

Designing portability into silicon IP

DESIGN FOUNDRY PORTABILITY INTO IP CORES, RATHER THAN APPLYING IT AFTER THE FACT.

Designing IP (intellectual-property) cores that you can port to various foundry protocols or process geometries is paramount in today's highly diverse standards-based-IP technologies. This technological advantage allows IP-core vendors to gain access to new customers and markets, enabling chip designers to be first to market with new products that ultimately consume less power and area. It is also imperative for IP-core vendors to offer products that can target multiple industry-standard protocols, such as PCIe (peripheral-component-interconnect express), SATA (serial-advanced-technology attachment), and GbE (gigabit Ethernet), which are the building blocks for switches, hard drives, chip-to-chip communications, optical communications, and networking.

For successful SERDES (serializer/deserializer)-IP development, you must design SERDES cores that you can easily and quickly port from foundry to foundry and into the latest and smallest-geometry processes. This approach ensures a comprehensive portfolio of IP, which results in a wide pool of potential deployments in customer applications. Customers all have their own reasons for wanting to go with one foundry and process technology rather than another. Going with the latest process technology can afford lower power dissipation and increases the number of transistors customers can fit onto a die, decreasing costs and increasing profit margins.

To design portable SERDES IP, analog- and digital-system designers must build a range of programmability into their circuits to make them robust over a range of process variations, and they must be able to rapidly verify circuits using various I/O-supply voltages. As with all development efforts, trade-offs arise between designing portability into IP and optimizing IP for a process. Designing SERDES IP with wide programmability allows for porting efficiencies and costs IP developers increased footprint for the IP on the chip.

Ideally, a design team can rapidly port a design without much redesign or modification of the physical design or layout. Verification, however, is always necessary because each foundry has its own set of model libraries, and designers must place and route all of the digital synthesized logic using each foundry's specifications. Even the same process technology can vary widely in performance among foundries. For example, one foundry's 65-nm, general-purpose process can differ from another foundry's 65-nm, general-purpose process in process and device parameters, as well as the typical variances in physical geometries and spacing rules. Therefore, efficient de-

velopment of IP cores always takes into account minimizing effort and time to address these differences.

ANALOG DESIGN AND PROGRAMMABILITY

Designing analog circuits with appropriate programmability in mind can vastly reduce the effort of porting from foundry to foundry or porting into new process technologies within the same foundry. SERDES IP typically includes a CDR (clock-and data-recovery) block in the receiving path, as well as a transmitter PLL (phase-locked loop), which generates the transmit high-speed clock. The CDR is responsible for recovering a clock from a serial binary-data stream that samples the high-speed data. The data and clock then go to a serial-to-parallel converter to deserialize the serial stream into bytes and words, which the converter then sends on to the core digital logic. The transmitter PLL is responsible for multiplying a reference-clock frequency into a high-speed bit clock that a parallel-to-serial converter then uses to serialize and transmit data words. The CDR and transmitter PLL contain VCOs (voltage-controlled oscillators) that provide a range of output frequencies. The chip designer tunes this circuit for the specified process-technology port.

Typically, digital control of the voltage inputs into the VCO tunes the circuit. For example, if a process technology operates at a slower rate than the chosen benchmark specifies, design-




Figure 1 An open-eye diagram is one indication that a port to a new process has been successful.

ers would tune the VCOs to a higher setting to provide the desired frequency of operation. This programmability allows the IP vendor to design and lay out CDRs and transmitter PLLs in a way that provides enough tuning range to achieve the correct output frequency for various process technologies. The IP designer must determine how much tuning range and resolution to build in to target all desired foundries and process technologies for a design.

Another parameter that typically varies among processes is the sheet resistance of polysilicon resistors. A polysilicon resistor with a given geometry may be more or less resistive than a polysilicon resistor of the same geometry in another process. To overcome these resistance mismatches between process technologies, you may need to allow for programmable resistances. Designers need resistors that they can switch into and out of a circuit. Although not the most area-efficient approach, this feature allows the designer to control the resistance without modifying the layout. An example of this usage is in transmitting/receiving-termination circuits.

