

Integration of Complex Mixed-Signal Solutions onto System on a Chip

Analysts predict that by the year 2005, over 70% of new system-on-a-chip designs will have mixed-signal content. To successfully meet this growing demand for integrated mixed-signal functionality on complex ASICs, semiconductor companies must address topics ranging from the process technology definition to tightly coupled tools and methodology for both ASICs and mixed-signal core development. This paper discusses the importance of mixed-signal methodology as a tool for SoC, the limitations and solutions, and provides a silicon example of mixed-signal CMOS system-on-a-chip.

INTRODUCTION

Growing demands for higher levels of integration onto increasingly complex ASICs have been driving the development of mixed-signal capabilities compatible with system-on-a-chip in the recent years. This combination of system-on-a-chip and mixed-signal functions presents a unique challenge to semiconductor manufacturers, affecting areas of chip manufacturing ranging from ASIC methodology, tools and verification flows, to process technology definition and modeling, system solution partitioning and analogue design capabilities and methodologies.

MIXED-SIGNAL AND SOC

Mixed-signal circuits, long thought of as a "black art", have been quietly making their way into mainstream CMOS ASICs.

The explosive growth in the communications market has driven demand for integrated mixed-signal solutions, as cost and power consumption push to reduce the number of discrete components and high-speed chip-to-chip connections required for these systems. Gigabit Ethernet PHYs, integrated baseband/IF/RF for Bluetooth and HomeRF, mega-pixel digital cameras, faster and larger capacity hard drives and optical drives - any of these applications alone provides a formidable challenge to the mixed-signal designer. Add in the task of integrating multiple mixed-signal interfaces together with digital logic, digital cores, and memories, and it's clear why mixed-signal SoC has not yet become an industry catch phrase. Given enough resources or time, mixed-signal SoC can consistently be achieved. But with time-to-market being the number one concern for semiconductor customers, it is necessary to look beyond the brute force approach and adopt an improved mixed-signal methodology.

MXS SOC METHODOLOGY

Integrating analogue functions into a large system-on-a-chip presents real challenges. A common scenario is that analogue functions are developed by one team and then transferred to the chip integration team which

is usually made up of digital designers in another location. This handover relies on a clear and complete specification, as well as implementation guidelines including information on placement, package bonding, power splits, isolation, shielding, external components, routing requirements and test requirements. Moreover System-on-a-chip tools are targeted at large digital designs, making integration and verification of the analogue blocks especially difficult.

One key to easing the integration of analogue and digital is to partition the development by encapsulating the analogue function within a digital shell. The idea is to include some digital logic within the analogue section that will be verified by both design teams. This is especially important for functions such as data converters, which have both digital and analogue content. The analogue team verifies the digital-to-analogue interface and provides a fixed layout of the entire block. All digital signals into or out of the function should pass through a digital library cell. The digital library cells are already characterized and models exist, thus negating the need to characterize and model the digital pins of the analogue section. The chip integration




