

Achieving Breakthrough Performance with Virtex-4, the World's Fastest FPGA

Webcast Date February 01, 2005

Presented by:

Peter Alfke, Director, Applications Engineering

Hitesh Patel, Sr. Product Marketing Manager, Design Software Division

Slide 1

Hi, I am Peter Alfke, applications engineering (not marketing) at Xilinx. This is the first in a series of talks on our Virtex-4 family. We are very enthusiastic about this family that has been shipping for 7 months now.

I will try to infect you with this enthusiasm.

Yes, we are #1 in 90 nm, as we will explain in detail.

Slide 2

Before starting the design of any new FPGA family, we always ask our customers about their problems and challenges, and we got many detailed suggestions.

We understand your shorter design times, the rapid obsolescence and the unrelenting cost pressure. We face the same problems.

Your designs get more complex and demand more performance, so do ours.

But, different from us, you face PC-board problems with reflections, ringing and crosstalk, and you struggle with power consumption and heat removal. And you want better tools, more intuitive, more powerful, faster and less buggy..

And you expect friendly and knowledgeable support.

We listened to you and we are trying to make your life easier.

Today we will talk about Virtex-4 **performance**.

Slide 3

90 nm. Why this obsession about 90 nm? You can't even see it, it's 5 times smaller than the wavelength of light.

The importance of high-volume 90 nm production is that it indicates the continuation of Gordon Moore's so-called Law. In 1965 he predicted an ever-increasing transistor count per chip, doubling every 1.5 to two years. This exponential growth has been the locomotive of our industry. In 40 years we have come from a dozen transistors to almost a billion, and the speed has been doubled every 5 years. And the reality of 90 nm technology now demonstrates that we can continue on that curve...

Compared to the previous 130 nm technology, 90 nm offers half the area, and thus almost half the cost per function. Smaller capacitance gives us faster circuits

and, together with lower supply voltages, significantly lower dynamic power consumption (per function and MHz)

While our competition rested at 130 nm, bemoaning the cost and the difficulty of further progress, Xilinx went ahead and geared up for 90 nm technology as early as 2001. As a result we have today shipped a hundred times more 90 nm devices than the rest of the PLD industry put together.

Not just more, but a hundred times more!

And shipping in volume to satisfied customer is more important than making inflated claims and producing doctored benchmarks.

Xilinx is #1 in 90 nm!

Slide 4

When we started designing Virtex-4, the goal was 500 MHz guaranteed operation of almost all sub-circuits.

Here are the four ingredients that got us to that goal:

90 nm is the foundation, we couldn't do without it.

But it alone does not assure superiority, since the competition, in due course, will have access to similar technology. We're all smart, and we all use the same silicon...We needed more than just technology.

Like clever circuit design, layout and testing methods. That is hard work, an ongoing evolutionary process of fine-tuning density and performance.

But the biggest performance boost comes from innovative architecture and novel functionality. That's where we are leaving the competition in the dust, and that's what we will describe in the next 30 minutes.

There is also continuous improvement in our design tools to make them easier to use, compile faster and achieve higher performance.

Slide 5

Here is the circle of architecture features, and you will see this picture again and again. We will work our way clockwise around the circle...

Slide 6

These are the performance highlights that we will describe in the next 30 minutes. And at the end we will use this same drawing to point out our performance advantages over Stratix-II. No, we're not defensive, we have something to be proud of.

Slide 7

The biggest performance boost came from improved functionality:

better clocking,
faster and wider parallel I/O,
multi-gigabit serial I/O,
on-chip memory with built in error correction and FIFOs,
versatile multiplier/accumulator circuits for DSP and other applications (like very fast counters).
We will talk about embedded microprocessors, both soft and hard,
and finally about combinatorial logic, flip-flops and routing in the fabric.

Slide 8 / 9

Let's start with clocking.

Every design, fast or slow, needs Global Clocks that can reach every one of the up to 200,000 flip-flops on a chip, with fast rise and fall times and minimal skew. In Virtex-4, the Global Clocks travel across the chip as differential signals which reduces duty cycle distortion.

We are all in favor of synchronous single-clock operation, but we do realize that real-world designs use multiple clocks. So we give you 32 Global Clocks, each capable of driving anywhere on the chip.

The relatively slow general-purpose interconnect lines are not a good choice for additional clocking. We give you localized Regional clocks, and also very fast I/O clocks for chip interfaces that go up to 1 Gbps DDR.

