PANEL

Fast, Cheap and Under Control: The Next Implementation Fabric

Chair: Abbas El-Gamal, Stanford University

Organizers: John Cohn & Andrew B. Kahng

Panelists: Ivo Bolsens, Xilinx, Andy Broom, AMI, Christopher Hamlin, LSI Logic, Philippe Magarshack, STMicroelectronics, Zvi Or-Bach, eASIC, Larry Pileggi, CMU

Overview

The semiconductor industry is caught on two horns of the economics dilemma: (1) the economics of technology - deep-subwavelength lithography and equipment cost, reticle enhancement technology and mask cost, and manufacturing variability and yield; and (2) the economics of design productivity - design turnaround time, availability of design skills, and portability of design effort. Standard-cells and the RTL methodology have taken us into the 90nm generation, but design is slow, expensive, and out of control. What we need is a next-generation fabric that will once again provide designers with "fast, cheap, and under control" implementation. The question: which fabric?

How deeply must the concept of regularity be engrained in the silicon implementation fabric to enable adequate yield and cost control in the face of CD variation and high mask NRE cost? Via-programmable fabrics such as eASIC (VPGA) provide an intermediate design point in the cost-density-performance space, but is this offering sufficiently attractive (let alone defensible)? Or, will traditional FPGAs continue to take up more of the market, starting from their foothold in low-volume and/or reconfigurable applications? On the other hand, regularity and programmability incur cost and performance losses as they abandon the leading edge of the scaling curve. Are such losses growing, and will we therefore always see viable ASIC and COT businesses? Finally, what are the views and needs of the platform SOC and pure-play foundry constituencies?

Panelist Statements

Ivo Bolsens, *Xilinx*

FPGAs ride the tide of Moore's Law. Process technology and architecture innovation are the two engines that have fueled a spectacular advancement in FPGAs over the past 10 years. During this period, price of FPGAs has been reduced by two orders of magnitude, logic capacity of FPGAs has been increased by two orders of magnitude, and performance has been increased by one order of magnitude. Whereas ASICs buck the tide of processing technology, FPGAs ride the tide. Deep submicron effects are breaking down the traditional modular design flow of traditional SOC architectures.

DAC 2003, June 2-6, 2003, Anaheim, California, USA.
ACM 1-58113-688-9/03/0006.

A growing portion of design time is spent dealing with deep submicron effects, at the expense of the creative process of design authoring. Surveys show that today, less than 20% of design time for complex SOCs is spent on design authoring.

Programmable FPGA platforms give designers the benefits of deep submicron but rather than focusing on getting the silicon to work, you can focus on getting the design to work.

Processor architectures have been trying to keep pace with Moore's Law by continuously increasing clock speeds to improve performance, and adding ever more complex memory hierarchies to cope with the communication bottleneck. Future technology nodes will make these techniques less viable. In contrast, FPGAs through their regular, parallel architecture and distributed memory organization can continue to benefit from dimensional scaling by adding more parallel hardware and distributed memory. Von Neumann architectures, originating from the days when silicon area was scarce, have been dictating sequential programming models. Spatial computing, exploiting massive resources of parallel hardware, will change the way we program future system platforms. We predict that FPGAs will become the heart of most systems over the next 5 years, replacing ASICs and processors as the fabric of choice.

Andy Broom, *AMI Semiconductor*

The industry has recently seen the emergence of a new breed of ASIC technologies designed to offer high performance and density, but with lower NRE costs and cycle times. Development of these technologies has been driven by spiraling NRE costs and continued increase in design cycle times with use of leading-edge processes. AMI Semiconductor is using this concept to address the growing need for low-cost FPGA replacements and cost-effective solutions for medium density ASICs. The XPA family of products offers users a design platform optimized for FPGA conversion while maintaining flexibility and ease of use for medium-density ASICs. Our unique hybrid process offers the advantages of leading edge technologies, but at a much lower cost of entry.

Christopher Hamlin, *LSI Logic*

The economics of very high complexity devices will be dominated by factors derived from opposite ends of the production chain: on one hand are fixed costs related to components such as masks, and on the other hand are skilled personnel and infrastructure required to specify and verify high-complexity device designs. The role of fabrics in ameliorating these otherwise debilitating economics is crucial: fabrics appear

to be essential to further progress in automation and abstraction, and fabrics permit otherwise unavailable economies of scale in production. Hence, the specific form assumed by fabrics is seen as a central problem for the industry. While it may be argued that the highest possible degree of regularity in fabric structures maximizes both their generality and their cost optimality, it is also true that traditional ASIC structures which tend to be irregular in key respects have often proven economically effective; similarly, genetic algorithms, for instance, can generate very irregular but efficient approximations to nearly optimal solutions. An alternative perspective offering possible paths to formal resolution of this tension between regular and random fabric structures is seen in the symmetry-breaking patterns and processes common in natural systems, which deeply embody aspects both of regularity and asymmetry in their functional organization. How this may apply to both the spatial and temporal aspect of fabric architecture is suggested.