Figure 1. Triple video DAC

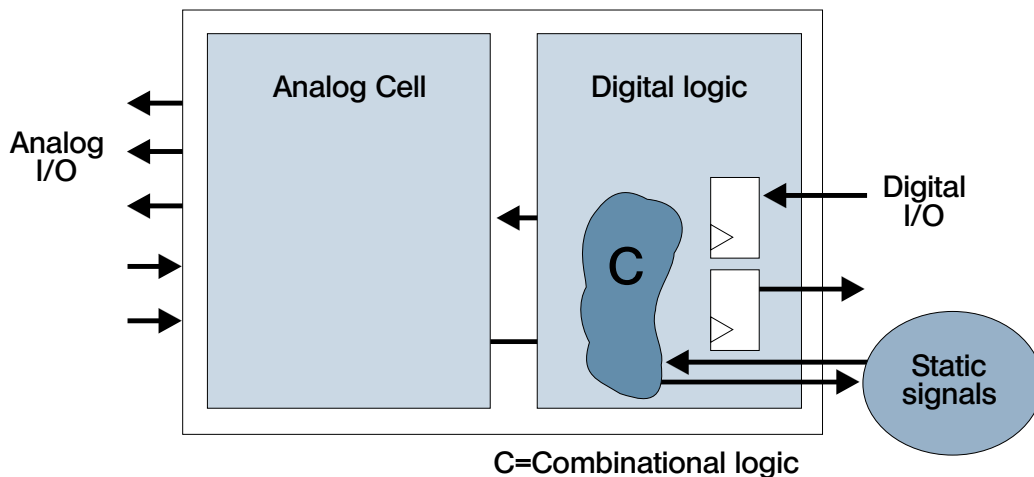


Figure 2.

team verifies the interfaces at chip level using standard digital simulators and timing analyzers.

Once the design partitioning is defined, design of the analogue block may proceed. Creativity is still allowed when converting specifications into circuits: without analog synthesis available, designers have to do it manually.

Standard flow generally adopted today is schematic entry, which is preferred to the analog extensions of VHDL and Verilog languages. The flow is: schematic entry > circuit simulation > layout > parasitic extraction > simulation. The key to accurate results is a quality circuit model. Specific internal rules must be followed on naming conventions, schematic entry, file generation, design engineer and mask designer flows. The task of developing accurate models cannot be underestimated. Parasitic models extracted from silicon must be integrated within the simulation tools for post-layout re-simulation to verify the analogue block. The reward is a high rate of right-first-time-by-design.

The standard circuit simulator works well for a small analogue function, but for more than a few thousand transistors you need to move to more complex mixed mode simulators. These simulators usually use a circuit level simulator and a digital gate level simulator that pass data to each other via a backplane. These tools can be useful, but are difficult to set up and require knowledge of both the analogue and digital simulators and their languages (spice, verilog, vhd). The EDA industry is moving to address this with support of analogue extensions to VHDL and Verilog language (Verilog AMS and vhd AMS). The main obstacle to widespread adoption of these is currently the effort required to develop characterised models. The ideal scenario is a single kernel simulator that supports circuit models as well as verilog or vhd AMS models and is easily configurable to allow use of different levels of abstraction for different design elements.

An important area the EDA industry is not yet addressing is the link between the electrical and the physical. This is particularly relevant when considering noise coupling between signals, and between signals and

the substrate. Today's tools do not adequately model and analyze these effects, which are playing an increasing role in limiting the performance of mixed-signal functions in noisy, switching digital ASICs.

A hierarchical method of design, based on simple functions or blocks, may greatly help with the re-use purpose. As an example a programmable voltage reference cell, accurately designed, can be part of a mixed signal library and be reused as part of more complex functions like analogue converters. A Sample-and-Hold amplifier is another block in this puzzle, as is an analog-to-digital converter cell, which contains the base elements of the conversion. If these functions are included in a single hard macro, the outcome is a complete ADC.

The degree of redesign is mainly related to the required performances; a relatively complex function like a fast 10 bits ADC, assuming the power supply is 3.3V, has an LSB value of about 3.2mV. In an ASIC there are two different types of noise involved: the




Figure 3.

