MODELLING WITH GENERALIZED STOCHASTIC PETRI NETS

M. Ajmone Marsan, G. Balbo, G. Conte,
S. Donatelli and G. Franceschinis

WILEY SERIES IN PARALLEL COMPUTING

Modelling with Generalized Stochastic Petri Nets

M. Ajmone Marsan, *Politecnico di Torino, Italy*
G. Balbo, *Università di Torino, Italy*
G. Conte, *Università di Parma, Italy*
S. Donatelli, *Università di Torino, Italy*
G. Franceschinis, *Università di Torino, Italy*

JOHN WILEY & SONS
Chichester · New York · Brisbane · Toronto · Singapore

Baffins Lane, Chichester
West Sussex PO19 1UD, England

National Chichester (01243) 779777
International (+44) 1243 779777

Other Wiley Editorial Offices

John Wiley & Sons, Inc., 605 Third Avenue,
New York, NY 10158-0012, USA

Jacaranda Wiley Ltd, 33 Park Road, Milton,
Queensland 4064, Australia

John Wiley & Sons (Canada) Ltd, 22 Worcester Road,
Rexdale, Ontario M9W 1L1, Canada

John Wiley & Sons (SEA) Pte Ltd, 37 Jalan Pemimpin #05-04,
Block B, Union Industrial Building, Singapore 2057

Library of Congress Cataloging-in-Publication Data

Modelling with generalized stochastic Petri nets / M. Ajmone Marsan ...[et al.].
p. cm. — (Wiley series in parallel computing)
Includes bibliographical references and index.
ISBN 0 471 93059 8
1. Electronic data processing—Distributed processing. 2. Petri nets. I. Ajmone Marsan, M. II. Series.
QA76.9.D5M64 1995
004'.36—dc20 94-20916
CIP

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0 471 93059 8

Produced from camera-ready copy supplied by the authors using LaTeX
Printed and bound in Great Britain by Bookcraft (Bath) Ltd

This book is printed on acid-free paper responsibly manufactured from sustainable forestation, for which at least two trees are planted for each one used for paper production.

Modelling with Generalized Stochastic Petri Nets

WILEY SERIES IN PARALLEL COMPUTING

SERIES EDITORS:

Preface

The goal of this book is to introduce readers to the use of Generalized Stochastic Petri Nets (GSPNs) for the performance analysis of distributed systems of different natures. We hope that a unified presentation of GSPN theory, together with a set of illustrative examples from different application fields, will further extend the success that GSPNs have obtained in recent years as a modelling paradigm for the quantitative analysis of distributed systems.

The book is addressed to beginners, rather than experts in the field of Petri Net applications to performance analysis, and aims to be a useful tool for a first approach to the GSPN field, and a support for a graduate course on modelling and performance evaluation of computer and communication systems.

For this reason, the book is written in an informal style, trying to always explain concepts in plain language before giving a formal definition (often even omitting formal definitions that are not strictly necessary), and including as many examples as possible, to visualize the concepts being explained.

Before a short outlook on the book's content, some historical remarks and acknowledgments are necessary.

The origin of GSPNs goes back to the early 80s, and is due to the cooperation of three of the authors, who participated in the research group that was organized in the city of Torino to join a national research project on the design, implementation and performance evaluation of multiprocessor systems.

Gianni Conte was the initiator of the group: he was the one with the most experience in multiprocessor architectures, and he had known about the possibility of adopting the Petri net paradigm for the description and analysis of such systems from the papers of L. Dadda. Gianfranco Balbo was the one with the most experience in the performance analysis field: he had just obtained his PhD in Computer Science from Purdue University with a thesis on queuing networks and their applications to the performance analysis of computer systems. Marco Ajmone Marsan had recently returned from UCLA, where he had met Mike Molloy, and had learnt about stochastic Petri nets (SPNs), and their usefulness. Actually, the idea of introducing immediate

transitions into the SPN paradigm was discussed several times along the corridors of Boelter Hall at UCLA between the two.

The original description of GSPNs was submitted to the 1982 European Workshop on Petri Net Applications and Theory in Varenna, but was rejected with the longest list of negative comments that we ever saw! In fact, the original version of the theory was not quite clean, but after a number of changes, a revised GSPN presentation was submitted to SIGMETRICS'83, and selected as one of the two best papers of the conference, deserving publication in the ACM Transactions on Computer Systems. The official birth date for GSPNs is thus the May 1984 issue of the ACM TOCS.

The first students to write a thesis on GSPN theory and applications were Giovanni Chiola at the Electronics Department of Politecnico di Torino and Gianfranco Ciardo at the Computer Science Department of the University of Torino. They were also the first to implement the software for a GSPN analysis tool. G. Ciardo, after his graduation, went to the USA for a PhD at Duke University with Professor Kishor S. Trivedi. G. Chiola, instead, remained in Torino, first at Politecnico, then at the University, where he made some very significant contributions to the development of GSPN theory, he applied GSPNs to the performance analysis of many different types of systems, and he developed the GreatSPN software tool for the analysis of GSPN models. We all deeply regret the choice of Giovanni to participate only in the early discussions about the contents of this book. Many of the ideas and results presented in the chapters on GSPN theory are his own, and he surely is today one of the most active and creative researchers in the field of GSPNs.

