An Extended Weak Isomorphism for Model Simplification (extended abstract)

Graham Clark*

Department of Computer Science, The University of Edinburgh, Edinburgh, EH9 3JZ, UK.
E-mail: gcla@dcs.ed.ac.uk

Abstract

This paper will describe an extension to Hillston's original weak isomorphism relation for PEPA. Previously it catered for the reduction of sequences of silent activities performed by a process; the stochastic model used was a GSMP which allowed generally distributed activities. We extend the relation to allow for the reduction of more general structure which need not encompass only silent activities. This impacts the calculation of performance measures and these effects are studied. The motivation is that consideration of the performance measure in partnership with the process can allow for useful reductions in model complexity.

1 Introduction

Stochastic process algebras (SPA) are a popular and elegant way of producing performance models for real-life systems. However, when modelling systems of even modest complexity, such models may suffer from the problem of state space explosion. If SPAs are to prove genuinely useful for performance modelling, solutions to this problem must be found.

One approach to dealing with large and complex Markov chains is to exploit the regular structure induced by the compositionality of the SPA model; for example, Hillston and Mertsiotakis [HM95] use the property of *near complete decomposability* where an exact solution of an appropriate Markov chain may be generated by combining the solutions of a set of smaller "sub-models". Our approach differs in that we do not exploit properties of the model alone. Instead we consider that with knowledge of the performance measure of interest, it is possible to simplify the model such that structures which do not affect the performance may be reduced. Indeed it should be possible to reduce structures which do influence the performance measure such that the extent of the influence may be determined. Therefore large models may be reduced to manageable sizes in order to calculate a particular aspect of performance.

2 Background

PEPA models are constructed from components, which are able to interact with each other. Each of these components is capable of performing activities. Formally, an activity

*This work is supported by EPSRC grant 95306547

$a \in \mathcal{A}ct$ is described as a pair (α, r), where α is the type of the activity, and r is the activity rate. The rate r is the parameter of an exponential distribution and thus governs the duration of the activity.

In this paper, we will concentrate on Cyclic PEPA. By definition, a PEPA process P is *cyclic* if it is a derivative of all the components in its derivative set, i.e.

$$Q \in ds(P) \Rightarrow P \in ds(Q)$$

This ensures that behaviour will always be repeated. In general a Cyclic PEPA model may be considered to consist of two layers: a lower layer of sequential components, combined by an upper layer of model components.

In [Hil94], Hillston describes a simple and effective way of calculating performance measures from PEPA models. Values are associated with particular action types; then the reward associated with a particular derivative is the sum of the rewards attached to activities which the derivative enables. Performance measures for a process are then derived from the total reward based on the steady state probability distribution.

2.1 Weak Component Isomorphism

We choose to extend a simpler style of relation than, for instance, the well-studied bisimulation equivalence. *Isomorphism* is a very strong relation; components are considered equivalent if there is a one-to-one correspondence between their derivatives, and these derivatives are capable of performing exactly the same activities. Isomorphism is not a bisimulation style of equivalence, and is very strong—certainly it is a congruence relation. Hillston [Hil94] adapts the isomorphism relation in order to reflect the hidden nature of τ activities. The new equivalence is called *weak isomorphism*.

The motivation behind this equivalence relation is that it provides a model simplification technique. To an outside observer, a τ activity is unobservable in terms of the action type—however a process would still have to expend effort to perform such an activity, and as such the work would take time. Therefore a sequence of identical τ activities would simply take more time to an observer. It therefore seems reasonable to attempt to simplify a process by amalgamating a sequence of τ activities into one, "longer" τ activity. This is essentially the extension to isomorphism captured by weak isomorphism. By reducing these "uninteresting" areas of a performance model, the effort required to solve the model is reduced. Furthermore, since rewards may not be attached to τ activities, Hillston guarantees that although the Markov chains of weakly isomorphic processes may not be equivalent, any performance measures calculated will be identical. In fact the amalgamation of a sequence of exponentially distributed τ activities leads to an activity with a non-exponential rate, and thus this process cannot be expressed in PEPA. However, Hillston defines a mapping to an alternative stochastic model, the *generalised semi-Markov process* (GSMP); by applying insensitivity results the use of a τ activity with an appropriately modified exponentially distributed rate is justified.

2.1.1 Insensitivity

It has been established that for particular GSMPs, an active element with a lifetime determined by a general distribution may be replaced by an active element with an exponential lifetime of the same mean, such that the steady state behaviour of the process is unaffected. A GSMP is said to be *insensitive* if the equilibrium distribution is dependent upon a set of governing distributions only through their means. This means that processes which are identical, except that the lifetime of an insensitive element is determined by different distributions with the same mean, will exhibit the same steady state solution.

Insensitivity is the principle that allows for the model simplification technique outlined earlier. If a GSMP is insensitive to one of its generally distributed active elements, that may be replaced with an active element which is exponentially distributed with the same mean. If this can be done for all generally distributed active elements then the GSMP corresponds simply to a Markov chain with the same steady state solution. This Markov chain can be solved to deduce performance measures from the reduced PEPA model.

