

HIGH PERFORMANCE COMPUTING

(H11)



HIGH PERFORMANCE COMPUTING

Proceedings of the International Symposium on
High Performance Computing
Montpellier, France, 22–24 March, 1989

edited by

J.-L. DELHAYE

*Centre National Universitaire Sud de Calcul
Ministère de l'Education Nationale
Montpellier, France*

E. GELENBE

*Ecole des Hautes Etudes Informatiques
Université René Descartes
Paris, France*



1989

NORTH-HOLLAND
AMSTERDAM • NEW YORK • OXFORD • TOKYO

ELSEVIER SCIENCE PUBLISHERS B.V.
Sara Burgerhartstraat 25,
P.O. Box 211, 1000 AE Amsterdam, The Netherlands

Distributors for the United States and Canada:

ELSEVIER SCIENCE PUBLISHING COMPANY, INC.
655 Avenue of the Americas
New York, N.Y. 10010, U.S.A.

ISBN: 0 444 88043 7

©ELSEVIER SCIENCE PUBLISHERS B.V., 1989

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of the publisher, Elsevier Science Publishers B.V./Physical Sciences and Engineering Division, P.O. Box 103, 1000 AC Amsterdam, The Netherlands.

Special regulations for readers in the U.S.A. - This publication has been registered with the Copyright Clearance Center Inc. (CCC), Salem, Massachusetts. Information can be obtained from the CCC about conditions under which photocopies of parts of this publication may be made in the U.S.A. All other copyright questions, including photocopying outside of the U.S.A., should be referred to the publisher.

No responsibility is assumed by the publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein.

Printed in The Netherlands

PREFACE

The European Symposium on High Performance Computing was held in Montpellier, France, from March 22 to 24, 1989 at the new Congress Center in Le Corum. It was organized by CNUSC (Centre National Universitaire Sud de Calcul, Ministère de l'Éducation Nationale), in association with IBM and CERFACS, and was sponsored by the District de Montpellier, Département de l'Herault, Région Languedoc-Roussillon.

The development of large scale computing is now considered an essential factor in research productivity, in all scientific and technical domains. The purpose of this symposium was to offer an overview of the state-of-the-art through papers concerned with both methods and applications, and to promote the exchange of experience and information.

These proceedings contain most of the papers presented at the meeting. Some of the authors who made presentations about recent research in industry at the conference meeting did not submit written reports of their work to be included in this volume for obvious reasons.

The papers in these proceedings are grouped into seven subject areas: Numerical Algorithms and Tools, Parallel Architectures, Applications in Fluid Dynamics, Applications in Chemistry, Applications in Physics, New Application Areas, and Education and Training.

The first area is concerned with numerical algorithms and tools. New computer architectures, whether they are parallel or vector, require elaborate algorithmic work in order to enhance performance and numerical stability.

The second area is devoted to parallel architectures, either related to supercomputers (CRAY, IBM) or to computers with a high level of parallelism (Hypercubes).

The next three parts present experience in three areas which have been heavy supercomputing users for a long time: Fluid Dynamics, Chemistry, and Physics. A high proportion of these presentations is concerned with the use and effect of vectorization and parallelization on the IBM 3090/VF.

New application areas such as medicine and economics are discovering the contribution which can be made by High Performance Computing. The four presentations in Part VI are thus of great interest because they open new horizons to those subjects' specialists.

Finally, the last section contains several contributions made as part of an open discussion on the training of supercomputer users. Emphasis must be placed on rapidly educating and training research workers, engineers and scientific students.

I wish to express my gratitude to Michèle Batlle, Pierre Richard, Erol Gelenbe, Francesco Pavoncelli for their work on the Organizing Committee, and to John Connolly, Brian Davis, Ian Duff, Dieter Haupt, Jean-Claude Ippolito, Chris Jones, Domenico Laforenza, Steve Lavenberg, Larry Lee, and Piero Sguazzero for agreeing to be members of the Program Committee.

In addition, I would like to express my special thanks to a number of friends and colleagues who helped considerably in making this symposium a success. They include Isabelle Wurth the Symposium Secretary, Sylvie Flamand, Kathy Barbieri, Andree Balzac, Jean-Claude Ippolito, Philippe Falandry, Bernard Di Miceli, Philippe Gairard, and Herb Budd.

