

Panel: Looking Backwards and Forwards

Organizer: Marco Casale-Rossi, Synopsys, USA

Panel Moderator: Giovanni De Micheli, EPFL, Switzerland

Panelists: Antun Domic, Synopsys, USA; Enrico Macii, Politecnico di Torino, Italy

Domenico Rossi, STMicroelectronics, Italy; Joe Sawicki, Mentor Graphics, USA

Abstract

Ten years ago, at 90 nanometers, EDA was challenged and deemed inadequate in dealing with increasing complexity, power consumption, and sub-wavelength lithography, thus harming the progress of mobile phones. Today, at 10 nanometers, integration capacity has increased by two orders of magnitude, power consumption has been successfully “tamed”, and 193 nanometer immersion lithography is still relied upon. Also thanks to EDA, tools, methodologies, and flows that were originally devised for design enablement for the emerging technology nodes, have been successfully re-deployed at the established technology nodes, where they represent a critical design differentiation factor. However, the battleground is changing again: after the billions of phones, trillions of “things” lie ahead. Moving forward, emerging and established technology nodes, digital and analog, hardware and software will be equally critical. What is EDA doing and, more important, what should EDA do – and is not doing – in order for the next decade to be as great as the past one? This panel session, moderated by EPFL Professor Giovanni De Micheli, gathers academia, semiconductor, and EDA industry to discuss the challenges and requirements of the new era.

Panel Moderator’s Introduction

Prof. Giovanni De Micheli, EPFL, Switzerland

The fast growth of semiconductor products and sales in the last three decades has been enabled by the coupling of semicustom design methods with Electronic Design Automation (EDA) tools and methods. The continuous downscaling of CMOS Field Effect Transistor (FET) dimensions enabled the semiconductor industry to fit an increasingly larger number of devices per chip and also to achieve higher performance per watt spent in computation. To keep pace with technology, EDA tools are challenged to handle both digital designs with growing functionality and device models of increasing complexity. More than ever, physical design tools are challenged by 3-dimensional devices (e.g., FinFETs) and by emerging devices such as gate all around transistors in Silicon

NanoWires (SiNWs), Carbon NanoTube (CNTs) and in other materials.

Moreover, new emerging nano-technologies are providing us with devices that are no longer simple switches, but switches controlled by the combination of electrical signals. This has been shown to be the case with SiNW and CNT controlled-polarity transistors. The arrival of such technologies has brought the need of new logic abstractions and in turn the requirement of new logic synthesis models and algorithms. Overall, the objective of increasing computational density and the performance/power ratio is now achieved by combining functionality-enhanced devices with EDA synthesis methods and tools that capture their essential properties. It is obvious that achieving competitive design in the 10nm range and beyond can no longer be thought in terms NANDs, NORs and AOIs. This requires a deep rethinking of computational models as well as EDA flows.

When looking forward, the manufacturing and design technology challenges have to be measured against the backdrop of new electronic systems and services that will permeate our society, such as the Internet of Things (IoT) in its various incarnations, including systems for health management operating on and inside human bodies, distributed systems for environmental and terrestrial defense, personalized social media with revolutionary I/O interfaces. All such systems require extremely low-power consumption, high dependability and security as well as seamless integration with sensors and transducers. This motivates even more a new wave of design models, algorithms and tools for an ever-evolving landscape of our digital society.

Panelists’ Position Statements

Antun Domic, Synopsys, USA

Looking backwards, I’m proud of what the EDA industry has accomplished in the last decade. Back then, EDA was deemed inadequate under several metrics, and often perceived as the obstacle towards the progression of Moore’s law. Today, EDA is considered a critical partner, an enabler at the emerging technology nodes, and a differentiator regardless of the technology node. Not only we have dramatically improved our results, but we have

also made these better results available to a wider design community. We have demonstrated how advanced design and EDA can make everything better, not only at 22/20 nanometers and below, but also at 32/28 nanometers and above.

Area is a very interesting example: in the last ten years, we have improved advanced RTL synthesis results by 30% in terms of area – incidentally, we have also improved performance, and power by approximately the same amount. In physical implementation, advanced design planning, and the ability to deal with an order of magnitude bigger blocks versus ten years ago, have made advanced EDA a critical enabler when the form factor or the die size are a key metric. At emerging technology nodes, our semiconductor partners tell us that the flat implementation of a hierarchical design can save silicon real estate, and power consumption – due to the lesser amount of buffering. At established technology nodes, the ability to deal with the convoluted floorplan and rectilinear shapes of the digital portion of leading-edge A&M/S designs is definitely a good reason to go for advanced EDA solutions.

