# Driving both technology and cost/performance

## Q&A with Chris Fanning and David Lee Rutledge,

*Lattice Semiconductor Corporation*

*Two senior execs at Lattice give us their views on the Lattice strategy, with an in-depth look at non-volatile FPGA technology, in this exclusive interview and short article.*

**PL: Lattice is making pushes in a lot of areas. Why such diversity in the offering, and how does it tie together from a strategy standpoint?**

**DLR:** When Lattice entered the FPGA market, we chose to focus on two key differentiators – Non-Volatility (NV) and Low-Cost – and set about developing technology and products in each area.

NV is the technology of choice for programmable logic, providing system-level benefits such as single-chip solutions, instant-on capabilities, and in-system programmability. We knew that we could leverage our NV expertise to deliver these benefits to FPGA users.

Lattice saw an opportunity to fill a void in the market by developing a Low-Cost, Feature-Rich FPGA and CPLD product line. Our strategy has been to develop products that work very well for most applications, allowing us to substantially reduce cost with little impact on performance.

So, while our products may seem diverse, they have been developed based on a consistent strategy to deliver value-added products, with NV and Low-Cost, Feature-Rich as our key differentiators.

**PL: What's the biggest technology trend driving your FPGA implementations today, and how do you see that affecting what you can do in the near future?**

**DLR:** Achieving higher levels of functional integration and performance while maintaining acceptable power consumption is a major challenge. Historically, the scaling of supply voltages and transistor feature sizes has provided great benefits in functional integration levels and performance, while lowering power consumption. In the future, these techniques will still provide substantially higher levels of functional density and performance, but at the expense of increased levels of static power consumption.

At 45 nm there will be a more direct trade-off between speed and power. For example, a high-density FPGA (~500 K LUT) built on 45 nm technology could easily consume over 10 W of standby power, with no clocks running, at the commercial +85 °C junction temperature limit.

Innovation in power management to optimize this trade-off is a high priority. Innovation is required concurrently across multiple disciplines, ranging from optimizing the basic process technology through the development of new "power-optimized" product architectures and also through the development of "power-aware" design tools.

Also, there is increased use of embedded SERDES channels as high-bandwidth chip-to-chip communication links. An FPGA with 40 channels of soon-to-be-available 10 G SERDES will have an incredible processing capacity of 400 Gbps. This trend will accelerate the development of radically new FPGA-based High-Performance Computing (HPC) platforms.

**PL: On the tools side, I see some additions of synthesis and simulation capability, and I'm guessing you're expanding further. What's the latest technology you are working on, and what's the impact?**

**CF:** Lattice has made a very significant investment in its design tool, ispLEVER. Two initiatives we see having great impact are physical synthesis and incremental design.

Physical synthesis should enable more optimal results in improved device performance and help accelerate the debug process. Incremental design is a collection of technologies that ultimately provides

more rapid turnaround time for achieving timing closure. Both these enable customers to meet timing requirements more easily and with fewer design iterations, even as FPGA devices become even larger and more complex.

Other Lattice-driven design tool advancements include:

- Power Calculator allows specifying parameters such as voltage, temperature, process variations, airflow, heat sink, resource utilization, activity, and frequency, and then calculates static (DC) and dynamic (AC) power consumption.
- Reveal uses a signal-centric model for embedded logic debug. Signals of interest are user-defined, and the tool adds instrumentation along with the proper connections to enable the required in-system analysis can then be performed. Users can specify complex, multi-event triggering sequences that make system-level design debug smoother and faster.

**PL: What about embedding operating systems on an FPGA core? How is this changing the way people design?**

**CF:** uClinux support expands our commitment to the open source model, and it consistently appears at the top of designer surveys we've seen as the preferred RTOS for embedded design. uClinux and the LatticeMico32 core allow designers to implement control systems in a design flow that builds on Lattice's open source, embedded solutions approach.

The adoption of embedded products has increased among FPGA designers in the last two to three years as the processors provided by FPGA vendors have dramatically improved in functionality, increasing designer productivity and lowering design risk. Embedded processors increasingly include a robust assortment of middleware such as DDR, DDR2, SDRAM memory controllers, Tri-Speed Ethernet Media Access Controller, and PCI 33 MHz Target, which automatically integrate into Lattice's Mico System Builder. Middleware has enabled designers to quickly and confidently configure a microprocessor in their design at very low cost, or no cost at all.

**PL: What should designers be doing differently now to get an advantage with their next FPGA-based design?**

**DLR:** We think in terms of two fundamental types of FPGA applications: *control-oriented* applications and *data path* applications. Control-oriented applications really have not evolved much over the years, implementing numerous small finite state machines and utilize many parallel I/Os for system monitoring and control. There are not too many new issues to deal with in this area.

Data path applications, however, have continued to grow and evolve. It is more important than ever for a system designer to consider how to best architect systems to effectively leverage the new high-speed SERDES-based capabilities of FPGAs. This increased data processing bandwidth will allow for radically new system-level architectures that can provide dramatically higher levels of cost/performance.

**CF:** And there's increasingly sophisticated functionality in ispLEVER to help the architect. Also, FPGA IP is particularly important, offering customers time to market and risk management advantages by providing proven, hardware-validated solutions that are typically parameterizable. These pre-engineered IP cores are cost-effective, and can save countless hours of development and validation effort.