Jean-Louis Delhaye
Chairman of the Organizing Committee
Montpellier, April 1989

CONTENTS

Preface	v
J.-L. Delhayé	
NUMERICAL ALGORITHMS AND TOOLS	
<i>Parallel Algorithms for Scientific Computing</i>	
D.J. Evans	3
<i>Use of Level 3 BLAS Kernels in the Solution of Full Sparse Linear Equations</i>	
P.R. Amestoy, M.J. Daydé, and I.S. Duff	19
<i>A Parallel Method for Solving Tridiagonal Systems of Linear Equations on the IBM 3090 Multiprocessor</i>	
R. Reuter and V. Zecca	33
<i>Round-Off Errors of a Parallel Conjugate Gradient Method</i>	
M. Boulbrachène, P. Cortey-Dumont, A. Huard, A. Laouar, and J.-C. Miellou	45
<i>Convergence Acceleration for Linear Systems Iterative Resolution and Applications to Computational Fluid Mechanics</i>	
N. Benbouda, P. Ferrand, and F. Leboeuf	55
<i>Parallel Implementation of Fast Clustering Algorithms</i>	
M. Bruynooghe	65
<i>Parallel Computing Comes of Age: Supercomputer Calculations for Lattice QCD and Spin Models on Advanced Architecture Computers</i>	
C.F. Baillie and G.C. Fox	79
<i>Distributed Sensing and Feedback Controls of Distributed Parameter Systems</i>	
H.S. Tzou	95

A Massively Parallelizable Heuristic Method for the Graph Partitioning Problem	
F. Meunier	109

PARALLEL ARCHITECTURES

Multiprocessor Speed-Up and the Activity Set Model of Program Behavior	
E. Gelenbe	121
Multitasking on CRAY and IBM Multiprocessors: Concepts and Experiences	
W.E. Nagel and F. Szelényi	133
Analysis of the Fast Fourier Transform on a Hypercube Vector-Parallel Computer	
L. Desbat and D. Trystram	143
Synchronization and Load Unbalance Effects of Parallel Iterative Algorithms	
L. Brochard, J.-P. Prost, and F. Faurie	153
A Scheme for Allocating Task Graphs on Hypercubes	
W.-T. Lo, S.K. Tripathi, and D. Ghosal	167
A Highly Efficient Implementation of Back-Propagation Algorithm on SIMD Computers	
A. Corana, C. Rolando, and S. Ridella	181

APPLICATIONS IN FLUID DYNAMICS

CAD and Structure Computation	
C. Petiau	193
Computational Fluid Dynamics on IBM 3090VF	
V. Saxena	211
N3S, a Numerical Simulation Software for Incompressible Industrial Flows	
Ph. Beque and F. Lacombe	221

An Efficient Vector and Parallel Multigrid Method for Viscous Transonic Flow M. Jayaram and A. Jameson	225
Natural Convection in Porous Media, Parallel Aspect J.P. Caltagirone, S. Carta, P. Fabrie, M.L. Janotto, and P. Richard	239
An Integrated Supercomputing Environment on the IBM 3090/VF for an Industrial Fluid-Dynamics Code F. Piccolo and V. Zecca	249

APPLICATIONS IN CHEMISTRY

Adaptation and Vectorization of the Gaussian 86 Quantum Chemical Program for the IBM 3090 with Vector Facility W. Koch, B. Liu, A.C. Scheiner, and D.J. DeFrees	261
Vectorized Computation of Complex Chemistry Flames N. Darabiha and V. Giovangigli	273
Vector and Parallel Restructuring for Approximate Quantum Reactive Scattering Computer Codes A. Lagana, O. Gervasi, R. Baraglia, and D. Laforenza	287
Vectorizing Pair-Potential "AMYP" Program for the Study of Molecular Associations F. Torrens, R. Montafiana, and J. Sanchez-Marin	299

APPLICATIONS IN PHYSICS

Vectorization and Parallelization of a Simulation Code for Magnetofluid Dynamics G. Chanteur and E. Porteneuve	311
Nuclear Matter Calculations Using Supercomputers A. Lejeune and M. Nihon	327
Numerical Simulations of Coronal Loop Heating S. Poedts, M. Goossens, and W. Kerner	337
Dynamics of Non-Periodic Materials: Spectral Moments Method C. Benoit and G. Poussigue	347

**Supercomputer: the Gateway to Atomic Scale Simulation of
Heteroepitaxial Growth**
V.V. Pham, D. Estève, and J.J. Simonne 357

The Josephson Fluxon Dynamics
J.C. Fernandez and G. Reinisch 367

NEW APPLICATION AREAS

Computational Problems in Orthopaedic Biomechanics
D.L. Taylor and D.L. Bartel 381

