Performance: The Future of Embedded Processing

MIPS Technologies, Inc.

Electronic devices are becoming increasingly sophisticated, at a rate that is nothing short of astounding. The market has an insatiable appetite for new features, as long as they are easy to use. Devices continue to converge, while new, more complex digital technologies enter the market at a staggering pace.

As if the resulting pressures on SOC designers aren't bad enough, their job is made even more difficult by relentless economic and competitive challenges. Not only must they design ever more complex devices, they must deliver them to market more quickly. Their devices have shorter and shorter life expectancies, the market place is increasingly fickle and unforgiving, and intense competition makes differentiation—and profit margins—more and more elusive.

Semiconductor technology advancements offer some relief, but bringing these advancements to bear in a timely fashion is itself a nearly intractable problem. The age-old technique of pasting together patchwork blocks of rigid specialized logic will not meet this challenge in the future. What's needed is a more flexible, more scalable solution. It should allow product features to be bound late in the design cycle, allow designs to be easily modified and upgraded to extend product lifetime, and allow designs to be reused to amortize the rising fixed costs of design and manufacturing.

One alternative is to construct SOCs from more regular, reusable, programmable building blocks, like microprocessor cores. While many believe this approach to be "inefficient" compared with the old specialized custom-logic approach, that perception is based on yesterday's problems and bygone technology. The rapidly approaching era of 100nm silicon features, 300mm SOI wafers, and 10-layer copper interconnects with low-*k* dielectrics changes everything. In this era, the dominant concerns shift from the efficient use of individual transistors to the more perplexing problems of reducing design time (i.e., time-to-market), increasing design flexibility and reuse, and taking better advantage of the awesome leverage provided by the semiconductor-technology learning curve.

Key to this solution will be high-performance, high-throughput processor cores. These cores will have considerably more horsepower than those typically found in today's SOCs. They will have enough performance headroom to enable high levels of software abstraction, have enough speed to easily meet real-time signal-processing deadlines without tedious low-level optimization, and

have enough processing and throughput capacity to handle multiple simultaneous multimedia data streams, which will characterize tomorrow's embedded devices.

Although high-performance cores are often associated with high cost and high power consumption, the association is valid only when myopically considering the core itself. In the larger context of a system, performance can be translated into overall lower power consumption and lower cost. For example, mathematically complex coding algorithms can dramatically reduce the signal power that must be broadcast to transmit error-free data across a noisy bandwidth-limited channel. High processor performance can also be utilized, for example, to implement soft peripherals, reducing the cost of other components and the total cost of integration.

Designers will continue to come under increasing pressure to get more complex devices into the market more quickly and to have those devices live longer in the marketplace. In response, they will be forced to turn to softer SOC designs embodying higher performance, higher throughput microprocessor cores. The benefits of this trend will be enormous and will be apparent to designers and end users alike.

Copyright © 1998-2002, MIPS Technologies, Inc. All Rights Reserved.