# Research Statement

Ali Ibrahim

## Motivation and Vision

Advances in wireless communication systems combined with the high penetration of mobile subscribers is creating a continuing demand for increased bandwidth and better quality of service. Future mobile terminals need to support higher data rates, full motion video and multimedia applications, a variety of wireless standards, be energy efficient, flexible, have a low time to market and be low cost.

Following Moore's Law, processors have become powerful, mobile, and embedded. This improvement is insufficient for portable computing and communication devices that increasingly employ rapidly evolving, sophisticated and power hungry signal processing techniques. Power dissipation, in the last few years, has emerged as a major concern. Projections on power density increases due to CMOS scaling indicate that this is one of the fundamental problems that will ultimately restrict scaling gains in the future. As a consequence, new architectural approaches will be needed in order to meet the performance and flexibility needs within acceptable energy budgets.

Embedded systems are pervasive in all aspects of our lives, and they come in many varieties such as PDAs, cell phones, digital cameras, etc. To obtain the desired level of energy-efficiency, most of these systems use specialized components. Current mobile phones, for instance, have a specialized baseband processing unit, and hardware acceleration units for some applications such as Mpeg decoding, turbo decoder, and cryptography. The number of these specialized units will even increase in the future considering that state-of-the-art embedded technology is the basis for substantial innovation in many other domains. This product differentiation has a tremendous amount of importance when power, energy, flexibility, and cost are considered. Designing an energy aware system that delivers the required performance and quality of service is a very challenging task, automating this process is even more challenging.

The key opportunity in the design of low power, high performance embedded processor relies on the opportunity for customization. Application specific integrated circuit (ASIC) design is the ultimate customization. In the ASIC approach, the circuit is customized to completely match the application characteristics. This level of specialization makes it difficult for the implementation to support multiple applications, or even evolutionary algorithmic improvements. Technology reuse in embedded systems is an important factor in increasing productivity and reducing costs. ASICs are also expensive, and incur lengthy design cycles. These aspects are poor matches for the rapidly evolving cellular telephony arena.

Designing, validating, and integrating (under shrinking design-time schedules and costs) energy-aware embedded systems will be the focus of future research. Over the next few years I plan to attack this problem from a variety of angles which can be divided into the following categories: Design automation of domain specific processors, energy aware domain specific compilers, and system design methodology. This research agenda will yield a body of techniques and tools that allow

skilled but non expert systems designers to quickly create and integrate efficient domain specific processors.

## Domain Specific Processor Design

A domain specific design approach, in the ideal case, makes it possible to reuse the design for a number of applications that fall in a particular domain. My research into in this field will be divided into two areas: architectures and compilers.

From the architecture prospective, there are many parameters that can be taken into consideration in the customizing process, this is mainly due to the numerous components that form an embedded system. These range from the memory system to processing or execution cores, and includes interconnection networks. Design parameters can also be divided into different categories. For instance, customization in the execution core can be in the address generation units, functional units, register file organization, bypassing logic, compression techniques, etc.

Compilers have a tremendous importance in the success of these processors. To fully take advantage of the architecture undertaken, a new style of compiler scheduling techniques is intrinsic to the research. For instance, the scheduling algorithms should: 1) explicitly manage the interconnection network between the architecture various components, 2) consider the application constraints such as power and performance, 3) be a combination of general purpose processors scheduling algorithms (e.g. modulo scheduling), and FPGA place&route algorithms to account for the limitation in the interconnection network.

Most of the issues that are related to design, implementation, and measurement of the performance-energy characteristics of a domain specific processor were addressed in my Ph.D. work [1, 2, 3, 4], and other projects that I worked on [5, 6]. In this work, my focus was in the design of a domain specific processor for wireless applications called ACT (adaptive cellular telephony). Even though the focus of the work was on the wireless domain, similar techniques can be applied to other domains [7].

In ACT, a high energy-delay efficiency was achieved through software controlled distributed memories, modulo addressed distributed single ported register files, multi-level reconfigurable interconnects, semi-reconfigurable address generation units, SIMD-ALUs, compression techniques, and extra hardware to support special wireless operations. The processor is basically a fine-grain VLIW architecture. The fine-grained software control provides considerable generality, and efficiency in term of energy-delay since different pipelines can be dynamically reconfigured to support a new processing phase that resembles data flows found in an ASIC implementation.

Hardware optimization by itself may not completely improve the efficiency of a program if a naive mapping is applied. In [8] we have shown how the performance of loop-carried-dependency algorithms can significantly improve if a combination of context switching, data flow retiming, and modulo scheduling are applied.

The efficiency of the processor was studied by mapping several algorithms taken from 3G wireless baseband, voice over IP, typical DSP algorithms, and kernels taken from the MPEG decoder to a verilog implementation of the processor. In general, the processor is within one to two orders of magnitude of the energy-delay of an ASIC implementation and is three to four orders of magnitude more efficient than a low power embedded processor implementation.

The aims of my future work, based on the current high performance low power architecture, ACT,

are: 1) develop algorithms and methods to automate the design of a domain specific processor and 2) investigating compiler scheduling techniques for such domain processors. Design automation can be achieved by placing restrictions on the inter- and intra-cluster communications and by varying the degree of reconfigurability and organization of the address generation units, the register files, the specialized function units, the scratch-pad memories, etc. Compiler scheduling can be improved by iteratively assigning groups of operations to groups of function units and performing routing at the same time. This will improve the scheduling quality through the use of global information as compared to an algorithm which places and routes one operation at a time.

Once the above tasks have been completed, the developed automation tools will allow us to 1) identify potentially power hungry sections of an application early in the design process, 2) characterize the effect of the design choices on the overall performance and energy of the application, 3) investigate other tools that make the design of complex systems more manageable.

One technology that will significantly benefit from this research is software radio (SR). SR is the science of implementing radios using software. It is expected to be a key component in future 4G wireless standards. It aims to get the software as close to the antenna as is feasible. Configuration on the fly can be done depending on what these radios need. Also they should be able reconfigure themselves based on the environmental conditions. I would like to investigate the interactions and the effect of the physical layer (hardware) on the embedded software design. I would like also to investigate the privacy and information security issues of such system.

## References


[1] A. Ibrahim, M. Parker, and A. Davis. Energy efficient cluster coprocessors. *International Conference on Acoustics, Speech, and Signal Processing*, May 2004.

[2] A. Ibrahim and A. Davis. ACT: A cluster co-processor for 3G wireless baseband. *Submitted to IEEE Transactions on Signal Processing*, Dec 2004.

[3] A. Ibrahim, A. Davis, and M. Parker. ACT: A low power VLIW cluster coprocessor for DSP applications. *Submitted to EuroMicro Conference on Real-Time Systems (ECRTS05)*, 2005.

[4] A. Ibrahim and A. Davis. Address acceleration mechanisms for an adaptive cellular telephony processor. *Submitted to International Conference on Multimedia and Expo (ICME) 05*, 2005.

[5] http://www.cs.utah.edu/impulse.

[6] J. Mahovsky, B. Wyvill, A. Davis, and A. Ibrahim. Robust ray-bounding volume hierarchy traversal with reduced precision integer arithmetic. *Submitted to Journal of Graphics Tools*, 2005.

[7] B. K. Mathew, A. Davis, and A. Ibrahim. Perception coprocessors for embedded systems. *Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia)*, 2003.

[8] A. Ibrahim and A. Davis. Exploiting data context switching in a low power VLIW coprocessor. *Workshop on Optimizations for DSP and Embedded Systems (ODES), in conjunction with IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, 2005.