

Lecture Notes in Computer Science

616

K. Jensen (Ed.)

Application and Theory of Petri Nets 1992

13th International Conference
Sheffield, UK, June 1992
Proceedings



Springer-Verlag

Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis

616

Advisory Board: W. Brauer D. Gries J. Stoer

BIBLIOTHEQUE DU CERIST



List of Referees for Petri Nets 1992

M. Agoulmine	H.J.M. Goeman	E. Pelz
M. Ajmone Marsan	R. Gold	W. Penczek
M. Akatsu	M. Goldwurm	L. Petrucci
C. Anglano	D. Gomm	M. Pezze
M. Aoyama	A. Guia	M. Pinna
P. Azema	N. Guelfi	J. Pitrel
G. Balbo	N. Götz	A. Poigné
K. Barkaoui	J. Hall	L. Pomello
D. Barnard	C. Hanen	B. Pradin-Chezalviel
E. Battiston	N.D. Hansen	P. Pulli
F. Bause	H. Hasegawa	A. Ramirez
A. Beltrami	I. Hatono	L. Rapanotti
S. Ben Ahmed	X. He	M. Rauhamaa
L. Bernardinello	A. Heise	W. Reisig
G. Berthelot	B. Henderson	M. Ribaud
A. Bertoni	K. Hiraishi	B. Sanchez
E. Best	S. Honiden	V. Sassone
J. Billington	H.J. Hooeboom	M. Sereno
R. Bosworth	P.W. Hoogers	S. M. Shatz
O. Botti	H. Hußmann	M. Shields
A. Bourguet	J.M. Ilie	T. Shimura
J. Brown	C. Johnson	C. Simone
J. Campos	G. Juanole	A. P. Sistla
H. Carstensen	B. Keck	M. Silva
H. Chehaibar	P. Kemper	V. Sliva
A. Cheng	E. Kindler	E. Smith
P. Chrzastowski-Wachtel	T. Kobayashi	N. Speirs
S. Christensen	K. Komota	R. Stroud
R. Coelho	F. Kordon	I. Suzuki
J.M. Colom	M. Koutny	H. Tamura
G. Conte	P. Lee	P. Taylor
J.P. Courtiat	D.L. Lee	E. Teruel
G. de Michelis	M. Leuschel	M. Tiusanen
P. Degano	J. Lilius	N. Treves
R. DeLemos	J. Lobo	N. Uchihira
J. Desel	A. Mader	T. Ushio
S. Donatelli	T. Massart	R. Valette
C. Dutheil	T. Matsumoto	A. Valmari
J. Engelfriet	G. Mauri	K. Varpaaniemi
S. English	D. McCue	F. Vernadat
J. Esparza	A. Moslemie	W. Vogler
P. Estrail	M. Mukund	K. Voss
J. Ezpeleta	T. Murata	R. Walter
G. Findlow	D. Murphy	T. Watanabe
G. Franceschinis	J. Murphy	J. Whitworth
H. J. Genrich	H. Müller	A. Yakovlev
C. Ghezzi	M. Nielsen	B. Zouari
C. Girault	K. Parker	P. Østergaard

Table of Contents

Invited Papers

<i>G. Balbo</i> Performance Issues in Parallel Programming	1
<i>W. Reisig</i> Combining Petri Nets and Other Formal Methods	24

Submitted Papers

<i>C. Autant, Ph. Schnoebelen</i> Place Bisimulations in Petri Nets	45
<i>K. Barkaoui, M. Minoux</i> A Polynomial-Time Graph Algorithm to Decide Liveness of Some Basic Classes of Bounded Petri Nets	62
<i>C. Brown, D. Gurr</i> Refinement and Simulation of Nets – A Categorical Characterisation	76
<i>G. Bruno, A. Castella, G. Macario, M.P. Pescarmona</i> Scheduling Hard Real Time Systems Using High-Level Petri Nets	93
<i>S. Christensen, L. Petrucci</i> Towards a Modular Analysis of Coloured Petri Nets	113
<i>J. Desel</i> A Proof of the Rank Theorem for Extended Free Choice Nets	134
<i>S. Donatelli, M. Sereno</i> On the Product Form Solution for Stochastic Petri Nets	154
<i>G. Findlow</i> Obtaining Deadlock-Preserving Skeletons for Coloured Nets	173
<i>H. Fleischhack</i> P-Superfairness in Nets	193
<i>H.J. Genrich, R.M. Shapiro</i> Formal Verification of an Arbiter Cascade	205
<i>C.A. Heuser, G. Richter</i> Constructs for Modeling Information Systems with Petri Nets	224
<i>K. Hiraishi</i> Construction of a Class of Safe Petri Nets by Presenting Firing Sequences	244

