is anticipated to have two broad benefits:
  - Better models and simulations of living phenomena
  - New models of computations that are biologically inspired.
- Inclusion of quantitative/stochastic elements is essential.

Thank you!