On the Modelling of Epidemics
C. Castillo-Chavez, K. Cooke, and S.A. Levin 389

Supercomputing in Economics: a New Field of Research?
H.M. Amman 403

EDUCATION AND TRAINING

Vector Processing at Watson Research
A.R. Rossi 413

High Performance Computing: How to Educate Users?
M. Batlle and G. Urbach 431

Training Users: the CNSF Experience
R.L. Feldman 441

**Towards a System Independant Approach to the Support of Vector
Computer Users**
A. Emmen 451

**NUMERICAL ALGORITHMS
AND TOOLS**

PARALLEL ALGORITHMS FOR SCIENTIFIC COMPUTING

David J. EVANS

Department of Computer Studies
University of Technology
Loughborough, Leicestershire, U.K.

Although mini-supercomputers cost an order of magnitude less than mainframes or supercomputers, their use in typical mainframe applications has been expanding rapidly beyond all expectations. Initially, their performance advantage was fast scientific computation with vector processing and embryonic parallel processing capabilities resulting in their suitability for simulation, scientific and technological applications.

More recently, however, with the increasing usage of parallel processing strategies and techniques, mini-supercomputers are emerging as serious contenders in an expanding arena of information processing applications.

Finally, the increasing trend of solving large linear systems derived from scientific and engineering problems is reconsidered and a parallel algorithm presented which obviates the need for 'linear solvers'.

1. INTRODUCTION

The past few years have seen the accelerated growth of a new breed of computers - the *mini-supercomputer*. These incredibly useful machines have transformed the computer scenario by their sheer capability and adaptability in many traditional application areas for they offer 25% of the power of a supercomputer for only 10% of the cost.

Clearly, then although the mini-supercomputer cannot compete on peak supercomputer performance they can provide a cost effective, high performance machine or work station in scientific and engineering application areas such as simulation and modeling, finite element/difference and matrix analyses, numerical solution of differential equations and approximation of functions in CAD and graphics.

How and where did these machines evolve? Historically the initial methods of achieving mini-supercomputer performance was via the use of special purpose *attached processors* (array processors). This was achieved by the rapid development of ASIC (application specific integrated circuit) design tools at board-level and chip-level implementations. The 32 bit array processor with its high performance processing of streams of data by firmware was initially used widely in signal and image processing but more recently has entered the domain of SIMD computing whilst the 64 bit array processor with the ability to execute programs rather than sub-

routines and with vector processing capability has emerged as the fore-runner of the mini-supercomputer.

However the earlier attached processor designs although adding extra processing ability suffered from severe host communication problems which demonstrated the need for executing more code within the attached processor. This need for large fast memories and wide bandwidth communication led to a *full function attached processor* containing all the features of a C.P.U. being incorporated.

The final stage was the elimination of the host system by introducing an operating system and the realisation of the *stand-alone mini-supercomputer*. In the future, we can expect to see the mini-supercomputer performance coming to the workstation market while off the shelf VSLI chip sets will be incorporated into desktop configurations.

Finally, mini-supercomputer performance can also be attained effectively by *multiprocessor implementations*. In fact, the recently introduced Sequent Computer Symmetry multi-processor system attains the M0-MFLOP threshold of supercomputer performance when equipped with an optional Weitek arithmetic co-processor.

Multiprocessor systems are optimized for multiple job throughput. Whilst *parallel processors*, in contrast, focus on the execution of single jobs. Both, however, employ a number of processors, and the prospective buyer should be more concerned with performance and cost than with what often amounts to architectural semantics.

2. CURRENT TRENDS IN PARALLEL ARCHITECTURES

Let us look briefly at some parallel architectures - a procedure which one might go through when deciding what sort of parallel architecture is best for a given class of problem.

To achieve parallelism by extending sequential architectures with such machines like the Cray 1S, or CDC Cyber or the latest line of VAX machine involves quite *expensive technology*. So I think we all agree that parallel processing is the correct approach for the future.

This calls into question what kind of parallel architecture to consider. You could put together a few very expensive processors. Some of the supercomputer companies, in fact, are recognising this trend and moving that way with the Cray XMP and the Japanese FUJITSU computers. A better approach is to put together many low cost processors - *micro-processors* in fact. I think the most convincing argument for this is that the microprocessor is the most cost effective package of a computer instruction.

There are many advantages of using microprocessor technology in a parallel architecture. First, there is a tremendous amount of market forces associated with using a current chip design for the processor. Secondly, another advantage is to have a wide range of high quality compilers available and software tools such as the Kuck & Associates parallel pre-processor, etc. Finally, the importance of the communications software must not be forgotten when it comes to building a system around standard microprocessors and with an operating system like UNIX available.

