

Priming the multicore pump

How will engineers approach the challenge of designing chips with hundreds – maybe thousands – of cores? By **David Boothroyd**.

How do we maintain the constant improvement in processor performance that has characterised electronics for decades? Answer: use more processor cores. Multicore processors are widely seen as crucial for the future of electronics and they are already out there, working successfully.

But typical mainstream devices only use a small number of cores, typically fewer than 10. What about chips featuring hundreds, or even thousands – a form of chip architecture never dreamt of in the early days of the semiconductor revolution? Many feel such chips will pose radically different – and difficult – challenges for chip designers, especially in terms of energy consumption and reliability.

That is certainly the philosophy underlying a multi institute project called PRiME, standing for Power-efficient Reliable Many-core Embedded Systems. And embedded systems are seen as a major potential user of massive multicore chips – in such application areas as automotive, industrial automation, telecommunications, consumer electronics and health and medical equipment.

Director of PRiME is Bashir Al-Hashimi, ARM professor of electronics at the University of Southampton. The idea of the PRiME project started around March 2012, following Prof Al-Hashimi talking with colleagues from Imperial College London and Manchester and Newcastle universities about combining resources to solve some of the key challenges in many core processing.

Several companies are involved too – ARM, Imagination Technologies, Altera, Microsoft Research and Freescale. Facilitating this kind of multi institute project is the aim of Programme Grants, a new initiative run by the Engineering and Physical Sciences Research Council.

“If it wasn’t for that initiative, it probably wouldn’t have been possible for all the partners involved in PRiME to come together,” said Prof Al-Hashimi.

PRiME, which began in earnest in May 2013, is targeting two fundamental technical challenges facing many computing and electronic devices: energy efficiency; and hardware reliability due to transistors wearing out.

“PRiME’s basic premise is that, unless we find solutions to these challenges, even as many core becomes available, they will not be that useful, especially in mobile and embedded applications,” said Prof Al-Hashimi. “We have to provide new hardware and software technologies that overcome these fundamental bottlenecks.”

Prof Al-Hashimi believes these will start to have a major impact within five years.

“We are predicting that, by 2020, processors with more than 100 cores will become available commercially and enable embedded devices that have the power of a supercomputer, but with the battery life of a smartphone. But what is used today to solve the energy and reliability challenges will not be good enough.”

Today, a typical multicore embedded or mobile device will have up to eight cores, but it is not just a matter of increasing the sheer number of cores to achieve the required enhanced performance. A move towards more variety in the nature of the cores used – greater heterogeneity – will be equally critical.

“We think one effective approach to reducing energy is not to treat all cores equally. This is not a new concept – ARM already has a processor based on this approach – the big.LITTLE – in which the ‘little’ processor does the mundane tasks, such as handling email, while the ‘big’ one does more processing intensive tasks like video decoding.”

Addressing systems with 100 cores or more


PRiME is looking to expand this to systems with 100 and more cores and, as Prof Al-Hashimi says, a key question is how to distribute applications across so many.

“This is what PRiME will try to address by developing intelligent software. But one key question is how many different kinds of core might be needed.”

Exactly how many is impossible to predict now, but the team expects several.

One technique likely to be critical is the use of adaptive mechanisms that will allow highly parallel embedded systems to manage and optimise their behaviour dynamically to achieve the desired system trade offs between performance, energy and resilience to hardware failures, guided by feedback received from lower system layers.

“When you have hundreds of cores of different kinds, many decisions must be made – for example, about which core a part of an application should run on and at what speed that core should run,” said Professor Michael Butler, who is leading the research on runtime management. “Tasks that need to communicate a lot need to be on cores close to each other. To do this, you need a smart runtime system – an extension of the operating system – that is making those decisions.”



Cover photos: Chris Sharp

For this to happen, the runtime system has to monitor what is happening, and some existing hardware does not facilitate this. So one PRiME target is to develop ways of giving the runtime more fine grained information to enable it to make the decisions.

"We will do this through smart software that works on information gathered from the hardware," Prof Al-Hashimi added.

Another key PRiME target is to achieve better 'energy proportionality' than current systems manage. Many systems today do not scale energy proportionately with the performance requirement. Typically, you reach a certain minimum level, after which it does not fall any further, even if the performance requirement does. With hundreds of cores, this will be a critically important factor. Achieving better power proportionality will require designers to change strategies towards embedded systems.

"Power proportionality thus becomes an issue affecting the performance across the layers of abstraction," Prof Al-Hashimi continued.

One of PRiME's themes is therefore to investigate the ways in which an embedded many core system consumes and distributes energy whilst maintaining its dependability in a wide dynamic power range. This should allow for

"We think one effective approach to reducing energy is not to treat all cores equally."

Prof Bashir Al-Hashimi

multiple modes, as well as graceful degradation of performance and Quality of Service.

