

A Guide to Modern FPGAs

FPGAs are growing fast in technology and density. The latest advances have generally helped the design engineer to handle complex designs in FPGAs with better accuracy in lesser time. Let us see how it is done



Microsemi IGLOO2 FPGA

ANAGHA P.

Design architects are now more comfortable selecting a field-programmable gate array (FPGA) as the central piece on the board. Given the quantum of vendor and community resources available at their disposal, FPGAs makes it easier for them to complete their design within stipulated timelines and budget.

FPGAs are now capable of serving as suitable alternatives to application-specific integrated circuit (ASIC) or application specific standard product (ASSP) implementations in many systems—so much so that some vendors are marketing their FPGAs as ASIC-likes.

Neeraj Varma, director of sales at Xilinx explained in an interview, “The whole belief is that the FPGAs in UltraScale generation have a lot many more features that are ASIC-like. The

key here is the ASIC-like class architecture. There are several changes we made in the architectures that are like ASICs. ASICs are always good in power and performance.

“FPGAs, traditionally, have an overhead, but we are trying to bridge this gap, and we did quite well in 28nm. But in 20nm, we are going to do even better.”

Earlier, FPGAs had few global clocks. As the density grew, skew problem also increased. Vendors, such as Xilinx, added more clocking regions, resulting in more clocking resources placed at different areas of the chip. This is how exactly it is done in ASICs to solve skew problem, and this is the reason why it is called an ASIC-class clocking.

Here, the software used is also an ASIC-strength design suite. The design suite has many built-in features where you are not allowed to make any mis-

takes, similar to how it is done for ASIC design tools. The tool itself takes care of so many things, saving your valuable time in repeating the design process.

Denser FPGAs

One of the points I came across in a white paper from Altera was that, because of shrinking of design cycles and design teams, cost-conscious designers were trying to move away from multi-chip solutions that were too costly and lacked flexibility. They were also moving away from single-chip solutions that could not meet power or performance objectives, or ASIC systems on chips (SoCs) which are too slow to get to market.

ARM-based SoCs combine a hard ARM processor, memory controllers and peripherals with customisable FPGA fabric in a single SoC. Initially, FPGAs were not dense enough and did not have hard cores on them.

Now, even for a single family of product, the user has quite a few options to select from, such as ARM-based SoCs that are user-customisable. This helps reduce system power and board size while increasing system perfor-



Virtex C2 AE-5 TXT FPGA from Xilinx

mance by integrating discrete FPGAs, digital signal processing (DSP) and microprocessor devices into a single, user-customisable SoC.

The built-in hardcore IP cores, which were not present earlier, are available now, providing more predictable performance figures. Every year we see shrinking of the process node in line with the industry growth and Moore's Law. Now 45nm, 28nm, 20nm and 16nm FPGAs are available. Also, now we are moving towards 14nm and 10nm nodes. This favours better density and reduction of power.

Adaptive look-up tables

Over the years, the basic building block of the FPGA logic fabric—the look-up table (LUT) based function generator—has remained almost the same. Experimentations were done with the number of inputs for LUT such as three, four, six and eight. Now they have adaptive LUTs which allow two outputs per LUT with two function generators to be implemented, sharing some of the inputs. The FPGAs mainly evolved by providing the same set of resources with different number of resources per family.

In addition, today it is all about providing ecosystems to the designer—evaluation boards, reference designs, design communities and partner IPs, ensuring FPGAs are able to meet the stringent time-to-market goals of the customer.

Things to consider while selecting an FPGA

If we presume that a device meets the basic design requirements, the next four things an engineer needs to consider are power consumption, board space, solution cost and design time. Whatever the application may be, these four factors can bias an engineer's decision across the vendors.

When it comes to low power consumption, iCE40 FPGA family from Lattice Semiconductor and IGLOO series of low-power FPGAs from Microsemi are some great choices in most low-power and battery-operated applications. It is particularly favourable in the Internet of Things (IoT) space where ultra-low-power programmable devices with ultra-small packaging are what a design engineer would be looking for.

Available resources, such as look-up tables (LUTs), registers, I/Os, phase-locked loops (PLLs) and memory, are considered before selecting an FPGA.

Another point to look for is the dedicated hard blocks that could be interfaced with the FPGA, such as the physical coding sublayer (PCS), peripheral component interconnect express (PCIe) end point and double data rate (DDR) memory controller. These could be added to the periphery for high-speed interface with the external world.

If the application requires a processor and an FPGA, then user-customisable SoCs are the best choice.

Selecting FPGAs for an application

High-end FPGAs feature:

- High performance
- High logic density
- High-speed interfaces
- UltraScale and 3D architecture
- Support complex features, architectural blocks and intensive computations
- High power consumption
- High cost
- Example: Virtex family by Xilinx, Stratix family by Altera, ProASIC3 family by Microsemi and Speedster 22i family by Achronix

Low-end FPGAs feature:

- Low power consumption per chip
- Low cost
- Low complexity
- Low speed of operation
- Low logic density
- Low-speed interconnects
- Do not support some features and architectural blocks
- Example: Spartan family by Xilinx, Cyclone family by Altera, Mach XO and ICE40 families by Lattice Semiconductor and Fusion family by Microsemi

Mid-range FPGAs


The mid-range FPGAs are an optimal solution that provide the users with a balance of cost and performance.