<i>P. Kemper, F. Bause</i> An Efficient Polynomial-Time Algorithm to Decide Liveness and Boundedness of Free-Choice Nets	263
<i>G. Klas</i> Hierarchical Solution of Generalized Stochastic Petri Nets by Means of Traffic Processes	279
<i>A.V. Kovalyov</i> Concurrency Relations and the Safety Problem for Petri Nets	299
<i>J. Lilius</i> High-Level Nets and Linear Logic	310
<i>V.M. Savi, X. Xie</i> Liveness and Boundedness Analysis for Petri Nets with Event Graph Modules	328
<i>E. Teruel, P. Chrzastowski-Wachtel, J.M. Colom, M. Silva</i> On Weighted T-Systems	348
Project Papers	
<i>G. Cutts, S. Rattigan</i> Using Petri Nets to Develop Programs for PLC Systems	368
<i>K. Lemmer, E. Schnieder</i> Modelling and Control of Complex Logistic Systems for Manufacturing	373
<i>J.C. Lloret, J.L. Roux, B. Algayres, M. Chamontin</i> Modelling and Evaluation of a Satellite System Using EVAL, a Petri Net Based Industrial Tool	379
<i>W.W. McLendon, Jr., R.F. Vidale</i> Analysis of an Ada System Using Coloured Petri Nets and Occurrence Graphs	384
<i>K. Varpaaniemi, M. Rauhamaa</i> The Stubborn Set Method in Practice	389
<i>L. Wilkens, J. Canning, P. Krolak</i> Modeling Fine Grain Computation via the Fusion of Two Extended Petri Nets	394

Performance Issues in Parallel Programming

Gianfranco Balbo

Dipartimento di Informatica, Università di Torino,
corso Svizzera 185, 10149 Torino, Italy
e-mail: balbo@di.unito.it

Abstract. The development of parallel applications requires the availability of tools that support their debugging and tuning. GSPN represent a formalism that is well suited for the construction of formal models of parallel programs that can be used for both validation and evaluation purposes. The analysis of GSPN models of parallel programs provides the information that is needed for deciding whether the objectives contained in the specifications of an application are met and for distributing the computation on a parallel architecture. In this paper we discuss a methodology for directly constructing a GSPN model of an application from its code and for deriving the parameters that are needed for obtaining the optimal allocation of the components of a parallel application on the computational units of a parallel architecture. A simple example is used throughout the paper to illustrate the different steps of the methodology and to show how these GSPN models can be used to check the efficiency of a parallel application.

1 Introduction

Parallel computers are widely recognized as the equipments capable of meeting the demands of high performance computing posed by new scientific and real time applications. Parallel programming is however still difficult because of the lack of tools that help in developing and debugging new implementations. Computer architectures and language characteristics restrict the class of applications that can be easily implemented [15] with parallel programs; indeed, concurrency, communication, synchronization and nondeterminism make the manual assessment of the correctness and of the efficiency of parallel programs particularly difficult.

The study of the characteristics of an application both from the point of view of correctness and performance, can be done at different stages of the software life cycle [27]. For instance, an analysis can be performed at the specification stage to ensure that the implementation will meet certain real time constraints [19] or to support the results of rapid prototyping [8]; alternatively, an evaluation can be carried on after the completion of the implementation to verify whether the results conform to the original specifications or to assess its efficiency through the computation of performance indices such as resource consumption indicators [28,16].

* This work has been supported in part by Ministero dell'Università e della Ricerca Scientifica e Tecnologica - 40% Project - and by the Italian National Research Council - Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo, Grant N. 91.00879.PF69.

In order to perform these tasks efficiently, tools must work on models of the real application that differ depending on the goals of the analyses. Different representations allow to characterize the behaviour of a program with different levels of detail [34,30,20]. It is however important that these representations be compatible so that abstract models can be augmented with more detailed descriptions of specific components to allow a modular and efficient analysis of large programs. In any case, the choice of the modelling formalism that is used throughout the software life cycle must easily integrate within the programming environment and must allow the characterization of both the static and dynamic behaviour of the program by means of analytic as well as simulation techniques.