3 Extensions to Weak Component Isomorphism

We now illustrate the extensions to be made to the weak isomorphism relation. Weak isomorphism caters for the reduction of what are called *reducible sequences*—in essence, these are sequences of states connected by τ activities, at the derivation graph level. The chain of states and connecting transitions is reduced to two states, and one τ transition with an appropriately modified rate. The performance of the unreduced part of the model is left unaffected. An example of a reducible sequence is given in Figure 1.

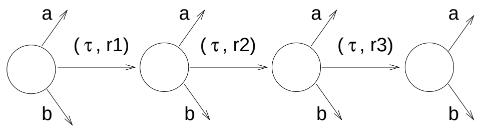

Figure 1: A reducible sequence

The motivation for making this reduction was that τ activities do not correspond to observable computation except that they take time. We extend the allowable reducible structure in order to cater for the simplification of a wider range of PEPA models. The extension is general in that we do not restrict the reduction to states connected by silent τ activities.

A first example would be a process capable of performing a choice between several silent activities, all of which would cause a transition to the same state. According to the operational rules of PEPA, an observer would simply witness less passing of time. However at the moment, the weak isomorphism relation would not capture this obvious extension. This would not lead to a simpler Markov chain which discards activity names and such multiple arcs anyway. However, it suggests much larger, more arbitrary structures which could be reduced. For example, consider the following PEPA process:

$$\begin{aligned} P &\stackrel{\mathrm{def}}{=} (\mathtt{a}, r_1).P' + (\mathtt{d}, r_4).P'' \\ P' &\stackrel{\mathrm{def}}{=} (\mathtt{b}, r_2).P'' + (\mathtt{e}, r_5).Q \\ P'' &\stackrel{\mathrm{def}}{=} (\mathtt{c}, r_3).Q \end{aligned}$$

If this process was to play no direct part in the calculation of performance measures, it should be possible to reduce it in the same style as the reduction of sequences of τ activities. Hillston's weak isomorphism relation is thus extended to cope with more general structure. The motivation is that the relation will give much greater scope for model simplification when particular performance measures are sought.

The relation will be extended more generally to allow for the amalgamation of non-τ activities. Weak isomorphism only caters for the reduction of silent sequences. Associating rewards with such silent activities is disallowed—in this way, the integrity of the final performance measure is maintained. We study the effect when rewards are associated with non-silent activities which are then reduced via the extended weak isomorphism relation. We seek bounds on the performance measure we may obtain from the reduced model.

We work with Cyclic PEPA—certainly it is necessary for the irreducibility of the Markov chain models. In this way we also aim to "rationalise" the weak isomorphism relation and give a syntax directed definition of the new extended relation. We aim for a congruence relation, which weak isomorphism is not.

The efficacy of this work will be proven by demonstrating a translation of a reduced Cyclic PEPA model to a GSMP, following Hillston. It will be illustrated that the GSMP is insensitive to its generally distributed active elements which represent the reduced structure in the PEPA model. Then, provably, these active elements may be replaced by others which have the same mean but are exponentially distributed. This leads to a simple correspondence with a Markov chain. This chain can be solved and performance measures calculated. The proof will be similar in style to that given by Hillston in [Hil93].

The work will be illustrated with an example showing the added benefits of the relation, and finally we will discuss some ideas for future work.

References

[Ber97] M. Bernardo. An Algebra-Based Method to Associate Rewards with EMPA Terms. In *to appear in 24th Int. Colloquium on Automata, Languages and Programming*, July 1997.

[Cla96] G. Clark. Formalising the Specification of Rewards with PEPA. In *Proceedings of the Fourth Process Algebras and Performance Modelling Workshop*, pages 139–160, July 1996.

[Her90] U. Herzog. Formal Description, Time and Performance Analysis: A Framework. Technical Report 15/90, IMMD VII, Friedrich-Alexander-Universität, Erlangen-Nürnberg, Germany, September 1990.

[Hil93] J. Hillston. PEPA: Performance Enhanced Process Algebra. Technical Report CSR-24-93, Department of Computer Science, The University of Edinburgh, 1993.

[Hil94] J. Hillston. *A Compositional Approach to Performance Modelling*. PhD thesis, The University of Edinburgh, 1994. CST-94-108.

[HL89] W. Henderson and D. Lucic. Applications of Generalised Semi Markov Processes to Stochastic Petri Nets. *Performance of Distributed and Parallel Systems*, pages 315–328, 1989.

[HL93] W. Henderson and D. Lucic. Aggregation and Disaggregation Through Insensitivity in Stochastic Petri Nets. *Performance Evaluation*, (17):91–114, 1993.

[HM95] J. Hillston and V. Mertsiotakis. A Simple Time Scale Decomposition Technique for Stochastic Process Algebras. *The Computer Journal*, 38(7):566–577, 1995.

[HP93] P. G. Harrison and N. M. Patel. *Performance Modelling of Communication Networks and Computer Architectures*. International Computer Science Series. Addison-Wesley, 1993.

[LS91] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, September 1991.

[Mil89] R. Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice Hall, 2nd edition, 1989.