


**THE GENESIS AND
SUPRENUM PROJECTS**

Oliver A. McBryan

CU-CS-719-94



University of Colorado at Boulder

DEPARTMENT OF COMPUTER SCIENCE

**THE GENESIS AND
SUPRENUM PROJECTS**

Oliver A. McBryan

CU-CS-719-94 1994

**Department of Computer Science
University of Colorado at Boulder
Campus Box 430
Boulder, Colorado 80309-0430 USA**

ANY OPINIONS, FINDINGS, AND CONCLUSIONS OR RECOMMENDATIONS
EXPRESSED IN THIS PUBLICATION ARE THOSE OF THE AUTHOR(S) AND DO NOT
NECESSARILY REFLECT THE VIEWS OF THE AGENCIES NAMED IN THE
ACKNOWLEDGMENTS SECTION.

The GENESIS and SUPRENUM Projects^{*†}

Oliver A. McBryan
Dept. of Computer Science
University of Colorado
Boulder, CO 80309.

Email: mcbryan@cs.colorado.edu

Abstract

The SUPRENUM project was a major effort in Germany to develop a national expertise in massively parallel processing both at the hardware and software level. We review the development of SUPRENUM and introduce the twelve later papers in this proceedings that deal in more detail with SUPRENUM. Taken together these papers present an overview of the project goals, its historical development, successes and failures of the project, and indicate ways in which the project has influenced MPP developments elsewhere.

GENESIS was a European ESPRIT project focused on massively parallel processing. It included both hardware and software architecture aspects. GENESIS was a natural development from SUPRENUM and used the Suprenum-1 machine as its primary hardware platform. We briefly discuss GENESIS and introduce the nine papers in this volume that represent the GENESIS effort. GENESIS has had a significant effect on many other MPP efforts. Several of the papers address these influences in some detail.

Keywords: SUPRENUM, GENESIS, MPP, parallel, supercomputer, portability, architecture, software, distributed memory, MIMD.

*Research supported in part by NSF Grand Challenges Applications Group grant ASC-9217394 and by NASA HPCG Group Grant NAG5-2218.

† To appear in Parallel Computing, October 1994.

TABLE OF CONTENTS

1. Introduction.....	1
2. The SUPRENUM Project	2
2.1. SUPRENUM Overview.....	2
2.2. SUPRENUM Papers.....	3
2.2.1. SUPRENUM Perspective.....	3
2.2.2. Computer Architectures.....	3
2.2.3. MPP Software and OS.....	3
2.2.4. Benchmarking and Applications.....	4
3. The GENESIS Project	6
3.1. GENESIS Overview.....	6
3.2. GENESIS Papers.....	7
3.2.1. GENESIS Perspective.....	7
3.2.2. Architecture Design.....	7
3.2.3. MPP Software	7
3.2.4. Benchmarking.....	8
3.2.5. Applications.....	8
References.....	9

The GENESIS and SUPRENUM Projects^{*†}

Oliver A. McBryan
Dept. of Computer Science
University of Colorado
Boulder, CO 80309.

Email: mcbryan@cs.colorado.edu

1. Introduction

This paper provides an overview of two major European projects involving massively parallel processing (MPP) - GENESIS and SUPRENUM. The two projects were closely related and in fact overlapped in time. This volume is entirely devoted to papers dealing with these projects. In this paper we will briefly introduce each project and then provide an introduction and associated context for each of the subsequent papers related to that project.

The SUPRENUM project was established as a consortium of German companies, university departments and research centers, with the goal of designing a Gflops MPP supercomputer for scientific applications. It resulted in the Suprenum-1 prototype computer which had a peak rate of over 5 Gflops.

The GENESIS ESPRIT project was established in order to design a second generation European MPP system. It utilized the developments already available from SUPRENUM, Supernode and several other European projects. It resulted in major advances in network interconnections, node design, MPP operating systems and languages, and benchmarking.

SUPRENUM, while strictly a German supercomputer project, led naturally to the broader European GENESIS project. The projects were in fact closely related and many of the key participants in SUPRENUM were also participants in GENESIS. For this reason it is difficult to strictly separate papers into GENESIS or SUPRENUM categories. This has been done here primarily to indicate the nature of the efforts, but at least one paper[6] clearly represents work equally attributable to both projects.

Section 2 introduces the SUPRENUM project and then briefly discusses each of the SUPRENUM papers in this volume. Section 3 similarly introduces GENESIS and follows with short introductions to the GENESIS papers in the volume.

