Building Up In 3D: the steps from SOC to 3D-ICs

I-Micronews

Stacked die are expected to begin showing up in volume in late 2012 and in 2013, turning what has been a science experiment into a mainstream way of designing and manufacturing SoCs.

This magnitude of this shift cannot be overstated, and clearly all of the pieces are not in place to make it all happen immediately. There also are significant technology challenges to overcome, as with any new technology. Nevertheless, for the semiconductor industry to continue building chips at a reasonable cost, chipmakers will be required to buy or re-use all types of IP, some of which was not developed at the most current process node, and to integrate that with the most advanced digital technology for processors, logic and memory.

“This opportunity is massively wide and very appealing to a lot of customers,” said **Simon Segars**, executive vice president and general manager of ARM’s Physical IP Division. *“But it also has a bearing on how you develop IP, which will have thousands of connectors up to memory. We have R&D work going on with IP, bandwidth issues, and the physics of how to put this all together.”*

ARM isn’t alone. Across the semiconductor supply chain there has been a monumental focus on getting this to work because it can mean the difference between an industry that shrinks to a handful of chipmakers and one that can encompass huge numbers of customers, which in turn consume more system-level tools, IP, platforms, memory and spur grow in a number of new sectors. *“This potentially will change what goes on a die,”* said Segars.

Defining 3D

What exactly constitutes 3D is likely to evolve, as well. The classic system-in-package idea is getting a facelift as better interposer technology between two or more chips becomes available. In addition, memory makers have been stacking die for the past two process generations.

At the most advanced end of the spectrum, 3D chips will be designed from the architectural level and laid out so that signals travel the shortest distance possible with the least amount of signal interference using the lowest amount of power necessary to drive signals. Companies such as IBM, STMicroelectronics, Samsung and Intel have been working on these kinds of structures for the past couple of years primarily for performance and power reasons.

The majority of chips, however, will fall somewhere in between. Instead of designing from scratch, many will use a base platform that has been approved for manufacturability by the major foundries. That will include the processor and the

logic, and at possibly some of the memory. On top of that will be chips that may be developed at a different process node and which are known to work, and which may contain specialized IP developed by a company that will separate their chip from others using similar components.

Sumit DasGupta, senior vice president of engineering at standards group Silicon Integration Initiative (Si2), said the first steps into 3D for fabless companies will likely be *"face on face with an interposer. The second phase will be three layers, which will include I/O, the processor and memory."*

While the IDMs will be the first to turn out these chips, the rules will have to be developed for everyone else—rules that involve how to model these chips, how to lay them out, what thickness the substrate will be, how close together certain components can be and business rules about who's responsible for what.

"The foundries and the technology providers will be driving the rules on this," said **Amit Marathe**, manager of reliability and modeling at GlobalFoundries. "There will be integration stacks and rule stacks. How that gets translated and implemented in the overall flow will need to become interactive."

GlobalFoundries is working with the major EDA companies—Synopsys, Mentor Graphics and Cadence—to develop those rules. Marathe said their customers may play a major role in defining those rules, as well. *"They need to set the guidelines of what can be competitive and what is not competitive,"* he noted.

From an EDA standpoint, what will become essential is integration of tools that can understand everything from place-and-route to thermal effects in multiple dimensions. *"In many ways the solution is just like multicore,"* said **Shay Benchorin**, director of marketing for Mentor Graphics' embedded software division. *"There are very high speed hardware lanes and multiple lanes out of the processor."*

What's missing

Other pieces are missing from the 3D flow, as well, most notably an understanding of the thermal effects of stacking on silicon stress. Putting chips together can result in a non-uniform expansion and contraction of the silicon. At older nodes and in two-dimensional chips this isn't particularly noticeable. In high-density 3D structures it can have a significant impact on performance.

