EE|Times
BETA
EE Life

SEARCH   ADVANCED SEARCH

Home   News & Analysis   Business   **EE Life**   Embedded.com   Design   Products   Education & Training   Events   Video   Sign In   Join

EE Life
Blogs
Message Boards
Start a Message Board

EE Times Home > EE Life

Blog

Xilinx multi-FPGA provides mega-boost re capacity, performance, and power efficiency!

Clive Maxfield
10/27/2010 1:36 PM EDT

I just returned from a trip to visit Xilinx in San Jose, California. While I was there, they showed me one of the most exciting things I've seen recently in programmable logic space (where no one can hear you scream).

The underlying idea seems simple enough – stick multiple FPGA die into a single package. It's the implementation that's so clever, and – as we shall discuss – the ramifications are truly enormous!

Xilinx 7 Series FPGAs
Before we plunge into the fray, let's take a step back to remind ourselves that – in June of this year – the folks at Xilinx announced their forthcoming FPGAs to be implemented at the 28 nm technology node. They're collectively calling these *"Xilinx 7 Series FPGAs."*

The folks at Xilinx are presenting these new devices as three ... hmmm ... I hesitate to say "three families" because that implies that they are functionally different ... I personally prefer to think of them as three branches of the same family, where these branches are called Artix-7, Kintex-7, and Virtex-7. All three of these family branches share a unified architecture that allows for ease of design, migration, and IP portability.

And where did these three names come from?

- **Artix** is rooted in Arctic, suggesting cool, low-power. The Artix-7 branch of products will offer 50% lower power and 35% lower cost than the Spartan-6 family, making this branch ideally suited for the cost-sensitive, high-volume markets served by ASSPs and ASICs.
- **Kintex** gets its roots form the word kinetic for movement and energy. Kintex offers the best combination of price and performance. The new Kintex-7 branch will deliver the performance of the existing 40 nm Virtex-6 family at half the cost. It not only addresses aggressive power and cost requirements with significant price/performance improvements over Virtex-6 and Spartan-6, but will also deliver on the emerging need for insatiable bandwidth in applications that include next generation broadcast systems and wireless networks.
- **Virtex** represents the summit and highest capability. The Virtex-7 branch will deliver the highest system performance and capability. The Virtex-7 branch will provide 2M logic cells and deliver 2X the system performance over previous generations. Virtex-7 is designed to meet the extreme performance needs of wired infrastructure, high performance computing (HPC) systems, and aerospace and defense among others.

Stacked Silicon = "More than Moore"
What the folks at Xilinx didn't say when they announced their 7 Series FPGAs was exactly how they intended to achieve the 2M logic cells in the Virtex-7 devices. I guess that (like most folks) I simply assumed that they were going to make bigger and bigger die containing more and more transistors. And, of

Comment

Paul Dillien
11/1/2010 2:15 PM EDT
In reply to "goafrit" on the likely selling price of the 2M LUT device, you can ...

Les_Slater
10/30/2010 11:55 AM EDT
My first experience with Xilinx was in the early stages of the XC2064; I may ...

More Comments >

Navigate to related information

News: Cadence, Xilinx form FPGA IP ecosystem microsite

Product: Smarter auto system designs with Xilinx FPGA development platforms

Product: Latest Xilinx ISE Design Suite supports plug-and-play ASIC-FPGA design

News: Xilinx demos Intel QuickPath technology for FPGAs

Product: Xilinx FPGA development platform for 3D TV

Most Popular

1 Is there a future for hybrid vehicles? Part 2
2 LED improvements extend automotive applications
3 USB Explained: An Introduction to USB and Its Future
4 Digital Signal Processing: A Practical Guide (Part 1)
5 Current Sensing Circuit Concepts and Fundamentals
6 Fundamentals of Wireless
7 Digital Signal Processing: A Practical Guide (Part 2)
8 The Inefficiency of C++, Fact or Fiction?
9 Digital Signal Processing: A Practical Guide (Part 3)
10 Digital Signal Processing: A Practical Guide (Part 4)

Product Parts Search

Enter part number or keyword

SEARCH

course, they will be doing this, but there's much more to the story...

