# MPSOC Design: The logical approach to IP reuse and embedded SOC development

By Chris Rowen, PhD, Tensilica, Inc.

# MPSOC Design: The logical approach to IP reuse and embedded SOC development

The rapid evolution of silicon technology is bringing a new crisis to system-on-chip (SOC) design. To be competitive, new communication, consumer, and computer product designs must exhibit rapid increases in functionality, reliability, and bandwidth and rapid declines in cost and power consumption.

At the same time, there is intense pressure on chip designers to develop increasingly complex hardware in decreasing amounts of time. Unless the industry adopts a more efficient and flexible approach to SOC design, the return-on-investment hurdle will simply be too high for many products. Both semiconductor design starts and the global rate of electronics innovation will decline.

The SOC design team faces a set of daunting challenges:

- **Design effort:** The design effort for large SOCs is exploding. As the blocks get more complex, Verilog- and VHDL- based logic design runs out of steam.
- **Verification difficulty:** The complexity of typical logic blocks – hence the number of potential bugs - grows much more rapidly than gate count. Design teams report spending 70% of their development time verifying their designs.
- **Cost of fixing bugs:** The larger design teams, bigger NRE fees, and lost profitability and market share make show-stopper design bugs intolerable.
- **Late hardware/software integration:** Software integration, the last step in the system-development process, routinely gets blamed for overall program delays. Late hardware/software validation is a critical risk for new product-development projects.
- **Complexity and change in standards:** The explosion in the number, complexity and rate of change of industry standards is obsoleting existing methods and building blocks. Complex new standards demand much greater computational throughput.

Although general-purpose processors can handle many tasks, they usually lack the bandwidth needed to perform complex data-processing tasks such as network packet processing, video processing, and encryption. Chip designers have long turned to hardwired logic to implement these key functions.

**Moore's Law = Opportunity + Crisis**

In 1965, Gordon Moore prophesized that integrated circuit density would double roughly every one to two years. Today, it is perfectly viable to build SOCs with more than 100 million transistors. Within a couple of years, we'll see billion-transistor chips built for complex applications. Unfortunately, the design task associated with these gargantuan chips is correspondingly terrifying. This phenomenon is captured in the Semiconductor Research Corporation's comparison of logic complexity and designer productivity shown in Figure 1.

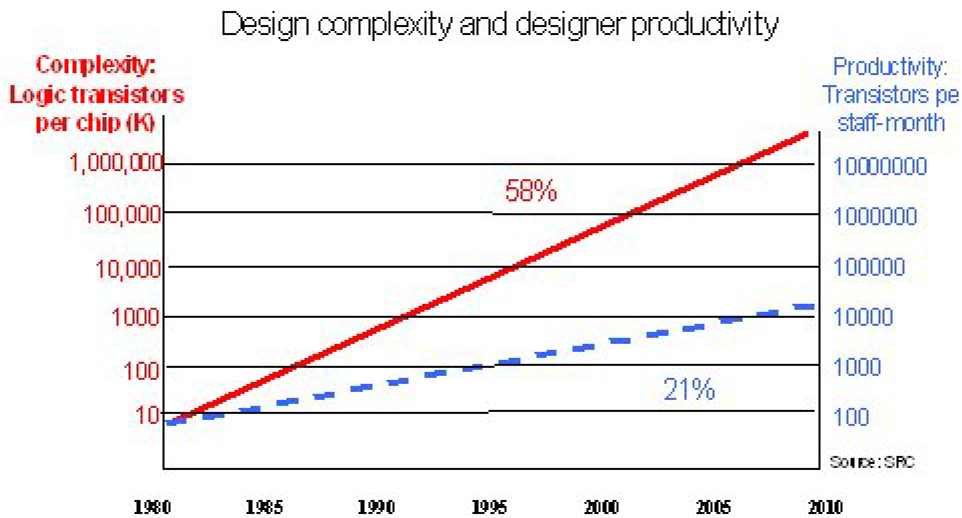


**Figure 1: A growing gap between achievable silicon complexity and designer productivity means that the industry needs a new, more efficient way to design SOCs**

The path to more effective SOC design is the multi-processor system-on-chip (MPSOC) design approach. The MPSOC approach gives designers flexibility to produce chips that are right the first time (lowering development cost) and remain right over time (increasing manufacturing volume and revenue).

Using this approach, SOC engineers can make a much fuller and more detailed exploration of the implementation possibilities early in the design cycle. They can better understand the design's hardware costs, application performance, interface, programming model, and other important characteristics.

**Programmability is Central**

SOC design methods that permit more programmability are a key enabler for delivering flexibility, which allows one chip to efficiently serve several different systems designs, while giving up few of the benefits of integration. The most popular ways to achieve system programmability is by using classical microprocessors or field-programmable gate arrays (FPGAs).