"At advanced nodes the die is thin so the TSVs through the die can add to the complexity of stress," said GlobalFoundries' Marathe. *"Initially there will have to be restrictive design rules to make sure that stress does not interfere with device performance. You need sufficiently long distances between the devices on a chip."*

CEA-Leti, the French consortium, is working on a 3D design flow in conjunction with companies such as STMicroelectronics and a number of startups out of the Grenoble ecosystem. Sylvian Kaiser, CTO of startup DOCEA Power, said the entire supply

chain will need some restructuring, as well.

“The relationship between suppliers and buyers will change,” he said. *“If you’re looking at package-on-package integration, the people who provide the memory chip will no longer sell to the equipment supplier. They will now collaborate with the chip developer.”*

That’s one of the primary focal points for standards group Silicon Integration Initiative, or Si2. **Steve Schulz**, president and CEO of the standards organization, said the first things that need to be addressed are common definitions for the supply chain and economic risk factors.

“We’re envisioning several stages of 3D,” Schulz said. *“Next will be 2.5D, which is an interposer between two chips, which will be followed by 2.9D, which will include some algorithmic calculations of some parts, but not actually designed that way. The design will be conceptual, but not all the analysis will be there.”*

That will be followed by process design kits that are completely 3D aware.

“What’s holding this all back is the market needs to be ready,” he noted. *“We might be able to get enough players together for a dictionary of terms in the next half year, although we’re not committing to a time table. That’s the foundation for the low-hanging fruit standards. After that we start looking at connecting thermal and packaging and yield.”*

Advantages and disadvantages

One advantage to using advanced processes in conjunction with older processes is that there is room for error—literally. Guard-banding will be almost essential in the first iterations of 3D stacking, and at 22nm there is plenty of available real estate.

“Until the models and prototypes are validated we will proceed with caution,” Marathe said. *“We will build up margin in the initial implementations, and as confidence grows we will get more aggressive about cutting out that margin. But the big advantage here is that you don’t have to scale everything in the same manner. The motivator in all of this is that it will put relief into the process challenges so you can put pieces together in the most meaningful manner.”*

Another great advantage is that more companies will be back on the Moore’s Law road map, even if they don’t develop at the most advanced node. If 3D design takes off as many industry experts expect it will, there will likely be a handful of players producing standardized platforms that can be connected to more innovative structures. That also means less time spent on what doesn’t add value—including verification of those platforms—and far less time spent integrating advanced digital designs and developing analog for nanoscale processes.

This isn’t an advantage for all companies, however. Those chipmakers that previously developed all the pieces may find themselves in a price/performance

war, relegating their platforms to commodity status instead of those platforms being the real differentiator in a chip. It also will shift much of the competitiveness in semiconductors to the software and IP side, with an emphasis on system-level tools for integration, verification and manufacturability. That will put many startups at a disadvantage at a time when startups are particularly needed to help solve some of the integration issues in 3D.

Conclusions

There is almost universal agreement that 3D will be important, game-changing and necessary. What isn't well understood are all the ramifications of this move, from the technology side as well as the business side. Who will be responsible, for example, if one layer of a multi-layer chip doesn't work with another layer produced by another company?

Some of the technology has been tested in limited ways, most notably through-silicon vias, but how chips will work when there are thousands of TSVs running through all layers remains to be seen. Interconnects need to be modeled, tested and verified to provide a better understanding of effects such as heat, noise and electromigration and electrostatic discharge. And while some very effective tools exist for automating chip design, there will be huge new opportunities for tools that can account for multiple effects at every step of the design process.

This is hard stuff and there are no simple answers. But if the cost of chips is to be tamed once again and brought back under control, it appears this is also necessary work—and work that will spawn many new opportunities.

SOURCE [1]

Source URL (retrieved on *04/23/2014 - 5:25am*):
http://www.mdtmag.com/news/2010/11/building-3d-steps-soc-3d-ics?qt-video_of_the_day=0

Links:
[1] http://www.i-micronews.com/lectureArticle.asp?id=5842