PANEL: Reconfigurable SoC--What Will It Look Like

Bryan Lewis, Gartner/Dataquest, US
Ivo Bolsens, Xilinx Corp., US
Rudy Lauwereins, IMEC, BE
Chris Wheddon, Quicksilver Technology, US
Bhusan Gupta, ST Microelectronics, US
Yankin Tanurhan, Actel Corp., US

Abstract

The argument against ASIC SoCs is that they have always taken too long and cost too much to design. As new process technologies come on line, the issue of inflexible, unyielding designs fixed in silicon becomes a serious concern. Without the flexibility of reconfigurable logic, will standard cell ASICs disappear and go the way of gate arrays? Will ASIC manufacturers lose their edge in providing intellectual value and become mere purveyors of square die area?

The argument in favor of FPGAs is that they have always provided great design flexibility because they were configurable. The argument against FPGAs is that compared to ASICs they have always been larger, slower and more expensive. Will FPGAs ever become efficient enough to replace ASICs in volume production applications? ASSPs can be designed with partial reconfigurability. Will they become the norm? Or, will new reconfigurable logic cores change the SoC game completely?

The answers to these questions will clearly impact system designers throughout the world and shape the future of the electronics industry. A panel of key industry executives each coming from a different area of the market with unique views will debate these highly controversial topics.

1. Current global market data (B. Lewis)

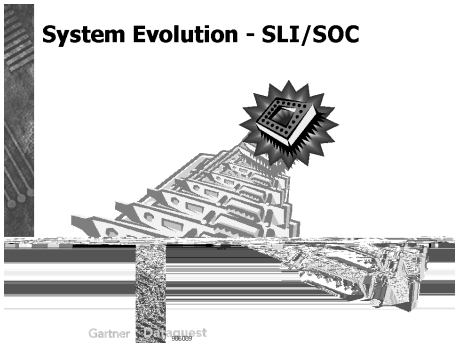

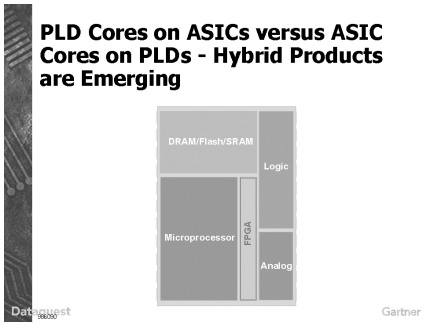

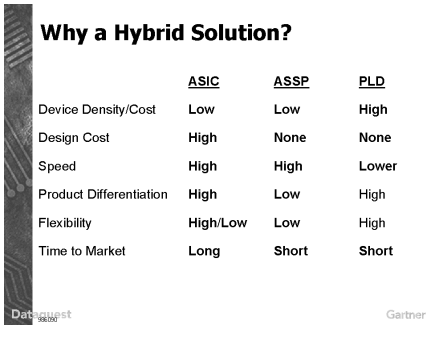

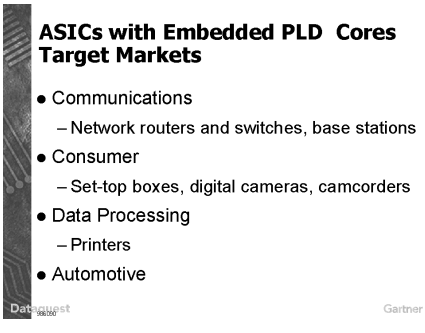

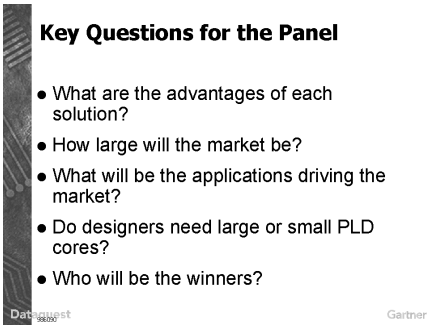

2. I. Bolsens, Xilinx Corp.

Technical limitations will force us to re-think the design of SoC platforms in the future. Providing the design community with a field-programmable hardware platform that contains millions of gates, running at clock frequencies of several hundred MHz, communicating with integrated processors and distributed memories at GByte/s data rates is definitely breaking down barriers to build tomorrow's products.

These great platforms are creating headaches for the EDA teams that are trying to capture this ever-growing design complexity, flexibility and heterogeneity in new modeling languages, co-design flows, verification techniques and IP-re-use strategies. Research teams from different fields are joining forces to master the problem of such large complexity.

But probably even more important, the future challenge will be in coping with the physical details. As we are reaching the capabilities to capture complete systems in a package, these systems will communicate with the rest of the world at multi-Gbit/sec data rates, creating challenges for signal integrity, packaging technology and testing. Parasitic effects such as substrate noise coupling, chip-to-package effects, high-speed transmission line effects will limit the performance of the products that make use of these powerful platforms. Moreover, these future platforms will draw tens of amps causing board and chip distribution problems as well as reliability issues.

SoC platforms will have to be accompanied by high-level design tools that allow designers to capture the growing complexity and flexibility (in time and space) of these components. But next to this, SoC platforms will have to integrate solutions that free the system engineer from dealing with the effects that result from smaller line-widths, higher speeds and higher power.