The problem is that when you first move to a new process technology there are issues with yield. Smaller devices have higher yield as illustrated in the graphic below. This explains why FPGA vendors typically come out with their mid-range devices first, because the larger FPGAs only become viable much later in the life-cycle once the process has been fine-tuned.
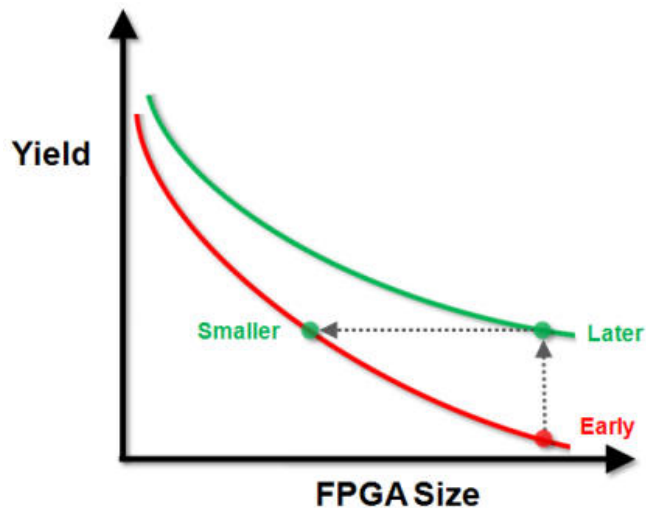

So how can this issue be addressed? Well, the folks at Xilinx have come up with something rather clever, which is to place multiple smaller die in the same package. Take a look at the graphic below. The four gold rectangles in the middle represent FPGA die. The large green square represents the main chip package. And we will return to consider the light blue square surrounding the four FPGA die in just a moment.
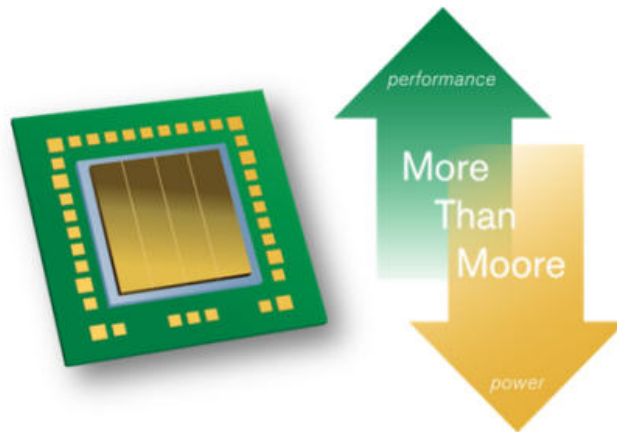
Now, having multiple chips in one package has been done many times before. Back in the 1990s we used to call them *Multi-Chip Modules (MCMs)*. More recently, we started to use the term *System-in-Package (SiP)*. We might think of the new Xilinx solution as *SiP TNG* (the next generation).