The presence of G. Ciardo at Duke helped make the GSPN paradigm known, and one of the first groups in the USA to adopt GSPNs was indeed the one led by Professor Trivedi. In particular, when the idea arose of organizing a first workshop to bring together the researchers active in the development and application of GSPNs and similar modelling paradigms, the support of Professor Trivedi was the key factor to give us sufficient confidence to organize the International Workshop on Timed Petri Nets in July 1985 in Torino. The name of the workshop was originally meant to be the International Workshop on Stochastic Petri Nets, but the fear of characterizing the topic too narrowly suggested the actual name. However, given the success of the Workshop in Torino, the subsequent editions adopted a different name: International Workshops on Petri Nets and Performance Models were held in Madison, Wisconsin, in 1987; in Kyoto, Japan, in 1989; in Melbourne, Australia, in 1991; in Toulouse, France, in 1993; and the next one will be hosted by Duke University in Durham, North Carolina, in October 1995.

Susanna Donatelli and Giuliana Franceschinis were the first students to choose GSPN theory and applications as a subject for their doctoral theses at the Computer Science Department of the University of Torino. S. Donatelli studied techniques for

the analysis of large GSPNs. G. Franceschinis developed coloured extensions of the GSPN paradigm, to exploit symmetries in the GSPN model definition and analysis. After the completion of their doctoral studies, both Susanna and Giuliana continued their research on GSPNs at the Computer Science Department of the University of Torino.

The continuing success of GSPNs, and the growing number of researchers interested in them, suggested the collection of the main results of more than ten years of research in this book, with the intent of providing newcomers to the GSPN field with a useful tool for their first approach with the formalism.

The book is divided into two parts. Part 1 comprises Chapters 1–6, and summarizes the main results in GSPN theory. Part 2 is made up of Chapters 7–11, which present examples of application of the GSPN methodology.

Chapter 1 contains an informal introduction to Petri nets, mainly intended for those readers who have had no previous experience with them. This chapter, however, is also useful for a first presentation of some Petri net models that will be used in later chapters.

Chapter 2 provides a formal definition of Petri nets, Petri net systems, and Petri net models, and contains a summary of some of the definitions and results of classical Petri net theory that are most useful for the development of the theory in the rest of the book.

Chapter 3 contains a first discussion about the introduction of temporal concepts into Petri nets, and provides an intuitive justification for the importance of priority in timed Petri nets. This discussion naturally leads to the problem of priority in Petri nets, which is addressed in Chapter 4, where formal definitions and theoretical results for Petri nets with priority are presented, extending the results of Chapter 2 to this environment.

After these four introductory chapters, the reader is ready to enter the GSPN arena. Chapter 5 provides the definitions of SPN (stochastic Petri net) and GSPN, as well as a description of the dynamic behaviour of these models. Chapter 6 illustrates the techniques for the analysis of SPN and GSPN models. Both chapters contain illustrative examples, to help the reader understand the concepts that are explained in the text.

The second part of the book is entirely devoted to application examples, to show how the GSPN methodology can be used in different fields.

Chapter 7 shows how GSPN models can account for activities with generally distributed durations, using as an example the classical central server model. In particular, the chapter shows how it is possible to develop a central server model in which the execution times are characterized with an Erlang-2 distribution, rather than with the usual negative exponential distribution adopted in SPNs and GSPNs.

The possibility of generalizing the approach to any phase-type distribution is also discussed.

Chapter 8 provides examples of the application of the GSPN approach to the performance analysis of flexible manufacturing systems. A model of the classical "Kanban" system, and two models of a push production system fed by either a continuous transport system or an automated guided vehicle are presented.

Chapter 9 illustrates the construction of GSPN models of polling systems, which find a variety of applications in the fields of telecommunications, manufacturing, and computer engineering. Rather than looking at the classical cyclic polling systems, we study *random* polling systems, since this choice allows the discussion of the approaches that can often be followed for the construction of *compact* GSPN models, with the advantage of a significant reduction in the cost of the model solution. An example of the reduction process that transforms a GSPN model into an SPN model is also contained in this chapter.

Chapter 10 provides an example of how GSPNs can be used to validate and evaluate concurrent programs. An automatic translation from the software to a GSPN is illustrated, and the techniques for the identification of faults in the software are outlined.

Chapter 11 provides examples of the application of the GSPN approach to the analysis of massively parallel computer architectures.

All the numerical results presented in the examples are obtained with the GreatSPN software package [Chi91a].

The book is completed by two appendices. Appendix A collects some of the results of the theory of stochastic processes that are necessary in the development of GSPN theory. Appendix B contains a glossary with most of the notation used throughout the book.

It should be noted that prerequisites for the study of GSPNs are the knowledge of some simple results of the theory of stochastic processes, and of the theory of Petri nets. Whereas the latter are presented in Chapter 2, the former are discussed in Appendix A. The reason for giving more space to the theory of Petri nets in Chapter 2, and for extracting the presentation of the basic results of the theory of stochastic processes from the main body of the book, lies in the fact that we imagine that a majority of the readers may already be familiar with the material in Appendix A, and in the wide variety of good books devoted to the theory of stochastic processes.

Before leaving the reader to the technical discussions, we wish to acknowledge the support and patience of Gaynor Redvers-Mutton, our Editor at John Wiley in England, and of her staff, and to thank all of our colleagues who read the early versions of the manuscript and provided useful comments and suggestion; in particular we wish to thank Giovanni Chiola and Gianfranco Ciardo for their very constructive remarks.

We also wish to thank L. Brino for providing some of the figures. We acknowledge the financial support of several agencies, which over the years have sponsored our research in the field of GSPNs. In particular, we are grateful to the Italian National Research Council, to the Italian Ministry for University and Research, and to the European Economic Community for their support.

Torino, Italy
June 1995

M.A.M.
G.B.
G.C.
S.D.
G.F.

Contents

List of Figures