*Research supported in part by NSF Grand Challenges Applications Group grant ASC-9217394 and by NASA HPCC Group Grant NAG5-2218.

† To appear in Parallel Computing, October 1994.

2. The SUPRENUM Project

2.1. SUPRENUM Overview

The SUPRENUM project evolved from a desire in Germany in the mid 1980's that the country remain at the forefront in computing technology. Two key research groups at the GMD were the primary proponents. The Numerical Group at the GMD, St. Augustin, led by Dr. U. Trottenberg, was primarily interested in methods for the numerical solution of partial differential equations (PDE) and their implementation on massively parallel processors (MPP). This group had pioneered the development of multigrid libraries for PDE solution and was particularly interested in MPP that could solve multigrid algorithms efficiently. The Numerical Group with its application orientation provided much of the simulation, specification, programming, benchmarking and evaluation expertise for SUPRENUM. The second group, at the GMD FIRST in Berlin, was led by Dr. W. Giloi and was involved in microprocessor and MPP architecture, and in operating system (OS) design. This second group provided the core expertise in SUPRENUM for hardware and OS design.

From concepts developed by these two core groups, the SUPRENUM project was constituted as a major German MPP initiative and came to include researchers at 15 institutions throughout Germany. The project was designed with both a research and a commercial aspect in mind. The commercial side was represented by the inclusion of several major German companies who had responsibility for hardware fabrication. The research side was represented by numerous university and research center groups who contributed to various aspects of design and evaluation. Overall coordination was handled by a new company, SUPRENUM GmbH, which was located in Bonn. The project began in 1985 and was completed in 1990.

The SUPRENUM Project resulted in the development of a prototype supercomputer, the Suprenum-1. There were considerable delays due to the complexity of the design problems involved, a situation typical of many MPP projects. As a result the Suprenum-1 appeared near the tail end of its opportunity window. However for a short time the 256-node Suprenum-1 was the fastest MPP computer available. Five of the systems were delivered to various research institutions in Europe and a substantial body of research was accomplished with those systems, some of which is reported on in this volume. Because no successor to the initial architecture prototype was funded, the Suprenum-1 was not a commercial success. However the SUPRENUM project was clearly extremely successful in the primary goal of developing MPP expertise within Germany. Today no one doubts that German researchers are among those at the forefront of MPP research world wide.

2.2. SUPRENUM Papers

The SUPRENUM papers naturally fall into four groups. First are three papers [1-3], by two of the key principal investigators, U. Trottenberg[1] and W. Giloi[2], and by an outside independent researcher, O. McBryan[3], which analyze the successes and failures of the project. The second [4] and third [5-7] groups of papers deal respectively with MPP architectures and software. The fourth, and by far largest, group of papers [8-12] are those dealing with benchmarking, performance evaluation and applications on SUPRENUM and related systems.

2.2.1. SUPRENUM Perspective

Dr. U. Trottenberg was involved with SUPRENUM from its inception, and headed the SUPRENUM GmbH company as its Scientific Director. In his insightful paper [1] *"Some remarks on the SUPRENUM project"* he reviews the origins of SUPRENUM, its organization and its major successes. Dr. Trottenberg also addresses some of the controversies that surrounded SUPRENUM, and discusses lessons to be learned from SUPRENUM for application to future projects.

Dr. W.K. Giloi was also involved in SUPRENUM from near the start and was the principal architect of the Suprenum-1 prototype. Dr. Giloi discusses the design issues involved in the Suprenum-1 hardware and operating system in the paper [2] *"The SUPRENUM Supercomputer: Goals, Achievements, and Lessons Learned"*. This paper provides the essential technical basis for understanding the remaining papers in the volume.

We have been familiar with the SUPRENUM project since its inception, and at the same time we were involved directly with several MPP projects in the USA. Our paper [3] *"SUPRENUM: Perspectives and Performance"* summarizes our views of the project and analyzes the reasons behind some of the difficulties experienced. This paper also presents performance measurements of Suprenum-1 that in fact might well be part of the fourth set of papers, on applications and benchmarks.

2.2.2. Computer Architectures

The SUPRENUM project has been critical in the development of several more recent MPP software and architecture projects. In the paper [4] *"Parallel Supercomputer Architectures and Their Programming Models"*, Dr. W.K. Giloi discusses some of these recent developments and relates them to the SUPRENUM roots.

