Abstract

Assertions based on the knowledge of the application can be very effective at detecting corruption of critical data or design faults that transparent techniques are unable to deal with. Particularly, high coverage assertions are very promising because of their low-overhead and high error coverage.

This thesis starts by characterizing the behavior of assertions-based error detection mechanisms under faults injected according to a quite general fault model. The main problems are identified as being the lack of protection of data outside the section covered by assertions, namely during input and output, and the possible incorrect execution of the assertions.

To handle those weak-points the Robust Assertions technique is proposed, whose effectiveness is shown by extensive fault injection experiments on a realistic control application. With this technique a system follows a new failure model, that is called Fail-Bounded, where with high probability all results produced are either correct or, if wrong, they are within a certain bound of the correct value, whose exact distance depends on the output assertions used.

We claim that this failure model is very useful to describe the behavior of many systems with low hardware and software redundancy.

Any kind of assertions can be considered, from simple likelihood tests to high coverage assertions such as those used in the two paradigms investigated in this thesis, Algorithm Based Fault Tolerance (ABFT) and Result-Checking (RC).

ABFT is the collective name of a set of techniques used to determine the correctness of some mathematical calculations. The basic idea is to apply some encoding to the input data of the calculation, execute the algorithm on the encoded data, and check that the encoding is preserved at the end. A less well known alternative is called Result Checking (RC) where, contrary to ABFT, results are checked without knowledge of the particular algorithm used to calculate them. These two error detection methods based on high coverage assertions assure a nearly complete error coverage for most of the matrix operations that are the basis of all scientific computation.

A comparative study between ABFT and RC shows that RC is a good alternative to ABFT for detecting errors due to physical or design faults. Being the first time that RC is evaluated by fault injection, it is shown that although both methods exhibit equivalent error coverage, RC has lower overhead and has the important advantage of being independent of the underlying algorithm.

To substantiate the practical value of RC, we have developed a result checker for a problem for which no such checker existed: the computation of eigenvalues and eigenvectors, the so-called eigenproblem. The proposed fault detection mechanism based on the new RC, implemented on top of routines from the LAPACK library using the Robust Assertions structure, is simultaneously very efficient and very effective. It has less than 2% performance overhead for medium to large matrices and it exhibited an error coverage greater than 99.7% with a confidence level of 99%, when subjected to extensive fault-injection experiments.

We conclude that high coverage assertions can be used to build highly reliable low-overhead systems.

Keywords:

Error detection, Algorithm Based Fault Tolerance (ABFT), robust ABFT, assertions, robust assertions, Result-Checking (RC), fault injection, failure models, fail-bounded, matrix operations, eigenproblem.

Paula Prata, Dep. de Informática da UBI