Routing is another good example: starting at 20 nanometers, it has become impossible to draw the copper interconnects of an IC without double-, triple-, or even quadruple-patterning. Without EUV, 5 nanometers could require octuple-patterning; while “multi-patterning” has allowed going beyond the minimum single-patterning pitch of approximately 80 nanometers, advanced EDA has made “multi-patterning” automated, hiding and waiving its complexity, thus allowing a transparent use of this sophisticated technique. The availability of advanced EDA tools suites has proven critical: if one uses an advanced EDA solution, one can “do more with less”. Moreover, more efficient “line-search” routing algorithms have resulted in much better routers under “simpler” design rules, making it possible to reduce layers at 28 nanometers and above. Our semiconductor partners tell us that moving from a 6-layer 130 nanometers A&M/S process variant to a 4-layer slashes 15-20% from the cost.

Finally, let’s look at power consumption. Voltage scaling use increased at 130 nanometers, when the dynamic power reduction started to be offset by the static power increase. At 90/65 nanometers, it became virtually impossible to design an IC without using sophisticated power reduction techniques. “Design for power” was an enabler that prevented massive amounts of “dark silicon”. Today, advanced EDA has made much of “design for power” techniques automatic and part of “standard” design. It has allowed a seamless use of a wide catalogue of techniques, making sure that the power intent is always correctly implemented and consistently verified throughout the design flow. In this respect, the availability of advanced EDA tools suites is a differentiator: again, if

one uses an advanced EDA solution, one can “do more with less”. Power reduction techniques that back then were meant for survival at the emerging technology nodes, are now fully exploited at the established technology nodes. Literally, scores of voltage/supply/shutdown domains even at 180 nanometers are common, providing incredibly power savvy solutions using the most cost effective process technology.

Looking forwards, we are certainly not done: complexity is the challenge, which manifests itself in many facets. At the dawn of the “Internet of Everything”, the efficient use of silicon resources remains imperative, regardless of whether we are designing 100 billion transistors at 10, 7, 5 nanometers, or 1 million at 180. Area and routing resources are an optimization target from RTL synthesis down to place and route. Today designs are both analog and digital, high performance and low power, etc. This poses, simultaneously, challenges that until very recently could be addressed only separately. The breadth of active technology nodes widens; the pace at which emerging technology nodes are adopted is getting asymmetric, as more than 90% of design starts are happening at 32/28 nanometers and above, and 180 nanometers is by far the most “designed” technology node, with more than 25% of the total design starts every year. This won’t change significantly over the next decade. So we will require that EDA tools, methodologies, and flows are tuned for a wider spectrum of manufacturing processes and application requirements. Software design tools and solutions will complement hardware ones, and these will be more and more critical for the realization of future systems as “electronification” and “siliconization” gain ground in new markets and applications.

Enrico Macii, Politecnico di Torino, Italy

Design automation has been one of the key enablers of the electronic revolution. Since the first wave of algorithms and tools for logic optimization (e.g., Espresso, Mini, MIS, SIS, etc.), innovation in EDA has gone hand-in-hand with technology progress. In the past decade, the focus of the tools has shifted from traditional optimization metrics, such as area, delay, power and testability, to more complex targets, such as manufacturability, temperature, ageing and process variation. And the entire design hierarchy has been supported by the most advanced tools, from system level all the way down to physical design. So far, though, one assumption has never been questioned: The circuits and systems to be designed were supposed to be manufactured using pure CMOS technologies. The biggest challenge for the EDA community has thus been the need of coping with feature size scaling and the effects of CMOS nanometer processes on the design flow.

The advent of heterogeneous systems, e.g., the “smart systems” which are at the basis of the IoT paradigm,

opens up a new era for EDA. In fact, smart systems represent a broad class of systems that can be defined as intelligent, miniaturized devices incorporating functionalities like sensing, actuation, and control; they are usually energy-autonomous and ubiquitously connected. In order to support these functions, they must include sophisticated and heterogeneous components and subsystems such as: Application-specific sensors and actuators, multiple power sources and storage devices, intelligence in the form of power management, baseband computation, digital signal processing, power actuators, and subsystems for various types of wireless connectivity.

Smart components and subsystems are designed and produced with very different technologies and materials. Then, the challenge in the realization of such smart systems goes beyond the design of the individual components and subsystems (an already difficult task by itself), but rather in accommodating a multitude of functionalities, technologies, and materials; as such, they involve solving problems of different nature. The widely acknowledged keyword in smart systems design is therefore integration.

