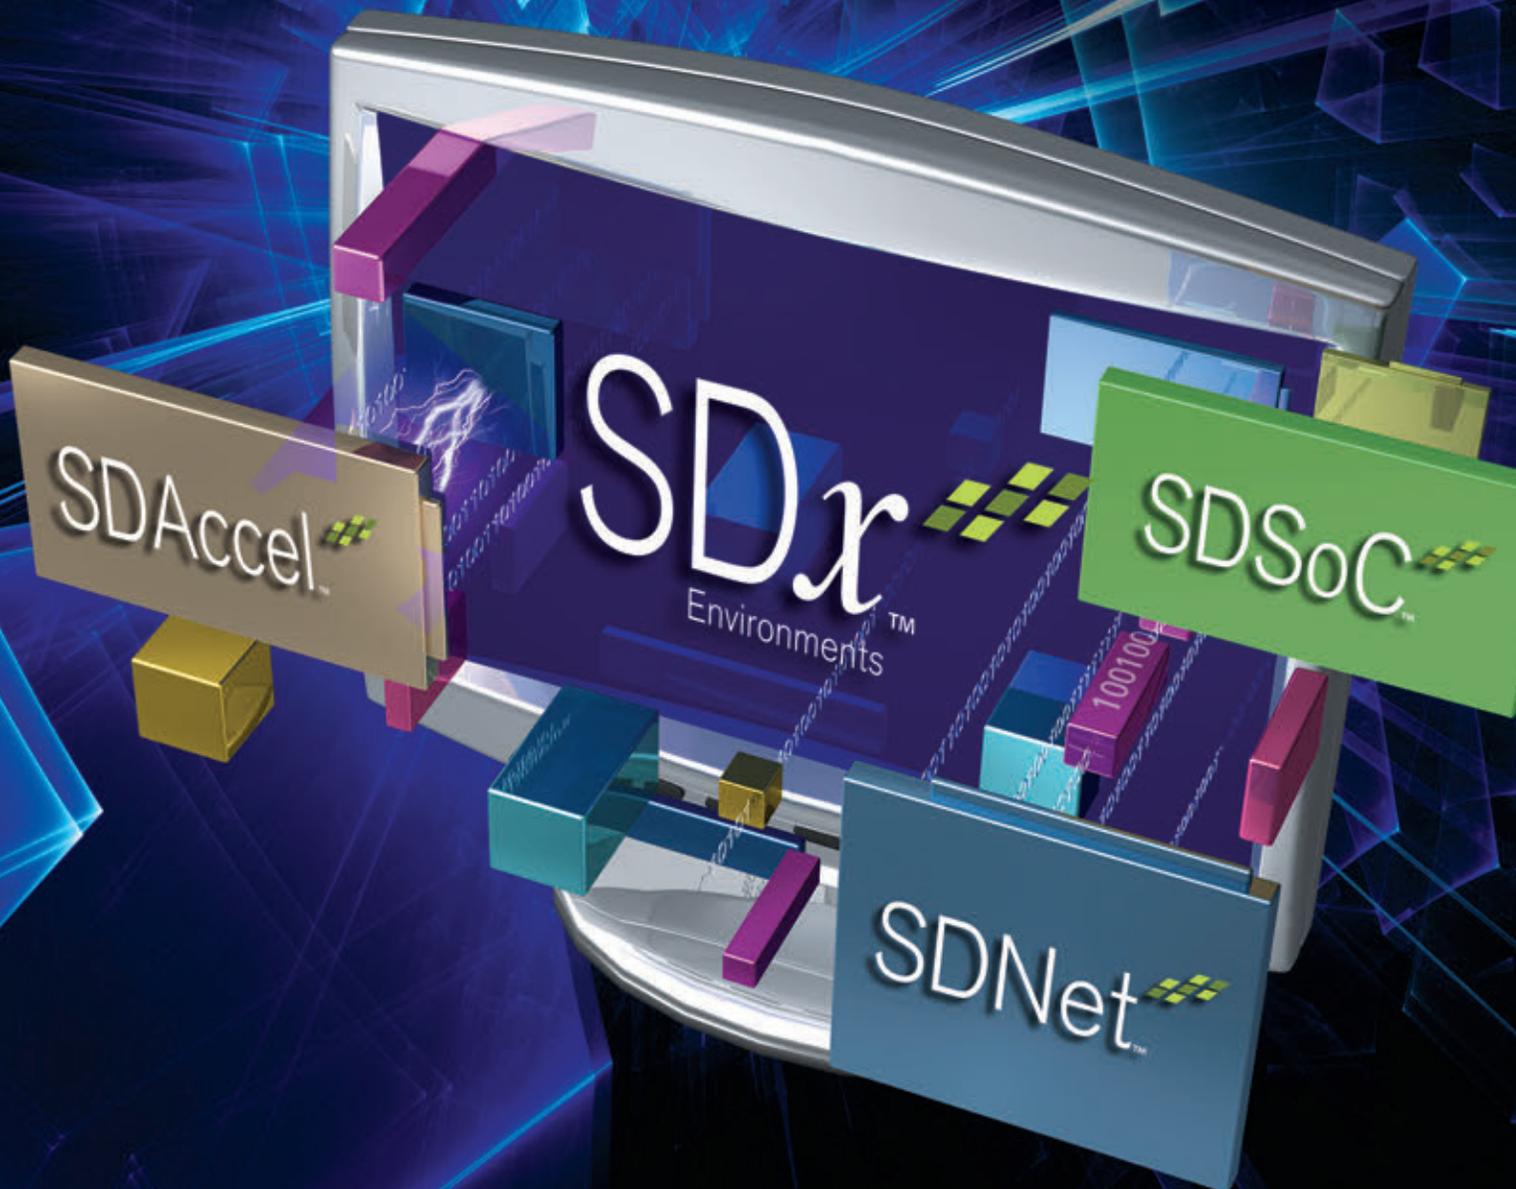


# All Programmable Abstractions: Programming Your Way

Xilinx's new SDx software-defined environments complement Vivado IPI, HLS and popular system-level design tools.



by Mike Santarini

Publisher, Xcell Journal

Xilinx, Inc.

mike.santarini@xilinx.com



To enable new levels of design team productivity and expand the reach of its All Programmable FPGAs, SoCs and 3D ICs to a much larger user base of software engineers, Xilinx® Inc. recently unveiled two new additions to its SDx™ development environment family. The new SDAccel™ development environment enables data center equipment programmers—with no FPGA experience—to program Xilinx FPGAs for data center and cloud computing infrastructure equipment using OpenCL™, C or C++. The resulting FPGA-based equipment will offer far superior performance per watt (performance/watt) than GPU- and CPU-based equipment. Xilinx has also unveiled the SDSoC™ development environment, which enables software developers—again, with no FPGA experience—to create systems in C or C++ targeting Zynq®-7000 All Programmable SoC and UltraScale+™ MPSoC platforms from Xilinx and third-party platform developers.

The SDx environments are the latest deliverables from Xilinx's All Programmable Abstractions campaign. The initiative is designed to enable software engineers and system architects to easily program Xilinx devices with development environments tailored to their design needs.

“The combination of SDNet, SDAccel and SDSoC will provide familiar CPU-, GPU- and ASSP-like programming environments to system and software engineers,” said Steve Glaser, senior vice president of the corporate strategy and

marketing group at Xilinx. “For the first time, those engineers will be able to derive the unique benefits of All Programmable devices for customized acceleration, 10x to 100x performance per watt and any-to-any connectivity with the required security and safety for the next generation of smart systems. Xilinx is enabling these next-generation systems to be more connected, software defined and virtualized. They must also support software-based analytics and enable more computing in the cloud, often driven by pervasive video and embedded vision. This requires SDx software-defined programming environments and heterogeneous multiprocessing with new UltraScale FPGAs and MPSoCs.”

The launch of the new SDAccel and SDSoC environments follows closely behind the spring 2014 release of the SDNet™ development environment, detailed in the cover story of [Xcell Journal issue 87](#).