In conventional SiPs the die are attached directly to the package substrate. In this case, compared to the tracks on the die, the tracks on the package substrate are relatively large, slow, and driving signals onto them consumes a lot of power. What Xilinx are doing is to use a special layer of silicon known as a "silicon interposer" combined with Through-Silicon Vias (TSVs) as illustrated below:
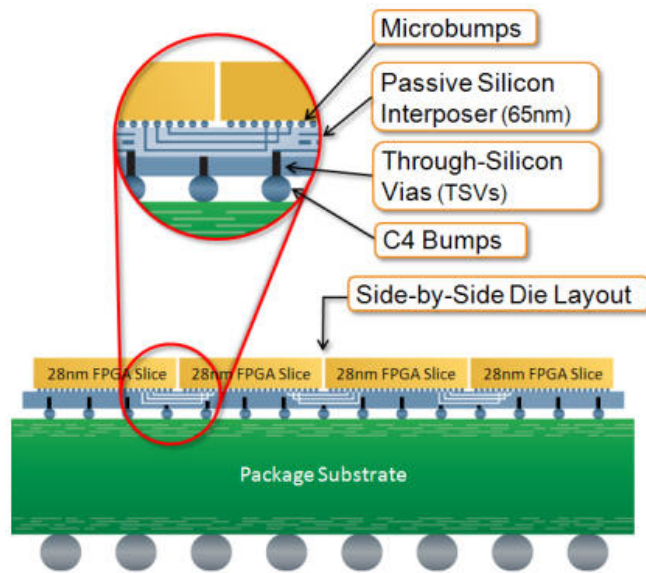
This technology may be referred to as *"Stacked Silicon Interconnect"* by some and *"2.5D integrated circuits"* by others. Depending on who is doing what to whom, the silicon interposer may be purely passive (that is, contain only tracks) or it may be active (it may also include devices like transistors and logic gates ... all the way up to complex macros and cores).

In this first Xilinx incarnation, the four FPGA die are implemented at the 28 nm technology node, while the *passive* silicon interposer is implemented at the 65 nm technology node. Implementing the large silicon interposer at this higher node reduces costs and increases yield without significantly degrading performance.

One way to think about this is that the silicon interposer essentially adds four additional tracking layers that can be used to connect the FPGAs to each other. And how many connections are we talking about here? Well, I bet you'll be surprised when I tell you that there are more than 10,000 connections between each pair of adjacent die!

On top of this, Through-Silicon Vias (TSVs) are used to pass signals through the silicon interposer to C4 bumps on the bottom of the interposer. These bumps are then used to connect the interposer to the package substrate.

Compared with having to use standard I/O connections to integrate two FPGAs together on a circuit board, this stacked silicon interconnect technology provides over 100X the die-to-die connectivity bandwidth-per-watt, at one-fifth the latency, without consuming any of the FPGAs' high-speed serial or parallel I/O resources.

Furthermore, by having the die sit adjacent to each other, Xilinx can avoid the thermal flux and design tool flow issues that would be introduced had a purely vertical die-stacking approach been adopted.

With regard to the design tools, the tracks on the silicon interposer are – to a large extent – seen as simply being long lines. The folks at Xilinx say that designers can simply "Press the Big Red Button" for the entire design to be automatically implemented across all four FPGA die as though they were a single large die. Alternatively, if the users wish to partition the design across the four die by hand, they can obtain 8 to 10% performance improvement on top of the staggering performance that is already offered by this technology.

Proven technology and supply chain
Of course, it's easy for folks to jump up and down and wave their arms around and tell you all sorts of things that sound wonderful, but when you come to look closely they aren't really there. There's a world of difference talking about this stuff and actually doing it. But I think it's safe to say that Xilinx have actually succeeded here, not the least that they've been working on it for a long, long time.

