

Integrating Inference with Stochastic Process Algebra Models

Jane Hillston

School of Informatics, University of Edinburgh

11th January 2018



Microsoft
Research

quanti.col

Outline


- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Outline

- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Process Algebra


- Models consist of **agents** which engage in **actions**.



- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.

Process Algebra

- Models consist of **agents** which engage in **actions**.



- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.

Process algebra model $\xrightarrow{\text{SOS rules}}$ Labelled transition system


Stochastic process algebras

Stochastic process algebra

Process algebras where models are decorated with quantitative information used to generate a stochastic process are **stochastic process algebras (SPA)**.

Stochastic Process Algebra

- Models are constructed from **components** which engage in **activities**.




- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.



- The CTMC can be analysed numerically (linear algebra) or by stochastic simulation.

Stochastic Process Algebra

- Models are constructed from **components** which engage in **activities**.




- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.



- The CTMC can be analysed numerically (linear algebra) or by stochastic simulation.

Stochastic Process Algebra

- Models are constructed from **components** which engage in **activities**.



- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.



- The CTMC can be analysed numerically (linear algebra) or by stochastic simulation.

Integrated analysis


Qualitative verification complemented by quantitative verification.

Integrated analysis

Qualitative verification complemented by quantitative verification.

Reachability analysis

How long will it take
for the system to arrive
in a particular state?




Integrated analysis

Qualitative verification complemented by quantitative verification.

Specification matching

With what probability does system behaviour match its specification?




Integrated analysis

Qualitative verification complemented by quantitative verification.

Model checking

Does a given property ϕ
hold within the system
with a given probability?




Integrated analysis

Qualitative verification complemented by quantitative verification.

Model checking

For a given starting state
how long is it until
a given property ϕ holds?



Benefits of integration

- Properties of the underlying mathematical structure may be deduced by the construction at the process algebra level.
- Compositionality can be exploited both for model construction and (in some cases) for model analysis.
- Formal reasoning techniques such as equivalence relations and model checking can be used to manipulate or interrogate models.
- For example the congruence Markovian bisimulation, allows exact model reduction to be carried out compositionally.
- Stochastic model checking based on the Continuous Stochastic Logic (CSL) allows automatic evaluation of quantified properties of the behaviour of the system.

Benefits of integration

- **Properties of the underlying mathematical structure** may be deduced by the construction at the process algebra level.
- **Compositionality** can be exploited both for model construction and (in some cases) for model analysis.
- **Formal reasoning techniques** such as equivalence relations and model checking can be used to manipulate or interrogate models.
- For example the congruence **Markovian bisimulation**, allows exact model reduction to be carried out **compositionally**.
- **Stochastic model checking** based on the Continuous Stochastic Logic (CSL) allows automatic evaluation of **quantified properties** of the behaviour of the system.

Benefits of integration

- **Properties of the underlying mathematical structure** may be deduced by the construction at the process algebra level.
- **Compositionality** can be exploited both for model construction and (in some cases) for model analysis.
- **Formal reasoning techniques** such as equivalence relations and model checking can be used to manipulate or interrogate models.
- For example the congruence **Markovian bisimulation**, allows exact model reduction to be carried out **compositionally**.
- **Stochastic model checking** based on the Continuous Stochastic Logic (CSL) allows automatic evaluation of **quantified properties** of the behaviour of the system.

Benefits of integration

- **Properties of the underlying mathematical structure** may be deduced by the construction at the process algebra level.
- **Compositionality** can be exploited both for model construction and (in some cases) for model analysis.
- **Formal reasoning techniques** such as equivalence relations and model checking can be used to manipulate or interrogate models.
- For example the congruence **Markovian bisimulation**, allows exact model reduction to be carried out **compositionally**.
- **Stochastic model checking** based on the Continuous Stochastic Logic (CSL) allows automatic evaluation of **quantified properties** of the behaviour of the system.

Benefits of integration

- **Properties of the underlying mathematical structure** may be deduced by the construction at the process algebra level.
- **Compositionality** can be exploited both for model construction and (in some cases) for model analysis.
- **Formal reasoning techniques** such as equivalence relations and model checking can be used to manipulate or interrogate models.
- For example the congruence **Markovian bisimulation**, allows exact model reduction to be carried out **compositionally**.
- **Stochastic model checking** based on the Continuous Stochastic Logic (CSL) allows automatic evaluation of **quantified properties** of the behaviour of the system.