While the new SDx environments enable software engineers and system architects to program the FPGA portion of Xilinx devices, the environments will also enable design teams with hardware engineering resources to be more productive and quickly converge on an optimized system to improve time-to-market. With a working system design in hand, hardware engineers can focus on optimizing FPGA layout and performance for system efficiencies, while software engineers further refine application code.

### ALL PROGRAMMABLE ABSTRACTIONS

When Xilinx began planning its 7 series All Programmable device family of FPGAs, 3D ICs and Zynq-7000 All Programmable SoCs back in 2008 under new CEO Moshe Gavrielov, it became evident that the rich capabilities of each of the members of the 7 series and future product lines would enable customers to place Xilinx's devices at the heart of their newest, most innovative products.

These All Programmable devices were far more sophisticated than the

glue-logic FPGAs of Xilinx's earlier years and enabled system functionality and end-product differentiation not achievable with any other architecture. To maximize the value of these new devices and jump ahead of the competition, management knew it was imperative for Xilinx to develop tools and methodologies that would enable system architects and even embedded-software developers—not just FPGA experts—to program Xilinx's newest devices. Further, the company would have to develop design environments for software engineers targeting key growth markets, and tailor those environments to the tools and flows these designers were accustomed to using. It would also be imperative to strengthen alliances with companies like The MathWorks and National Instruments, which already have established environments that enable nontraditional-FPGA users to leverage the power efficiency, performance and flexibility of Xilinx All Programmable devices.

For established customers, providing design environments for every member of the design team would guarantee efficiency and shorten time-to-market. If

sophisticated enough, the environments would essentially “democratize” All Programmable FPGA and Zynq SoC design, enabling individual architects and software engineers who have no experience designing with FPGAs to program these devices without help from hardware designers. Worldwide, software engineers outnumber hardware engineers 10 to 1. Thus, Xilinx and Alliance Program partners providing such development environments (or the hardware platforms supporting them) would expand both their user base and their revenue.

That strategy, along with a subsequent development effort to enable software engineers and system architects to program Xilinx devices with environments tailored to their design needs, is what Xilinx calls All Programmable Abstractions (Figure 1).

### VIVADO HLS AND IPI: FIRST STEPS UP IN DESIGN ABSTRACTION

The first serious step up in design abstraction began in 2011 with Xilinx's acquisition of privately held high-level-synthesis (HLS) tool vendor AutoESL. The merger was followed in 2012 by Xilinx's public release of the

HLS technology integrated into its ISE® Design Suite and Vivado® Design Suite tool flows. Xilinx selected the AutoESL technology after an exhaustive study by Berkeley Design Automation demonstrated that AutoESL's HLS tool was the easiest to use and offered the best quality of results of all the HLS tools available from the EDA industry. Its origins in the EDA world also meant that the use model for the tool was targeting ASSP and system-on-chip (SoC) system architects and well-rounded design teams who had experience in C and C++ programming along with a working knowledge of hardware-description languages (Verilog and VHDL) and chip design requirements.

With Vivado HLS, such broadly skilled architects and design teams can create algorithms in C and C++ while using Vivado HLS to compile and convert those algorithms into an RTL intellectual-property (IP) block. Then, an FPGA designer can take that block and others in RTL that they've created or licensed, and assemble the IP into a design using Xilinx's IP Integrator (IPI) tool. Next, the FPGA designer takes the assembled design through the Vivado flow, performing HDL simulation, timing and power optimization, placement and routing. Ultimately, the designer generates a netlist/bit file and configures the hardware of the targeted All Programmable device. If the design includes a processor (either the Zynq SoC or a soft MPU core), the configured device is then ready for embedded-software engineers to program.