Parallel programs are developed to obtain high-performance computing and is thus a major aspect of their implementations that of allocating their components on the computational units of parallel architectures in order to maximize their efficiency. Real parallel architectures are however characterized by a limited number of processors and by a limited degree of connectivity (not every computational unit can directly communicate with any other unit of the architecture) that constraint their capabilities. Intuition suggests that processes that are concurrently active should be allocated on different nodes of a parallel computer in order to exploit parallelism. Communication among processes can however modify this picture. Indeed, processes allocated on the same processor communicate through common memory in a very fast manner. Communications among processes allocated on different processors, on the other hand, take place through relatively slow links. It follows that when mapping a parallel program on a parallel architecture, several counteracting effects must be taken into accounts. For instance, processes that are simultaneously active and that interact very strongly may be better allocated on the same processor trading the loss of parallelism with the reduction of communication latency. On the contrary, processes with very loose interactions can be easily allocated on separate processors.

These considerations are usually formalized as an optimization problem whose objective function accounts for the communication and processing costs. The form of the objective function depends on the structure of the parallel program and the coefficients depend on the amount of data exchanged among processes, on their mutual distance, and on the amount of local processing performed by each processor [23].

To solve this problem, many different models of parallel programs have been presented in the literature, typical examples being the *task graph* [18] and the *communication graph* [33] in which nodes represent processes and arcs correspond to communications and/or synchronizations. A great effort has been devoted to devise methods for obtaining the optimal allocation strategy [7,31], but little or no effort has been devoted to the problem of constructing these models starting from real programs.

In this paper we show how Generalized Stochastic Petri Net (GSPN) [3] models can be used to estimate characteristic parameters of parallel programs, and thus to construct the corresponding communication graphs, with the possibility of exploiting all the classical analysis results based on GSPN for validating and evaluating these applications. In particular, we address the problem of constructing a GSPN model of the program flow, the possibility of doing this automatically, and the trade-off between automatic generation of the model, efficiency of the analysis and usefulness

of the results. We consider the impact of including control variables in the models and the way of representing communication among processes. Finally we discuss the choice of which features to include in the model, i.e. the abstraction level of the representation, that strongly depends on the type of analysis we want to perform.

To apply these techniques, we focus our attention on parallel languages that allow applications to be organized as sets of cooperating tasks using a message passing paradigm of the rendez-vous type. Major examples of languages of this type are Occam, CSP and Ada. We shall also take into account the case of communications of the non-blocking type, like those allowed by CsTools which is a programming environment available on Meiko's [25] parallel computers.

The work reported in this paper is part of a project for the definition and the implementation of an integrated programming environment for the development of parallel applications organized as sets of parallel processes that interact by message passing following the CSP [21] paradigm. A single formalism based on GSPN is used within this environment for specifying, designing, implementing, testing, and evaluating parallel programs. Detailed information on this project can be found in [5,6,22].

The balance of the paper is the following. Section 2 discusses the possibility of using static analysis techniques for characterizing the behaviour of parallel programs. Section 3 describes the transformation steps that must be undertaken to produce the desired model starting from the code of parallel programs. Section 4 overviews the graph models that are used to solve the problem of mapping parallel programs on parallel architectures. Section 5 describes how the communication graph of a parallel program can be constructed starting from the solution of the GSPN representation of the same application. Section 6 indicates how the model of a parallel program together with the indications of its allocation on a parallel architecture can be used to check the efficiency of the implementation. Section 7 concludes the paper with indications on the problems that are still open and with a discussion of future research efforts that will be undertaken in this field.

2 Static Analysis of Parallel Programs

Two types of strategies may be followed to infer the properties of parallel programs and to obtain their optimal execution. Programs that exhibit very dynamic behaviours may be run on representative sets of input data and under the control of dynamic allocation policies. The results obtained from these sample executions are used to identify the properties of these programs and the observed balance of workloads and communications may be interpreted as a measure of the quality of the allocation policies. Programs that are instead characterized by an internal structure may be analyzed independently of their inputs to identify their properties and may be optimized a priori (i.e., statically) in order to make the best use of the capabilities of the architecture.

Although appealing, the first strategy may be impractical because of the difficulty of choosing representative test cases, of measuring the behaviour of a parallel application, and of the necessity of devising fast decision policies that allocate resources and restructure applications without forcing the programs to wait for long