

The Reinvention of the Microprocessor

Chris Rowen
Tensilica, Inc.
3255-6 Scott Blvd
Santa Clara, CA 95054
+1 408 986 8000

rowen@tensilica.com

ABSTRACT

The microprocessor has been with us for more than 30 years, and has evolved in response to available silicon technology and to electronic system requirements. Now, shifts in basic silicon scaling and embedded systems are forcing a significant shift in the architecture, the use and the design process for microprocessors.

Three fundamental trends in the design of embedded systems are driving a basic redefinition of the microprocessor. First, total system complexity, especially software complexity, is growing dramatically. Second, the cost, power, and size constraints of high-volume systems (e.g. digital televisions, mobile phones, games, digital cameras, printers, etc.) make high-integration silicon implementation not just attractive, but mandatory. Third, Moore's Law transistor scaling is reaching limits in circuit power density. These trends underlie important shifts in electronic design, including

- *greater importance in finding and exploiting system **concurrency**,*
- *increased use of **processors**, not only in traditional control and DSP roles, but also in roles previously reserved for hard-wired multi-media, security, and protocol-processing tasks, and*
- *a basic shift in the architecture of processors, away from complex one-size-fits-all architectures running at high frequency, to much more energy-efficient processors optimized for each **system-on-chip application platform**.*

This talk addresses the basic changes underway in the electronics environment. It highlights how "Moore's Law" silicon scaling continues to give designers more transistors, but no longer offers the power improvements of the past. In fact, processor clock frequency is now constrained by battery and package limitations, so continued system-level performance gains must come primarily through architectural, rather than circuit or silicon improvement.

This talk also outlines the evolution of architectural and micro-architectural features of processors over the past three decades and sketches the (often modest) performance

benefit and the (often severe) silicon cost of major innovations, such as microcode, pipelining, superscalar, out-of-order execution, caches, coherency, SIMD, and simultaneous multi-threading. When processors were developed separately from the rest of the system, this generic optimization served system needs well. However, the emergence of, automatic generation of complete processor hardware and software dramatically improves processor efficiency. And in the era of application-specific chips, creating a new processor for each new chip family no longer carries a cost or risk penalty.

State-of-the art processor generation offers four dimensions to the configuration of the microprocessor – instruction set, memory system, processor interface, and processor control functions. Together, these enable quick creating of an enormous range of optimized processors at various degrees of application specialization. Moreover, these processors incorporate versatile inter-processor communications channels that permit order-of-magnitude improvement in communications bandwidth and energy efficiency. Automatically generated processors change the design of electronic systems in two basic ways:

1. *Tuned processors are smaller, faster and more energy efficient than the generic processors that preceded them.*
2. *Tuned processors also serve as an alternative to hard-wired logic blocks, bringing full programmability from high-level languages to high-performance functions.*

These two effects make it both possible and desirable to build application-oriented chips using processors as a basic building block – 5 or 6 processors per chip has become routine, even in consumer systems and chips with almost 200 32-bit processors are used in production network routers. The optimization of processor configuration and processor-to-processor communication has been a central theme in system-on-chip development. This shift from generic processors, often with hardwired logic accelerators, to configurable processor-based system design appears fundamental to improved design productivity and end-product efficiency.