

$$\nabla \times \vec{H} = j\omega\epsilon\vec{E} + \vec{J}$$

TECHNICAL BRIEF

cadence®

$$\dot{X}(t) = f(X(t), u(t))$$

$$\nabla \cdot (\epsilon\vec{E}) = \rho$$

$$\nabla \times \vec{E} = -j\omega\mu\vec{H}$$

$$\nabla \times \frac{1}{\mu} \nabla + j\omega\sigma\vec{E} - \omega^2\epsilon\vec{E} = -j\omega\vec{J}$$

$$\Delta t_k = \frac{-\dot{X}_k}{\ddot{X}_k}$$

$$X(t) = \frac{\partial f}{\partial X} X + \frac{\partial f}{\partial u} u$$

Computational Digital Software

Cadence has been using the term “computational software” to unify many of the algorithms that underlie EDA tools. This is apparent in the digital full flow, which starts at synthesis and runs through physical design and signoff. Cloud vendors sometimes casually say things like “run for a week” but they seem surprised when EDA algorithms literally take a week or longer using 32 or more CPUs during that period. Maybe some machine learning on the largest datasets take similar periods, but the algorithms used have improved a lot. Resnet-50 (a standard neural net trained on Imagenet) only takes 15 minutes on 32 CPUs with 32 GPUs. The ImageNet Large Scale Visual Recognition Challenge contains 1,281,167 images for training, 50,000 for validation, and 100,000 for testing. By EDA standards, these are small numbers.

When engineers explain that a chip has tens of billions of transistors, it is not quite as challenging as it sounds. Firstly, a gate is usually four transistors, and most standard cells contain multiple gates, so the number of “placeable objects” is an order of magnitude less. Most SoCs contain significant amounts of memory, which are generated by a memory compiler. The biggest chips of all, BPUs and multicore CPUs, have a lot of regularity—a core is created and then reproduced to fill the chip. Nonetheless, the computation required to take a SoC from RTL to complete layout is one of the most intense in all of EDA and includes dealing with perhaps hundreds of macro blocks and tens or hundreds of millions of standard cells.

Digital Full Flow

There are multiple core engines involved in the digital flow: engines for synthesis, placement, clock-tree, global routing, detailed routing, timing, extraction, and power. These all need to be best in class to get a best-in-class result. In the past, these engines would run one after the other (synthesis, then placement, then clock tree...and so on). That is not accurate enough for all the physical details that need to be correct in a modern leading-edge process like 16nm and below. There also needs to be world-class integration so that, for example, placement can be considered and partially completed during synthesis, or that synthesis can be used during physical design to restructure logic.

Integration between synthesis, placement, and optimization is so important that we call this iSpatial technology. Interconnect layers in a modern process all have very different capacitance and resistance profiles, so deciding how to use them needs to be early in the design process, which involves techniques such as layer assignment, useful clock skew, and via pillars. The iSpatial technology allows a seamless transition from the Cadence® Genus™ Synthesis Solution to the Cadence Innovus™ Implementation System using a common user interface and database.

Runtime is very important when a design runs for a week. Saving 10% saves more than a day. A savings like this can

result from several reasons: better scalability into the cloud (be able to use 48 machines), better core algorithms (a faster placer), and/or better integration between engines.

Over the past several years, Cadence has created a digital full flow with common integrated engines, as illustrated in Figure 1, resulting in technology leadership. While development is never over, this flow gives Cadence the platform to move into the future. Advanced node (5nm, 3nm, and beyond) is still very important to our most advanced customers.

