

Summary of the Workshop on **Energy-Aware COmputing (EACO)***

Wednesday, 25th May 2011, 2-5pm

University of Bristol, Merchant Venturers Building, Room 4.01

EACO Workshop Coordinators: Kerstin Eder and David May

Please send any feedback to Kerstin.Eder@bristol.ac.uk.

July 7, 2011

1 Purpose

To bring together researchers and engineers with interests in *energy-aware computing* for discussions to identify intellectual challenges that can be developed into collaborative research projects. We strive to go significantly beyond the state of the art.

2 Participants

Steve Kerrison	Microelectronics Research Group
Daniel Schien	Sustainable Computing
James Hanlon	Microelectronics Research Group
Leon Atkins	Algorithms Research Group
Jose Nunez-Yanez	Microelectronics Research Group
Lorna Colquhoun	RED
Roger Shepherd	ST-Ericsson
Zuner Bortolotto	School of Physiology and Pharmacology
Chris Preist	Computer Science
Kerstin Eder	Microelectronics Research Group
David May	Microelectronics Research Group
Rob Hunt	Microelectronics Research Group
George Papaharalabos	BRL

*The EACO workshop series is supported by the Institute for Advanced Studies (<http://www.bris.ac.uk/ias>) and the Merchant Venturers School of Engineering.

3 EACO Workshop 1: “Intellectual Challenges”

3.1 Overview

The following is a list of *Intellectual Challenges* (ICs) identified at the first EACO workshop. The individual ICs are to be labelled with an estimated timeline, indicating **S** - short, **M** - medium or **L** - long term challenges. The labelling should be established based on discussions in consultation with the EACO community.

An interesting “classification” of ICs would be to distinguish between “incremental” approaches that advance the state of the art in a step by step fashion as opposed to “radically new paradigm shifting” approaches. Again, this classification provides scope for discussion within the EACO community.

3.2 Strategy to take ICs forward

It is expected that student projects can be set up to advance **S** challenges so that we can obtain early results on case studies. These can later be used to underpin research funding applications. We should make every effort to try and get these early results published. Challenges labelled **M** are prime candidates to put forward for internal funding at FEN or UoB level. Challenges labelled **L** are much longer term and we need to watch out for funding opportunities at national and international level, gradually building up or profile and track record making use of **S** and **M** challenges.

At this stage the challenges are listed in no particular order.

IC1 Learning from biology: More Pre-Processing. An approach based on approximate pre-processing with on demand more detailed focused computations should be more energy efficient than traditional fully accurate computations. Where can this be exploited?

In this context: **Selected Connectivity.** Can we design systems that “know” what processing will be required and that adjust accordingly?

IC2 The Dark Silicon Challenge. Radically new approaches are required to address the “Dark Silicon Challenge”, i.e. the fact that engineers can design systems so dense that power supply only permits supplying small fractions of them with power at a time. Can we turn the “Dark Silicon Challenge” into a feature inspired by biology? After all, the the brain only uses selected parts at a time.

One potential solution is increased heterogeneity in the design at the SoC level. With a surplus of transistors that can not all be powered at the same time, it will increasingly make sense to use some of the transistor budget for specialised accelerators. These accelerators may only be used a fraction of the time, but they can be highly optimised for performance and energy efficiency for these specific tasks. Increased heterogeneity of design presents significant challenges in terms of design trade-offs and for system-level programming.

IC3 Massively Parallel Extremely Low Power Processing. Build a board with many extremely low power processors and experiment with ways to use them inspired by biology (brain).

XMOS devices may be ideal to start this locally, but there may be other options to be explored with industrial collaborators who come forward.

IC4 From “Always On” To “By Default Off”. A paradigm shift is needed from gates to the application. For instance, demonstrate that event-driven computation is more energy efficient than traditional polling loops. A paradigm shift is needed towards a more energy efficient concurrent programming model.