The choice now is whether to have a tightly coupled or a *shared memory* architecture, or a loosely coupled or *distributed memory* architecture? Both types of systems are now the subject of extensive competition in the marketplace.

More recently it has become evident that the shared memory architectures with a moderate number of processors have a programming model that is much more closely aligned with the traditional uni-processor approach. So ultimately the machine is easier to use and program. On the other hand, loosely coupled architectures, i.e. hypercube, or mesh-type topologies do have the advantage of using a few hundred or even a few thousand processors. However because of the communication problems that are involved, trade-off's are made to decrease the number of interconnects between the processors. In addition, the mapping of the algorithms on to such architectures has to be seriously reconsidered which can often involve extensive reprogramming overheads.

Thus midrange multi-processors, sometimes called multi-microcomputers achieve super-minicomputer and even mini-supercomputer performance by combining 16 to 64 processors with memory and I/O subsystems, usually in tightly coupled configurations. Like low-end multi-processors, they use proprietary system buses, standard microprocessor I/O buses and variations on the UNIX theme.

Parallel processing systems are specialized multi-processor systems. They fall into two categories: common backplane implementations that achieve parallel operation through the use of computer technology, and architectures like hypercubes and systolic arrays in which processors communicate with a limited number of cohorts. Such systems are the focus of much activity in terms of both research and publicity.

Few technology issues remain unresolved in the multi-processor field. Loosely and tightly coupled systems of minicomputer and microprocessor-based APU's are now available from several suppliers. Dynamic load balancing is now common, and the capability to parallelize tasks is rapidly emerging.

Differences between implementation strategies usually revolve around the specific tasks for which the systems are best suited. In general, low-end multi-processors are used in the workstation, network server and multi-user applications, while multi-microcomputers and large multi-processors compete with super-minicomputers.

The forecast for the computer industry as a whole for the remainder of the decade predicts a growth of roughly 10% per year. If the forecast is right, growth in the multi-processor segment would presumably have to come at the expense of (and be resisted by) the established firms. New market entrants such as mini-supercomputer and RISC suppliers will provide fierce competition.

Clearly, a pitched battle will rage (with a high casualty rate) between super-minicomputer, mini-supercomputer and multi-microcomputer/parallel-processor suppliers for what is now the high end of the super-minicomputer market. The flexibility of multi-processor systems (variable granularity, incremental growth capability and fault-tolerance, etc.) suit them well for specialised markets and midrange multitasking applications. Finally

parallel computers can be expected to replace an entire segment of mid-range machines, including those currently categorized as mini, supermini and main-frame. Single processor-based equipment will hold down the low end, with massive supercomputers on the other end of the price and performance range, (Fig.1).

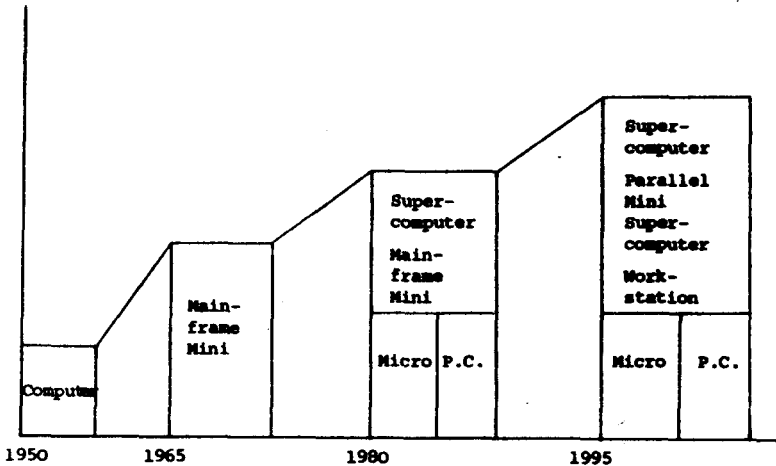


FIGURE 1: Changes in computer markets due to the impact of parallel architectures.

3. FUNDAMENTAL ISSUES IN PARALLEL ALGORITHM DESIGN FOR LINEAR SYSTEMS

Research into parallelism indicates that often we cannot directly extend the techniques used in serial algorithms to parallel computers because efficient serial algorithms do not necessarily lead to efficient parallel algorithms. Our example illustrates the 'divide and conquer' strategy. This (very old) idea is probably the first that would occur to anyone designing an algorithm for parallel machines. One first partitions the problem into several sub-problems and then the parallel machine solves the subproblems simultaneously. Unless the subproblems are independent, the data must be communicated between the subproblems. For the success of a divide and conquer technique, it is necessary that the gains in solving the subproblems in parallel are not outweighed by the extra work required to obtain the solution to the whole problem from that of the subproblems. An essential ingredient of the strategy is that the subproblems are rearranged in 'explicit' form so that a parallel algorithm can be achieved.