Philippe Magarshack, *STMicroelectronics*

The next-generation fabric will in fact be a heterogeneous 'mosaic' of fabrics, consisting of:

- Reusable HW IP for well-defined standards (MPEG4, JPEG, etc.)
- Standard I/O (I2C, FireWire, USB, PCI, SATA, SPI4, etc.)
- Value-added RF and mixed-signal components
- Embedded fine- and coarse-grain reconfigurable HW
- Large-scale usage of multi-processors and associated memory
- All connected by a network-on-chip

Whether special-purpose fabrics such as embedded FPGAs or embedded Standard Cells will make it in mainstream design flows depends on the availability of specific design tools. These tools should allow the SoC designer to map part of a platform architecture seamlessly to these fabrics, and stitch the results together so as to allow the designer to run SoC-level analysis and prediction of timing, performance and yield.

The fabric that is already mainstream and is likely to become more so in the future is embedded memory, including classical 6T-SRAM, 1T-SRAM and embedded DRAM. As more configurability goes into software, and the software has to be as close as possible to the processor(s), the case is simple to make for more Mbits on-chip. The question really is what to make of the non-memory transistors.

Here, despite the extremely intricate issues of timing closure, electromigration, voltage-drop, on-chip variability and negative-bias temperature instability, we believe that close partnership between process developers, circuit designers and EDA developers will allow us to produce robust HW design flows in 90nm and below. The resulting tools and methodologies will allow those that master them to attain leadership in implementing key dedicated HW, be it for high-speed or low-power applications.

Zvi Or-Bach, *eASIC*

The future of logic IC cannot be an extension of the past. Cost of product development in standard-cell is getting too high and cost of silicon, performance and power penalties of FPGA are too severe for applications to which ASIC and ASSP are being applied. FPGA vendors would like to claim that they could take over. Yet they are developing new fabrics themselves, and they incorporate ever-increasing amounts of standard-cell logic into their existing fabrics.

The 'Next Implementation Fabric' must answer these challenges and enable cheap, fast, cost-effective nanometer implementation. Programmable interconnects are dominating FPGA area, power and performance, and will do so even more in future technology generations. By replacing programmable interconnections with via-configurable interconnections, an ultimate fabric can be provided. Combining such a fabric with direct-write e-beam lithography provides the Holy Grail of the silicon business: zero mask cost and no minimum quantity as in standard-cell, without the huge die area, performance, and power cost of FPGA. eASIC has taken this idea all the way through silicon implementation. The eASIC fabric is an all- layers regular fabric that is configured by a bit-stream for the logic definition, and has an e-beam-defined via layer for routing finalization, providing:

- No mask costs
- Very short turnaround time
- No minimum volume
- Ease of design
- Reconfigurable – ECO, debugging, ePLD, logic BIST
- Volume price ~ standard-cell ASIC (+10%)
- Performance ~ standard-cell ASIC (-10%)
- Power ~ standard-cell ASIC (+10%).

Lawrence T. Pileggi, *CMU*

As evidenced by the decreasing number of ASIC design starts, non-recurring engineering costs for standard-cell designs are becoming prohibitively expensive due to design time and mask costs that are incurred by the multiple design spins required to produce working silicon. There are no "yet-to-be-discovered" EDA solutions that will provide single-pass, affordable implementation for complex sub-100nm ICs based on standard cells in their present form. Instead, we must employ new forms of regular logic and circuit fabrics that can: (1) amortize portions of the design costs over larger volumes; (2) provide greater predictability to accommodate robust, single-pass implementation; and (3) facilitate customization by way of reduced mask costs for medium-volume applications. The application-dependent question becomes: "how much regularity can we afford?" An FPGA is at the regularity extreme, with no manufacturing customization, but with a high penalty in terms of performance, power and area. Exploring the middle-ground of new regular logic fabrics that lie in between standard cells and FPGAs can permit more optimal performance-cost tradeoff solutions. Such regular logic fabrics that rely on customization by way of a simplified mask set, however, must be integrated as part of an application-domain-specific IC (ADSIC) platform that includes memory, I/O, analog, etc. In contrast to existing FPGAs and standard products which try to be general and all-encompassing, ADSICs could enable improved optimization of performance-cost tradeoffs.