Outline

- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Modelling in a Data Rich World

There are many situations in which we wish to model and analyse behaviour of **complex systems**, which are **operational** and generate **data**, but which many not be completely transparent to us.

I will use the example of **systems biology**, where particular biological phenomena are observed and wet lab experiments can typically collect data on some parts of the system, but the basic mechanisms, or the parameters governing their behaviour, are unknown.

Modelling in a Data Rich World

There are many situations in which we wish to model and analyse behaviour of **complex systems**, which are **operational** and generate **data**, but which many not be completely transparent to us.

I will use the example of **systems biology**, where particular biological phenomena are observed and wet lab experiments can typically collect data on some parts of the system, but the basic mechanisms, or the parameters governing their behaviour, are unknown.

Molecular processes as concurrent computations

Concurrency	Molecular Biology	Metabolism	Signal Transduction
Concurrent computational processes	Molecules	Enzymes and metabolites	Interacting proteins
Synchronous communication	Molecular interaction	Binding and catalysis	Binding and catalysis
Transition or mobility	Biochemical modification or relocation	Metabolite synthesis	Protein binding, modification or sequestration

A. Regev and E. Shapiro *Cells as computation*, Nature 419, 2002.

Formal modelling in systems biology

- Formal languages provide a convenient interface for describing complex systems, reflecting what is known about the components and their behaviour.
- High-level abstraction eases writing and manipulating models.
- They are **compiled** into **executable models** which can be run to deepen understanding of the model.
- Formal nature lends itself to automatic, rigorous methods for analysis and verification.
- **Executing** the model generates data that can be compared with biological data.

- ... but what if parts of the system are unknown?

Bio-PEPA modelling

- The **state of the system** at any time consists of the **local states** of each of its “species” components, describing biochemical entities.
- The local states of components are **quantitative** rather than functional, i.e. biological changes to species are represented as distinct components.
- A component varying its state corresponds to it varying its **amount** through reactions modelled as interactions between components.
- The effect of a reaction is to vary the parameter of a component by a number corresponding to the **stoichiometry** of this species in the reaction.

Bio-PEPA modelling

- The **state of the system** at any time consists of the **local states** of each of its “species” components, describing biochemical entities.
- The local states of components are **quantitative** rather than functional, i.e. biological changes to species are represented as distinct components.
- A component varying its state corresponds to it varying its **amount** through reactions modelled as interactions between components.
- The effect of a reaction is to vary the parameter of a component by a number corresponding to the **stoichiometry** of this species in the reaction.

Bio-PEPA modelling

- The **state of the system** at any time consists of the **local states** of each of its “species” components, describing biochemical entities.
- The local states of components are **quantitative** rather than functional, i.e. biological changes to species are represented as distinct components.
- A component varying its state corresponds to it varying its **amount** through reactions modelled as interactions between components.
- The effect of a reaction is to vary the parameter of a component by a number corresponding to the **stoichiometry** of this species in the reaction.

Bio-PEPA modelling


- The **state of the system** at any time consists of the **local states** of each of its “species” components, describing biochemical entities.
- The local states of components are **quantitative** rather than functional, i.e. biological changes to species are represented as distinct components.
- A component varying its state corresponds to it varying its **amount** through reactions modelled as interactions between components.
- The effect of a reaction is to vary the parameter of a component by a number corresponding to the **stoichiometry** of this species in the reaction.

The semantics

The semantics is defined by two transition relations:

- First, a **capability relation** — is a transition possible?;
- Second, a **stochastic relation** — gives rate of a transition, derived from the parameters of the model.