There are two dimensions of integration that represent the main obstacle towards mainstream design of smart systems: Technological and methodological. As already experienced in other domains (e.g., digital and analog design), a solution has been found first for the technological issues. Advanced packaging technologies, such as system-in-package (SiP) and chip stacking (3D IC) with through-silicon vias, allow today manufacturers to package all these functionalities more densely, combining various technological domains in a single package. SiP technology works nicely because it allows merging of components and subsystems with different processes, and mixed technologies using state-of-the-art advanced IC packaging capabilities with minor impact on the IC design flow.

Design methodologies, however, are falling behind: Current smart system design approaches use separate design tools and ad-hoc methods for transferring the non-digital domain to that of IC design and verification tools, which are more consolidated and fully automated. This solution is clearly sub-optimal and cannot respond to challenges such as time-to-market and request of advanced sensing functionalities. A big step towards effective, large-scale design of smart systems would then be that of changing the design of such systems from an expert methodology to a mainstream (automated, integrated, reliable, and repeatable) design methodology, so that design cost is reduced, time-to-market is shortened, design of the various domains is no longer confined to teams of specialists inside IDMs and system miniaturization can be achieved with limited risks.

This objective can be met by defining and implementing a structured design approach that explicitly accounts for integration as a specific constraint, thus minimizing manual hand-off. The ability of exchanging a wide range of complex design parameters between components and subsystems from different technologies, packages, and architectural templates in a holistic co-design framework is an extraordinary challenge, which requires closing several technical and cultural gaps by means a multidisciplinary approach.

The task of constructing a flexible modeling, simulation, design and integration EDA platform for miniaturized smart systems represents, in my opinion, one of the major challenges that the EDA industry, and the EDA scientific community at large, will have to face in the coming decade.

Domenico Rossi, STMicroelectronics, Italy

Accessing the most advanced CMOS platforms still represents one of the most critical aspects to innovate and stay competitive in digital products. New nodes are introduced every 18 months (while the 14nm is taking off, the 10nm will be soon brought up while the 7nm is already in the R&D POR of most of the major players) while the R&D costs and the complexity of the produced to be developed are both dramatically.

One way not to be trapped in the so called “innovation death spiral” and exploit across what these new technologies are capable of across all the different markets, relies on the timely availability of “robust since the early adoption” of the EDA ecosystems.

While, in case of very high volume application (Wireless and high end CPUs), the brute force can be applied and paid back, “design efficiency” is indeed the only possible, technological and financial solution applicable in most of other cases.

We are referring to “EDA ecosystem”, as multivendor tools chains, rather than a suite from a single vendor, can represent the best and fastest solution. From this standpoint, as experienced even in the recent past, sometimes this went very well, sometime less.

One thing which, by sure, went well has been the support given to comply with the adoption of layout rules/constraints exploding and the masks composition/manufacturing. From this standpoint, the improvement in the hardware mask print resolution has been efficiently complemented by the EDA capabilities and RET, OPC and multi-patterning techniques have made possible the bring up of 14nm and 10nm without introducing alternative, but not yet mature lithographic technologies (EUV).

From the other hand, other aspects perhaps trivial, but not from the user perspective, should have deserved a

better attention and a more structured approach. All of us remember the VHDL Verilog dualism and the relevant costs involved. The same happened with UPF and CPF for the description of the power intent, with the associated ambiguity in the case of a multi-vendor flow. We cannot also forget the approach used by CCS-ECSM for library description: as a technology provider, we had to duplicate the effort for our IP deliveries, internally and to the customers.

This said, a quantum leap in bringing the design ecosystem up (including the bring up of very demanding IPs) and quickly getting through the physical implementation of very complex digital ICs is fundamental to fully exploit the potential of the new most advanced technologies.

For sake of simplicity, ST experience in designing ASICs for Networking (very simple indeed from the architectural standpoints, predictable in terms of roadmap evolution, but very demanding in terms of technology adoptions) is reported

Analog, but not only IP

Even if not evident at all, the time spent in designing, developing and integrating analog IPs into an ASIC design flow analog IPs define the time a new technology is used for ASICs for Networking. These are the cases of High Speed Links SERDES, High Speed ADC and DAC and, to different extend, TCAM memories. From this standpoint boost the design productivity is fundamental.

Physical Implementation

A lot has been already done (but only recently!!). Taking (almost full) the opportunity given by the multiple cores sitting in the farms, engineers can today run a place-and-route job for a 5-6M instance sub-chip with a throughput approaching the 1M instance per day, but there is still a lot to do.