Personally I am amazed that they've managed to keep this secret. It seems that they've actually been working on this for the last four or five years. They created their first test vehicle at the 90 nm node in 2008; the second test vehicle at the 40 nm node in 2009; and the third test vehicle at the 28 nm node this year in 2010 (check out the picture of this latter test vehicle below – I actually held this little beauty in my sweaty hands):
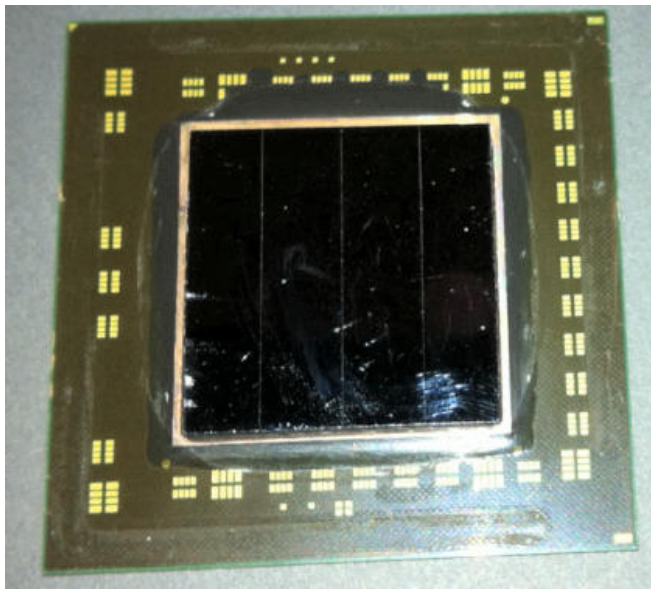
Now, the silicon graveyard is littered with technologies that seemed to be a good idea at the time, but which never succeeded because their originators failed to ensure that all of the players were in place. Not this time – as illustrated in the graphic below, Xilinx have fully solved the big infrastructure supply chain problem, which has proved to be the show-stopper for other folks in the past.
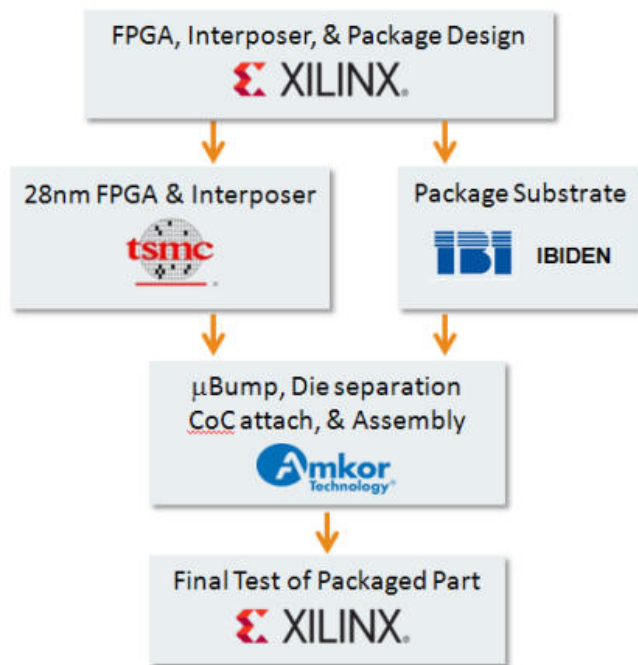
The current state of play is that the Xilinx design tools are already geared up to take full advantage of this new technology starting with the ISE 13.1 Beta release, and we should be seeing the first engineering samples of this technology around the middle of 2011. I cannot wait!

Mind-Boggling Implications
As I mentioned at the beginning, the ramifications of this new technology are truly enormous! Here are just a few thoughts off the top of my head...

First, as I noted above, FPGA vendors typically come out with their mid-range devices first, because the larger FPGAs only become viable much later in the life-cycle once the process has been fine-tuned. Well, by gathering four medium capacity die (and remember that the term "medium" is relative – these are actually honking big die whichever way you look at them) into a single package as described here, the effect is as though we had immediate access to the largest members of the family.

Another way to look at this is that we are getting next-generation density in this

generation's technology. And, of course, as the process becomes fine-tuned and the yield improves, the folks at Xilinx can boost the capacity even further.

Another consideration is that we don't have to limit ourselves to four FPGA die in a package. There could be fewer (two or three) or more (six, eight...). Also, the FPGA die don't have to be homogeneous. Although my understanding is that all four die will be identical in the initial releases, there's no reason why Xilinx may not decide to "mix-and-match" in the future – for example, combining two DSP-intensive die with two SERDES-intensive dia. Or how about replacing one or two of the die with pure memory die, or .... I tell you, the more you think about this the more exciting it becomes.