Example



```
k_s = 0.5;
```

```
k_r = 0.1;
```

```
kineticLawOf spread : k_s * I * S;
```

```
kineticLawOf stop1 : k_r * S * S;
```

```
kineticLawOf stop2 : k_r * S * R;
```


```
I = (spread,1) ↓ ;
```

```
S = (spread,1) ↑ + (stop1,1) ↓ + (stop2,1) ↓ ;
```

```
R = (stop1,1) ↑ + (stop2,1) ↑ ;
```

```
I[10] ⊗ S[5] ⊗ R[0]
```

Example



$$k_s = 0.5;$$

$$k_r = 0.1;$$

$$\text{kineticLawOf spread} : k_s * I * S;$$

$$\text{kineticLawOf stop1} : k_r * S * S;$$

$$\text{kineticLawOf stop2} : k_r * S * R;$$

$$I = (\text{spread}, 1) \downarrow ;$$

$$S = (\text{spread}, 1) \uparrow + (\text{stop1}, 1) \downarrow + (\text{stop2}, 1) \downarrow ;$$


$$R = (\text{stop1}, 1) \uparrow + (\text{stop2}, 1) \uparrow ;$$

$$I[10] \quad \boxtimes \quad S[5] \quad \boxtimes \quad R[0]$$


* *

Optimizing models

Usual process of parameterising a model is iterative and manual.




Alternative perspective



Model creation is data-driven

Machine Learning: Bayesian statistics



- Represent belief and uncertainty as probability distributions (prior, posterior).
- Treat parameters and unobserved variables similarly.
- Bayes' Theorem:

$$P(\theta | D) = \frac{P(\theta) \cdot P(D | \theta)}{P(D)}$$

posterior \propto prior \cdot likelihood

Modelling

Thus there are two approaches to model construction:

Machine Learning: extracting a model from the data generated by the system, or refining a model based on system behaviour using statistical techniques.

Mechanistic Modelling: starting from a description or hypothesis, construct a formal model that algorithmically mimics the behaviour of the system, validated against data.

Comparing the techniques

Data-driven modelling:

- + rigorous handling of parameter uncertainty
- limited or no treatment of stochasticity
- in many cases bespoke solutions are required which can limit the size of system which can be handled

Mechanistic modelling:

- + general execution "engine" (deterministic or stochastic) can be reused for many models
- + models can be used speculatively to investigate roles of parameters, or alternative hypotheses
- parameters are assumed to be known and fixed, or costly approaches must be used to seek appropriate parameterisation

Comparing the techniques

Data-driven modelling:

- + rigorous handling of parameter uncertainty
- limited or no treatment of stochasticity
- in many cases bespoke solutions are required which can limit the size of system which can be handled

Mechanistic modelling:

- + general execution "engine" (deterministic or stochastic) can be reused for many models
- + models can be used speculatively to investigate roles of parameters, or alternative hypotheses
- parameters are assumed to be known and fixed, or costly approaches must be used to seek appropriate parameterisation

Developing a probabilistic programming approach

What if we could...

- include information about uncertainty in the model?
- automatically use observations to refine this uncertainty?
- do all this in a formal context?

Starting from the existing process algebra (Bio-PEPA), we have developed a new language **ProPPA** that addresses these issues.

Outline

- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Probabilistic programming

- A programming paradigm for describing incomplete knowledge scenarios, and resolving the uncertainty.
- Programs are probabilistic models in a **high level language**, like software code.
- Offers **automated** inference without the need to write bespoke solutions.
- Platforms: IBAL, Church, Infer.NET, Fun, Anglican, WebPPL,....
- Key actions: **specify** a distribution, specify **observations**, **infer** posterior distribution.

Probabilistic programming workflow

- Describe how the data is generated in syntax like a conventional programming language, but leaving some variables uncertain.
- Specify observations, which impose constraints on acceptable outputs of the program.
- Run program forwards: Generate data consistent with observations.
- Run program backwards: Find values for the uncertain variables which make the output match the observations.

Probabilistic programming workflow

- Describe how the data is generated in syntax like a conventional programming language, but leaving some variables uncertain.
- Specify observations, which impose constraints on acceptable outputs of the program.
- Run program forwards: Generate data consistent with observations.
- Run program backwards: Find values for the uncertain variables which make the output match the observations.

Probabilistic programming workflow

- Describe how the data is generated in syntax like a conventional programming language, but leaving some variables uncertain.
- Specify observations, which impose constraints on acceptable outputs of the program.
- Run program forwards: Generate data consistent with observations.
- Run program backwards: Find values for the uncertain variables which make the output match the observations.

Probabilistic programming workflow

- Describe how the data is generated in syntax like a conventional programming language, but leaving some variables uncertain.
- Specify observations, which impose constraints on acceptable outputs of the program.
- Run program forwards: Generate data consistent with observations.
- Run program backwards: Find values for the uncertain variables which make the output match the observations.