The impact of moving from “always on” to “by default off”. It would be interesting to estimate the savings this could bring on a big scale, i.e. re-calculate existing energy consumption predictions for 10 to 20 years into the future based on the assumption that systems are “by default off” rather than based on what we currently have.

What is the impact of shifting from desk top PCs to tablets for consumption of electronic media? Is there a net gain or reduction in energy consumption?

IC5 Power/Energy Characterization of Processor Designs. Early estimation of power or energy consumption of processors/SoCs is needed to enable hardware design engineers to better understand the properties of the design as it is being developed.

The key challenges to achieve this are:

- **Early Provision of Use Cases.** These provide the typical workload for which a power or energy profile is required.
- **Fast Estimation of Power/Energy Profile of Hardware.** Current techniques are computationally expensive and rely on information that only becomes available late in the hardware design flow. Moving to higher levels of abstraction should help provide earlier profiles at the cost of losing accuracy.

IC6 Power/Energy Characterization of Software. What is the power or energy behaviour of key software use cases running on selected embedded systems, e.g. mobile platforms?

Can software power modelling be done independent of the target processor, i.e. without having to model the target processor (in detail)?

If software developers could understand the relative savings/costs of different software design options and algorithms then that would help them explore the design space wrt energy efficiency of software applications which adds a new dimension to the traditional software design space.

IC7 Parallel Design of Complex Software Systems. Re-design complex applications such as a web browser targeted at parallel execution with energy efficiency as a primary design goal. How does it differ from current solutions?

In this context: **Can we design efficient methods to distribute computation over resources?** How?

IC8 Low Power Multi-Voltage Design and Verification. Voltage domain crossing impacts functional design and hierarchy. Better understanding of the impact of design decisions is required to enable more informed decisions and to ease verification.

IC9 Dynamic Power Control of Mobile Computer Systems. When and by how much should a computer system change its voltage and frequency? This is currently a hard challenge. Prediction is needed to enable more informed decisions, particularly in environments where voltage/frequency scaling impact real-time behaviour and timing constraints.

IC10 More “Power” to Programmers. Programmers need to be able to design software with energy efficiency as a primary design goal.

One approach is to provide constructs in programming languages to explicitly express energy budgets within code e.g. in terms of assertions similar to how functional properties are being expressed in hardware design.

Another approach is to provide programmers with constructs that enable them to control the degree of accuracy of the computation at program level, explicitly trading accuracy for energy efficiency.

Static analysis should enable more informed decisions during design space exploration so that programmers better understand the tradeoffs and consequences of their decisions.

IC11 Raising Energy Consumption Awareness. Can we provide user feedback for transactions or computations? Raising end-user awareness may help change end-user behaviour towards reducing energy consumption.

IC12 Energy Autonomous Robots. Some robots, e.g. the EcoBot, generate their own energy. It is essential that accurate estimates become available so that such systems can assess the feasibility of a task in term of there being enough energy to complete it before engaging in its execution. These estimates go beyond estimating the cost of computation and require energy consumption modelling at system level including mechanical parts of the robot.

IC13 Communication vs Computation. For different architectures and applications, show the trade-off between communication vs computation in terms of energy consumption and other design goals such as performance and code size. This can be explored at all levels of abstraction in the system stack, from algorithms down to architecture and micro architecture.

IC14 Re-evaluating Algorithms for Energy Efficiency. Finding the best “hardware fit” for an algorithm is a challenging task. This is further complicated by the performance focused optimizations of compilers and at microarchitectural level. Some of these are counter-productive when it comes to optimizing for energy consumption. Research questions like *Can we find an architecture or hardware configuration where Bubblesort outperforms Quicksort wrt energy consumption?* provide interesting material for discussion and for student projects.

In this context we also need to invest effort into the: **Re-evaluation of Traditional Metrics for Computational Complexity.** The metrics traditionally used to determine the computational complexity of an algorithm in terms of the input data size are not precise enough to capture energy consumption. For instance, the number of data items does not capture the amount of switching in the hardware caused by the data stream.