To help these embedded-software engineers with the programming chore, Xilinx provides an Eclipse-based integrated design environment called the Xilinx Software Development Kit (SDK), which includes editors, compilers, debuggers, drivers and libraries targeting Zynq SoCs or FPGAs with Xilinx's 32-bit MicroBlaze™ soft core embedded in them. Introduced more than a decade ago, the SDK has evolved dramatically as Xilinx's integration of processors on its devices has transitioned

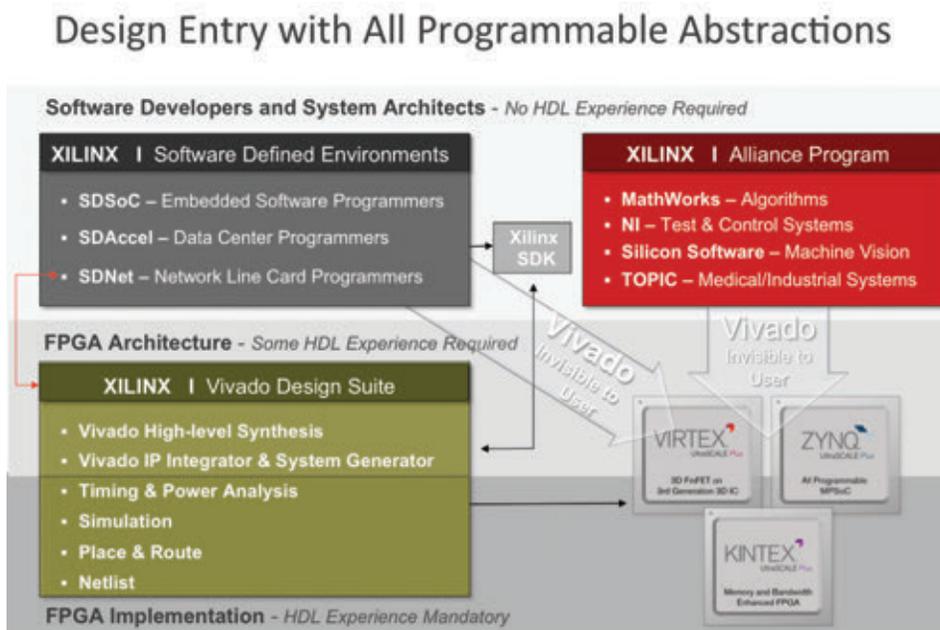


Figure 1 – With the new SDx design environments and those from Alliance members, Xilinx is enabling more innovators to leverage Xilinx All Programmable devices to add greater value to their next-generation products.

from hardened DSP slices and soft-core MCUs and MPUs (8, 16 and 32 bits) to a hardened 32-bit PowerPC® in Virtex®-4 and Virtex®-5 FPGAs; hardened 32-bit ARM® processors in the Zynq SoC; and 64-bit ARM processors in the upcoming Zynq UltraScale+ MPSoC.

### ALLIANCE MEMBERS BRING XILINX VALUE TO MORE USERS

For well over a decade, Xilinx has had very close partnerships with National Instruments and The MathWorks. Both of these companies offer unique high-level development environments tailored to the specific audiences they serve.

[National Instruments](#) (Austin, Texas) offers hardware development platforms fanatically embraced by control and test system innovators. Xilinx's FPGAs and Zynq SoCs power the NI RIO platforms. National Instruments' LabVIEW development environment is a user-friendly graphics-based program that runs the Vivado Design Suite under the hood so that National Instruments' customers need not know any of the details of FPGA design. Some perhaps don't even know a Xilinx device is at the heart of the RIO products. They can simply program their systems in the LabVIEW environment and let NI's hardware speed the performance of designs they are developing.

In the case of [The MathWorks](#) (Natick, Mass.), more than a decade ago the company added FPGA support to its MATLAB®, Simulink®, HDL Coder and Embedded Coder with Xilinx's ISE and Vivado tools running under the hood and completely automated. As a result, the users—who are mainly mathematician algorithm developers—could develop algorithms and speed algorithm performance exponentially, running the algorithms succinctly on FPGA fabric.

Over a decade ago, Xilinx added an FPGA architecture-level tool called System Generator to its ISE development environment and, more recently, the Vivado Design Suite, in order to enable teams who had FPGA knowl-

edge to further tweak designs for additional algorithm performance gains. This combination of MathWorks and Xilinx technologies has helped customer companies produce thousands of innovative products.