A Probabilistic Programming Process Algebra: ProPPA

The objective of ProPPA is to retain the features of the stochastic process algebra:

- simple model description in terms of components
- rigorous semantics giving an executable version of the model...

... whilst also incorporating features of a probabilistic programming language:

- recording uncertainty in the parameters
- ability to incorporate observations into models
- access to inference to update uncertainty based on observations

A Probabilistic Programming Process Algebra: ProPPA


The objective of ProPPA is to retain the features of the stochastic process algebra:

- simple model description in terms of components
- rigorous semantics giving an executable version of the model...

... whilst also incorporating features of a probabilistic programming language:

- recording uncertainty in the parameters
- ability to incorporate observations into models
- access to inference to update uncertainty based on observations

Example Revisited



$$k_s = 0.5;$$

$$k_r = 0.1;$$

$$\text{kineticLawOf spread} : k_s * I * S;$$

$$\text{kineticLawOf stop1} : k_r * S * S;$$

$$\text{kineticLawOf stop2} : k_r * S * R;$$

$$I = (\text{spread}, 1) \downarrow ;$$

$$S = (\text{spread}, 1) \uparrow + (\text{stop1}, 1) \downarrow + (\text{stop2}, 1) \downarrow ;$$

$$R = (\text{stop1}, 1) \uparrow + (\text{stop2}, 1) \uparrow ;$$

$$I[10] \quad \boxtimes \quad S[5] \quad \boxtimes \quad R[0]$$

* *

Additions

Declaring uncertain parameters:

- `k_s = Uniform(0,1);`
- `k_t = Uniform(0,1);`


Providing observations:

- `observe('trace')`

Specifying inference approach:

- `infer('ABC')`

Additions



```

k_s = Uniform(0,1);
k_r = Uniform(0,1);

```

```

kineticLawOf spread : k_s * I * S;
kineticLawOf stop1 : k_r * S * S;
kineticLawOf stop2 : k_r * S * R;

```

```

I = (spread,1) ↓ ;
S = (spread,1) ↑ + (stop1,1) ↓ + (stop2,1) ↓ ;
R = (stop1,1) ↑ + (stop2,1) ↑ ;

```

```

I[10] ⊗ S[5] ⊗ R[0]

```


```

observe('trace')
infer('ABC') //Approximate Bayesian Computation


```

Semantics


- A Bio-PEPA model can be interpreted as a CTMC; however, CTMCs cannot capture uncertainty in the rates (every transition must have a concrete rate).
- ProPPA models include uncertainty in the parameters, which translates into uncertainty in the transition rates.
- A ProPPA model should be mapped to something like a distribution over CTMCs.







model




CTMC







model



set
of CTMCs



model



μ

distribution
over CTMCs

Constraint Markov Chains

Constraint Markov Chains (CMCs) are a generalization of DTMCs, in which the transition probabilities are not concrete, but can take any value satisfying some constraints.

Constraint Markov Chain

A CMC is a tuple $\langle S, o, A, V, \phi \rangle$, where:

- S is the set of states, of cardinality k .
- $o \in S$ is the initial state.
- A is a set of atomic propositions.
- $V : S \rightarrow 2^{2^A}$ gives a set of acceptable labellings for each state.
- $\phi : S \times [0, 1]^k \rightarrow \{0, 1\}$ is the **constraint function**.



Constraint Markov Chains

In a CMC, arbitrary constraints are permitted, expressed through the function ϕ : $\phi(s, \vec{p}) = 1$ iff \vec{p} is an acceptable vector of transition probabilities from state s .

However,

- CMCs are defined only for the discrete-time case, and
- this does not say anything about **how likely** a value is to be chosen, only about **whether** it is acceptable.

To address these shortcomings, we define **Probabilistic Constraint Markov Chains**.



Probabilistic CMCs

A **Probabilistic Constraint Markov Chain** is a tuple $\langle S, o, A, V, \phi \rangle$, where:

- S is the set of states, of cardinality k .
 - $o \in S$ is the initial state.
 - A is a set of atomic propositions.
 - $V : S \rightarrow 2^{2^A}$ gives a set of acceptable labellings for each state.
 - $\phi : S \times [0, \infty)^k \rightarrow [0, \infty)$ is the **constraint function**.
-
- This is applicable to continuous-time systems.
 - $\phi(s, \cdot)$ is now a probability density function on the transition rates from state s .