3. R. Lauwereins, IMEC vzw

The cost of deep sub-micron semiconductor technology drives us away from ASICs and toward moderately reconfigurable SoC platforms.

The NRE cost of design as well as mask fabrication for deep sub-micron integrated circuits with hundreds of millions of gates will become exceedingly large. Hence, the design and production of a custom IC for a single application will only be cost efficient for very large volumes, e.g. hundreds of millions of units per year. It can already be observed that the number of different design starts is decreasing by roughly 20% per year, although production volume is still increasing.

This leads us to flexible, reconfigurable platforms that can be tailored after manufacturing to all applications in an application domain. The NRE cost of ASICs of several hundred million gates can indeed hardly be amortized over the volume of a single application. Therefore, tomorrow's platforms will not be high complexity ASICs. For mobile embedded signal processing applications, the platforms will not be scaled-up versions of today's fine-grain reconfigurable field programmable gate arrays (FPGAs), because the unnecessary amount of gate-level reconfigurability will lead to excessive power consumption. Also, the lack of application specialization in a sea-of-microprocessors approach would also lead to too much power consumption. I foresee that tomorrow's platforms will consist of an application domain-specific mixture of microcontrollers, digital signal processors, embedded memory blocks, hardwired accelerators and embedded coarse-grained, reconfigurable logic. The amount of software and hardware reconfigurability should be kept to a minimum; just enough to be able to map any application of the targeted application domain efficiently onto the platform.

Many challenges will have to be solved. It should be possible to select an application domain and specify its properties. A platform architecture should be deducted from these specifications such that the right balance between flexibility and power consumption is obtained. Once such a platform is built, an application should be mapped onto the platform, thereby optimally employing the opportunities offered by the platform's flexibility. When the application shows a substantial amount of dynamic behavior, the platform's flexibility could even be used at run-time.

4. C. Wheddon, Quicksilver Technologies

The adaptive computing machine could be the in-system reconfigurability solution for next-generation wireless applications.

Third-generation (3G) and fourth-generation (4G) wireless terminals will be required to provide exceptionally higher levels of service than their second-generation (2G) counterparts. This would not create a cause for concern if it was not for the ever-increasing, simultaneous demand for increased mobility, which

introduces several physical requirements (e.g., longer battery life, reduced size, lighter weight, etc). Wireless terminals are fast approaching a point in the product development roadmap where, without a paradigm shift in the basic design architecture, they will no be able to simultaneously meet demands for both service and mobility. The ideal solution would be to take advantage of the ASIC, while retaining the flexibility of the DSP. This is very essence of the adaptive computing machine (ACM), which is being developed as a platform comprising: a hardware device, micro OS kernal, development tools and application software. The ACM platform will allow software downloads at exceptionally fast speeds to give the user "applications on demand." Benchmark performance comparisons have shown that the ACM is an order of magnitude improvement over conventional technologies.

5. B. Gupta, ST Microlelectronics

The inevitable addition of flexibility and reconfigurability to complex system-level ICs must also include corresponding levels of software and system tool support.

The ASIC and SoC industry is undergoing a series of challenges that are intensifying the debate over future design options. The first debate is regarding the design and implementation of increasingly hard-to-build ICs. Issues of EDA productivity gaps, design closure and verification play an important role here. The second axis of debate is how best to amortize the NRE costs and design cycle time of new products. In other words, the question is how to add enough flexibility to enable faster design and more reuse without compromising too much in area or power.

The ASICs and ASSPs that are going to be built in the future are going to be more complex and harder to design regardless of the underlying implementation technology. The design solutions should be ones that address both sets of challenges in a unified fashion. Simply using an embedded FPGA is not enough to address the design productivity gap, rather it tends to delay facing the problem. The need for flexibility and reconfigurability will have to be pervasive and at all levels (HDL, software, system specification, etc). The amount of flexibility will be determined by the application domain and, idealistically, by battery size. So, the thesis that we will build flexible ASICs and ASSPs is inevitable from an economic viewpoint, but the nature of the flexibility and configurability will extend far from just using FPGAs and must be accompanied by software and system tool support.

6. Y. Tanurhan, Actel Corp.

The embedded FPGA core adds reconfigurability to SoC applications to help get designs to market faster and keep them in the market longer

New larger and feature heavy FPGAs tend to be expensive and make a big footprint in system-level board designs, but they put time-to-market competitive pressure on ASICs. An embedded reconfigurable FPGA core presents an alternative that offers system-level architects on-chip design flexibility and risk avoidance that can the speed time to market of an SoC and help maintain the cost efficiencies of keeping high-volume system applications in entirely ASIC silicon.

Large, complex SoC designs in ever diminishing process geometries will be accompanied by very expensive NREs with mask sets approaching US$1 million. Yet to remain competitive, SoC designs will still be required to get to market quickly. More than ever, design re-spins must be avoided—keeping costs from getting out of hand and maintaining a time-to-market edge over alternative solutions. Having the ability to add a reconfigurable logic core makes sense and will become widely adopted. This is because such a maneuver can provide a comparatively small on-chip reprogrammable die area, a better performance/die area ratio than dedicated off-chip FPGAs and a better performance/power ratio when compared to standard software IP alternatives. A reasonably small, but optimized reconfigurable logic core embedded in an SoC provides several significant advantages for system-level designers, not the least of which are design flexibility, savings in time-to-marketing and cost and extended product lifecycles.