More recently, other companies have begun contributing to the Xilinx environment as well. Xilinx recently welcomed two European companies, TOPIC Embedded Systems and Silicon Software, to its Alliance Program, specifically for their high-level development environments focused on medical and industrial markets.

[TOPIC Embedded Systems](#) (Eindhoven, the Netherlands) has a unique development environment called DYP-LO, which was featured in [Xcell Journal issue 89](#). Targeting the Zynq SoC and, soon, the Zynq MPSoC, the environment takes a singular approach to system-level design, enabling system architects to create a system representation of their design in C or C++ and run it on the Zynq SoC's dual-core ARM Cortex™-A9 processing system. When users find sections of the design that are running too slowly in software, they can drag and drop the sections to a window that converts the C into FPGA logic (running Vivado HLS under the hood) and places it in the Zynq SoC's programmable logic. They swap code snippets back and forth until they converge on an optimal system performance. TOPIC has seen first use in medical-equipment development and is starting engagements with manufacturers of industrial equipment.

[Silicon Software](#) (Mannheim, Germany) is the newest addition to the quartet of Alliance members enabling software engineers to leverage Xilinx All Programmable devices. The company's VisualApplets graphical image-processing design environment allows system architects and software engineers targeting Zynq SoC platforms to build new innovations in industrial machine vision applications—and do so without assistance from a hardware engineer. In Xilinx's booth at the SPS Drives trade show in 2013, Silicon

Software demonstrated an optical-inspection system it developed with the VisualApplets environment. The demo showed a system performance speed-up of 10x by using the VisualApplets environment to move image-processing tasks from the Zynq SoC's processing system to the device's FPGA logic.

### DEMOCRATIZING FPGA AND SOC DESIGN WITH SDX

With the SDx software-defined development environments, Xilinx is building a series of higher-level design entry environments targeting software developers and system architects in key markets. Both SDAccel and SD-SoC run the entire Vivado flow automatically under the hood, and require no direct use of Vivado tools and no hand-off to a hardware engineer. For its part, SDNet does not access Vivado HLS and has a unique two-step use model targeting line card architects.

In the first step, line card architects use an intuitive, C-like high-level language instead of arcane microcode to design the requirements and create a specification for a network line card. The SDNet development environment generates an RTL version of the design based on that specification. The flow then requires a hardware engineer to implement the RTL into the targeted FPGA.

In the second step, SDNet also allows network companies to test and update protocols in the high-level language and upgrade the functionality of the cards—even after deployment in the field—without hardware design intervention. The flow enables companies to quickly create and update cards, flexibility that's ideal for software-defined networks.

The two new environments, SDAccel and SDSoc, take the SDx philosophy into new application areas.

### SDACCEL OPTIMIZES DATA CENTER PERFORMANCE/WATT

According to a March 2014 article in the *Data Center Journal*, huge data centers

that form the heart of companies like Google, Facebook, Amazon and LinkedIn “consume upwards of 3 percent of the world’s electric power, while producing 200 million metric tons of CO<sub>2</sub>.” That enormous power consumption costs data centers more than \$60 billion a year in electricity fees. The power consumption cuts deeply into the bottom-line profitability of even the biggest dot-com companies but also has an untold cost on the environment.

As more companies look to leverage cloud computing and analysis via Big Data; and as video and streaming media become more pervasive worldwide; and as more people join wireless networks and upgrade toward 5G, there is a relentless, exponentially increasing demand for more data centers with even better performance. The current trajectory, according to the *Data Center Journal* article, will bring data center traffic to 7.7 zettabytes annually by 2017. That means that data center power consumption will rise astronomically if left unchecked. Today the main culprit in this power consumption is the Intel x86 processors that form the bedrock of most data center equipment. MPUs today provide good, but not optimal, performance and high power consumption.

