HPCS : Client-Server Support for High Performance Computing

Inauguraldissertation
zur
Erlangung der Würde eines Doktors der Philosophie
vorgelegt der
Philosophisch-Naturwissenschaftlichen Fakultät
der Universität Basel

von
Gérald Jean-Pierre Prétôt-Eckenstein
aus Le Noirmont, JU

Basel, im Februar 1999

Genehmigt von der Philosophisch-Naturwissenschaftlichen Fakultät auf Antrag der Herren

Prof. Dr. Helmar Burkhart

Prof. Dr. Béat Hirsbrunner

Basel, im April 1999

Prof. Dr. S. Schmid, Dekan

Research Reports in Computer Science

Band 5

Gérald Jean-Pierre Prétôt

HPCS: Client-Server Support for High-Performance Computing

Shaker Verlag
Aachen 1999

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Prétôt, Gérald Jean-Pierre:
HPCS: Client-Server Support for High-Performance Computing /
Gérald Jean-Pierre Prétôt. - Als Ms. gedr. - Aachen : Shaker, 1999
(Research Reports in Computer Science ; Bd. 5)
Zugl.: Basel, Univ., Diss., 1999
ISBN 3-8265-6700-5

Printed in Germany.

ISBN 3-8265-6700-5
ISSN 1436-6967

Shaker Verlag GmbH • P.O. Box 1290 • D-52013 Aachen
Telefon: 0049/2407/9596 - 0 • Telefax: 0049/2407/9596 - 9
Internet: www.shaker.de • eMail: info@shaker.de

Abstract

Distributed and parallel computing are well-established and successful fields. In both certain levels of standardization have been achieved and reliable tools are available for software development and deployment. However, so far hardly any effort has been made, to *integrate* the two worlds in a generic way. This is the goal of the work presented here.

This thesis introduces a new model for *client-server support for high-performance computing* called *HPCS*. The model offers a generic approach to making parallel, high-performance implementations of *compute services* accessible to sequential clients in the form of software *components*. It exploits the abstraction provided by *parallel data structures* and *parallel procedures* to allow *transparent* access from remote clients. In HPCS parallel data structures are presented as *abstract data objects* and parallel procedures as *services*. Access is provided using standardized client-server *interfaces* such as the ones specified by the *CORBA* standard. This is done in a way that does not require the service programmer to be concerned with the client-server interface nor the client programmer with the parallel nature of the service implementation.

Call-persistent data objects and a minimal set of methods for their manipulation are proposed as a mechanism to work with parallel data structures that remain on the server between different service calls. A *parallel request handler* split into two components, the *Multiplexer* and the *Dispatcher*, provides transparent access to parallel services.

ACS, a concrete implementation of HPCS targeting the parallel programming language *ALWAN* and *CORBA* implementations is presented. Thus, ACS enables writing high-performance services in ALWAN which are translated into executable parallel programs and interfaces that allow access to the services from CORBA clients.

Sources of *overhead* are discussed and some measurements are presented that provide a feeling for the possible behavior of an HPCS application. Examples from the fields of *volume rendering*, *signal processing* and *computational fluid dynamics* demonstrate how ACS is used in applications.

Preface/Roadmap

The presentation of our work is divided into five parts:

Introduction The general ideas and motivation for our approach to client-server support for high-performance computing, are discussed. We have a brief look at the two technologies, the integration of which is the goal of this work: distributed and parallel computing, *CORBA* and *ALWAN* in particular. Part I concludes with an overview of work that is related to ours.

Design and Implementation We introduce an abstract approach to the integration of high-performance computing and client-server infrastructure (HPCS), and present *ACS*, an implementation for the ALWAN parallel computing system and CORBA. The ACS server and client programmers' views are described. Finally, we discuss the overhead incurred by HPCS and some rough measurements made in the ACS environment to give the reader a general feeling for the possible behavior of HPCS applications.

Demonstration Applications Three applications using parallel servers are presented: our main example, a 3D volume rendering application from the medical domain (RDVOL), but also a signal processing library/application (LIBSIG/IDAHO), and a simple, educational computational fluid dynamics code (CFD). These applications were used as test cases for ACS and to demonstrate how it works.