Example: Artix-7 and Kintex-7 series by Xilinx, Arria series by Altera, ECP5/ECP3 series by Lattice Semiconductor and IGLOO2 series by Microsemi

Selecting FPGAs

Over the past few years, a distinct segmentation of FPGA vendors has occurred depending on the device capabilities: distinct high-end and low-end product lines. If the application requires intensive computation, the FPGAs preferred are those with complex DSP blocks, high-speed inputs/outputs (I/Os) and configurable memories. They are seen serving consumers looking at ultra-high-definition televisions, 3D televisions and medical applications, such as 265-channel ultrasound. Examples of FPGAs delivering these requirements are Virtex, Stratix, ProASIC3 and Speedster 22i.

For low-cost applications, where a trade-off has to be made for perfor-



Lattice ECP3 low-power FPGA

mance, FPGA families, such as Spartan, Cyclone ECP5 and Fusion, are used. Typical applications include emerging communication products, such as heterogeneous networks (HetNet), gigabit passive optical network optical line terminals (GPON OLTs) and small form-factor pluggable (SFP) modules.

Choice of process node

Process nodes are available in 45nm, 40nm, 32nm, 28nm, 22nm, 20nm and 16nm values. The 10nm nodes are already under development.

From the user's perspective, a process node is not everything and the user does not have much control on it. Provided a device meets the design specifications, power and cost budgets and availability, designers often tend to overlook the process node.

However, because of the reduction in process node values, the newer

dyes are much denser and hence, for the same dye size of an older process node, more logic can fit in on the lower process node devices. Thus, lower process node certainly provides higher performance if the chip is fully utilised.

Up to 130nm, as the process node is lowered, the power and cost are also reduced. Beyond that, any decrease in size of the transistor will cause leakage to increase. This results in higher power consumption. Due to this reason, relatively higher process nodes are more preferred for low-power consumption applications.

When we consider the higher power consumption due to leakage, the higher clock rates that new nodes provided did not seem as much of a value proposition for every designer. But now with the introduction of FinFET technology, the semiconductor fabs and so FPGA firms are able to solve this problem to an extent.

New and exciting fields of application

While the bread-butter market for FPGAs has always been wired/wireless communications, FPGA products have forced their way into a variety of markets, such as industrial, medical, automotive and consumer electronics.

Nowadays we come across FPGAs in almost every field. They are used to power various devices ranging from miniature, wearable and handheld devices to smartphones, auto infotainments and medical equipment. They are also used in complex rugged systems, such as military and defence equipment and satellite systems. Additionally, FPGAs are being used in small glue logic implementations to high-speed computing applications.

High-performance video applications

FPGAs are playing a huge role in video applications. With large on-chip memories, DSP blocks and high-performance logic fabric, the FPGAs are being used to implement entire image pipe with multiple video IPs, implementing the compression algorithms and high-speed

Major contributors to this report

- **Neeraj Varma**, director of sales, Xilinx India
- **Rahul Prakash Joshi**, technical manager, elnfochips
- **Rajesh Chakkingal**, lead architect, Mistral Solutions
- **Rajiv Nair**, regional sales manager - India, Lattice Semiconductor
- **Sheik Abdullah**, project leader - FPGA design, iWave Systems Technologies Pvt Ltd
- **Shubha S.**, senior tech lead, Dexcel Electronics Designs

interfaces for videos.

The video applications require frame or line buffering at multiple points in the path. With HD video, the video resolution and the frame rates have increased. With high-performance devices, DDR memory interfaces, readily-available DDR controller IPs and high-speed I/Os, the FPGAs are meeting the required high bandwidth for video applications.



Altera Stratix HC IV FPGA

FPGAs in other fields

The wearable electronics market has been touted as the next big wave in consumer electronics industry. Low-density and ultra-low-density FPGAs are well-positioned to ride this wave. In addition, we see success in complementary markets such as human-machine interface (HMI), office printers and imaging sensor characterisation.

Surveillance cameras, consumer electronics and medical imaging are the domains where FPGAs are being used extensively. They are also used in some domains where there are stringent security standards for the design to be implemented, such as avionics. The cockpit display systems, data acquisition and control mechanisms are all powered by FPGAs. Defence is yet another sector where the FPGA usage is high for a variety of applications.

FPGA as a hardware accelerator

High-performance computing and hardware accelerators are some of the new areas where FPGAs are getting introduced. For example, in stock exchange, where decisions are to be de-

rived based on fast computations, possibilities of platforms having multiple FPGAs with design comprising register-transfer level (RTL), soft-core processor and high-speed serial interfaces with host system are being explored.

Computation-intensive algorithms, such as encryption, decryption and compression, are getting off-loaded from software applications running on processors and are getting implemented as hardware accelerator IPs in FPGA.

Why an FPGA

For products that are not sold in huge volumes, it makes sense to provide the solution on FPGA. This helps you reduce the number of overall devices on the board, making the board compact and design more secure. This makes reverse engineering difficult and brings up the design much quicker than earlier days, to reach the market faster.

The industry is all about being smarter these days. Now what we need more is not just higher bandwidth but also the ability to transmit more intelligent data in the given bandwidth. FPGAs have come a long way to help the designers meet such requirements. ●

The author is a technical correspondent at EFY