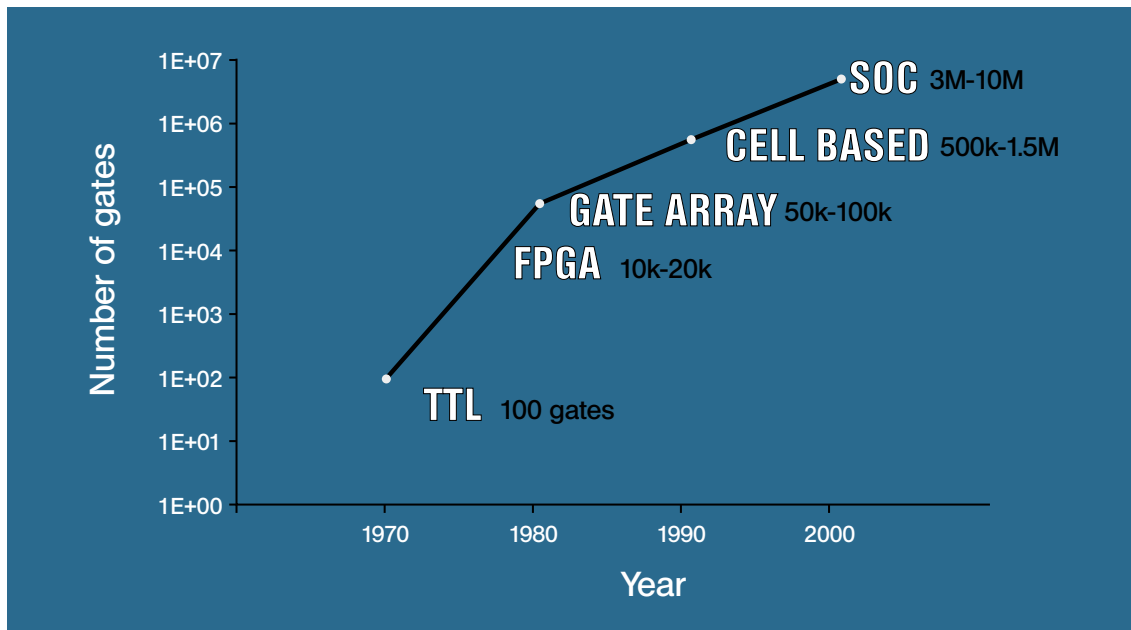


Figure 4.

inherent noise of the function and the interference noise. The first is generated by intrinsic sources like thermal or flicker noise and can be reduced but never eliminated; the second is generated by the switching digital part and fast processors, always included in SoC, and has several sources like crosstalk, substrate coupling, parasitic elements, power supplies and others.

There are some techniques to shield the analogue functions, but almost everything is done manually, following specific rules and the experience of the designer. In practice, the re-use of a mixed signal element could require a time from a week to several weeks of redesign, lay out and new simulations, depending on the application.

One IP represented by a triple video digital-to-analog converter is shown in figure 1. This converter has been successfully re-used in both ASSP and full custom ASIC designs with minimum effort.

EXAMPLE OF SOC SILICON

Despite the barriers already discussed, it is feasible to integrate complex mixed-signal systems into SoC. In February 1999, LSI Logic demonstrated one of industry's first single-chip CMOS CDMA baseband processors at CTIA in New Orleans. This chip contained a 14-bit voice codec, transmit DAC/filter, receive ADC, bandgap, 32kHz oscillator, DSP, microprocessor, RAM and ROM for a total of 7 million transistors on the chip. The mixed-signal portion alone comprised of approximately 1 million transistors. It was developed in LSI Logic's G11ä 0.25um drawn single-poly digital CMOS process.

Partitioning was used at the system-level to cleanly delineate between the analogue subsystems and the digital processing of the baseband chip. Within the analogue subsystems, hard macros were developed for key analogue blocks, each with a distinct digital

interface to the rest of the chip. This facilitated verification of each subsystem, while enabling re-use of individual analogue cells. Digital netlists provided accurate timing information for these interfaces and were used to simulate system performance at the chip-level. Simple analogue behavioural models were created to enable automated verification of connectivity between the hard macros and the rest of the chip ■

Marco Mariani is mixed signal manager of ASIC product marketing with LSI Logic Europe.

Marco, who has been with the company since 1997, is responsible for ASIC mixed signal product marketing across Europe. Based out of LSI Logic's offices in Milan, Marco is also involved in helping to define the requirements for future analogue developments within LSI Logic. In his present role Marco is supporting customers on refining requirements for analogue functions that are part of SOC (System On a Chip) solutions. Marco has a Master Degree in Electronics and his hobbies are archery, kayaking and photography.