The solution of banded and tridiagonal linear systems occurs frequently in the solution of partial differential equations in Computational Mathematics. Here we introduce a new parallel algorithm which can be readily extended to cope with multi-dimensional boundary value problems.

Consider the following second order parabolic equation,

$$\frac{\partial U}{\partial t} = \frac{\partial^2 U}{\partial x^2}, \quad 0 \leq x \leq 1, \quad 0 < t \leq T, \quad (3.1)$$

subject to the initial-boundary conditions,

$$\begin{aligned} U(x,0) &= f(x), \quad 0 < x < 1, \\ U(0,t) &= g(t), \quad 0 < t \leq T, \end{aligned} \quad (3.1a)$$

and $U(1,t) = h(t).$

A uniformly-spaced network whose mesh points are $x_i = i\Delta x$, $t_j = j\Delta t$ for $i=0,1,\dots,m,m+1$ and $j=0,1,\dots,n,n+1$ is used with $\Delta x=1/(m+1)$, $\Delta t=T/(n+1)$ and $\lambda = \Delta t/(\Delta x)^2$, the mesh ratio.

A weighted approximation to the differential equation (3.1) at the point (x_i, t_{j+1}) is given by,

$$\begin{aligned} -\lambda\theta u_{i-1,j+1} + (1+2\lambda\theta)u_{i,j+1} - \lambda\theta u_{i+1,j+1} &= \lambda(1-\theta)u_{i-1,j} + (1-2\lambda(1-\theta))u_{i,j} + \\ &\lambda(1-\theta)u_{i+1,j}, \quad i=1,2,\dots,m. \end{aligned} \quad (3.2)$$

This approximation can be displayed in a more compact matrix form as,

$$\begin{pmatrix} a & b & & & \\ c & a & b & & \\ & c & a & b & \\ & & & \ddots & \\ & & & & c & a & b \\ & & & & & c & a \end{pmatrix} \begin{matrix} u_1 \\ u_2 \\ \vdots \\ \vdots \\ u_{m-1} \\ u_m \end{matrix}_{j+1} = \begin{matrix} f_1 \\ f_2 \\ \vdots \\ \vdots \\ f_{m-1} \\ f_m \end{matrix}_j, \quad (3.3)$$

i.e., $Au = f$, (3.4)

where, $c = -\lambda\theta$, $a = (1+2\lambda\theta)$, $b = -\lambda\theta$;

$$\left. \begin{aligned} f_1 &= \lambda(1-\theta)(u_{0j} + u_{2j}) + \lambda\theta u_{0,j+1} + (1-2\lambda(1-\theta))u_{1j}; \\ f_i &= \lambda(1-\theta)(u_{i-1,j} + u_{i+1,j}) + (1-2\lambda(1-\theta))u_{ij}, \quad i=2,3,\dots,m-2,m-1; \\ f_m &= \lambda(1-\theta)(u_{m-1,j} + u_{m+1,j}) + (1-2\lambda(1-\theta))u_{mj} + \lambda\theta u_{m+1,j+1} \end{aligned} \right\}$$

and

$$u = (u_{1,j+1}, u_{2,j+1}, \dots, u_{m,j+1})^T \text{ and } f = (f_1, f_2, \dots, f_m)^T. \quad (3.4a)$$

We note that f is a column vector of order m consisting of the boundary values as well as known u values at time level j while u are the values at time level $(j+1)$ which we seek. We also recall that (3.4) corresponds to the fully implicit, the Crank Nicolson, the Douglas and the classical explicit methods when θ takes the values 1 , $\frac{1}{2}$, $\frac{1}{4}$, $-1/12\lambda$ and 0 with accuracies of the order $O([\Delta x]^2 + \Delta t)$, $O([\Delta x]^2 + [\Delta t]^2)$, $O([\Delta x]^4 + [\Delta t]^2)$ and $O([\Delta x]^2 + \Delta t)$ respectively.

Let us assume we have an *even number of intervals* (corresponding to an odd number of internal points, i.e. m odd) on the real line $0 < x < 1$. A similar analysis holds for m even. We can then perform the following

