

IQ Magazine Interview with NXP's Geoff Lees:

NXP Launches the LPC1100

By Glenn ImObersteg, IQ Magazine

NXP

NXP announced the release of the LPC1100 on November 16. The LPC1100,

based on the ARM Cortex-M0 processor core, will be available from distributors in mid-December. Billed as the "lowest priced 32-bit mcu solution in the market", the LPC1100 offers higher value and ease-of-use than existing 8/16-bit microcontrollers. The following interview is between Geoff Lees, the Vice President and General Manager, Microcontroller Product Line, NXP Semiconductors, and Glenn ImObersteg, the publisher and editor of IQ.

IQ: *I've heard from industry sources that you worked closely with ARM in the creation of this new Cortex processor core. How were you involved?*

Geoff: Yes, we were involved closely. The first requirements were that we were looking for a Cortex-based core with a smaller dynamic power and static power profile than the current M3 or the existing ARM7, and we talked with ARM for a while about a number of ways to implement that idea. The end result was the proposal to go ahead with the project code-named "Swift", which later became the Cortex-M0 processor core.

IQ: *When was the project initiated?*

Geoff: It started a couple of years back, and really reached fruition with the announcement of the Cortex-M0 processor by ARM at the 2009 Embedded World in Nuremberg, last February. At that show we were introduced as a lead partner in the development of a series of microcontrollers based on Cortex-M0.

IQ: *Why the M0? What was the rationale behind having one more Cortex core?*

Geoff: Well really it's nothing more or less than a return to the original RISC concept. It's all about the basic Thumb instruction set, which is at the heart of all ARM processors since the ARM7TDMI, the ARM926EJ, and the ARM11 family, as well as all the Cortex series. This basic RISC instruction set is what we wanted to bring to market for the first time in stand-alone form. For our purposes, we view the pure 35 instructions of the original Thumb instruction set as absolutely perfectly suited to compete in the 8- and 16-bit markets.



Geoff Lees, the Vice President and General Manager, Microcontroller Product Line, NXP Semiconductors.

IQ: *The 8-bit and 16-bit market has historically belonged to microcontrollers like the 8051 and the PIC, and now you're targeting those applications for the Cortex-M0?*

Geoff: Yes, and the idea is, for 32-bit architectures to appeal to the 8-bit user, it's not just enough to match the price. Over time we've seen customers view 32-bit as more performance than they need, possibly taking more power than they need, and they're pretty convinced that

they are going to have to use more code to achieve what they want. And what Cortex-M0 has the potential for, is to actually remove all of those obstacles and show that Thumb can be more efficient than the architecture they're working with.

"The idea is, for 32-bit architectures to appeal to the 8-bit user, it's not just enough to match the price."

Geoff: For the launch of our new LPC1100 Cortex-M0 family, we've utilized the Core-mark benchmark, which is an open source benchmark suite available from EEMBC, to demonstrate that code size from the LPC1100 can be typically half the code size of an existing 8-bit solution.

Even peripheral code and small software drivers won't be any larger in Cortex-M0 than in a typical 8-bit architecture. There's even likelihood that it will be reduced in size.

However, any significant application code is going to be smaller by as much as 30-50 percent. We believe that it's a true value proposition for current 8-bit users, who pretty much never saw any benchmarks in the past between Thumb and 8-bit. That's because Thumb was usually benchmarked only against other ARM instruction sets, x86 or MIPS, or some other 32-bit type architectures--but it was never really compared with 8-bit architectures.

IQ: *So I imagine it's really quite a surprise to see that an ARM derived instruction set from over ten years ago actually has higher efficiency and smaller code size than today's 8-bit and 16-bit architectures.*

Geoff: Exactly, and maybe we shouldn't be surprised at this, because 8-bit cores have their origins in the beginnings of the Semiconductor industry. They were all designed around a time when engineers were handcrafting designs, even before the arrival of EDA tools. Silicon real estate was in short supply, and as a consequence, cores had limited register sets, typically some general-purpose accumulators, and one or two dedicated registers for math-- if you were lucky.