***Chris Fanning*** *is Corporate Vice President of Enterprise Solutions and manages Lattice's software, intellectual property, and technical support businesses. He previously worked at The Boston Consulting Group. He holds an MS in Computer Science from Worcester Polytechnic Institute, and an MBA from The University of Chicago.*

***David Lee Rutledge*** *is Corporate Vice President of FPGA Product Development, with two decades of programmable logic experience including development of the first in-system programmable PLD. After working for Harris Semiconductor, Rutledge joined Lattice in 1983 as its 31st employee. He holds a BSEE and MSEE from Purdue University.*

Lattice Semiconductor

503-268-8000

www.latticesemi.com

# Non-volatility: the choice is clear

*By David Lee Rutledge*

Non-Volatility (NV) offers the same compelling benefits to FPGA users as it always has to CPLD users: a single-chip solution, instant-on capabilities, design security, and more. In fact, these features are arguably more valuable to FPGA applications than to CPLD applications.

So, why isn't everyone making NV FPGAs?

NV FPGAs require much more sophisticated process technology and circuit design techniques than their SRAM-based cousins. The key to NV success is providing sophistication without significantly impacting cost or performance, compared to an equivalent SRAM FPGA.

## Anatomy of a true NV FPGA

A true NV FPGA must employ high-performance, embedded flash memory technology (Figure 1). Historically, this approach had relatively low performance due to the additional wafer processing steps required to embed flash memory in a logic process. Lattice has worked closely with our foundry partner Fujitsu to develop proprietary embedded Flash technologies that enable exceptionally high performance for minimal additional cost.

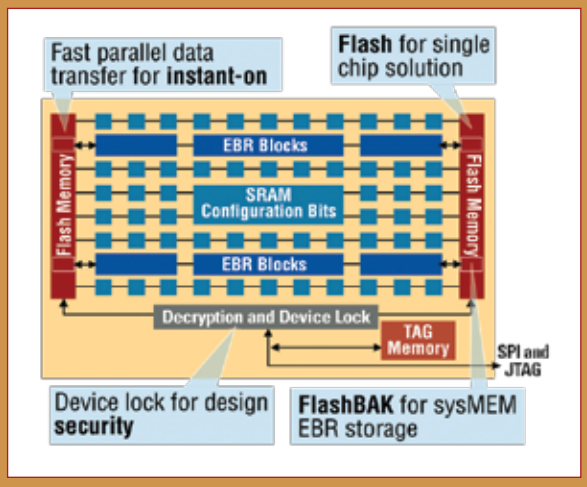


Figure 1

*Instant-on*, a feature that allows an FPGA to be used immediately upon startup, is essential in many applications. Typical SRAM FPGAs require 10s to 100s of milliseconds at startup using external boot PROMs. This prevents the use of SRAM FPGAs in certain critical control logic sections of systems that must be "live" immediately upon the application of power. Lattice has developed proprietary circuit design techniques that allow its NV FPGAs to become active virtually instantly after startup (within 1-2 mS).

Design security is also a key advantage for our NV FPGAs, in two distinct ways. First, our embedded flash memory allows the entire user design to be stored in on-chip Flash memory. The configuration bitstream is never exposed, as it is when using SRAM FPGAs with boot PROMs. It is virtually impossible for a hacker to access the configuration data of a NV FPGA.

Second, our embedded flash memory allows support of true AES encryption of the configuration bitstreams (Figure 2). This is a new feature that has been added to the latest 90 nm LatticeXP2 family. It supports the storage of a user-programmable, on-chip 128-bit encryption key. Once this key is stored on-chip, the user can encrypt the configuration bitstream for a given design and transmit it to the FPGA (over unsecured transmission lines). Once on-chip, the bitstream is decrypted using the stored AES key, and is never exposed off-chip.

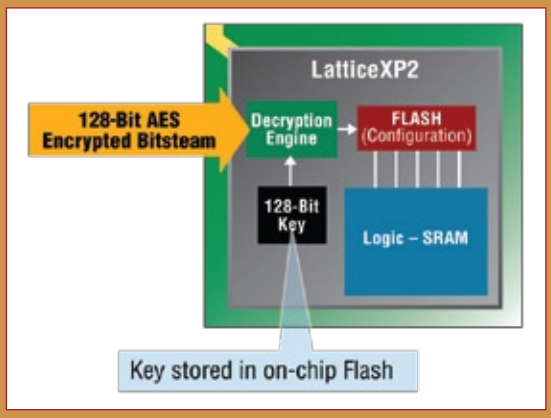


Figure 2

## Alternative NV technologies

FPGAs have been developed that are called NV, but they are not true single-chip, embedded flash solutions. Instead, they combine two independent chips (a flash memory and an SRAM-based FPGA) in a single package using a stacked die technology. Compared to true NV FPGAs, these hybrid products have severe limitations.

For example, stacked die technology cannot support instant-on because essentially it is identical to two-chip, SRAM-based technology, and so has the same lengthy boot time as an SRAM FPGA (Figure 3).

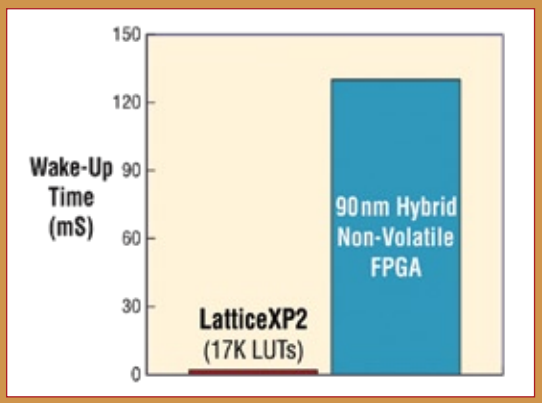


Figure 3

Stacked die technology also does not enable ultimate design security because any stacked-die implementation must transmit the unencrypted configuration data from the boot PROM to the SRAM FPGA over bond wires within the package. A hacker need only perform minor surgery to gain access to the wires that transmit the configuration data.