

High level architectural modelling for early estimation of power and performance

Koen Claessen, Carl Seger, Mary Sheeran, Emily Shriver and Wouter Swierstra
Chalmers University of Technology and Intel

1 Introduction

We present a work in progress report on yet another project that aims to use programming language technology to help raise the level of abstraction at which microprocessor and system-on-chip design is done. We are fully aware that inside Intel alone there have been many such projects, none of which has succeeded. In this talk, we explain why we are trying again, and clarify the differences between our approach and the previous ones.

2 The problem we are trying to address

There is general agreement about the need to raise the level of abstraction at which complex integrated circuits are designed and verified. The increasing use of SystemC and SystemVerilog is evidence of this need. Yet, we and others feel that these languages and associated tool-sets are failing to deliver the expected benefits. For example, a move to SystemVerilog inside Intel has not given the promised speed-up in simulation. We see, once again, a window of opportunity for more advanced programming language technology, and particularly functional programming languages.

3 Our plan of action

We think that the key to making progress in raising the level of abstraction without losing too much control over the final low level details is to perform a se-

rious study of the abstractions involved. The simple abstractions that cleanly separated the various levels of design have served us well, but are now breaking down. Our aim is to build a domain specific language, embedded in Haskell, in which abstraction is a first class citizen. We plan to support fast (preferably parallel) simulation, early estimation of both power and performance, visualisation of the results of those estimations, and finally refinement from specification to implementation. In addition, we want to be able to do all the simulations symbolically, as well as concretely. You might say that we want to see how far you can get by pushing the limits of symbolic simulation, to cover areas like architectural power modelling, where it has not (to our knowledge) been much used before.

This project has started in February 2009, but it is based on earlier work by the authors. In the talk, we will survey related work and discuss some initial design decisions. Furthermore, we would like to discuss the many remaining open problems that we see. We aim to provoke a discussion about the reasons why previous attempts have failed. Perhaps one of those reasons is the fragmentation in our community – with too many individual efforts and too little real collaboration? A discussion of ways in which this community might work *together* to increase the chances of success this time would be a welcome result of this talk.