A Linear Time Algorithm for Sequential Diagnosis in Hypercubes *

Sanjeev Khanna†
Department of Computer Science
Margaret Jacks Hall
Stanford University
Stanford, CA 94305-2140

W. Kent Fuchs
Center for Reliable and High-Performance Computing
Coordinated Science Laboratory
1308 W. Main Street
University of Illinois
Urbana, IL 61801

April 27, 1994

Abstract

This paper describes a new sequential diagnosis algorithm for hypercubes. The algorithm is based on the PMC model and it assumes the existence of a central observer for syndrome decoding. If we denote the total number of processors in a given hypercube by N, then the algorithm achieves $\Theta(\sqrt{N \log N})$ degree of diagnosability using only $O(N)$ tests over all iterations of diagnosis and repair. The aggregated syndrome decoding time is also shown to be $O(N)$ for this algorithm. The number of iterations of diagnosis and repair needed by the algorithm is $O(\log N)$.

Keywords : Diagnosability, Fault-Tolerance, Hypercubes, Multiprocessor Systems, Sequential Diagnosis, System-Level Diagnosis.

1 Introduction

In fault-tolerant multiprocessor systems, an important aspect is the diagnosis of faulty processors. The process of identifying faulty processors in a system by conducting tests on various processors and interpreting the test results, is known as *system-level diagnosis* [19, 11].

A widely-studied model for system-level diagnosis is the PMC model [19]. This model assumes that a system can be partitioned into units capable of testing each other along available communication channels. When a unit tests another unit, it declares it to be fault-free or faulty based on the test responses. This evaluation is considered reliable iff the testing unit is fault-free. The faults are assumed to be permanent in this model. This model has been extended by several authors to certain probabilistic models where each processor may fail with a certain probability and a fault-free processor is able to determine the status of a faulty processor correctly only with a certain probability [7, 13, 4, 6, 20, 9, 10, 18].

*This research was supported in part by the Semiconductor Research Corporation under Contract 91-DP-109 and in part by the Department of the Navy and managed by the Office of the Chief of Naval Research under Contract N00014-91-J-1283.

†This work was done when the author was at the Center for Reliable and High-Performance Computing, University of Illinois at Urbana-Champaign.