Q&A: Embedded FPGA Facilitates Reconfigurable SoC/MCU RTL Blocks

Electronic Design
William Wong
Fri, 2016-10-14 12:46

Field-programmable gate arrays (FPGAs) have provided developers with flexibility. Recently, FPGAs have incorporated system-on-chip (SoC) microprocessors such 64-bit, ARM Cortex-A57 cores. This simplifies programming and integration, but there are other alternatives as well, such as Flex Logix's embedded FPGAs that make it possible to integrate an FPGA within an ASIC.

I talked with Flex Logix's Cheng Wang about the firm's embedded FPGA technology. The two-and-half-year-old year old company already has embedded FPGA IP available in TSMC 28 nm and TSMC 40 nm, and will soon be available in TSMC 16 nm.

Wong: I hear Flex Logix won the Outstanding Paper Award at ISSCC, which means you must have some impressive technology. Tell me a little bit about the company and what you have developed?

Related

Q&A: Embedding FPGAs into Custom SoC Designs

Understanding FPGA Processor Interconnects

The Principles of FPGAs

Wang: FPGAs today are about 75% to 80% programmable interconnect and only 20% to 25% programmable logic. We invented a new high-density interconnect that makes the FPGA about half the size, with higher utilization and fewer metal layers. We realized that this would be very attractive for enabling embedded FPGAs in SoCs, so we developed IP cores that we prove in silicon and use to construct embedded FPGA arrays from 100 to >100K lookup tables (LUTs), with optional MACs and RAM. We also provide the software to program them. Our IP is proven on TSMC 28HPM/C, is in fab on TSMC 40ULP, and in design on TSMC 16FF+/c.

Using our platform, customers are able to design their next-generation MCUs and SoCs with reconfigurable RTL that can be quickly, easily, and cost-effectively updated or changed at any time after fabrication, even in-system. As every chip designer knows, having to change the RTL blocks at any point in the design process could easily cost multiple millions of dollars and add three to six months to the design schedule. Flex Logix has solved that problem.

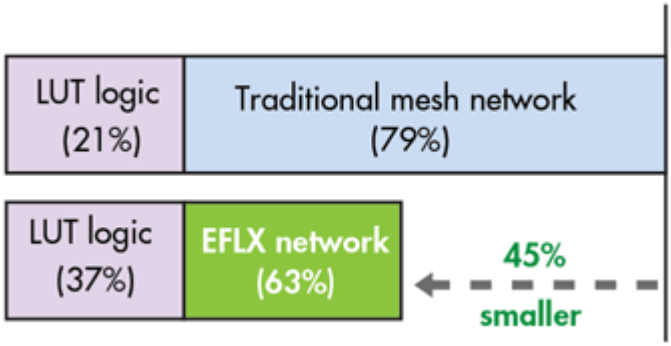

ng: Processors have been a popular embedded IP since the 1990s. Why has it taken so long for embedded GA to become available?

Wang: A few chip companies have some programmable logic on their chips, and we expect this trend to continue. The Flex Logix platform is a complete, scalable, silicon-proven embedded FPGA offering with high density; any SoC or MCU company can integrate as little or as much as they need.

Wong: Why do people want to embed an FPGA in their chip anyway?

Wang: Two reasons. First, some customers want to be able to update critical RTL, such as protocol logic, to keep up with changing industry standards and customer requirements. With embedded FPGA, the RTL can be updated quickly at anytime and in the field in-system, like a firmware update for an embedded processor.

Second, some customers want to be able to customize a single-chip design for multiple customers or market segments by changing the RTL in a certain part of the chip. For example, MCUs at 90 nm have dozens of variations with small differences in the I/O. At 40 nm, masks cost $1 million, so being able to offer the dozens of variations with a small amount of embedded FPGA can save a significant amount of money and speed time to market for the custom versions. Customers might not even be told an embedded FPGA is being used to provide the customization.

Wong: How do customers use embedded FPGAs?

Wang: We see a wide range of applications, from very large networking chips down to small MCU/IoT chips. In 40-nm applications such as MCU/IoT, the emphasis is on power, so we optimize our offerings to have more power-management modes, low-voltage state retention, and other features. In 28-/16-nm applications, the emphasis is on performance, so we optimize for that.

The highest performance requirement is typically where we operate in the control path or data path and we need to clock at the frequency of the surrounding hardwired RTL ASIC. In this case, customers typically use us in blocks of 1000-or-less LUTs implanting fast control logic with 1 or 2 LUT stages between flops.

I/O requirements tend to be very large, especially on inputs. A relatively lower performance requirement is I/O control, such as in a MCU or the IoT. Here, we can enable local processing of I/Os to reduce the overall system power by not having to activate the MPU, or implement additional serial I/O functions as needed. An intermediate application is where we are a block of reconfigurable RTL on a processor bus.