Semantics of ProPPA

The semantics definition follows that of Bio-PEPA, which is defined using two transition relations:

- Capability relation — is a transition possible?
- Stochastic relation — gives **distribution of the rate** of a transition

The distribution over the parameter values induces a distribution over transition rates.

Rules are expressed as state-to-function transition systems (FuTS¹).

This gives rise the **underlying PCMC**.

¹De Nicola *et al.*, *A Uniform Definition of Stochastic Process Calculi*, ACM Computing Surveys, 2013

Semantics of ProPPA

The semantics definition follows that of Bio-PEPA, which is defined using two transition relations:

- Capability relation — is a transition possible?
- Stochastic relation — gives **distribution of the rate** of a transition

The distribution over the parameter values induces a distribution over transition rates.

Rules are expressed as state-to-function transition systems (FuTS¹).

This gives rise the **underlying PCMC**.

¹De Nicola *et al.*, *A Uniform Definition of Stochastic Process Calculi*, ACM Computing Surveys, 2013

Simulating Probabilistic Constraint Markov Chains

Probabilistic Constraint Markov Chains are open to two alternative dynamic interpretations:

- 1 Uncertain Markov Chains:** For each trajectory, for each uncertain transition rate, sample once at the start of the run and use that value throughout;
- 2 Imprecise Markov Chains:** During each trajectory, each time a transition with an uncertain rate is encountered, sample a value but then discard it and re-sample whenever this transition is visited again.

Our current work is focused on the **Uncertain Markov Chain** case.

Simulating Probabilistic Constraint Markov Chains

Probabilistic Constraint Markov Chains are open to two alternative dynamic interpretations:

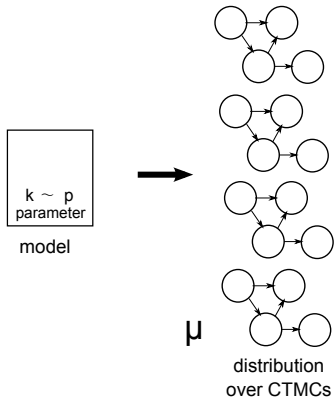
- 1 Uncertain Markov Chains:** For each trajectory, for each uncertain transition rate, sample once at the start of the run and use that value throughout;
- 2 Imprecise Markov Chains:** During each trajectory, each time a transition with an uncertain rate is encountered, sample a value but then discard it and re-sample whenever this transition is visited again.

Our current work is focused on the **Uncertain Markov Chain** case.


Outline

- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Inference



Inference



Inference

$$P(\theta | D) \propto P(\theta)P(D | \theta)$$

- Exact inference is impossible, as we cannot calculate the likelihood.
- We must use approximate algorithms or approximations of the system.
- The ProPPA semantics does not define a single inference algorithm, allowing for a [modular approach](#).
- Different algorithms can act on different input (time-series vs properties), return different results or in different forms.

Inferring likelihood in uncertain CTMCs

Transient state probabilities can be expressed as:

$$\frac{dp_i(t)}{dt} = \sum_{j \neq i} p_j(t) \cdot q_{ji} - p_i(t) \sum_{j \neq i} q_{ij}$$

The probability of a single observation (y, t) can then be expressed as

$$p(y, t) = \sum_{i \in \mathcal{S}} p_i(t) \pi(y | i)$$

where $\pi(y | i)$ is the probability of observing y when in state i .

The likelihood can then be expressed as

$$P(D | \theta) = \prod_{j=1}^N \sum_{i \in \mathcal{S}} p_{(i|\theta)}(t_j) \pi(y_j | i)$$

Inferring likelihood in uncertain CTMCs

Transient state probabilities can be expressed as:

$$\frac{dp_i(t)}{dt} = \sum_{j \neq i} p_j(t) \cdot q_{ji} - p_i(t) \sum_{j \neq i} q_{ij}$$

The probability of a single observation (y, t) can then be expressed as

$$p(y, t) = \sum_{i \in \mathcal{S}} p_i(t) \pi(y | i)$$

where $\pi(y | i)$ is the probability of observing y when in state i .

