Performance Metrics For Configurable Computing

Henk Spaanenburg
Honeywell Technology Center
3660 Technology Drive
Minneapolis, MN 55418

Metrics determine what kind of conclusions may be drawn from benchmark results, and also affect how benchmarks must be performed. The metrics in use in scientific computing benchmarks address mainly four questions:

1. how can a given algorithm be characterized
2. how good is an algorithm to solve a problem on a given implementation
3. how good is a machine across multiple problems, and
4. how effectively does an implementation-algorithm combination scale?

When evaluating architectural and packaging options for an architecture, one commonly encounters the problem of meeting performance requirements within the constraints of weight, volume and power envelope as well as the amount of computation performance that can be realized with a given physical envelope. This assessment process can be guided by a metric that we have found to be relatively consistent in past applications. It incorporates throughput in million operations per second (MOPS), weight (and implicitly volume) in kilograms, and power in watts. The MOPS/(kg.watt) ratio has been used to evaluate technology and packaging tradeoffs.

The claim is that given a particular technology (pre-VHSIC, VHSIC phase 1 and 2) and a particular packaging approach (representing various die size per real estate area ratios), the selected combination will produce a system where the MOPS/(kg.watt) is known to be within a certain order of magnitude. A commercial supercomputer, such as the Intel Paragon (Gamma), assuming 7.7 GFLOPS, 3000 lbs and 116 kW of power, would represent a 0.0005 MOPS/(kg.watt). A Honeywell militarized Touchstone (Sigma) avionics supercomputer, assuming 7.7 GFLOPS, 82 lbs and 2.9 kW power, would represent a 0.71 MOPS/(kg.watt). For example, future avionics systems need to have a MOPS/(kg.watt) on the order of several hundred (209) for a 1.8 GFLOPS/20 GOPS Touchstone enhanced radar preprocessor realized on one double sided (2 lbs), liquid cooled (200 W) SEM-E form factor board (1 GFLOPS ~ 10 GOPS).

In our experience, FPGA-based systems for any function perform at an order of magnitude better MOPS/(kg.watt) than a similar DSP-based implementation but still at an order of magnitude less than a full ASIC implementation. Our evaluation will confirm that observation, but in addition, will provide insight into the mechanism of reconfiguration and its related timing expense.

In addition to the above metric, others have been defined to measure the effectiveness especially of reconfiguration aspect of configurable computing devices. A direct function-for-function evaluation, especially relative to ASICs, is not the proper way to evaluate configurable computing. Of interest are some of the metrics as proposed by DeHon at MIT as part of their reinventing computing program.

Proceedings of The Thirtieth Annual Hawwaii International Conference
on System Sciences ISBN 0-8186-7862-3/97 $17.00 © 1997 IEEE

1060-3425/97 $10.00 (c) 1997 IEEE