Process Placement for Distributed Computations

by

Harry Ivor Rubin

B.A. (University of Minnesota) 1980

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy
in

Computer Science

in the
GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:
Professor Domenico Ferrari, Chair
Professor Chittoor V. Ramamoorthy
Professor Charles A. O'Reilly

1992

The dissertation of Harry Ivor Rubin is approved:

____

Professor Domenico Ferrari, Chair — Date

____

Professor Chittoor V. Ramamoorthy — Date

____

Professor Charles A. O'Reilly — Date

University of California at Berkeley

1992

Process Placement for Distributed Computations

ABSTRACT

Process Placement for Distributed Computations

by

Harry Ivor Rubin

Doctor of Philosophy in Computer Science

University of California at Berkeley

Professor Domenico Ferrari, Chair

The processes of a distributed computation can be assigned to the processing nodes of a distributed computer system in many different ways. The choice of how many processes to execute on each node (the *arrangement*), and which specific processes to execute on each node (the *configuration*), usually have a significant effect on the execution time of the distributed computation. The goal of this research is to find an effective and efficient placement procedure for homogeneous, loosely-coupled distributed systems. We simulated all possible combinations of program characteristics, and studied how the characteristics influence the performance of various arrangements and configurations.

We show that the numbers of arrangements and configurations quickly grow very large as the numbers of processes and processing nodes increase; the numbers are too large to handle on any ad hoc basis. Since the differences in response time between configurations can be very large, it is worthwhile to find and use good placement procedures.

Over a variety of different approaches to selecting the fastest configuration, we consistently found that the CPU demands of the processes determine which configuration will be fastest, and that these demands can be used to make fairly good placement

decisions.

Processes spend far more time waiting for other processes to send them messages than they spend waiting for those messages to travel through the communication network, even in extremely communication-intensive programs. This means that speeding up computation is much more important in improving response time than is speeding up communication.

In all the situations we have studied, spreading out the CPU demand of the processes as much and as evenly as possible, using as many processing nodes as are available, leads to the fastest response time. While this has been widely believed, it had not previously been demonstrated or quantified. We believe this is the best approach to making placement decisions for any numbers of processes and processing nodes.

Approved by:

____

Professor Domenico Ferrari Date

Acknowledgements

Thanks to Ethan Munson for help with the statistical analyses; Ethan gave most generously of his time and expertise. Thanks also to Jim Propp for recommending *Enumerative Combinatorics* by Richard Stanley.

My graduate advisor, Domenico Ferrari, has stood by me through it all. He deserves many thanks for his help, advice, and support.

I would like to thank all those in the U.C.–Berkeley Computer Science Division who have made my time here educational and enjoyable. I want to especially thank my fellow graduate students and friends.

Thanks, too, to the communities of which I have been a part.

I could not have accomplished this without the love and support of my parents and my sisters. They have my deepest thanks for everything.

My thanks to Wendy, for what we shared together, and for caring about me.

To Lael, who has meant so much to me, many thanks for the support, for the encouragement, for so much else, and most of all, for caring about me.

Abridged Table of Contents

Complete Table of Contents