And one last point to ponder is that Xilinx currently have a push to use SRAM-based FPGAs in space applications. (Creating radiation-tolerant SRAM-based FPGA designs is something of a "hot-button" for me at the moment). Well, how about using three of the die to implement triple-modular redundancy (TMR) at the die level, and then use the third die to perform house-keeping tasks like implementing the voting circuits and constantly reading the configuration data for the other die, performing CRC checks on that data, and reloading as required. (The functions on this fourth die could themselves be implemented in TMR fashion and this die could also be monitoring and reloading its own configuration data as necessary.)

New ideas are popping into my head as I pen these words. I tell you, we certainly do live in interesting times...

print   email   rss   ⊕ SHARE   💬 post comment   16 Comments

Comments

10/27/2010 1:54 PM EDT

I am really excited about this announcement. My head is buzzing with regard to the implications. Can you think of anything I've missed?

Max the Magnificent

Sign in to Reply

10/28/2010 11:07 AM EDT

This really is exciting, especially the potential to use dissimilar die. Xilinx should license the interconnect technology and publish die requirements for interface. No limit to what one could do here. There's lots of natural applications. One that comes to my mind is an image sensor between a couple of these die.

Les_Slater

Sign in to Reply

10/28/2010 11:08 AM EDT

Hi Max

Very interesting. As you mentioned, somewhat similar technology has been tried before. Don't I recall wafer scale integration which held the promise that it would yield humungous great memories?

Paul Dillien

Did the Xilinx guys make any mention of on-chip redundancy? This tried-and-trusted technique has been used for years in regular structures like DRAM, and the columnar structure of a Virtex FPGA looks very similar. Altera don't talk about it as much these days, but I'd be very surprised if they are not still using redundancy to boost their yields. It's most beneficial on large die, which is just where your yield graph needs a lift.

The "SiP TNG" could be combined with a NV-memory to hold the configuration (as per Spartan-3AN) to provide a single chip solution. It could also hold the program code for those embedded ARM processors that are expected next year. Adding a DRAM chip

would provide system storage that is typically "off-chip". What about analog? A couple of high performance ADC's and the product starts looking attractive to the base station guys. Perhaps an RF chip? Couple this with a low-cost Artix and you could address more wireless applications. The list goes on...

Paul

10/29/2010 7:27 AM EDT

irrelevant point: a pedant would like to say that the plural of 'die' is 'dice'. I mention it because for example 'these three die' brings me up short, to a shuddering halt, and disrupts the flow of the article. But what do I know...

DickH

10/29/2010 9:33 AM EDT

What do any of us know? Why is the plural of herring herring? If you look at the Wikipedia is says:

"The wafer is then cut into rectangular blocks, each of which is called a die. Each good die (plural dice, dies, or die) is then..."

Max the Magnificent

If you visit Dictionary.com the plural is *dies* with regard to machinery or an engraving stamp and *dice* with regard to architecture (also *dice* with regard to the small cubes with numbers used in board games).

This a tricky one -- I know what you mean when you say it brings you to a shuddering halt -- there are other things that do the same to me -- but I think it's what you are used to -- when I was new to the industry I was told that the plural of die was die (in the context of unpackaged silicon chips) and it sort of stuck -- sorry :-)

10/29/2010 8:05 AM EDT

Good one from Xilinx. They lead this industry and certainly the industry leader. Xilinx is competing very well and innovating. But who knows what the next Gen of these devices will cost with all these new features and technologies. Hope it can still be affordable.

goafrit

Did I miss it? When is this new idea coming to the market?

10/29/2010 9:22 AM EDT

I think they will have engineering samples mid-2011

Max the Magnificent

10/29/2010 11:13 AM EDT

Thanks Max, I dropped out of high school because I was flunking English and wouldn't have had enough credits to graduate anyway. English usage is not

Les_Slater

always consistent and that's true in the technical realm also. Also 'dice' as plural for 'die' grates on me. Hope nobody's loosing any sleep over this. For a tendency to shudder over trivia, see a physician.