Slide 10

Digital Clock Management circuits use a clever feedback scheme to insert the appropriate amount of clock delay, to completely compensate the chip-internal clock delay, or even the PC-board delay.

This means that large chips can achieve the same performance as smaller ones.

The DCM can also be used for frequency synthesis, simultaneously multiplying and dividing the incoming clock frequency by any integer up to 32.

The DCM outputs can also be phase shifted with very fine granularity, one 256th of the clock period, or 35 picoseconds. Dynamic phase alignment can be used to adjust the clock with respect to data, but we will also mention a different mechanisms that work inside the I/O, independent of the global clocks.

"Picosecond" has become an important new word in our vocabulary, and in yours.

Slide 11 / 12

Let's now talk about inputs and outputs

Traditionally, even large systems were based on one central clock, distributed across the pc-board to its many destinations. The clock was supposed to arrive everywhere at the same time, so that all ICs would live in the same time zone and could easily communicate with each other.

This pretty picture was destroyed by today's faster clock rates, where pc-board delays can no longer be ignored.

The solution is source-synchronous design, where a clock is sent together with every data wire or bus. Clock delay and data delay are then (almost) identical, and pc-board routing delays have become irrelevant.

At the receiving end, clock and data must be aligned appropriately. Source-synchronous clocking adds clock lines, and it adds synchronization complexity.

But we'll see how Virtex-4 helps you solve these problems.

Slide 13

Each Virtex-4 input pin (yes: each and every pin) has its individual programmable precision delay line called IDELAY, and also has a serial-to-parallel converter with its own counter, register, and clock divider. As output, each pin also has its own parallel to serial converter.

The I/O data rate is anything up to 1 Gbps.

Slide 14

This is a simplified block diagram of the IDELAY circuit, which consists of a 64-tap delay line. The total delay is servo-controlled by a 200 MHz clock, giving each tap a fixed delay difference of exactly 78.125 picoseconds, unaffected by temperature, voltage or chip parameter changes.

Slide 15

Here you see how IDELAY can be used either to move the clock or strobe into the center of the data eye, or to move data optimally with respect to the fixed clock.

Slide 16

Here is the input serial-to-parallel converter that can adapt fast inputs to slower structures on the chip.

Slide 17

And here is the opposite, a parallel-to-series converter that changes slow parallel data into a faster serial output bitstream.

Slide 18

Virtex I/O can interface to many memory types. Since all pins have identical capabilities, up to 432 data lines can be used at up to 600 Mbps using a double data rate protocol. This large number of identical I/O pins sets Virtex-4 apart from the competition.

Slide 19

Xilinx offers a memory evaluation board ML461 for sale or loan to interested customers. It can be used as demonstrator testbed or as a model for many different memory interfaces.

Slide 20

There is also ML450, the evaluation board that demonstrates many communications interfaces.

These are just two of the many Virtex-4 based evaluation boards. Explore the list, and take a good look at the extremely low-cost, 500 dollar versatile ML401 board.

Slide 21

Now let's switch to the ultra-fast multi-gigabit transceivers, faster and more versatile than the Rocket I/O transceivers in our previous Virtex families. Output bit rate is now any value between 600 Mbps and more than 10 Gbps. There is the traditional complement of functions: parallel interface to the fabric, FIFOs, CRC, 8B/10B coding, and selectable pre-emphasis on the transmitter output. New features are 64B/66B coding and a sophisticated equalization capability on the serial input. This makes the MGT compatible with different trace lengths on pc-boards and backplanes.

Slide 22 / 23

So much about I/Os.

Let's move to features that enhance the performance in digital signal processing applications.

It may come as a surprise to some that FPGAs can easily outperform even the newest and fastest dedicated digital signal processor chips that run at outlandish clock rates.

FPGAs derive their superior DSP performance from massive parallelism, where data is being manipulated in up to 512 cascadable multiplier/accumulators simultaneously. A couple of numbercrunchers can never compete with that kind of performance.

Slide 24

Such circuits are not new, our competitor designed them as shown here. Nice circuit, but when you want to cascade more than four of them, you need to dive into the fabric and implement wide adder structures. That reduces performance drastically. Simple math tells us that one extra nanosecond reduces a potential 500 MHz speed to 333 MHz.

Slide 25

The Virtex-4 DSP slice, (that's what we call our MAC) has all the required hooks and pipelines for unlimited expansion from the bottom to the top of the chip, at full speed. We can maintain 500 MHz operation through 32 MACs in the smallest chip, and 96 MACs in the biggest chip. No external logic at all.

Slide 26

Here is a more detailed look at the interaction between multipliers, adders, and registers.