Software engineers, in abundant supply worldwide, find these MPUs the easiest devices to program. To solve the performance portion of the data center problem, many companies have been creating equipment with graphics processing units (GPUs) or CPU systems accelerated by GPU cards. GPUs have performance that’s far superior to that of CPUs in data center applications but unfortunately, far worse power consumption. Together, the performance is extremely fast but the power consumption is abysmal.

To get the best of both worlds, many companies are turning to an FPGA-centric approach in which they pair FPGAs with other processors to maximize data center equipment performance.

A number of data center equipment vendors have demonstrated that a dis-

crete FPGA paired with a discrete CPU raises power per card minimally but improves performance dramatically, yielding a significant improvement in performance/watt. Others believe that performance/watt can be further improved with a chip that combines an x86 processor core interconnected to FPGA logic on a single SoC. Some think that a seemingly even lower-power and equally high-performance solution would be to integrate FPGA logic with ARM 64-bit processor IP on a single SoC.

The main deterrent to using FPGAs in the data center has been programming. Data center developers are accustomed to programming x86-based architectures and typically come from a pure software-programming background. A first step designed to help developers move CPU programs to faster GPUs was the industry’s open development of the OpenCL language. Over the last two years, OpenCL has evolved even further to enable customers to target FPGAs. This development is opening up new possibilities for future data center equipment architectures and even ubiquitous networks.

By launching the SDAccel environment, Xilinx has bridged that programming gap and paved the way to enabling data center engineers to use OpenCL, C or C++ to pro-

gram FPGA-based platforms without requiring design intervention from hardware engineers. Tom Feist, senior director of design methodology marketing at Xilinx, said the SDAccel development environment for OpenCL, C and C++ enables data center programmers to build equipment with up to 25x better performance/watt than CPU- and GPU-based systems.

Feist said that in the SDAccel flow (see Figure 2, the SDAccel development environment demo), which targets x86 CPUs paired with acceleration cards running Xilinx’s 20nm Kintex® UltraScale FPGAs, software developers identify applications that need acceleration, code and optimize the kernels in OpenCL, and then compile and execute the applications for the CPU. They then cycle to estimate kernels and debug them to come up with cycle-accurate models. Next, they use SDAccel to compile the code and implement it automatically in the FPGA (possible because Vivado runs under the hood). They can then run the application and validate the performance of the applications accelerated with the card vs. performance when running solely on the CPU. “They can run iterations in this cycle until they find the optimal balance between performance

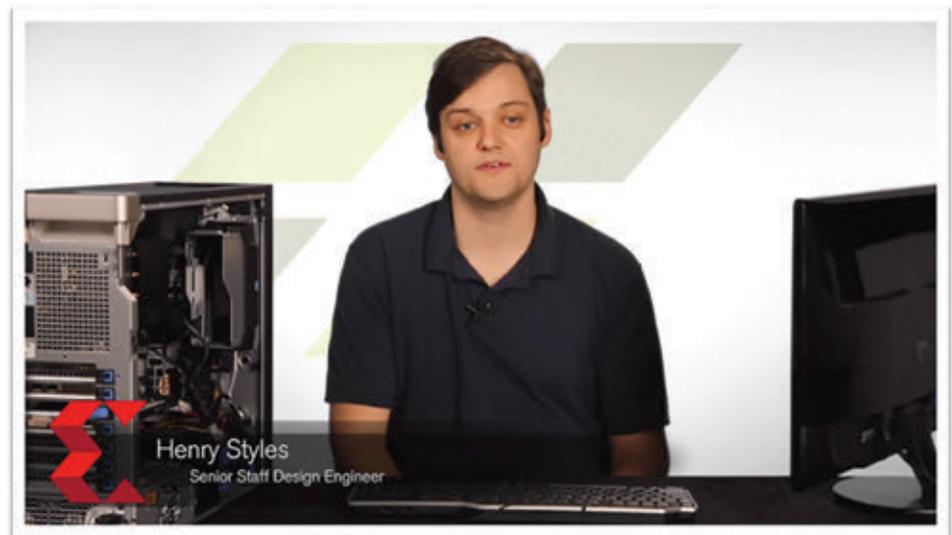


Figure 2 – In this [SDAccel development environment demo](#), engineer Henry Styles shows how to use the SDAccel development environment for acceleration using a standard x86 64-bit workstation containing an Alpha Data ADM-PCIE-7V3 accelerator.