10/29/2010 12:04 PM EDT

There are many reasons why English isn't consistent -- I'm actually planning on writing a series of articles on English (grammar, history, other stuff) explaining things like how to set about writing articles and papers and suchlike -- I'm hoping it will be a lot more interesting than it sounds here (grin) -- I'll try to post the first one next week -- Max

Max the Magnificent

10/29/2010 2:34 PM EDT

Max, it might be useful to re-read Warren Weaver's introduction to the '63 edition of Claude Shannon's 'A Mathematical Theory of Communication'.

Les_Slater

10/29/2010 2:59 PM EDT

At the moment I am so back-logged with stuff to read that I cannot think of adding anything else to the pile (my head hurts :-)

Max the Magnificent

10/29/2010 3:11 PM EDT

It's fairly short, sixteen pages, and not very technical. A short word of introduction from Warren Weaver:

"The word communication will be used here in a very broad sense to include all of the procedures by which one mind may affect another. This, of course, involves not only written and oral speech, but also music, the pictorial arts, the theatre, the ballet, and in fact all human behavior."

http://academic.evergreen.edu/a/arunc/compmusic/weaver/weaver.pdf

Les_Slater

10/29/2010 3:17 PM EDT

Oh, OK, I can squeeze that in ... thanks for the link -- Ill read it this weekend -- Max

Max the Magnificent

10/30/2010 1:38 AM EDT

This does not look very good implementation. Xilinx tools - ISE and EDK - are not so friendly or stable. This design implemntation will using them more difficult. Also, as new process geometry evolves, this may soon become obsolete. We would also like to know the reliability data from Xilinx.

hm

10/30/2010 11:55 AM EDT

Les_Slater

My first experience with Xilinx was in the early stages of the XC2064; I may have gotten some of their first product. My design was a 3-chip CRT controller. The three chips handled I/O for independent CRTs, the timing, memory addressing and I/O and serializing the data into 2-bit pixels for each monitor. All three parts were synchronous from a single clock source. The biggest difficulty was matching propagation delays, never mind min/max. The most difficult was the I/O (remember data was going through the FPGA to the memory, back to the FPGA and serialized out while the RAM addressing was being kept within timing constraints). Like I said I/O timing was the most difficult to accommodate. This was 25 years ago but relatively speaking we have some of the same problems. Back then we had problems that wouldn't fit into a single die; we still do.

11/1/2010 2:15 PM EDT

Paul Dillien

In reply to "goafrit" on the likely selling price of the 2M LUT device, you can figure out a ball-park.

I analysed the prices of existing products in my "FPGA Market Report" (a shameless plug...) based on the cost of 1k off 4-input logic elements.

Using the largest device currently offered by Xilinx (XC6VLX760-1FF1760C) with a claimed equivalent capacity of 759k LEs and the sample quantity pricing of $15,622, this translates to a price of $20.59 per 1,000 LEs. The proposed XC7V1500T will have have 1,954,560 4-input equivalents. Assuming that the price of the largest 28-nm Virtex-7 device is set, say, at 70% of the existing 40-nm Virtex-6, suggests that the price per 1k LEs would be around $14. Multiplying this price by the capacity suggests a ball-park selling price of around $28,000.

That level of pricing would not be an issue for customers, say, who are prototyping ASIC devices.

The other way to think about it is that Xilinx has announced a device that is significantly bigger than anything proposed by Altera, and can price accordingly.

Perhaps I should not discount the recently announced Achronix device with 2.5 M LE equivalents to use the Intel 22-nm because samples of the first device are promised for Q4 of next year, compared to XC7V1500T in 2H of 2011.

Please sign in to post comment

Submit Comment   Follow Comments

More EE Times   **EE Times Network**