Slide 27

And these are the innards of the DSP slice, with its wide inputs and outputs, note especially the two sets of wide expansion inputs and outputs. They are the secret behind the constant 500 MHz performance.

Slide 28

DSP slices are not just for digital signal processing, you can also use four slices to form a 6-to-1 multiplexer, 36 bits wide, and capable of 500 MHz data rate.

Slide 29

And there are more ideas. The accumulator can obviously be used as a 32-bit synchronously loadable counter, or as a non-loadable 48-bit synchronous counter. Both can run at 500 MHz, thanks to the fast carry structure that is required for DSP applications.

It is well known that any multiplier can also be used as a barrel shifter.

But it might be surprising that the accumulator can also be used to build a low-jitter direct digital synthesis phase accumulator. The 500 MHz clock rate would normally cause 2 ns of jitter, but a BlockRAM and 16 IDELAY circuits can create a virtual 8 GHz clock rate, and keep the jitter below 200 picoseconds.

The smallest Virtex-4 chip has 36 such DSP slices, the largest has 512. It is safe to assume that some will be left over for such non-conventional applications.

Slide 30 / 31

Let's now look at BlockRAMs.

The size and basic structure is retained from previous generations, but the speed has been increased to 500 MHz clock rate, using the pipeline register incorporated in the data output. The two ports still have individual width control, and in write mode the user can choose between automatically reading the previously stored data or the new data. Two neighboring BlockRAMs can now be combined to form a 32K x 1 RAM without loss of speed, or they can be combined to form a 512 deep 64-wide RAM with automatic Hamming error correction, without using any extra logic.

Each BlockRAM also contains its own FIFO controller, as we will see in a few minutes.

Slide 32

First we show how the synchronous nature of the BlockRAM makes it very efficient to build state machines running at up to 350 MHz. Here is a hint at various solutions, from 64 states to 256 states

Slide 33

with up to 45 outputs and up to 4 control inputs.

Slide 34

For more flexibility, add one or two CLBs.

Slide 35

Now we look at First-in-First-Out memories, which are very popular for moving data from one clock domain to another. Most FIFO applications use different clock for writing and reading. That's why they use a FIFO.

Conceptually FIFOs are very simple, using just a dual-port RAM, two counters and some comparators. The devil is in the details. When write and read clock are unrelated (as they often are) the FULL and EMPTY decoder must compare

counter values in both clock domains, which requires Gray-coded counters. That in turn makes it difficult to detect partial full-ness, the so-called dipstick. All this consumes additional circuitry, but there is a bigger problem: synchronizing FULL and EMPTY to their respective relevant clock domains means crossing clock boundaries and living with metastability issues. This become really tricky at multi-hundred megahertz clock rates.

Slide 36

Luckily the original designer of the world's first FIFO IC (the 1971 Fairchild 3341) now works at Xilinx and helped us design the 500 MHz dual-clock FIFO in Virtex-4. So we can claim 34 years of experience in FIFO design.

We also developed a method for thorough self-testing the FIFO, which cannot be done with conventional timing analysis. We used a 500 MHz read clock and let the FIFO go empty 200 million times per second. We then monitored the asynchronous handshake. There had not been a single failure when we stopped the test after ten-to-the power of 14 cycles. Users can trust this FIFO controller, even at 500 MHz asynchronous operation. It is a complete solution, available in each and every Virtex-4 BlockRAM.

Slide 37 / 38

Microprocessors embedded in FPGAs have become popular, since many of the larger (and even of the smaller) FPGA-based designs can benefit from the versatility and flexibility of an embedded microprocessor. Most FPGA-based designs contain a mixture of speed requirements. There is fast logic in the fabric, in counters, DSP slices, BlockRAM and the I/O. Let's call that **nanosecond** logic. But there usually also are slower and more complex function, often in control, that must be completed in milliseconds or perhaps microseconds. Such slower functions benefit from the versatility, flexibility and ease of design of an embedded microprocessor.

Slide 39

There are two ways to embed a microprocessor in an FPGA, soft or hard. Soft microprocessors are implemented in the fabric, using BlockRAM to store data and instructions. PicoBlaze and MicroBlaze are the two soft microprocessors offered by Xilinx. They use from 100 to 500 slices of logic plus a few BlockRAMs. Their performance can reach 120 DMIPS, but that is no match for the hard implementation of a PowerPC, as available in Virtex-4 FX.

