



## Not Bad Die

Xilinx EasyPath Explained

by Kevin Morris, FPGA Journal

We always thought we knew how it would go down.

Under cover of darkness, our black-clad insertion team would rappel down the walls of the super-secret Xilinx fortress in the desert. With the kind of precision timing and teamwork found only in movies and editorial feature introductions, we'd scan the perimeter and locate the vulnerable point. A diamond-tipped drill bit driven by a silent motor would bore a hole just large enough for our fiber-optic viewing tool, and the telling video would be immediately beamed back to FPGA Journal headquarters. At the same time, however, we'd inadvertently trigger the alarm system, and our squad would fragment, running for cover. In a panic, our leader would end up like Charlton Heston, running for his life through the streets shouting, "EasyPath is Bad Die! EasyPath is Bad Die!"

It didn't work out that way.

Instead, we went in the front door at Xilinx wearing our regular, security-issued name badges (with snappy photos of ourselves in case anyone we met needed to see what we looked like). We were escorted to a conference room where a group of Xilinx personnel calmly explained EasyPath to us – no big secrets, no dramatic reveal -- in truth, FAR more boring than the previous paragraph.

We'd always heard stories and made assumptions about EasyPath, and it turns out most of them were completely wrong.

First Wrong Assumption: EasyPath is just a press-release/marketing program

Truth: Nope. EasyPath hasn't gone away, and Xilinx actually earns money on it.

Explanation:

EasyPath was announced somewhat hurriedly way back at the time when Altera first announced their HardCopy FPGA-to-ASIC conversion program. In case you don't remember, in those days, Xilinx and Altera would play press-release ping-

pong with a cadence that resembled a third-grade playground argument.

Press release 1 – Monday, Company A: “We have the biggest FPGAs!”

Press release 2 – Tuesday, Company B: “We have the biggest FPGAs IN PRODUCTION!”

Press release 3 – Wednesday, Company A: “Oh yeah? Well, we have the biggest FPGAs with SerDes I/O!”

Press release 4 – Thursday, Company B: “No you don’t!”

Press release 5 – Friday, Company A: “Yes we do!”

Press release 6 - “Nuh-uh”

Press release 7 - “uh-huh”...

Designers around the world were both informed and amused by these press releases. BusinessWire and PR Newswire made a fortune. Technology editors nodded off to sleep.

A few years ago, Altera announced HardCopy as a cost-reduction strategy for their high-end FPGAs. If you designed a system with one of their FPGAs and wanted a low-cost version of that design in an ASIC for volume production, they could (and still can) whip out an ASIC in short order with the same functionality, lower power, faster logic, and much lower unit cost.

Right after that, predictable as pie, here came the Xilinx announcement of EasyPath - “Hey, look, we have a cost reduction strategy too.” OK, that isn’t really what it said, but that’s what we read with our specially-coated editorial press-release reading glasses. We figured Xilinx had just cooked up something to put out in a press release so that when customers came to them and said, “but, Altera has a cost reduction program...,” they could calmly reply, “yep, we’ve got one too...,” and the customer would go on to purchase FPGAs, and they’d probably never get to the stage of needing to actually cost-reduce anything. It would make the cost-reduction issue disappear from the normal FPGA-selling landscape.

It turns out that Xilinx actually put engineering effort into the EasyPath program, coming up with a way to test FPGAs specifically for a customer’s design. By dispensing with the usual full-blown testing of every single LUT, multiplier, and I/O, considerable cost could be cut. Also, by just testing what was used by the customer’s design, net yield was probably higher on EasyPath wafers than on normal FPGA wafers.

Second Wrong Assumption: No customers would actually use EasyPath; they just wanted it as a check-in-the-box competitive security blanket.

Truth: A large number of EasyPath designs are currently in production, and more are coming on every quarter. Xilinx says EasyPath revenues have more than doubled.

Explanation: We assumed that people would go in with the idea of cost-reducing “later on,” when their design was stable and when their device went into volume production. We believed that “later on” would never come, as the design would always be changing, and many products would never reach volume production. It turns out that many do – and they want the significant (30%-70%) unit cost reduction they can get from EasyPath. Also, since the EasyPath program is exactly the same FPGA device and package they were already using (not an ASIC replacement), many design teams gained an advantage from not needing to re-qualify the EasyPath version of their designs.

Third Wrong Assumption: EasyPath is bad die.

Truth: Xilinx pinky-swears that EasyPath gets its own wafers, and that those wafers go through the separate, specific EasyPath testing rather than the normal FPGA testing.

Explanation: It made sense at the time. We believed that Xilinx would test dies to see if they worked for general FPGA use – if so, they’d be binned as regular FPGAs. If they failed (with certain types of failures), they’d be re-tested as EasyPath candidates. We developed all sorts of elaborate theories based on this bit of misinformation. We were visualizing the yields going up for regular FPGAs and pushing EasyPath parts out...

Apparently we didn’t know enough about IC testing.

Wrapping it up:

EasyPath is a completely different cost reduction strategy from Altera’s HardCopy. While you should make your own comparison list for your own application, ours will help you get started. The pros and cons as we see them are:

EasyPath parts can be produced faster and more reliably from the time you order them – because they’re really just the same FPGAs you were already using. Even though HardCopy has an excellent track record for delivery, it is an ASIC spin.

EasyPath parts are exactly the same FPGAs you were already using, so there

should be no re-validation of your design required. In some cases, replacing an FPGA with even a pin-compatible, exact-functionality-match ASIC like HardCopy might require some re-qualification.

EasyPath has every function of the corresponding FPGA except arbitrary re-programmability. It should even be bug-compatible, exhibiting the same quirks as the original FPGA. (Because, did we mention, it's exactly the same FPGA? We did? OK, never mind.) HardCopy has historically not supported every single feature available in the company's FPGAs (SerDes I/O comes to mind.)

EasyPath should have shorter turnaround time (just setting up the test) and lower NRE charges than HardCopy. (Of course, we're not negotiating for you, so your mileage may vary.)

HardCopy is an ASIC, is metal-programmed, and thus doesn't require configuration logic. You can save the BOM cost of that circuitry when you swap the FPGA for the HardCopy device.

HardCopy should have noticeably better performance than the corresponding FPGA. Convert your FPGA to a HardCopy ASIC and get an almost-instant performance boost.

HardCopy should have significantly lower power consumption than the corresponding FPGA. EasyPath will have exactly the same power consumption as the FPGA version because... oh, you know.

HardCopy will most likely have better radiation tolerance than an FPGA. This is just our theory, and we have to point out as a caveat that neither strategy is qualified as radiation-tolerant as far as we know. SRAM FPGAs require special radiation mitigation strategies, and the registers and memories of HardCopy devices would still be vulnerable unless specific strategies like TMR were employed in the design. HardCopy should be more stable in the face of the rare but occasional low-altitude single-event-upset, however, because of the metal interconnect instead of SRAM-based configuration logic.

Well, there you have it. Apparently our crack team of cost-reduction researchers can take off their black overalls and ditch the night vision goggles. Dang! That was gonna be fun.

*by Kevin Morris, FPGA and Structured ASIC Journal*

*May 27, 2008*