ARCHITECTURE SYNTHESIS FOR EMBEDDED SYSTEMS

A vision statement for multidisciplinary research in microelectronics

Nigel Topham

Vision Statement

I believe the research communities in Microelectronics, Computer Architecture, and Optimising Compilers need to engage in multidiscplinary research to address four critical challenges facing embedded systems over the next ten years and beyond. These challenges are **Complexity**, **Energy Efficiency**, **Flexibility** and **Tolerance of on-chip delays**. The over-arching vision of this proposal is "the creation of tools that will learn how to build future systems on silicon".

Justification and Relevance

Microelectronics is an implementation technology for complex digital and mixed signal systems. To exploit it effectively, those who design microelectronic systems must be able to create designs quickly, correctly, and with a good understanding of the expected performance, power consumption and cost. I am sure we are all familiar with the major trends in microelectronics: transistor densities are growing exponentially; wire delays increasingly dominate system timing; static power consumption in CMOS circuits is growing alarmingly; and, the cost of developing a chip using a leading-edge technology is also growing alarmingly. Together, these problems are stifling growth in the SoC industry, with a direct effect on the UK economy.

How can we translate these gross trends, and their associated design problems, into a coherent programme of research? I believe we must look at the effect they have on the design process, and how this in turn affects design cost.

A New Approach - Systems that Learn How to Design

As silicon geometries shrink below 90nm the effects of on-chip signal delays and static power consumption will change the landscape of system-on-chip design. Whereas in the past the performance characteristics were largely dependent on logical complexity and logic partitioning, future performance will depend on wire lengths, routing congestion, and physical partitioning. These aspects of design must be considered at a high level, and at an early phase in the design process. They can no longer be considered a mechanical, back-end, process. Physical effects will have a profound impact on architectures and the compiler technologies needed to overcome the problems of on-chip signal transmission delays.

Future computing systems will be severely constrained by these challenges unless new ways of designing complex system-on-chip devices are found. We need to find new ways of automating the design of low-power, high-performance, embedded systems based on tools that ***learn*** how to optimise the three principal degrees of freedom - architecture, microarchitecture and compilation.

Previous research in automating the construction of processors and compilers from architecture descriptions has not taken physical factors into account. Companies such as ARC, CoWare, Stretch, Tensilica, and others, have led the field in configurable and flexible architectures for embedded systems.

Today, some of these systems are able to perform a crudely-automated search for architectural improvements, but have no knowledge of how the systems will really behave when implemented in silicon. Architectural optimisation based on physical constraints remains a manual process that is both time-consuming and expensive. Based on current trends, it will become more time-consuming and more expensive.

When physical factors are taken into account the design space becomes hugely complex and difficult to navigate. The growth in design complexity, coupled with the increasing effects of wire delay, renders manual optimisation of embedded system design an almost intractable problem. This vision statement proposes a significant program of research to develop ideas, techniques and tools that are able to synthesise architectures, micro-architectures and compilers, and to learn how to optimise those systems based on information fed back from the physical domain.

A Research Programme in Intelligent Design Synthesis

The issues highlighted here suggest an extensive research programme to investigate fundamental techniques in the co-synthesis of micro-architectures, architectures and compilers. It would be multidiscplinary, in the sense that it would involve UK researchers spanning microelectronic circuit design, processor microarchitecture, embedded systems design, and high performance compilers. This programme would not seek to impose novel computational structures on a market that is unprepared or unwilling to accept radical change. Instead, this is an attempt to recognise the fundamental issues in SoC design and propose far-reaching and ambitious solutions based on advanced automation and machine learning.

Significance

This research would have a far-reaching impact on industry and academia by introducing learning mechanisms into the SoC design process. If successful, this research would lead to better SoC design, achieved with lower R&D cost, resulting in wider deployment of SoC technology.

Success or failure would ultimately be measured by whether the research achieves its primary goal of creating tools that learn how to optimise architectures, microarchitectures and compilers.

Communicating the aims of this research to the wider public could be challenging. The goals are somewhat technical, but the concept of an intelligent system that is able to design the next generation of electronic devices could capture the imagination of the general public.

A research agenda such as this would bring together knowledge and research from iterative compilation to low power circuit design, and this would benefit many areas of academic research. From a commercial point of view solutions to the four fundamental problems highlighted in this proposal would be highly valued.

Scale

A research programme with the goal of automating embedded system design using machine learning clearly has a very broad scope. I would anticipate that several research groups in the UK could bring their particular expertise to bear in distinct but complementary areas. The application of machine learning to compiler optimisation is already an active research area in the University of Edinburgh, and already involves international collaboration.

Timeliness

In the UK commercial arena there are several companies producing or developing configurable SoC components for embedded systems. Around the world there are many competitors, some with a distinct technical advantage.

The challenges addressed by this research program are already with us, and in the next few years will become increasingly problematic. A programme of research to address these challenges is therefore required urgently. However, the scale of the problem and the ambitious nature of a solution based on systems that “learn how to design”, makes this a 10-year objective.

What are the risks? Firstly, it’s worth pointing out an important risk that is *not* present. This research vision has a clearly defined goal, so there is no risk of wasting research funding to produce a solution to a problem that does not exist. I believe the main risk is associated with the multidisciplinary nature of this unified solution to a set of problem that are typically dealt with in different domains, namely physical design, microarchitecture, architecture and compiler optimisation. A managed programme of research, involving key academic groups and UK industry will minimise this risk.

Nigel Topham

Professor of Computer Systems
School of Informatics
University of Edinburgh

npt@inf.ed.ac.uk