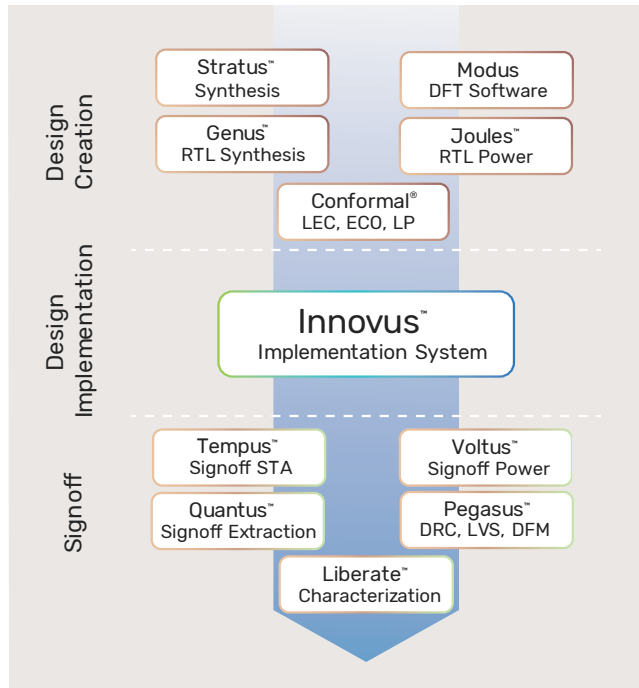


Figure 1. The Cadence digital full flow integrates all of these critical engines

Another important trend is advanced packaging, often called “More than Moore”. Integrating everything onto the most advanced node is not always the most cost-effective approach, and it makes a lot of sense to use the most advanced node only for the heart of the SoC that requires it, and then use a less aggressive node to interface to the outside world. Not only does this approach yield better, but it also can result in faster time to market since the current chip can use the I/O designs from the previous generation, and the current generation I/Os can be developed off the critical path to be ready for the next-generation chip, perhaps a year later.

Machine Learning

Further improvement can be achieved using modern machine learning techniques. One of the limiting factors in any advanced SoC design is the availability of enough designers. In some ways, running EDA tools has something in common with running a nuclear power plant: a lot of it is very repetitive, but only the best engineers are required. Machine learning allows computation, particularly in the cloud, to substitute for routine human interaction and, thus, a major increase in productivity.

The goal for digital full flow is eventual automation, sometimes called “no human in the loop.” But, as with automobiles moving toward autonomous driving, this goal must be approached through incremental steps. Perhaps this can be called “fewer humans in the loop.” When an engineer is running a tool, looking at the result, and then tweaking some parameters before running the flow again, there are a lot of opportunities to tweak the parameters automatically, at least some of the time.

There also is opportunity to improve predictability. At some level, each tool in the flow has a standard of “goodness” that is tied up in how well it integrates with the next stage. A good placement is one that the global router handles well, and so on. Almost all of the algorithms in EDA are computationally intractable, in the sense that getting the exact optimal solution is not possible. Instead, heuristics are used. But this is another area where machine learning can be used, with heuristics at one level using machine learning to better predict how things will be downstream.

Better prediction leads to both a better result and, at least potentially, a faster runtime due to the reduction in the amount of iteration required. Diagrams of a flow always look much more linear than they really are since there is so much iteration going on under the hood. It is like the famous description of a swan as serene on top but paddling furiously underneath. EDA tools often run smoothly from start to finish, but it takes a lot of furious paddling to make everything that smooth.

Mediatek

Dr. SA Hwang, the general manager of the computing and artificial intelligence technology group at Mediatek, said:

“We spend a significant effort tuning our high-performance cores to meet our aggressive performance goals. Using the new machine learning capabilities in the Innovus Implementation System’s GigaOpt Optimizer, we were able to automatically and quickly train a model of our CPU core, which resulted in an improved maximum frequency along with an 80% reduction in total negative slack. This enabled 2X shorter turnaround time for final signoff design closure.”

Summary

The evolution of EDA tools can be broken down into three phases since the mid-1980s, resulting in today’s computational software. By integrating and co-mingling previously independent design, analysis, and implementation tools, optimum results can be achieved for SoC designs. By partitioning and scaling computation to thousands of CPU cores and servers, order-of-magnitude faster computation can be achieved. By introducing machine learning to improve and harness design heuristics for system optimization, productivity can be greatly increased.

The Cadence digital full flow is an excellent example of modern computational software that contains best-in-class engines together with best-in-class integration to make everything work cleanly together. Adding machine learning improves productivity even more.



cadence[®]

Cadence is a pivotal leader in electronic design and computational expertise, using its Intelligent System Design strategy to turn design concepts into reality. Cadence customers are the world’s most creative and innovative companies, delivering extraordinary electronic products from chips to boards to systems for the most dynamic market applications. <http://www.cadence.com>.

© 2020 Cadence Design Systems, Inc. All rights reserved worldwide. Cadence, the Cadence logo, and the other Cadence marks found at www.cadence.com/go/trademarks are trademarks or registered trademarks of Cadence Design Systems, Inc. All other trademarks are the property of their respective owners. 15331 11/20 SA/LL/PDF