Thoughts on the "Future Directions of Design Automation Research"

Andreas Kuehlmann, Cadence Berkeley Labs
October 24, 2006

Some of the main challenges we will face in the next 10-15 years in the design of electronic systems can be characterized as follows:

- The traditional CMOS process will continue to be used extensively. However, our current design flow based on horizontal sign-off models will further deteriorate and break down on many levels from device level, logic level, to behavioral levels. This will have dramatic ramifications on the utilization of libraries, IP components, and sign-off interfaces and have a significant impact on the algorithmic flow for optimizing and verifying systems and system components.
- Pushing the integration density for CMOS processes to its limits will likely force us to abandon current methods for process and device characterization and require modeling approaches based on the actual device physics possibly going down to atomistic levels. This would soften the traditional hard boundaries between design CAD and TCAD and foster a new interface contract between design and manufacturing.
- System solutions will increasingly integrate electronic and non-electronic components on one substrate or into one package and also utilize non-CMOS materials. This will require comprehensive DA solutions that model, analyze, and optimize closed-loop systems in multiple-domains such as mechanical, biological, micro-fluids, etc., involving sensors, data processing elements, and actors.
- The challenges driven by increasing system complexity will push true support of higher modeling levels for exploring implementation options, validating functional correctness as well as analyzing and optimizing manufacturing robustness and reliability. The modeling must be done across implementation domains from hardware to software, digital to analog and include the environment in which the system is to be used.

Specifically, the challenges from increased system complexity will require new approaches to model, analyze, and optimize components across traditional narrow modeling boundaries. To combat complexity, classical DA flows for electronic systems relied on clean sign-off models between different layers of abstraction. For example, the application of standard cells with crisp logic and static timing models facilitated the use of logic and timing abstraction in early logic synthesis flows. Such clean boundaries increasingly disappear on all levels. For instance, starting in the late nineties, logic and physical synthesis merged to address the increasing importance of interconnect delay. In a similar manner we can now observe that the compact models for devices, interconnect structures, standard cells, functional macros etc. become increasingly bloated and require the exposure of more details from lower levels to accurately account for implementation

details. In general, these detailed characteristics become dominated by the physical and electrical context in which the components are placed, thus "uniqifying" the individual instances. Two important challenges caused by this trend can be summarized as follows:

- In the absence of a practical monolithic formulation of the overall design optimization problem, iterative approaches are required to handle global optimization goals. Such methods iterate between higher level design decisions and lower level analysis steps from which the results are used to readjust the design objectives. A fundamental challenge for such global optimization approach is convergence, i.e., ensuring that the iterations improve the design goals monotonically and eventually terminate. For this, the individual optimization steps must be "stable", i.e., small changes at the input cause only small changes at the resulting output. Many traditional algorithms developed in DA, including placement, routing, logic synthesis, do not have this property. Non-stable algorithms will increasingly cause difficulties, as more effects of the implementation must be considered on both sides of the design flow: On the manufacturing side effects such as CMP, litho, resist, and etching play a growing role. On the system side, any intelligent architectural exploration requires quick synthesis runs all the way down to layout to realistically evaluate architectural design choices.
- Typically, much of the computational complexity in DA has been addressed by abstraction, i.e., hiding as much detail as possible to simplify the computation. This reflects a particular "slicing" of the overall design problem in a manner that makes the computation feasible. As outlined above, abstractions are traditionally done "horizontally" using "lean" modeling cuts at the different steps of the design flow. Given that the complexity of the models grows dramatically and more low-level information needs to be considered at higher design levels, this slicing will become increasingly infeasible. New abstraction approaches such as dynamic adaptation of modeling details based on criticality may provide alternative slicing methods that offer more detail when required and less detail when crude estimations are sufficient. Such alternative slicing could have the potential to address the complexity issues in a more adaptive manner but would also have dramatic ramifications on tool architecture, individual algorithms, as well as the overall design flow.