Overall, an interesting research programme may be a general “**RE-ASSESSMENT OF SYSTEM DESIGN METHODS FOR ENERGY EFFICIENCY**” covering all aspects of system design from applications / algorithms down to the hardware.

IC15 **Education.** Software developers need to gain an insight into what consumes energy when a program is being executed. They need to develop an intuition to be able to anticipate the impact their design decisions have on energy consumption during execution. This requires tools that enable energy-aware design space exploration and enhanced curricula that promote energy efficiency to a first class design goal for software developers.

4 Rationale for the EACO Initiative

Energy efficiency is now a major (if not the major) concern in electronic systems engineering. The European Commission's 2011 Work Programme on ICT has a strong focus on making systems energy efficient. The next funding call, Call 8, is expected in July 2011 with a deadline in January 2012. The 2011 Work Programme includes "Minimising Energy Consumption of Computing to the Limit (MINECC)" for Call 8. The aim is to "lay the foundations for radically new technologies for computation that strive for the theoretical limits in energy consumption while maintaining or even enhancing functionality and performance." In particular, one target outcome is focused on "software models and programming methodologies supporting the strive for the energetic limit (e.g. energy cost awareness or exploiting the trade-off between energy and performance/precision)." In the UK, the government and the research councils have made Energy Efficiency one of the top priorities of their programmes, for details see <http://www.rcukenergy.org.uk>. This is the right time to identify and address the intellectual challenges that lead to energy-aware computing. With a modest amount of investment at this stage we could be in a very strong position to take a lead in this field.

5 Background

Traditionally, researchers and engineers work within one or perhaps two layers of the *system stack* with very limited overlap, e.g. software engineers, computer architects or hardware designers. However, energy-aware computing is a challenge that requires investigating the entire system stack from application software and algorithms, via programming languages, compilers, instruction sets and micro architectures, down to the design and manufacture of the hardware, or alternatively bottom up. This is because energy is consumed by the hardware performing computations, but the control over the computation ultimately lies within the software and algorithms, i.e. the applications running on the hardware. It has recently become clear that, while hardware can be designed to save a modest amount of energy, the potential for savings is far greater at the higher levels of abstraction in the system stack. The greatest savings are expected from energy-consumption-aware software. Hence, addressing the challenge of energy-aware computing requires collaboration between researchers from all the above named areas and a good understanding of the applications that will drive software and hardware development in the future.

6 Opportunity

The current activities on energy efficient computing within the Microelectronics research group already cover almost all layers traditionally related to microelectronics in the system stack. This is a unique strength of the Microelectronics research group and puts us into an excellent position to make significant advances by collaborating. However, we must reach out to the upper layers of the system stack, including algorithms and applications, ranging from pervasive computing in health care, via robotics to the large scale complex IT systems that enable cloud computing. We must understand the wider impact on sustainability and end-user behavior, and also alternative models for energy efficient computation such as those inspired by biology.

7 Aim and Objectives

A series of three workshops on Energy-Aware COmputing has been funded by the Institute for Advanced Studies in collaboration with the Merchant Venturers School of Engineering to bring together researchers from diverse backgrounds with a common interest in pushing the boundaries of Energy-Aware COmputing. The aim of the EACO workshop series is to develop a new research initiative on Energy-Aware COmputing at the University of Bristol. Our initial and longer term objectives are to:

- bring together researchers for discussions with the aim of finding common ground,
- take stock and position ourselves,
 - Where are we now? What is required in the future?
 - What is the state of the art? Where are the barriers?
 - What are the intellectual challenges?
- formulate a joint vision,
 - How can we advance the state of the art?
 - How can we address the intellectual challenges?
- decide on a research and collaboration strategy,
- network nationally and internationally (both within the EU and further abroad) with leading experts in the field (both in academia and in industry) to investigate the potential for research collaboration, and
- prepare to compete for research funding to kick-start this new initiative e.g. by
 - addressing funding bodies at national (EPSRC) and international (EC) levels as well as
 - raising industrial sponsorship for our research.