2.2.3. MPP Software and OS

The PEACE operating system was designed within the SUPRENUM project to provide an efficient scalable massively parallel OS. Many difficult design decisions were involved, particularly those related to minimizing the OS overheads for message passing subroutine calls. PEACE has continued to evolve beyond the SUPRENUM project, and is the OS for the Manna architecture discussed in the

previous paper [4] by W. Giloi. Dr. W. Schröder-Preikschat reviews all of these topics in his paper [5] "*PEACE - A Software Backplane for Parallel Computing*".

SUPRENUM was the first MIMD supercomputer to support Fortran 90 array extensions in each node. The SUPRENUM architecture, with its complex vector nodes, posed a serious challenge for compiler development. However there was no body of pre-existing Fortran 90 programs. The SUPERB project developed an automatic parallelization facility capable of generating fully parallel SUPRENUM Fortran (including message passing insertions) from an initial F77 application program. H.P. Zima, P. Brezany and B.M. Chapman describe SUPERB in their paper [6] "*SUPERB and Vienna Fortran*". Vienna Fortran is a related project developed under the GENESIS project.

SUPRENUM supported a comprehensive array of facilities for providing user access to parallelism. At the lowest level, SUPRENUM Fortran provided a Fortran 77 extension that included a) Fortran 90 array extensions, b) tasks as a Fortran datatype and c) message passing integrated into Fortran I/O statements. At an intermediate level, PARMACS provided a portable message passing environment for SUPRENUM. At the highest level, an applications library of grid operations called COMLIB provided a powerful facility for scientific computing applications. R. Hempel reviews each of these areas in his paper [7]: "*Applications Programming Interfaces for SUPRENUM*".

2.2.4. Benchmarking and Applications

The remaining SUPRENUM papers all deal with benchmarking, performance evaluation and applications on SUPRENUM. Taken together these papers show that Suprenum-1 was a successful prototype which delivered as much as 25% of its peak rate of 5 Gflops on real application programs. Most of the measurements described in these papers relate to grid oriented computations, an area where SUPRENUM software provided particularly strong support.

The paper [8] "*Performance Modeling of Grid Problems - A Case Study on the SUPRENUM System*", by H. Mierendorff, H. Schwamborn and M. Tazza, discusses the problem of understanding Suprenum-1 performance measurements through detailed performance modeling. The paper shows that for an important class of grid algorithms, performance can be predicted accurately provided sufficiently detailed parameterizations of the architecture are used.

The later sections of our previously mentioned paper [3] "*SUPRENUM: Perspectives and Performance*" presents results of a series of Suprenum-1 benchmarks we developed for the well-known Shallow Water Equations, frequently used in the atmospheric and oceanographic modeling communities.

Multigrid methods were an initial focus point in SUPRENUM planning because of the strong multigrid research program at the GMD. In the paper [9] "*Implementation of a Multigrid Algorithm on SUPRENUM and Other Systems*", M. Alef describes the implementation of a multigrid algorithm on the Suprenum-1 and reports the measured performance.

In contrast with the previous papers which by and large report results specific to SUPRENUM, the paper [10] "*LiSS - An environment for the parallel multigrid solution of partial differential equations on general 2D domains*", by H. Ritzdorf, A. Schüller, B. Steckel and K. Stüben, describes a very general package developed in the SUPRENUM project for the solution of non-linear PDE on MPP machines. The package utilizes the portable PARMACS communication library (also developed within SUPRENUM) and as a result LiSS runs on most MPP systems.

O. Pätzold, A. Schüller and H. Schwichtenberg present quite diverse measurements of Suprenum-1 performance in their paper [11] "*Parallel Applications and Performance Measurements on SUPRENUM*". In contrast to the previous papers that discussed only grid applications, this paper includes also measurements for two other types of algorithm: Monte Carlo methods in statistical mechanics, and chemical transport computations. Suprenum-1 is shown to do well on all three types of simulation.

The final paper [12] in this section, by G. Fleischmann, M. Gente, F. Hofmann and G. Bolch, entitled "*Performance Analysis of Parallel Programs Based on Model Calculations*", describes a general framework for performance analysis of MPP programs.

3. The GENESIS Project

3.1. GENESIS Overview

GENESIS was a 3 year European ESPRIT project that extended from 1989 to 1992. The primary goal was to pool the technical resources of a number of key European research organizations and companies in an effort to design a second generation massively parallel supercomputer. GENESIS had many partners, including several MPP manufacturers, national computer science research centers, and university departments.