The problem with programmability is that there's a tremendous gap in silicon efficiency and performance between creating a hard-wired design and implementing the same design using programmable technology. The tradeoff for programmability can be as large as a factor of 100.

**Domain-Specific Flexibility**

System developers need not use completely universal silicon because of the economics. For example, a designer of digital cameras doesn't need to use the same chip that's used in a high-end optical network switch. The difference in benefit derived from a chip shared by 100 similar designs, versus one shared by 10,000 designs, is relatively modest, as shown in figure 2. Designers can easily afford a chip-level design platform that is appropriate to their application domain, yet flexible within it.

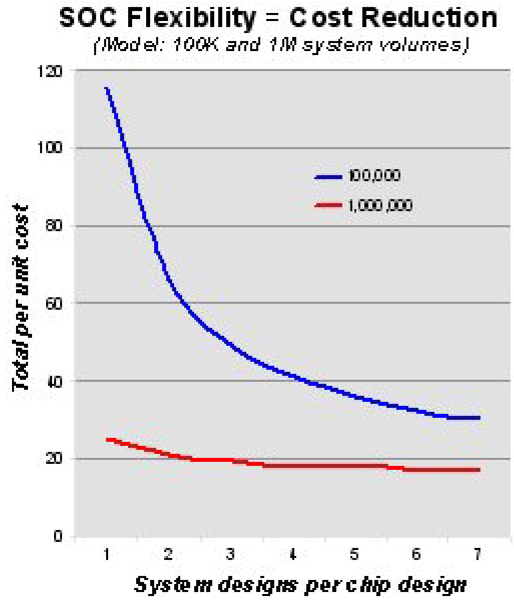


**Figure 2: System designs per chip design ($10M design cost, $15 manufacturing cost, 5% premium for programmability)**

Today, designers generally develop most SOC subsystems using hard-wired logic instead of off-the-shelf microprocessor cores because general-purpose microprocessor architectures are usually not fast enough to achieve design goals.

**The Processor as SOC Building block**

The basic building block of this MPSOC design methodology is the configurable, extensible microprocessor core, created by a processor generator that uses high-level specifications of application-domain requirements in the form of instruction-set descriptions or even examples of the application code that need to run to generate small, efficient, application-specific, programmable microprocessors.

Configurable processors can perform traditional microprocessor tasks quite efficiently. But because these configurable processors are able to incorporate the data paths, instructions, and register storage for an application's own natural data types, they also support virtually all of the functions that chip designers have historically implemented as hard-wired logic.

The introduction of configurable, extensible processors changes the SOC design equation. These processors can now deliver very high performance. The performance per gate, per square millimeter of silicon, per watt, or per clock delivered by these processors often rivals or exceeds the performance of hard-wired logic blocks that they replace.

The true leverage of the configurable, extensible processor arises from the way it allows designers to move tasks more easily between hardware and software. Because a much wider variety of embedded subsystems all fit within the capabilities of a configurable, extensible processor, the effort to move code running on a generic processor to an application-specific processor is very low because the functional specification remains in software, generally written in high-level languages such as C or C++.

Some of the design problems the MPSOC design methodology addresses include:

1. **An inadequate reuse model:** Semiconductor Intellectual Property (SIP) reuse has been a watchword in the industry for the last decade because building million-gate designs

from scratch isn't practical. Unfortunately, most RTL blocks are difficult to reuse. However configurable, extensible processors are much easier to reuse because critical functions are implemented in software.

2. **Inefficient use of memory blocks:** With the MPSOC approach, most of the memories in the system can be tested, initialized, managed, and controlled by the associated processors. This creates greater opportunities for flexible sharing and reuse of these on-chip memories.

3. **Difficult system modeling:** Because MPSOC systems are based on processors, processor-based instruction-set simulators (ISS) can be used to simulate these systems. Instruction-set simulators are much faster than RTL simulators so it is easier to instantiate individual chip models into a system model and you can run very large numbers of simulation cases and very long test sequences through the system simulation.

**Transition to MPSOC Design**

The MPSOC design methodology often takes hold in situations where existing design methodologies have hit the wall. This design methodology rapidly grows to take over other functions that are traditionally implemented in hardware and functions that have not yet hit the wall in terms of complexity because it's a faster, easier path to creating mega-gate SOCs than RTL design. The methodology can further grow to subsume even functions already implemented in software on traditional processors because application-specific processors also function like general-purpose processors. The MPSOC design methodology offers significant benefits to all members of the SOC design team in terms of simplifying the system design, shortening the development cycle, unifying system design from the hardware- and software-development perspectives, and increasing the reusability of both subsystems and SOC platforms built with this methodology. Together, these benefits represent a fundamental shift in the return-on-investment (ROI) equation for the chip or system builder. The volume and profitability of the SOC increases, while the cost of development decreases. This shift in ROI heralds both economic recovery and accelerated innovation in the electronics industry.