It is impractical to always build working chips to investigate such behaviour deeply, so much of PRiME's work – at least initially – is the development of formal models, mathematical models and algorithms that enable the simulation and analysis of the behaviour of many core systems.

"These are being used to reason about energy reliability interplay, enabling potential trade offs between energy efficiency and resilience in many core architectures. For example, we have been looking at ways of measuring the power consumption at runtime and it is clearly fundamental to our approach. One way of doing it is by using performance counters on processors. Effectively, these are very low level counters that look at things like cache misses and how they correlate to the power consumption of the chip."

But not all of PRiME's work is theoretical. Demonstrators of its research results have already been developed – so far, mostly multimedia applications – but more will come as the project continues. One aim is to benchmark the new technologies it develops against systems that do not actively manage the energy reliability trade off.



This involves two different approaches. The first will use FPGAs to implement RISC soft cores for energy/reliability management. Existing examples without active management include the BERI FPGA soft core and the RAMP Blue platform: an FPGA system containing 1008 Xilinx MicroBlaze cores (32bit, 90MHz) running off the shelf applications and scientific benchmarks.

“Existing platforms, soft cores and many core OSs such as these will be used to evaluate energy efficiency, fault tolerance parameters and their interplay,” said Prof Al-Hashimi. “The pioneering work of ClearSpeed and systems such as the XMOS Xcore are beginning to emerge which consider these parameters.”

Cross layer collaboration

Using this understanding, and as outputs from PRiME’s other research themes become available, an innovative FPGA based platform will be developed to support benchmarking. This platform will incorporate architectures and cross layer collaborative mechanisms developed from PRiME’s formal models.

“It is envisaged the platform will contain heterogeneous and scalable processors with up to 1024 cores, representing the requirements of high performance embedded systems and future applications over the next 5 to 10 years,” Prof Al-Hashimi said. “The platform will also take advantage of new technologies that become available during the project, using IP from our industrial collaborators (for example ARM’s big.LITTLE, Altera’s NIOS II) where available.”

The second approach will investigate novel computer architectures, such as those inspired by biology and the human brain – one being SpiNNaker, developed over the last few years at Manchester University. This is based on a custom multiprocessor SoC featuring 18 ARM968 processor cores. Its novelty lies in the interprocessor communication mechanism that enables large numbers of

The PRiME Project group at Southampton.

Foreground: Prof Al-Hashimi

Background, from the left : Dr Geoff Merrett; Dr Tom Kazmierski; Dr Anup Das; Gery Scott; Luis Maeda-Nunez; and Dr Rishad Shafik.

very small packets, each representing a neural ‘spike’, to propagate across the machine in much less than 1ms – the requirement for biological real time.

“The machine will ultimately scale up to a system with more than 1million ARM processors and, at this scale, fault tolerance and energy efficiency are significant engineering concerns,” Prof Al-Hashimi contended. “As an existing state of the art platform, PRiME will use SpiNNaker to identify and develop the new runtime approaches to reliability and energy efficiency.”

SpiNNaker could help the PRiME project in a number of ways. One potential option for reliability is to leave some cores uncommitted to provide fault tolerant redundancy. Unused cores could also be power gated to minimise power dissipation and protect against the more ‘permanent’ faults.

“While inter chip links are fault tolerant, more pertinent to PRiME is the system’s ability to recover from software crashes. In order to permit run-time recovery, we will consider forms of checkpointing; for example, storing redundant data in non local memory to reduce the chance of corruption. State restoration will also be investigated, migrating tasks to another core or, in the worst case, a different chip.”

Highly parallel complex systems

Prof Al-Hashimi admits the principle underlying the design of SpiNNaker is different to the way one would design a high performance embedded system – for example, it doesn’t have what would, by today’s standards, be considered an operating system. But the link with PRiME and SpiNNaker is that it brings great expertise over many years in the design of highly parallel, complex computer systems.

The spinoff benefits of PRiME will not just be technological, but personal, Prof Al-Hashimi believes.

“There are more than 40 people working on PRiME, most of whom will make their way into UK companies. So projects like this are great training grounds.”

A future challenge for many core systems is that application developers will face architectures that are changing all the time and it is unrealistic to expect them to alter their code constantly to deal with this.

Dr Geoff Merrett, who is leading PRiME’s application theme, said: “We are trying to protect them – for instance, the runtime system will take care of jobs like deciding at what speed the processor will run. But we will also need some information from them, like what quality of service is needed in areas like video decoding.

“Generally, applications for many core systems are evolving and can be expected to be vastly different in five years; they will be selected and refined throughout PRiME’s duration,” he said.

Pushing embedded systems up to hundreds, even thousands, of cores will be one of the most demanding challenges the electronics industry has ever faced. If it succeeds, PRiME will have been a true pioneer of that development.