

The Challenge Of Re-Use

By Kalar Rajendiran
Senior Director of Marketing, eSilicon Corp.

Overview

Current statistics show that up to 70% of a system on chip is now comprised of third-party IP and homegrown IP that was re-used from previous chips.

There are several very good reasons for this trend. First, it's far quicker to integrate and modify existing IP blocks than to build them from scratch. Second, it's often a safer bet that IP, which is tested in the market and approved for manufacturability by foundries, will work as planned. And third, it usually costs less to buy or license a functional IP block than to build one.

On a whiteboard this all makes perfect sense. Reality is quite different, however, particularly at advanced process nodes. Not all the IP goes together as planned for a variety of reasons. And not everyone has the skill set to make it all work together.

This white paper will examine the factors involved in the re-use of existing IP and the addition of third-party IP.

Area and performance factors

At 130nm and 90nm it was possible to re-use much of the IP inside a chip without concern about whether it would function at the newer process nodes. The biggest concern was that the area would be too large and that it would impact the layout for additional functionality on a chip. More functionality requires more logic, more memory, and an understanding of how those additional features will impact other performance.

At 65nm and 45nm, area has become far less of a consideration because of the amount of real estate available at those process nodes. The shrinkage of wires and transistors has opened up much more room for SoC designers. Now the issue has become whether those re-usable blocks will perform well enough in the newer nodes.

At 28nm and 22nm, there will be so much space available on a die that many chip companies don't even know what to do with it all. That makes physical size of IP far less of an issue, and it also makes it easier to add many more new functions. A mobile Internet device, for example, can easily include everything from complex games to a video camera and global positioning system and still have plenty of room left over for memory, logic, I/O blocks and multicore processors.

The bigger challenge is whether the IP being re-used from one chip to the next can be optimized for the same level of performance—or increased performance—given the other components with which it needs to be integrated. It may not matter as much for a camera function as for an Internet search capability, but too many functions working together can slow the overall device.

Power factors

A far bigger issue than performance is whether all the IP can fit into a given power window.

That becomes particularly complicated as IP is re-used from former designs because it often introduces multiple voltages into an SoC. That not only makes power spikes likely, which can do everything from affecting signal integrity to damaging components, it also makes it harder to reduce leakage with a single architectural approach. And it makes it orders of magnitude harder to verify that all the power domains are working properly and that signal integrity is maintained.

Consider, for example, a simple two-core processor with two power modes—on and off. Each core needs to be verified in each mode, and with both cores in the on and off states so that when both cores are running at peak performance the chip doesn't overheat or draw too much power. Now consider the same chip with eight cores, multiple voltages, multiple modes—on, off, and various states of sleep—and with complex signal interactions. In these cases, re-using IP may be less optimal than buying new IP at the same voltage with built-in verification IP.

Skill factors

While all these problems can be solved, not all engineers are capable of solving them because the complexity has escalated well beyond the scope of training for most engineers. In addition, there is a certain level of trial and error and skill that comes from doing many chips, and fewer companies can afford to do many chips because of the rising cost of development.

While there are no hard numbers, developing an SoC at advanced nodes will cost at least tens of millions of dollars, even when derivative chips are factored into the equation. There are very few companies capable of designing and developing those kinds of chips, and there are even fewer capable of doing multiple chips a year. And while some great advances have been made in tools automation, engineers get better the more they do a certain job or work out a certain problem.

Does an ARM processor have the optimal cache, for example? Can an engineering team put all the pieces together in a way that works right the first time, and which can be optimized for area, power and performance? Is IP from Cadence or Synopsys better for a specific chip or a specific part of a chip? And does that change depending upon the foundry processes?

These are not simple questions to answer, and sometimes there is more art than science in the answer. But getting a feel for a manufacturing process and the quirks of tools isn't something you learn in one chip implementation, which is why the value chain producer model has gained so much traction over the past couple of process nodes.

Business bias

One other consideration is whether the IP being chosen for re-use is the best IP. In the case of large integrated device manufacturers, which have spent significant time and effort creating IP, throwing away that investment in favor of third-party IP would be a political mistake for engineering managers. At the same time, keeping the IP in a design can weaken the final product and ultimately the entire company.

The problem is that what has been developed isn't always the best IP, particularly for a new application. It can compromise performance or strain a power budget, limiting what else can be built into the product. What exists in the arsenal of an IDM isn't always the best-in-class IP for a particular design or a particular market. Ultimately, it can add to the

cost of the design because of the workarounds that are necessary to make it work sufficiently well. It also can limit the overall competitiveness of the product. And in a worst-case scenario, it can do both.

Compare this with a value chain producer's perspective, where there is no bias toward one vendor's IP vs. another's. In this case, it's up to the open market to produce the best IP it can, and the choice is based on a best-in-class analysis for each design, market or specific situation, such as manufacturability or power or environmental conditions, rather than what's been developed in the past.

Conclusions

Disintermediation in the SoC design chain coupled with unprecedented complexity are turning the bulk of the work in chip development into a job for teams of specialists. Just as foundries became the specialists in manufacturing, there is a need for specialists in everything from layout to verification, from system modeling to power modeling, and from writing RTL code to writing firmware and middleware all the way up the applications.

There is massive movement across the supply chain. Development teams may create differentiating IP, but increasingly they are leaving the job of developing more standard IP and for re-using other IP to teams of experts who can best capitalize on it.

Each new SoC is an amalgamation of challenges left over from past implementations of IP, plus new IP, process variations, power issues, performance considerations, and even more complex place-and-route and synthesis. It is a place where the most complex technology issues on the planet need to be solved as efficiently as possible, both financially and with great speed. It will be up to everyone involved in this process to find the place where they can add the most value into a broadening ecosystem—and to shift their focus from doing everything to doing what they can do best.



By Kalar Rajendiran
Senior Director, Marketing
eSilicon Corporation

Kalar Rajendiran is the Senior Director of Marketing for eSilicon Corporation, where he is responsible for the company's corporate marketing, branding and marketing operations. He joined eSilicon in 2000, prior to the company's public launch and plays a key role in eSilicon's market positioning and growth. He established many of the business processes, forward cost models, pricing models and other sales and marketing tools used by eSilicon during its customer engagement phases. He also negotiated and put in place most of eSilicon's strategic IP agreements with its value chain partners.

Rajendiran has held several senior engineering and marketing management positions with National Semiconductor, LSI Logic and Larsen & Toubro. He holds a BE (Honors) in Electrical Engineering from Anna University, an MS in Computer Science from Texas Tech University, and an MBA from Santa Clara University.