The likelihood can then be expressed as

$$P(D | \theta) = \prod_{j=1}^N \sum_{i \in \mathcal{S}} p_{(i|\theta)}(t_j) \pi(y_j | i)$$

Inferring likelihood in uncertain CTMCs

Transient state probabilities can be expressed as:

$$\frac{dp_i(t)}{dt} = \sum_{j \neq i} p_j(t) \cdot q_{ji} - p_i(t) \sum_{j \neq i} q_{ij}$$

The probability of a single observation (y, t) can then be expressed as

$$p(y, t) = \sum_{i \in \mathcal{S}} p_i(t) \pi(y | i)$$

where $\pi(y | i)$ is the probability of observing y when in state i .

The likelihood can then be expressed as

$$P(D | \theta) = \prod_{j=1}^N \sum_{i \in \mathcal{S}} p_{(i|\theta)}(t_j) \pi(y_j | i)$$

Calculating the transient probabilities

For finite state-spaces, the transient probabilities can, in principle, be computed as

$$\mathbf{p}(t) = \mathbf{p}(0)e^{\mathbf{Q}t}.$$

Likelihood is hard to compute:


- Computing $e^{\mathbf{Q}t}$ is expensive if the state space is large
- Impossible directly in infinite state-spaces

Basic Inference

- Approximate Bayesian Computation is a simple simulation-based solution:
 - Approximates posterior distribution over parameters as a set of samples
 - Likelihood of parameters is approximated with a notion of distance.


Basic Inference

- Approximate Bayesian Computation is a simple simulation-based solution:
 - Approximates posterior distribution over parameters as a set of samples
 - Likelihood of parameters is approximated with a notion of distance.




Basic Inference

- Approximate Bayesian Computation is a simple simulation-based solution:
 - Approximates posterior distribution over parameters as a set of samples
 - Likelihood of parameters is approximated with a notion of distance.




Basic Inference

- Approximate Bayesian Computation is a simple simulation-based solution:
 - Approximates posterior distribution over parameters as a set of samples
 - Likelihood of parameters is approximated with a notion of distance.




Basic Inference

- Approximate Bayesian Computation is a simple simulation-based solution:
 - Approximates posterior distribution over parameters as a set of samples
 - Likelihood of parameters is approximated with a notion of distance.



Basic Inference

- Approximate Bayesian Computation is a simple simulation-based solution:
 - Approximates posterior distribution over parameters as a set of samples
 - Likelihood of parameters is approximated with a notion of distance.



Approximate Bayesian Computation

ABC algorithm

- 1 Sample a parameter set from the prior distribution.
- 2 Simulate the system using these parameters.
- 3 Compare the simulation trace obtained with the observations.
- 4 If distance $< \epsilon$, accept, otherwise reject.

This results in an approximation to the posterior distribution.

As $\epsilon \rightarrow 0$, set of samples converges to true posterior.

We use a more elaborate version based on Markov Chain Monte Carlo sampling.

Inference for infinite state spaces

Various methods become inefficient or inapplicable as the state-space grows.

How to deal with unbounded systems?

- Multiple simulation runs
- Large population approximations (diffusion, Linear Noise, . . .)
- Systematic truncation
- **Random truncations**

Expanding the likelihood

The likelihood can be written as an infinite series:

$$\begin{aligned} p(x', t' | x, t) &= \sum_{N=0}^{\infty} p^{(N)}(x', t' | x, t) \\ &= \sum_{N=0}^{\infty} \left[f^{(N)}(x', t' | x, t) - f^{(N-1)}(x', t' | x, t) \right] \end{aligned}$$

where

- $x^* = \max\{x, x'\}$
- $p^{(N)}(x', t' | x, t)$ is the probability of going from state x at time t to state x' at time t' through a path with maximum state $x^* + N$
- $f^{(N)}$ is the same, except the maximum state cannot exceed $x^* + N$ (but does not have to reach it)

Using Russian Roulette truncation we can estimate the infinite sum with a random truncation.