Tools must be anyway “more” stable but also better supporting ancillary features making the task faster and by far more efficient. Just as an example, there is no real self-monitoring of the implementation tools able to generate information useful to the next runs. Today, this is up to the “user”, with this “user” figuring up how the algorithms work and how they can be then tuned. From this standpoint, a kind of built-in self-learning engine having access and greatly exploiting an exhaustive set of information could better drive for more consistent results.

This same concept of self-learning could also be applied to the floor planning activities. The tools are today supposed to support automatic power grid synthesis and floor plan but retrofits to get around problems of congestion, timing and current/power densities are, as a matter of fact, manual, and relying only on designer sensibility, feeling and experience. To make the story

short, we are missing, even in this case, the global approach that makes this retrofit fully automatic.

Changing the subject, we need tools and methodologies to address the power related aspects. In ASICs for networking we are used to face products with switching activities in excess of 5X if compared to most of standard processors: the management of the power density and the removal of hot spots cannot rely on any automatic tool. The identification of the most critical situations and the on-the-fly introduction of decoupling cells as well as the management of power crowding should be one of the key parameters the tool itself should take care.

Done for networking, this could help the ADAS (Advanced Driver Automotive Systems) asking for the adoption of advanced CMOS technology to a pace the Automotive market never witnessed, but compliant with zero PPM quality standards even when the ICs is asked to work in tough temperature conditions.

From this standpoint, we have also to reconsider what the role of the DFT will be. Usually and universally DFT is considered, even in my organization, a front end activity, but is this still true? Why is it needed to perform, later during the implementation, the scan chain reordering to alleviate the congestion, or addressing the same issue for the compression logic? Even in this case, a radical change in the approach is required.

All in all, this is what we need to make the most advanced technologies accessible to “dummies” and not only to a few high volume applications.

Joe Sawicki, Mentor Graphics, USA

The history of the semiconductor industry is defined by the pursuit of growth. The industry has become an engine of the modern economy by successfully targeting, exploiting and driving the needs of successive new application segments. Over the past decade mobile communications has been the main growth driver while the PC market also provided a last measure of growth as it moved to maturity. In 2015 we are at an inflection point. PC growth has flat lined, with replacement timelines extending and mobile platforms providing an alternative for many users. Mobile is moving to a more mature stage with both margins and growth under severe pressure. The Internet of Things is seen by many as the next opportunity for the semiconductor industry. It is an incredibly broad term, covering the Fitbit in my pocket, an internet gateway in my car, and an industrial manufacturing solution. All have in common a few elements: a radio to communicate, a processor to manage data, and, often, a sensor to collect data from the outside world. What makes this wave different from previous waves is that it does not require the next technology node to implement. I believe this will send us on two parallel development

paths moving forward. One to enable continued scaling (a broadly deployed IOT would require a massive, networking and server infrastructure), and one to enable IOT.

One of the first things to note about IOT devices is that they are systems, and the requirement for sensors makes it unlikely that they will be implemented as SOC's. Package techniques like the INFO solution from TSMC will be important, driving an approach to package integration that will raise the importance of integration between tools for PCB-oriented integrated-packaging and IC's. Systems aren't just hardware either. Numerous parts of the software stack can be made far more efficient by the availability of IP and tools focused on IOT applications. Here the elements like IOT gateway reference designs and real-time operating systems point the way to a new segment for EDA to serve.

On the hardware level, IOT designs will require low-power, low-cost implementations. Here technologies originally implemented to enable advanced node designs are easily reused and retargeted. Low-power design techniques move directly across. Other technologies will be retargeted. For example, high-compression DFT technologies will be targeted at low-pin-count test, helping to enable lower cost packaging. We are also already seeing established node variants that leverage the learning curve provided by the original advanced designs to hit a new point on the power/cost/performance curve, enabling even more attractive designs.

None of this means that scaling will end. On the technical side, we are already deep into the development phase for 7nm and can see a reasonable path to 5nm. The need is there. The amount of data that would be collected, transmitted, and processed will require an underlying infrastructure backbone that will drive increased transistor densities for years to come. Here EDA has been a critical contributor to progress as computational lithography has been one of the primary enablers of feature scaling in the absence of EUV. This will continue even after the eventual introduction of EUV as feature sizes at that node will be small enough to continue to require computational lithography to enable viable yield.

For years we've had a history of synergistic growth between the semiconductor industry and EDA, as EDA delivers on critical tools and IP to enable to effective exploitation of new markets. This trend can be expected to continue moving forward. Scaling will continue to drive innovations to implement these very challenging processes. IOT will leverage retargeted technologies, and add new tools and IP to accelerate this new market.