Future Work and Conclusions We attempt a look at how this work could be extended, conclude from the presented work the position and value of our approach and summarize our experience during the development of ACS.

Appendix The appendix contains the ALWAN-to-IDL mapping and the source code (IDL) of the demonstration applications, as well as source code of the program used for measurements. Raw data from the timing measurements is also listed.

Acknowledgement

I feel grateful to a large number of people who have accompanied and supported me during my research and while writing this thesis.

First of all I would like to thank my supervisor (and more) Prof. Dr. Helmar Burkhart for giving me the opportunity to work in this exciting field, for providing guidance when needed but also the environment and freedom to pursue my own ideas, which made the work a truly enjoyable experience. I would also like to thank Prof. Dr. Béat Hirsbrunner for refereeing the thesis and for his valuable and much appreciated feedback.

A very special thank goes to my family, to my wife, Sandra, and my son, Sébastien, who both had to suffer through some difficult times, especially towards the end of my work. Although I must have often been quite unbearable, they still kept supporting me as only a loving family can.

My thanks also go to my parents Leny and François Prétôt and my brother René, my dear parents in-law, Monique and Heinz Eckenstein and my brothers in-law, Sacha and Fabian. They provided the net that sometimes kept me from falling.

Doing this work would not have been half the fun without the valuable discussions with my colleagues at the Computer Science Department, my family's dear friend and author of ALWAN, Dr. Guido Hächler, Dr. Edgar Lederer, the discussions with whom I will miss as much as the famous chocolate cakes, Robert Frank, who always seemed to be there and know an answer even for the stupidest questions and who kept the computers and networks running (most of the time, anyway) and Dr. Birgit Westermann, who is the author of the parallel Rdvol code and had to endure many long discussions about ALWAN, ACS and the amount of diskspace available on the file server.

Dr. Gonzalo Travieso maintained the ALWAN tools and removed a number of idiosyncrasies in the code (aka *bugs*) without which work on ACS would have been impossible. Kay Ruchti wrote the core of the parallel CFD code and was always willing to discuss and give feedback.

Other important sources for feedback were Dr. Niandong Fang, Dr. Beat Hörmann, Dr. Robert Sinclair and Dr. Walter Kuhn. I am also grateful for Romeo Dumitrescu's technical support that was always available when needed.

Last but by no means least of the people at the Computer Science Department whom I owe, are our former secretary, Regula Kneubühler and our current secretary Karin Liesenfeld. The good services and sup-

port that they provided are innumerable. I really don't know how they put up with us sometimes and still managed to bring color and sunshine into our offices.

I would like to thank the Swiss Science Foundation and the Swiss Federal Office of Education and Science for sponsoring parts of this work[1] and some of our collaborators in the projects, especially Patrick Jacques, Jean-Luc Viellé and Chi Ngo Duc of Spacebel Informatique S.A., Délphine Délavaux of Aérospatiale and Pierrick Beaugendre of the IRISA.

Grzegorz Mysliwiec worked on the parallel code for the LIBSIG server.

There are a number of people at Hoffmann-La Roche where I continued to work part time during the full duration of my work on the thesis. Dr. U. Goetz who very much supported my first steps "going back to school".

Prof. Dr. Don Kaiser, Mark Peeters, Dr. Alfred Steinhardt, Nathalie Legrenzi and Claudia Bohnet, my managers, who were kind enough to tolerate my part time presence in the office and who provided a work environment that allowed me to wear more than two hats without going absolutely crazy.

Much of what is left of my sanity I owe to my dear friends Maja, Tobi and Patrick of the *MMM*, with whom I so often enjoyed the pleasures of music, from which I drew pure energy.

There are so many people whom I have the honor of calling my friends, who accompanied me through this time and who continue to be an important part of my life. Any list would be incomplete. Thank you all!

[1] Projects: SPP IF 5003-034357 and ESPRIT HPCN project 21029, *PACHA*

Contents

List of Figures

List of Tables