The address range of 8-bit architectures was typically under 64K, and some early cores had even less. Over the years, various 8-bit cores have been expanded through paging or through extended instruction schemes, all of which have led to inefficient increase in core area.

If you combine that, with the fact that 8-bit cores have not been supported by the latest design tool generations, latest process nodes, and advanced EDA tools for place & route and synthesis.

IQ: *You mentioned the EEMBC benchmark scores for the M0. How do they compare with traditional 8-bit architectures?*

Geoff: Well the advantage of the Cortex-M0 is that if you look at benchmark results against nearly all of the commonly available 8-bit architectures, it's actually significantly better than all of them. There's no one 8-bit architecture that is a particular target: I would say all 8-bit architectures are going to be put under pressure by our LPC1100 family--and also other Cortex-M0 products from other vendors in the future.

IQ: *So how does the pricing of the Cortex-M0 compare at 10,000 units against other devices, for example?*

Geoff: In general, 8-bit products with minimum feature sets in the 8K to 32K Flash range, are priced at the \$.70-\$1.00 range, but actually some of the 8-bit and 16-bit architectures are not that cheap. Some of them are in the \$1.00-\$2.00 range. A lot of it depends on the particular process node that the vendor is using. The more advanced the process, typically the more aggressive the pricing.

IQ: *To continue with pricing, I understand you're pricing the LPC1100 at \$0.65 each in the 10,000-unit range?*

Geoff: Yes, we are. That reflects one of the most efficient designs that we have ever been able to implement. In fact, when you look at all the designs we've done in 8-bit over the years, this product range is actually able to come in using less silicon than virtually all of the 8-bit designs we've ever done. And that's really a tribute to the "Swift" design team, (the Cortex-M0 design team at ARM), who we feel have implemented one of the most striking, small size, low-power core designs in the industry.

IQ: *How does the M0 compare in power efficiency to other Cortex cores like the Cortex-M3 or Cortex-R4?*

Geoff: Well compared to the Cortex-M3, the M0 core power is under 50%. So scaling up, for example, to the Cortex-R, its probably in the range of 5x lower.

"Instead of just competing in the main 32-bit embedded space, now ARM is competitive in the total embedded space."

IQ: *That's impressive, and the beauty is that it allows you to start with one single scalable architecture.*

Geoff: Exactly, and it means that instead of just competing in the main 32-bit embedded space, now ARM is competitive in the total embedded space. And that really increases the market potential from 4-5 billion dollars up to 15 billion dollars for the entire embedded market.

IQ: *We've compared the Cortex-M0 to 8- and 16-bit cores, but how does it compare to 32-bit cores like PowerPC or Coldfire?*

Geoff: It's hard to get specific data, but we know the approximate core sizes. For example, in the Coldfire range, the Coldfire V1 is somewhat competitive with the Cortex-M3. So the Cortex-M0 then is in the range of 1/3 of that size. But the V1 is the entry level of the Coldfire series. The V2 and the V4 are larger than that so really, from the existing 32-bit cores, there are no cores in the same range and scale as the Cortex-M0. Realistically, the wide scale acceptance of the Thumb instruction set and its' extensive tool support over the last ten years, puts M0 in a class of its own.

IQ: *What kind of applications are you going to go after? Are you going beyond the mobile and other markets that typically have been the realm of the ARM Core?*

Geoff: We're going after the whole 8-bit and 16-bit market. The 8-bit market is less verticalized than the 32-bit. You see a lot of SoC designs, and a lot of custom designs in the 32-bit market. The 8-bit space is much more about general-purpose products that have been optimized and expanded over time to suit particular customer requirements. But, they are all pretty much general-purpose products. So, the LPC1100 is already at what we feel is the entry level point in the mid-range 8-bit space, the 32 to 48-pin package range. Today we're not addressing the low-pin markets, ie, the 6-pin and 8-pin architectures-- those are typically very low-end single-function control applications. In the mid-range 8-bit space, Cortex-M0 is extremely competitive.

