Working paper, January, 1997.

Barriers for fun and profit

Thomas E. Warfel

Univ. of Pgh. Medical Center
Pittsburgh, PA xxxxx

David R. O'Hallaron

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Here is the abstract.

1 Introduction

To compute $\mathbf{y} = \mathbf{K}\mathbf{x}$ on a set of processors, we must consider the data distribution by which vectors and matrices are stored. The vectors $\mathbf{x}$ and $\mathbf{y}$ are stored in a distributed fashion according to the mapping of nodes to processors induced by the partition of elements among processors. If a node i resides in several processors (because i is a vertex of several elements mapped to different processors), the values x_i and y_i are replicated on those processors. The matrix $\mathbf{K}$ is distributed so that K_{ij} resides on any processor on which nodes i and j both reside. Figure 1 demonstrates this method of distributing data.

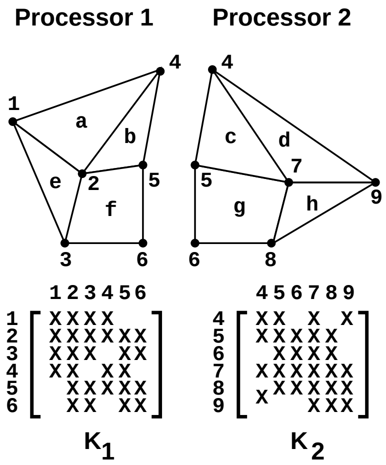

Figure 1: A finite element mesh

The processors of modern parallel and distributed computer systems are commodity microprocessors whose peak performance is doubling every 18 months or so. Clearly, communication

This research was sponsored in part by the Advanced Research Projects Agency/CSTO monitored by SPAWAR under contract N00039-93-C-0152, in part by the National Science Foundation under Grant CMS-9318163, and in part by a grant from the Intel Corporation.
The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either express or implied, of the U.S. government.

performance will have to improve as processor performance increases, but the question is how much and in what respects. It is crucial to understand communication requirements because some parts of the communication system will not be commodity, and are therefore expensive.

2 Conclusions

Here are the conclusions.