Figure 3 – In this [SDSoC development environment demo](#), principal engineer Jim Hwang uses SDSoC to create a simple image-processing pipeline to detect motion and insert motion edges into a live HD 1080p video stream running at 60 frames per second.

and power and achieve up to 25x performance/watt improvement over CPU and GPU implementations,” said Feist.

In this issue of *Xcell Journal*, Feist and colleagues contribute a detailed overview of the SDAccel environment and a second article showing SDAccel in action (see pages 38 and 42).

### SDSOC MULTIPLIES EMBEDDED-SYSTEM INNOVATORS

While SDNet enables line card developers to quickly create next-generation networks with a unique, “softly defined” approach, and while SDAccel allows data center equipment vendors to achieve the best performance/watt of their next-generation data centers, the new SDSoC development environment may well have the broadest impact among Xilinx’s user base. That’s because SDSoC targets the vast numbers of embedded-system design teams and specifically, their software engineers designing for the majority of markets Xilinx serves with its Zynq SoCs. The SDSoC environment enables users to now configure the logic—not just program the processor of embedded systems running Zynq SoC-based hardware platforms—in C and C++.

“Software developers are accustomed to programming motherboards, ASSP

platforms and ASICs without requiring a hardware engineer to do anything,” said Nick Ni, SDSoC product manager. “With the SDSoC development environment, they can program Zynq SoC and MPSoC platforms in the same manner as they have with ASSPs. But what’s unique is that with SDSoC, they can now create complete system designs in an Eclipse IDE environment using C or C++ targeting the Zynq SoC and Zynq UltraScale+ MPSoC platforms.”

Ni said that in using the SDSoC environment, embedded-software developers create their designs in C or C++ and test to see what portions aren’t running optimally on the Zynq SoC’s processing system. They highlight the suspect code and command the SDSoC environment to automatically partition that code into the Zynq SoC’s programmable logic to speed up the system performance. Ni said the SDSoC environment can move a software function to FPGA logic with the click of a button. And it doesn’t require a hardware engineer to do it. The compiler in SDSoC will generate the entire Vivado project as well as the bootable software image for the targeted hardware platform.

“We are essentially enabling embedded-software engineers to become

system engineers with the SDSoC environment for our Zynq SoCs,” said Ni (see Figure 3, the SDSoC development environment demo).

The SDSoC environment leverages a macro compiler that takes the C/C++ code users have designated to accelerate in the Zynq SoC’s logic. By running the Vivado Design Suite under the hood, the environment automatically turns the code into an IP block and configures that block into the device’s logic, automatically generating a driver in software.

SDSoC provides board support packages (BSPs) for Zynq All Programmable SoC-based development boards including the ZC702 and ZC706, as well as third-party and market-specific platforms, such as the ZedBoard, MicroZed, ZYBO, and video and imaging development kits.

“We’ll be adding more BSPs to SDSoC in the coming months, especially as more third-party platform companies develop systems with the Zynq SoC,” said Ni. “The SDSoC not only expands the user base for Xilinx but also for companies developing platforms leveraging the Zynq SoC.”

Ni is quick to note that while the main goal of the SDSoC environment is to enable the vast number of embedded-software designers to now create entire systems with Xilinx Zynq SoCs, users with traditional FPGA backgrounds and design teams can also greatly benefit from using the environment.

“It enables design teams to quickly design a system architecture in C and C++ and then try different configurations to get the optimal performance they require,” Ni said. “If they do have FPGA designers on their team, they can have them use Vivado tools to optimize the blocks and logic layout further.”

For additional news about Alliance member support for the SDSoC environment, see the Xpedite section in this issue (page 66). For further information on Xilinx’s All Programmable Abstractions, visit <http://www.xilinx.com/products/design-tools/all-programmable-abstractions.html>.