Slide 40

These one or two PPCs per chip run at max 450 MHz, achieving 700 DMIPS each. That is more than 3 times the performance of the best soft microprocessor cores. Moreover, the PPC has two 16 kilobyte caches, which in simple applications, can be used to store all the data and instructions. Larger applications use BlockRAMs or even external memory.

Compared to previous Virtex families, the new PPC in Virtex-4 is faster and adds an interface that supports high performance co-processing.

Slide 41 / 42

The last item in our circular tour is the fabric, represented by the CLB structure and its routing. Each CLB consists of four slices, and each slice has two variable look-up tables plus carry and other expansion logic, plus two flip-flops. None of the LUT inputs are shared, which makes it easy for the software to place unrelated functions into one CLB, for highest packing efficiency. The fabric and interconnect structure can support the 500 MHz clock rate of the blocks that we described before.

Not every design implementation will be able to run at 500 MHz system speed. Performance can easily suffer from extra delays. Just 1 ns reduces 500 MHz to 333 MHz. But the CLB and interconnect structure, using shallow combinatorial logic and pipelining when appropriate, can sustain the 500 MHz of the blocks we mentioned.

Slide 43

Now we come to the last subject, benchmarks.

Who needs benchmarks?

Benchmarks are there to evaluate and compare competing solutions. Ideally, benchmarks should guide and help the user to save time and effort.

Before you buy a washing machine or a car you probably consult Consumer Reports or Road and Track, to get expert and unbiased competitive information. If there were such independent institutions or publications for FPGAs, what would they do?

They would obviously compare equivalent speed grades, would use the most appropriate software options and timing constraint methods. They would compare meaningful designs and would utilize all available resources, including novel embedded functions. Such benchmarks would then be meaningful, trustworthy and fair.

No such luck with FPGA benchmarks. They are, as you know, generated and published by the competing manufacturers, and:

truth is the first casualty of the benchmark wars.

Think about it: No marketer would ever admit that his product is inferior to the competition !

You can be absolutely certain that every published competitive benchmark will declare itself the winner. That's unavoidable.

Caveat emptor!

You, the user, must separate the chaff from the wheat. You must detect for example, that something is fishy when one manufacturer declares himself the winner with a 39% performance superiority in structures that are similar, and using similar technology. If it sounds too good to be true, it probably is...

Slide 44

Nevertheless, Xilinx also does benchmarks, mainly for internal purposes, to inform our designers and our FAEs. And here is what a third party, Symplicity, says about the Xilinx benchmarking methodology,.....

Slide 45

.....and here are our benchmarks of 43 different designs. Relative performance has a wide spread, from 20% in Stratix-II favor to 40% in Virtex-4 favor, with an average 10 % advantage for Virtex-4.

10% is really about all one can possibly achieve when implementing similar structures and using similar technology. Both companies use the same silicon, we both have competent designers, and we both have excellent fabs...

Slide 46

But these primitive benchmarks are not the real measure of performance. Those archived old designs ignore the new capabilities and leave out the new and the really important features, that we talked about for the last half hour. When any designer fights for the highest performance, he or she will obviously take advantage of the best available stuff.

Slide 47

So let's forget all this benchmarking and just compare the performance-oriented parameters that have been published for Virtex-4 and Stratix-II. These numbers cannot be fudged or debated.

Virtx-4 wins by a factor 1.6 for general I/O bandwidth, and a factor 3 for memory bandwidth. The MHz speeds are not very different, since they are dictated by the requirements of existing standards, but Virtex-4 offers a much larger number device pins with greater flexibility.

Regarding multi-gigabit transceivers, there is no contest, 10 Gbps to nothing.

Virtex-4 has a 20% advantage in RAM clock rate and a 40% advantage with the expandable multiplier accumulators.

The Virtex-4 PowerPC is at least 3 times faster than any NIOS-II implementation, and Virtex-4 has a 10% speed advantage in the fabric, when fair and appropriate methodologies are used.

Slide 48

To summarize:

I hope I made it clear that

Xilinx **was, and is** first in the use of 90 nm technology, in design, in manufacturing and in sales.

Xilinx has shipped a hundred times more 90 nm circuits than all our competitors combined.

Virtex-4 is the world's fastest FPGA, as measured by benchmarks, and as demonstrated by its many innovative functions and new capabilities.

Slide 49

Appropriately, Electronic Products magazine voted Virtex-4 the Product of the year.

Slide 50

Please join us two weeks from now for the next seminar in this series, where we will describe the, perhaps surprising, Virtex-4 reduction in supply current and power consumption.

Thank you for listening, I hope it has been of some value to you.