Future of Design Automation

To determine the future direction of Design Automation, we need to take a careful look at the trends in design (systems and circuits) and technology. On the enterprise system side, customers demand virtualization, availability and serviceability. On the consumer side, customers' appetite for portability, 24/7 platform agnostic accessibility to information and data is insatiable. Ease of use for the customer invariably translates into almost intractable complexity for chip design, measured in terms of billions of transistors and multitudes of functions. At the same time, relentless dimensional and frequency scaling has finally met its match: increase in power due to deep pipelining limits frequency scaling and the increasing dominance of manufacturing variability takes a bite out of performance as well.

For chip designs to be manageable and affordable, the overall complexity of the chip must be reduced. At the same time, to maintain chip performance increase on the Moore's curve the overall chip performance can not be sacrificed. One way to resolve this seeming conflict is to raise the level of abstraction and adopting a hierarchical approach. In a nutshell, the overall chip complexity is reduced by lowering the complexity of the cores or components (performance is reduced) but increase the complexity at the chip level to maintain the chip performance. This is the classic complexity reduction from O(MN) to O(M) + O(N) by adopting a hierarchical approach. To elaborate, the chip will be made up of a large number of heterogeneous components (hundreds in 5 years, perhaps thousands in 15 years). Each component has relatively low architectural complexity and will run at a relatively low frequency to achieve power efficiency and reasonable power density for the chip. As a result, this approach also solves the power problem. The performance is to be recovered by the increased architectural complexity at the chip level and this presents the biggest challenge as well as the biggest opportunity to Design Automation. When there are thousands (and doubling every few years to stay on Moore's Curve) of components in a chip, the components are the new "logic gates". Just as we need a design tool to synthesize a microprocessor from logic gates, we will need a new design tool to synthesize these new chips from the components. The problem can be described as follows:
Given a set of components (e.g. processor cores of different types, optimized engines that accelerate specific tasks, caches at different levels, interconnect fabric of different types (could be asynchronous) and performance, autonomic management units, custom ASICs, FPGAs etc.), a set of software applications, and a set of acceptance criteria (e.g. performance target, power budget, cost, time to market etc.), find the "optimal" chip level architecture that will meet all the criteria.
I assert that it will be very difficult for even the most seasoned architect to come up with a good solution due to the sheer complexity of the problem at the chip level. Moreover, a system/chip is as only as good as its components. Therefore there is the added dimension of finding the "optimal" granularity and design point for each component. A complete solution must involve considering the design of the compiler, the operating system, the chip, and the components in a holistic fashion. I am not referring to existing work on developing chip-specific compilers to increase performance and expedite software development. I am proposing to "characterize and parameterize" the compiler, the hypervisor, and the components to enable exploration and optimization. In conclusion,

raising system level design automation to include the software, compiler, OS, hypervisor stack is one big challenge for DA in the future.

On the technology side, the situation is somewhat uncertain. There are a lot of speculations about the end of CMOS and about the potential candidates for post CMOS technology. Unfortunately, while there are some very promising candidates for memory device, the picture for logic device is much murkier. It usually takes a technology a good ten years to evolve from being a proven prototype to being high-volume production-ready. The current state of affair of these nano-devices led me to believe that CMOS will still be the predominant technology in the next decade. The question is: what is the roadmap for CMOS in the next decade? Here, there are two scenarios as well. One camp believes that CMOS will stop at around 25 nm the other believes that CMOS can scale to the single digits. I will focus on the former scenario for no other reason than that it is more disruptive and therefore more interesting. The role of design automation will be to extract the maximum value out of a technology node since there will not be a next generation to give you the scaling benefits. The value proposition and investment will shift to new materials, structures and devices, novel circuit families, and radical packaging. New materials, new structures and new device topologies require innovative modeling, probably in the quantum mechanical domain. Novel circuit families need new logic and physical synthesis techniques as well as new design methodologies. New packaging like 3D IC naturally requires substantial retooling of the whole suite of physical design tools, among other things. However, a more disruptive change is the following. The industry has always relied upon design migration to gain performance and cost benefit by mapping designs from a technology node to the next. If technology scaling stops, design automation will be a key ingredient in making up for the shortfall.

With increasing geopolitical and geological risks, the future is the age of information and data gathering and processing. In order to incorporate all these functionalities under one roof, chips of the future can be very heterogeneous. They will consist of MEMS, optoelectronic, bio-electronic, RF, analog and digital components. The functioning of some of these components is strongly dependent on the package; therefore chip-package co-design is necessary to achieve an “optimal” design point. The possibility of using a 3D package opens up even more challenges to design automation. A lot of innovations are required to address issues such as modeling (with the package), verification, chip level simulation, testing and so on. The fact that the energy domain stretches from mechanical, thermal, optical to electrical simply compounds the difficulty of the problem.

Design Automation in the future will expand into areas which are beyond its traditional scope, simply because the functions of chips and systems are rapidly evolving and expanding. Different automation disciplines will integrate and merge to create new and exciting fields just like logic synthesis has merged with placement and routing to become the important area of design closure. There is no doubt that automation will be critical to the survival, continuation and prosperity of leading edge technology.