IQ: So what are your plans for features such as CAN, Dataconverters (DAC), temperature sensors, etc.?

Geoff: We've announced in our press release that the next versions of the LPC1100 family are already in development. In fact, we're testing the CAN version right now, and we're including industry-standard CAN software on-chip ROM, so a customer is able to merge those drivers with their own application code. So the whole of the on-chip Flash memory is available to the customer for their application.

IQ: And when will that be available?

Geoff: Well that's actually in final validation now; we're planning on launching that in early 2010. The next derivatives are already in development, and those include 12-bit ADC and DAC, a temperature sensor, much more higher resolution timer system for motor control, and PWM systems.



LPCXpresso Development Kit.

IQ: In one of the more pro-active steps I've seen in the launch of a new product, NXP has announced the simultaneous release of new low-cost development tools to support the M0 and the M3, like the mbed and LPCXpresso. What can you tell us about them?

Geoff: LPCXpresso is really the culmination of efforts on our part to bring the tool chain to a wider audience. You know, we really appreciate the ARM development tools ecosystem; it's one of the strengths of ARM. But as we enter the 8-bit market, customer expectation is a little different than in the existing 32-bit markets. Customer expectations have really been driven the last years by the proprietary architecture vendors who have developed tool chains. They have to do that on their own because they use proprietary architectures, and they optimize the tool chains themselves and provide that ecosystem from within that company. And so, customers are used to tools coming for free, they're used to tools being part of the support package that they get from those vendors. The current ARM ecosystem is very strong-- it offers high performance compilers and a lot of benefits in the 32-bit space.

But there's a void in the market for entry-level tools for the ARM environment, and LPCXpresso is really at the forefront of that. What we've developed is a fully integrated tool chain, where we provide the hardware and ecosystem. We're able to offer an eclipse debug environment, IDE, compiler, and a target evaluation system all in one small low-cost small form-factor board. Over time we will work with more tools vendors who have new innovative solutions and can use this platform to get to market quickly.

One of the beauties of the design of this tool is that it's really an open platform that can host any software solution. We use our latest ARM9 based high-speed USB controller to interface to the development environment, the PC and the web. The board itself does not require installed firmware, as all necessary development firmware is downloaded at the start of the session. This controller contains everything needed to connect to the internet via any platform, and becomes the debug vehicle. The target system can be either one of our latest Cortex-M3 LPC1340 USB controllers, or the LPC1100 Cortex-M0 series.

In addition, we've adopted the same form-factor as the recently launched, popular mbed rapid prototyping tool. This LPCXpresso board had been designed to be pin-compatible with the popular recently launched mbed rapid prototyping tool. After doing evaluation projects or proof of concept prototyping, the user can move on to a more formal development tool based on the LPCXpresso system. This really provides a continuum between the mbed world, for the new users, and fast adopters, through to the formal software verification and production stage.

IQ: LPCXpresso is a collaboration between NXP, Code Red and Embedded Artists, isn't it?

Geoff: Yes, those are our first partners. However, the architecture, schematics, board design are not proprietary, and are all available on-line at nxp.com/lpcxpresso.

IQ: What's impressive about this is it comes in a plain envelope, its under \$30, its cheaper than almost any electronics you can buy anywhere else, and the user manual is a slip of paper that just says "To get started, go to www.nxp.com/lpcxpresso".

Geoff: Yes, that's my idea of a user manual.

IQ: It's a great idea. I'm sure the users opening this for the first time will be impressed. Is there anything else you would like to discuss?

Geoff: I think that pretty much covers everything. We are looking forward to the challenge, if you will, the charge of Cortex-M0 into the embedded markets.