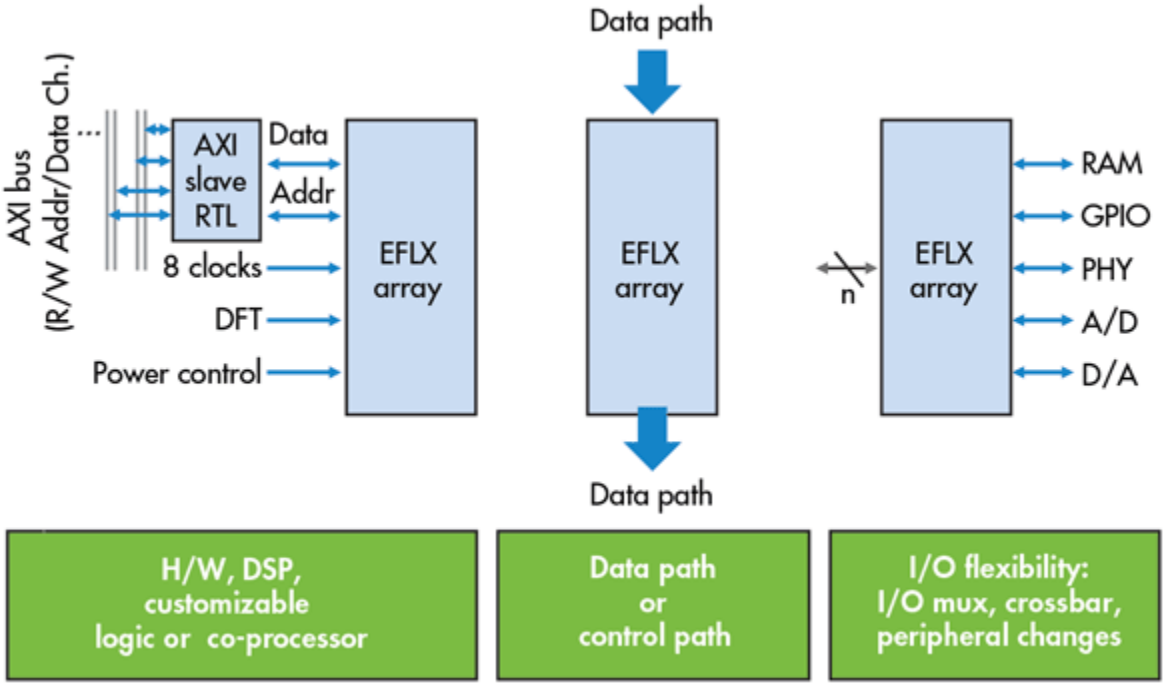

Wong: What is the difference between an FPGA chip and embedded FPGA?

Wang: The technology is very similar, except with embedded FPGA there's no need for SERDES and PHYs because on-chip signaling is very fast. The density is similar, at least for Flex Logix's embedded FPGA, in the same process generation. The real difference is the users. FPGA chips are used primarily by systems companies, some in high volume. Embedded FPGAs are used primarily by chip companies who need to integrate a small amount of FPGA-like flexibility into their chips.

Wong: Why don't they use an external FPGA chip instead of embedded FPGA?

Wang: It comes down to I/O. Chip-to-chip I/O is relatively slow and pins/packages are expensive. On-chip block-to-block data connections can be much wider and faster. If a customer needs a 512-bit bus coming into our embedded FPGA block at a 500-MHz to 1-GHz data rate, that's impossible to do with an external FPGA. Cost and power are other considerations that favor embedded FPGAs. This is analogous to SRAM, because on-board SRAM can have wide and fast buses, much wider and faster than if the SRAM were on a separate external chip.

Wong: What are your specific offerings today?

Wang: We offer two hard IP cores, the EFLX-100 and EFLX-2.5K. The 100 and 2.5K describe the number of LUTs (each of our LUTs is a 5-input LUT or dual-4-input LUTs with two flip-flops). We have an all-logic and DSP (embedded MAC = multiplier-accumulator) version of each. The cores are arrayed to make the EFLX array size the customer needs from 100 to >100K LUTs. To get density, we offer hard IP.

Today, the EFLX-2.5K Logic/DSP IP cores are available in TSMC28HPM/C, are proven in silicon, and are in customer designs. The EFLX-100 Logic/DSP IP cores are available in TSMC40ULP and a validation chip is in fab now. We have also begun design of EFLX-100 in TSMC16FF+ and will quickly follow that with EFLX-100 Logic and EFLX-2.5K Logic/DSP in TSMC16FFc. In addition, we have the EFLX Compiler, which we provide to customers to map their RTL into the EFLX array, determine timing, and generate a bitstream for programming the chip. Customers have been using the EFLX Compiler for a year now.

	EFLX-100 Logic	EFLX-100 DSP	EFLX-2.5K Logic	EFLX-2.5K DSP
T 40ULP/EF	In fab	In fab	–	–
T 28HPM/HPC	–	–	Silicon proven	Silicon proven
T 28HPC+	Roadmap	–	Roadmap	Roadmap
T 16FF+/FFC	In design	–	In design	In design

Wong: You've obviously done a lot of chip design and now work with a lot of different customers doing chip design. Do you have any advice for chip designers today?

Wang: If you focus on minimizing costs, you need to remember that cost is more than silicon area. There's a cost to mistakes that take mask spins, and there's a cost to not having the functionality a customer needs and losing the business. There's also a cost to not being able to handle changing protocols and shunting the task to a slow, overloaded processor. In any case, the focus should be on maximizing ROI (return on investment) and trying to extend chips' lives by adding flexibility for unforeseeable change.

Wong: What is your vision for Flex Logix in five years?

Wang: Geoff Tate, our CEO, was the founding CEO of Rambus, which had three people when he joined and became a multi-billion-dollar market cap public company. Our goal at Flex Logix is to do the same. We have recruited venture capitalists who share our goal, such as Peter Hebert of Lux Capital and Pierre Lamond of Eclipse Ventures. Pierre was co-founder of National Semi and chairman of many major chip companies including Cypress and Microchip.

To scale rapidly, we have hired executives with proven track records who have managed much larger groups so that they can help us build Flex Logix quickly as demand grows. This includes a new lead systems architect we hired from Intel to build a team that provides architecture and applications support for assisting customers in integrating embedded FPGA. We see very high interest in our technology at most companies we have met with, and believe embedded FPGAs will become a pervasive technology.

Source URL: http://electronicdesign.com/fpgas/qa-embedded-fpga-facilitates-reconfigurable-socmcu-rtl-blocks