Expanding the likelihood


The likelihood can be written as an infinite series:


$$\begin{aligned} p(x', t' | x, t) &= \sum_{N=0}^{\infty} p^{(N)}(x', t' | x, t) \\ &= \sum_{N=0}^{\infty} \left[f^{(N)}(x', t' | x, t) - f^{(N-1)}(x', t' | x, t) \right] \end{aligned}$$


where

- $x^* = \max\{x, x'\}$
- $p^{(N)}(x', t' | x, t)$ is the probability of going from state x at time t to state x' at time t' through a path with maximum state $x^* + N$
- $f^{(N)}$ is the same, except the maximum state cannot exceed $x^* + N$ (but does not have to reach it)

Using Russian Roulette truncation we can estimate the infinite sum with a random truncation.








Outline

- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Example model



```

k_s = Uniform(0,1);
k_r = Uniform(0,1);

```

```

kineticLawOf spread : k_s * I * S;
kineticLawOf stop1  : k_r * S * S;
kineticLawOf stop2  : k_r * S * R;

```

```

I = (spread,1) ↓ ;
S = (spread,1) ↑ + (stop1,1) ↓ + (stop2,1) ↓ ;
R = (stop1,1) ↑ + (stop2,1) ↑ ;

```

```

I[10] ⊗ S[5] ⊗ R[0]

```

```

observe('trace')
infer('ABC') //Approximate Bayesian Computation


```

Results


Tested on the rumour-spreading example, giving the two parameters uniform priors.

- Approximate Bayesian Computation
- Returns posterior as a set of points (samples)
- Observations: time-series (single simulation)



Results: ABC



Results: ABC



Results: ABC



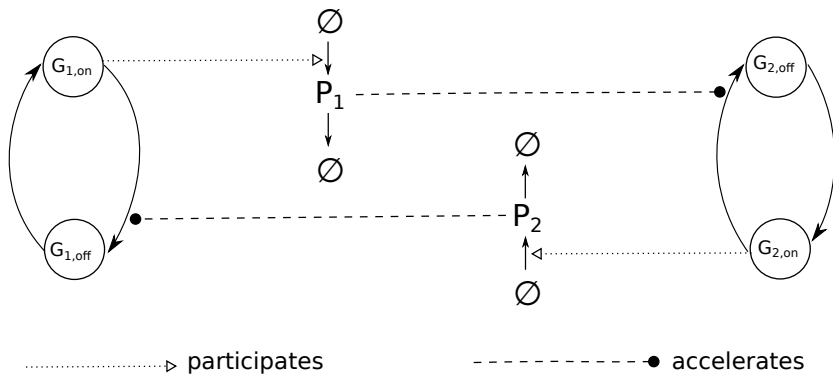
Genetic Toggle Switch

- Two mutually-repressing genes: promoters (unobserved) and their protein products
- Bistable behaviour: switching induced by environmental changes
- Synthesised in *E. coli*²
- Stochastic variant³ where switching is induced by noise

²Gardner, Cantor & Collins, *Construction of a genetic toggle switch in Escherichia coli*, Nature, 2000

³Tian & Burrage, *Stochastic models for regulatory networks of the genetic toggle switch*, PNAS, 2006

Genetic Toggle Switch



Toggle switch model: species

```
G1 = activ1 ↑ + deact1 ↓ + expr1 ⊕;
```

```
G2 = activ2 ↑ + deact2 ↓ + expr2 ⊕;
```

```
P1 = expr1 ↑ + degr1 ↓ + deact2 ⊕ ;
```

```
P2 = expr2 ↑ + degr2 ↓ + deact1 ⊕
```

```
G1[1] <*> G2[0] <*> P1[20] <*> P2[0]
```

```
observe(toggle_obs);
```


```
infer(rouletteGibbs);
```

```
 $\theta_1 = \text{Gamma}(3,5); //\text{etc...}$   
  
kineticLawOf expr1 :  $\theta_1 * G1$ ;  
kineticLawOf expr2 :  $\theta_2 * G2$ ;  
kineticLawOf degr1 :  $\theta_3 * P1$ ;  
kineticLawOf degr2 :  $\theta_4 * P2$ ;  
  
kineticLawOf activ1 :  $\theta_5 * (1 - G1)$ ;  
kineticLawOf activ2 :  $\theta_6 * (1 - G2)$ ;  
kineticLawOf deact1 :  $\theta_7 * \exp(r * P2) * G1$ ;  
kineticLawOf deact2 :  $\theta_8 * \exp(r * P1) * G2$ ;  
  
G1 = activ1  $\uparrow$  + deact1  $\downarrow$  + expr1  $\oplus$ ;  
G2 = activ2  $\uparrow$  + deact2  $\downarrow$  + expr2  $\oplus$ ;  
P1 = expr1  $\uparrow$  + degr1  $\downarrow$  + deact2  $\oplus$  ;  
P2 = expr2  $\uparrow$  + degr2  $\downarrow$  + deact1  $\oplus$   
  
G1[1] <*> G2[0] <*> P1[20] <*> P2[0]  
  
observe(toggle_obs);  
infer(rouletteGibbs);
```


