The Cedar System and an Initial Performance Study*

D. Kuck, E. Davidson†, D. Lawrie †, A. Sameh
C.-Q Zhu†, A. Veidenbaum, J. Konicek, P. Yew,
K. Gallivan, W. Jalby†, H. Wijshoff†, R. Bramley†, U.M. Yang
P. Emrath, D. Padua, R. Eigenmann, J. Hoeflinger, G. Jaxon†, Z. Li†,
T. Murphy, J. Andrews, S. Turner

Center for Supercomputing Research and Development
University of Illinois
Urbana, IL, 61801

Abstract

In this paper, we give an overview of the Cedar mutliprocessor and present recent performance results. These include the performance of some computational kernels and the Perfect Benchmarks® . We also present a methodology for judging parallel system performance and apply this methodology to Cedar, Cray YMP-8, and Thinking Machines CM-5.

1 Introduction

Tremendous progress in VLSI technology today has made it possible to build large-scale parallel systems with dazzling peak performances. Several such systems have even been commercialized over the last 5 years. However, the goal of building a general-purpose large-scale parallel system remains quite elusive. Many of the systems still have very little software support and are very difficult to program. The sustainable performance from various real applications running on those machines remains erratic and unpredictable. These phenomena show that we still do not know how to build parallel machines, how to program such machines, or how to characterize the performance of such machines for writing good application codes. The difficulties stem from the fact that in order to build a large-scale parallel machines that can deliver "practical parallelism," we need to understand the interactions among system architecture, system software and parallel application codes.

The Cedar project brought together a group of people in the areas of computer architecture, parallelizing compilers, operating systems, and parallel algorithms/applications to help solve the real problems associated with building a "complete" parallel system, and to study the effects of interaction among these components on such a machine [GKLS83, KDLS86]. The machine has been in full operation since late 1990. The Cedar experience includes the architecture, compiler, OS, and application perspectives and this paper attempts to summarize these for the architecture community.

We describe the machine organization in Section 2, concentrating on the unique aspects of Cedar. Programming and compilation for Cedar are discussed in Section 3. Performance measurements of the systems and the interpretation of those results are presented in Section 4.

2 The Organization of Cedar

Cedar is a cluster-based shared memory multiprocessor. The system consists of four clusters connected through two unidirectional interconnection networks to a globally shared memory (Fig. 1). Each cluster is a slightly modified Alliant FX/8 system with eight processors. In this section we first summarize the features of these clusters and then describe the unique features of Cedar. For a more detailed overall description of Cedar see [KTVZ91].

*This research was supported by the Department of Energy under Grant No. DE-FG02-85ER25001 and by the National Science Foundation under Grants No. US NSF-MIP-8410110 and NSF-MIP-89-20891, IBM Corporation, and the State of Illinois.

†Affiliated with CSRD for a portion of the project.