The project included all aspects of MPP design and use. Preliminary subtasks were defined in each of the following areas: Current State of the Art, Global Architecture, Abstract Machine Definition, Node Architecture, Interconnection Network, Distributed Operating System, Programming Language Environment, User Interface, and Application Software. After the second year all hardware development funding from GENESIS was terminated and portability of software became the dominant theme for GENESIS funded projects.

Initially all evaluation work focused on the Suprenum-1 prototype, but was soon widened to include most of the currently available MPP systems. This resulted in the need to define a portable benchmark suite which could allow useful comparisons between architectures. This suite, along with a portable numerical algorithm library to support it, were among the most useful and enduring outputs of GENESIS. This effort also indicated the importance of the portable PARMACS message passing environment developed under SUPRENUM.

Many of the GENESIS architectural results were incorporated into later products. For example the Meiko CS-2 MPP system is based on concepts developed in GENESIS. Similarly the Manna MPP system developed at GMD First in Berlin, has its origins in GENESIS. Language developments from GENESIS have played an equally important role in many areas, including the High Performance Computing Forum and its HPF language, as well as the Message Passing Interface Forum and its MPI specification. GENESIS partners have played an important role in both the HPF and MPI standardization efforts. Several current research projects are direct descendants of GENESIS - for example PPPE and RAPS.

The essential switch in GENESIS from a hardware oriented project to one oriented to portable software occurred during the second year. After that, portability (based on PARMACS) became the core goal of GENESIS. Initially SUPRENUM GmbH was the managing contractor for GENESIS. In the second year, Dr. U. Trottenberg founded Pallas GmbH (basically incorporating most of the software components of SUPRENUM GmbH) and thereafter Pallas became prime contractor for GENESIS.

3.2. GENESIS Papers

The ten GENESIS papers in this volume are collected into five categories. The initial paper [13] provides an overview of the project. Two papers [14-15] deal with architectural issues. A further two papers [5,16] focus on MPP software developments. The paper [17] discusses the portable GENESIS benchmark suite. Finally, four papers [18-21] describe the GENESIS applications effort.

3.2.1. GENESIS Perspective

A. Hey introduces the GENESIS section of this volume with his comprehensive overview paper [13] *"The Genesis Esprit Project - An Overview"*. This paper describes the origins of GENESIS, the initial consortium members and project plan and their evolution over time, and highlights the key results of GENESIS research. The paper also sets the stage for the remainder of the GENESIS papers in the volume.

3.2.2. Architecture Design

A major goal of GENESIS was to develop a second generation European supercomputer design that would build on the lessons already learned in SUPRENUM. Considerable effort was devoted to analysis of interconnection strategies. O. Kolp describes some of these network simulations in his paper [14] *"Performance Estimation for a Parallel System with a Hierarchical Switch Network"*.

While GENESIS has had an impact on many subsequent projects, perhaps the most significant of these is Meiko's CS-2 MPP computer. Meiko was a GENESIS participant and the CS-2 design was influenced by the GENESIS project. J. Beecroft, M. Homewood and M. McLaren describe the CS-2 interconnection network in detail in their paper [15] *"Meiko CS-2 Interconnect Elan-Elite Design"*.

3.2.3. MPP Software

Vienna Fortran is a Fortran extension designed especially for distributed MPP systems and represents one of the major accomplishments of the GENESIS project. Vienna Fortran has had a significant impact on the design of the HPF MPP language designed by the High Performance Computing Forum. A number of GENESIS partners were members of the HPF Forum. H.P. Zima, P. Brezany and B.M. Chapman describe the language in the paper [6] *"SUPERB and Vienna Fortran"*. Because SUPERB was developed within the earlier SUPRENUM project, and Vienna Fortran was a natural outgrowth of SUPERB, this paper appears in the SUPRENUM section of the volume.

A major goal of GENESIS was to develop portability strategies for MPP applications. The SUPRENUM PARMACS message passing system was adopted as a basis for communication, and GENESIS projects then developed higher level libraries using PARMACS. The GENESIS numerical library was based on a related

transputer project library, which was converted to PARMACS and designed to be applicable to most distributed MIMD systems. The library development is described in detail in [16] *"The Design and Implementation of a Portable Parallel Numerical Library"* by L.M. Delves, C.A. Addison and O.A. Aziz.

3.2.4. Benchmarking

GENESIS included a significant MPP benchmarking effort. The initial goal was to make a careful evaluation of SUPRENUM. Later the goal was to develop a portable benchmark suite that would allow meaningful assessment of systems and useful comparisons between systems. C.A. Addison, V.S. Getov, A.J.G. Hey, R.W. Hockney and I.C. Wolton discuss the GENESIS benchmark project in their paper [17] *"Benchmarking for Distributed Memory Parallel Systems: Gaining Insight from Numbers"*.

