

THE NEED FOR POWER BENCHMARKING OF RECONFIGURABLE ARCHITECTURES

Tobias Becker, Peter Jamieson, Wayne Luk
Department of Computing
Imperial College London
{tbecker,pjamieso,wl}@doc.ic.ac.uk

Peter Y. K. Cheung
Department of EEE
Imperial College London
p.cheung@imperial.ac.uk

Tero Rissa
Nokia Devices R&D
Finland
tero.rissa@nokia.com

Rapidly evolving standards, convergence of increasingly complex features and growing time to market pressure are pushing handheld consumer device manufacturers to consider alternatives to ASICs and microprocessors. There is a clear demand for power efficient circuits that are flexible, while capable of delivering performance through parallelism. Reconfigurable architectures, such as FPGAs, have potential in meeting the demand for flexibility and performance, but they often miss the power requirements by up to several orders of magnitude. To stimulate the development of more power-efficient device we propose a benchmarking suite for power and energy in reconfigurable devices.

The objective of this benchmarking suite is to allow a fair comparison of different configurable devices based on a number functionalities that are representative for the target product. It should prevent unreasonable optimisations specific to the benchmarks while stimulating general optimisations in the architecture.

Past attempts of FPGA benchmarks suffered from several shortcomings: The Prep benchmark was intended to measure the true logic capacity of devices but vendor tools were optimised to recognised test patterns and report favourable results. MCNC provides a range of simple circuits and state machines and has often been used for FPGA benchmarking. In MCNC, the circuits are specified as netlists which limits its applicability to a fine-grain configurable logic only and cannot take advantage of more efficient dedicated resources such as memories or DSPs. The circuits are also too simple and not representative of real applications and most importantly, lack input stimuli.

Benchmarking FPGAs hence confronts us with a number of challenges. Traditional processor benchmarks are often simply passed through a standard compiler to obtain executable code. Synthesisable RTL code would be a comparable input format for FPGAs, however this would be inadequate since too many design decisions are already made at this level and synthesis results are most likely suboptimal. This is due to the variation of architectural features inside different devices. These variations include LUT input size, special LUT modes such as shift registers or memories, LUT to flip-flop ratio or connectivity of the

routing fabric. Devices can also provide different dedicated blocks such as processors, DSPs or memories and last but not least, the low-power capabilities can vary widely from simple clock gating to advanced sleep modes and variable supply voltages.

In order to achieve an implementation that is optimal for a given device, the functionality of the benchmark needs to be specified on a higher level of abstraction. One example could be an encryption standard or an image processing function where the functionality is strictly specified but not its implementation. The benchmark user can then decide on the optimal way to map this functionality to a device given its available features. Hardend RAMs or multipliers for example are better choice in many cases than using fine-grain LUT-based logic.

Another crucial aspect of benchmarking are workloads that provide input stimuli to the benchmark circuit. The usual benchmarking scenario, where the circuit is tested under maximum throughput has limited significance for evaluating the suitability of the circuit for low power applications. Realistic stimuli have to create different levels of activation with varying active and inactive periods. Only this allows to assess the effectiveness of low power modes. The benchmark user can then for example measure the difference between processing chunks of data in short bursts with sleep states in between compared to continuous processing at lower clock frequencies.

To address the discussed issues we propose a benchmark suite with the following aspects: We provide a range of benchmarks that aims to cover the application space sufficiently. Each benchmark is specified on a functional level that allows user-optimised implementations but also a synthesisable version to allow quick tests and comparisons with less-optimal but possibly good-enough results. Benchmarks are combined with a range of workloads that can stimulate the circuit with long inactive periods or close to maximum throughput. This approach will require more effort from the benchmark user than traditional benchmarks, but will allow exploration and comparison of different architectures and features. We hope that this will ultimately lead to technological improvements enabling the use of FPGAs in mobile devices with strict power constraints.