Chapter 8
GPA:
A Multiformalism, Multisolution Approach to Efficient Analysis of Large-Scale Population Models

Jeremy T. Bradley
Imperial College London, UK

Marcel C. Guenther
Imperial College London, UK

Richard A. Hayden
Imperial College London, UK

Anton Stefanek
Imperial College London, UK

ABSTRACT

This chapter discusses the latest trends and developments in performance analysis research of large population models. In particular, it reviews GPA, a state-of-the-art Multiformalism, Multisolution (MFMS) tool that provides a framework for the implementation of various population modelling formalisms and solution methods.

1 INTRODUCTION

A decade ago Sanders (Sanders 1999) noted that in spite of research advances in the performance modelling and analysis field, the work of performance analysts has by no means become easier since both the expectation in their work as well as the complexity of systems to be evaluated have also grown considerably. To equip performance analysts with expressive modelling formalisms and efficient solution methods, Sanders suggested a multiformalism, multisolution (MFMS) paradigm that would not only allow modellers to model systems in a natural composite manner using different formalisms in the same model, but also provide them with different analysis options for such heterogeneous composite models.

While MFMS research and tools development have led to a better range of software products for performance analysts, there are still quite a lot of open research challenges in the performance

DOI: 10.4018/978-1-4666-4659-9.ch008

community with regards to improving formalisms and their solution techniques. Population models, which are the focus of this chapter, are one such area that has recently received a lot of attention in the literature due to the growing need to model and analyse crowd behaviour (Massink et al. 2011), biological systems (Ciochetta & Hillston 2009) as well as large, distributed communication systems with thousands of network participants (Stefanek et al. 2010). The challenge in analysing these models is the fact that the state-space of population models increases exponentially in the number of interacting individuals/agents, making it computationally expensive to solve such models using traditional Monte Carlo simulation techniques. Moreover, even moderate population models with only a few hundred agents often exceed the capabilities of traditional numerical state-space avoidance and largeness tolerance solution methods. As a consequence, novel mean-field analysable formalisms have been developed, which can efficiently handle models with large populations. The efficient mean-field/fluid analysis techniques look at the models from a macroscopic point of view instead of treating every component at an individual level. By aggregating the behaviour of individual components, it is often possible to derive a set of Ordinary Differential Equations (ODEs) whose solution expresses the evolution of probabilistic measures such as the means and higher order moments of populations.

Grouped PEPA Analyser (GPA) (Stefanek et al. 2010) is an advanced software solution for population modelling and mean-field analysis. Originally developed for the analysis of the Grouped PEPA process algebra (GPEPA) (Hayden & Bradley 2010, Hillston 2005) in 2009, the tool has since been extended to support a range of different population modelling formalisms and solution methods. As this change continues, the tool is slowly embracing more and more MFMS principles, with a strong focus on population modelling formalisms. In this chapter we give an overview of mean-field analysable population models and describe how the architecture of GPA facilitates the implementation of new formalisms and solution techniques for such models. The chapter is organised as follows; In Section 2 we define population modelling, introduce different classes of mean-field analysable formalisms in Section 2.1 and formally describe a Population CTMCs, the central intermediate representation used by GPA, in Section 2.2. In addition to this, Section 2.3 reviews related MFMS and population modelling tools. Section 3 and 4 describe the architecture of GPA and show how different formalisms and solution techniques were implemented. Finally, Section 5 discusses future extensions for GPA and we present our conclusions in Section 6.

2 BACKGROUND

Population models describe interactions between individuals, which are grouped into populations. Individuals can represent a number of different entities or agents such as people, telecommunication equipment or vehicles to name but a few. While the individual behaviour of agents can be described using a small set of rules, the simulation of population models becomes infeasible when looking at the interaction of thousands or millions of individuals. However, when grouping a large number of individuals into populations, it is possible to evaluate the effects of interactions using efficient mean-field analysis techniques rather than simulation.

2.1 Mean-Field Analysable Population Models

Before we discuss the population model classes that are currently supported by GPA, we would like to give a brief overview of different population model classes and measures that can be calculated with mean-field fluid techniques. To distinguish between different classes of population models, we consider different possibilities of

Related Content