3.2.5. Applications

GENESIS supported a major effort to parallelize a variety of scientific application codes. It is not possible to present all of these application efforts in this volume. However in their paper [18] *"GENESIS Application Software"*, K. Solchenbach, C.-A. Thole and U. Trottenberg review the wide range of applications included, the main research teams involved, and the influence of the effort on continuing MPP research projects. The following three short papers are then provided as representative of the overall applications research effort.

E. Gerteisen's paper [19] *"Preliminary Performance Results of the Massive Parallel Aircraft Euler Method"*, is typical of GENESIS applications research. It describes the parallelization and performance evaluation of a complex finite element code using the portable PARMACS message passing library. T. Kauranne discusses the parallelization of weather models in the paper [20] *"Summary of GENESIS work at the European Centre for Medium-range Weather Forecasts (ECMWF)"*. Finally J.J.H. Miller and S. Wang describe the parallelization of a 3-D semiconductor device simulator in their paper [21] *"On the Implementation of a 3-D Semiconductor Device Simulator on Distributed Memory MIMD/SIMD Machines"*.

References

- [1] U. Trottenberg, "Some remarks on the SUPRENUM project", Parallel Computing, this volume.
- [2] W.K. Giloi, "The SUPRENUM Supercomputer: Goals, Achievements, and Lessons Learned", Parallel Computing, this volume.
- [3] O. McBryan, "SUPRENUM: Perspectives and Performance", Parallel Computing, this volume.
- [4] W.K. Giloi, "Parallel Supercomputer Architectures and Their Programming Models", Parallel Computing, this volume.
- [5] W. Schröder-Preikschat, "PEACE - A Software Backplane for Parallel Computing", Parallel Computing, this volume.
- [6] H.P. Zima, P. Brezany and B.M. Chapman, "SUPERB and Vienna Fortran".
- [7] R. Hempel, "Application Programming Interfaces for SUPRENUM", Parallel Computing, this volume.
- [8] H. Mierendorff, H. Schwamborn and M. Tazza, "Performance Modeling of Grid Problems - A Case Study on the SUPRENUM System", Parallel Computing, this volume.
- [9] M. Alef, "Implementation of a Multigrid Algorithm on SUPRENUM and Other Systems", Parallel Computing, this volume.
- [10] H. Ritzdorf, A. Schüller, B. Steckel and K. Stüben, "LiSS - An environment for the parallel multigrid solution of partial differential equations on general 2D domains", Parallel Computing, this volume.
- [11] O. Pätzold, A. Schüller and H. Schwichtenberg, "Parallel Applications and Performance Measurements on SUPRENUM", Parallel Computing, this volume
- [12] G. Fleischmann, M. Gente, F. Hofmann and G. Bolch, "Performance Analysis of Parallel Programs Based on Model Calculations", Parallel Computing, this volume.
- [13] A. Hey, "The Genesis Esprit Project - An Overview", Parallel Computing, this volume.
- [14] O. Kolp, "Performance Estimation for a Parallel System with a Hierarchical Switch Network", Parallel Computing, this volume.
- [15] J. Beecroft, M. Homewood and M. McLaren, "Meiko CS-2 Interconnect Elan-Elite Design", Parallel Computing, this volume.
- [16] L.M. Delves, C.A. Addison and O.A. Aziz, "The Design and Implementation of a Portable Parallel Numerical Library", Parallel Computing, this volume.

- [17] C.A. Addison, V.S. Getov, A.J.G. Hey, R.W. Hockney and I.C. Wolton, "*Benchmarking for Distributed Memory Parallel Systems: Gaining Insight from Numbers*", Parallel Computing, this volume.
- [18] K. Solchenbach, C.-A. Thole, U. Trottenberg, "*GENESIS Application Software*", Parallel Computing, this volume.
- [19] E. Gerteisen, "*Preliminary Performance Results of the Massively Parallel Aircraft Euler Method*", Parallel Computing, this volume.
- [20] T. Kauranne, "*Summary of GENESIS work at ECMWF*", Parallel Computing, this volume.
- [21] J. Miller, S. Wang, "*On the Implementation of a 3-D semi-conductor Device Simulator on Distributed Memory MIMD/SIMD machines*", Parallel Computing, this volume.