Experiment

- Simulated observations
- Gamma priors on all parameters (required by algorithm)
- Goal: learn posterior of 8 parameters
- 5000 samples taken using the Gibbs-like random truncation algorithm


Genes (unobserved)




Proteins



Observations used



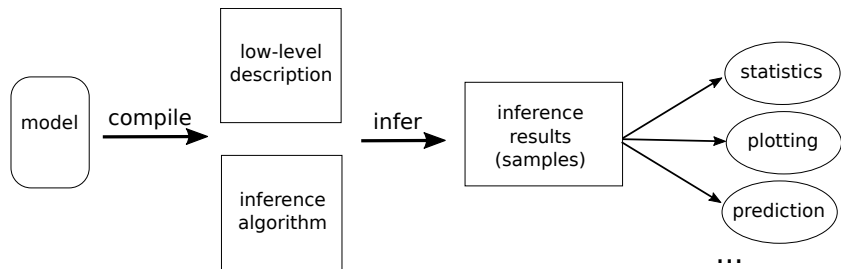
Results



Outline

- 1 Stochastic Process Algebras
- 2 Modelling in a Data Rich World
- 3 ProPPA
- 4 Inference
- 5 Results
- 6 Conclusions

Workflow

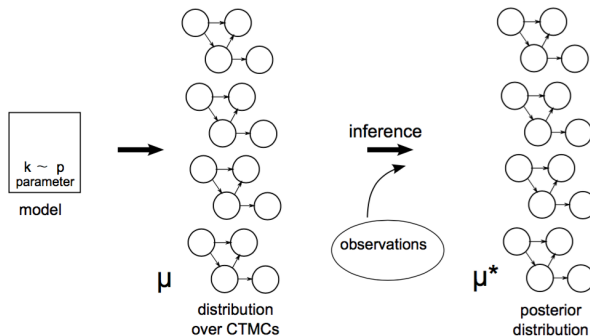


Summary

- ProPPA is a process algebra that incorporates uncertainty and observations directly in the model, influenced by probabilistic programming.
- Syntax remains similar to Bio-PEPA.
- Semantics defined in terms of an extension of Constraint Markov Chains.
- Observations can be either time-series or logical properties.
- Parameter inference based on random truncations (Russian Roulette) offers new possibilities for inference.

Challenges and Future Directions

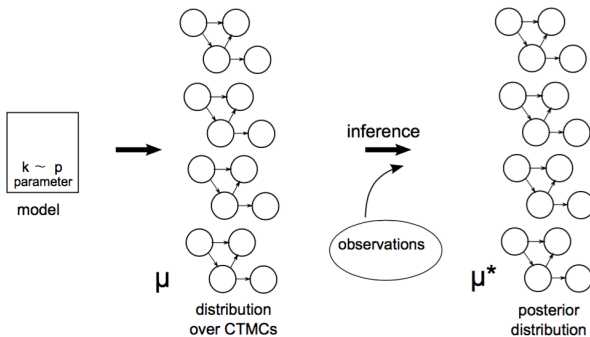
■ The value of observations



Can we reason about the “distance” between μ and μ^* ?

Challenges and Future Directions

■ Heterogeneous populations



What if we are seeking the “optimal mix” rather than the best individual representative?

Thanks

- Anastasis Georgoulas



- Guido Sanguinetti



Funding from ERC, Microsoft Research and CEC FET-Proactive Programme FoCAS.

Thanks

- Anastasis Georgoulas



- Guido Sanguinetti



Funding from ERC, Microsoft Research and CEC FET-Proactive Programme FoCAS.