When designing SERDES IP to enable quick porting to smaller process geometries, it is important to take note of the I/O-supply voltages that each process technology supports. For instance, some 65-nm processes support 2.5 and 3.3V I/O transistors, whereas others may support only 1.8V devices. To further complicate matters, overdriven, 2.5V devices may have a wider range of operation. IP-circuit designers should design the SERDES to accept a range of I/O voltages and lay out the circuits such that the IP's physical designer can easily swap in different I/O transistors. This step usually involves having the IP's physical designer develop a floorplan with room for the larger I/O devices in the layout and for the IP-verification engineers to verify operation of the circuits with all the devices. With careful design, this step can be as simple as swapping layers in the physical-design software based on the customer's preference. Another possibility is to use the overdriven I/O devices with extended widths, which can support a range of supply voltages larger than 10% that vendors typically specify for I/O-supply variations of SERDES. A designer must be mindful not only of multiple-foundry design and portability into new process geometries, but also of the process-technology options that are available, such as different I/O devices.

DIGITAL CALIBRATION OF ANALOG CIRCUITS

An effective technique for using programmability is automatic digital calibration. In this process, a designer puts a digital-calibration loop around an analog-circuit block. This technique is useful not only for minimizing design changes when porting, but also for extending a SERDES' ability to minimize effects due to process, voltage, and temperature variations. Automatic calibration and a wide enough tuning range on key analog circuits aid in getting better porting results without requiring extensive verification for different process technologies and options. The digital-calibration circuitry should take up minimal area and should become smaller than the analog block as process geometries get smaller. Plus, the benefits can be substantial. The calibration helps compensate for differences in performance among each foundry's technology and among newer technologies. For example, the calibration

loop can automatically tune the VCO's voltage inputs so that the VCO oscillates at the correct frequency even if the performance of a new foundry's process technology differs greatly from that of the process currently in use. Programmability and calibration are important tools for ensuring first-pass silicon success when porting to new technologies.

Using this design approach to achieve first-pass silicon success in a port to a different process technology, one company designed a test chip for a new process technology. The company had recently completed a test chip for a foundry's general-purpose process, and the goal was to quickly port to the equivalent geometry's lower-power node. The designers placed a wide tuning range on the VCOs, added programmability to their transmitter/receiver-termination resistors, and designed for multiple-I/O transistor devices. The team had time to verify only the major blocks, such as the master bias, and little time for physical-design changes. The team relied on the use of built-in programmability and calibration.

The design would need few changes to meet the specifications of the lower-power process technology. A few months later, the team received both the lower-power and the general-purpose test chips and began characterizing both. While testing, the team verified the calibrated values of the lower-power and the general-purpose test chips and observed similar performance. The end result was that the lower-power test chip performed as successfully as the general-purpose test chip, with minimal time and effort in porting from the general-purpose to the lower-power chips.

AUTOMATED VERIFICATION

The other missing link of designing portability into IP is the use of specially designed and optimized CAD (computer-aided-design) tools that streamline porting. The goal is to design IP that designers can quickly and easily port, and the CAD tools you use play a big part in achieving this goal. You can devise a customized verification-suite setup for rapid verification. Such a suite would comprise custom PERL (Practical Extraction and Report Language) scripts and Microsoft (www.microsoft.com) Office macros that automatically generate testbenches, run simulations, and extract formatted results. In such a setup, the highly automated verification process relieves designers of the task of extracting data so that they can focus on designing the circuits and analyzing the pertinent results. An automated verification suite that all the designers use also makes it easier to share data and for designers to easily and quickly begin new projects without long learning curves. Some companies also offer specification templates for each major analog-circuit block, and these templates allow quick, repeatable, and consistent verification and produce formatted results for easy review.

PHYSICAL DESIGN FOR PORTING

Supporting a range of I/O-device options inevitably requires allocating a large enough area to support the largest devices a user might drop into the circuit. Similar ideas apply to supporting multiple foundries; the layout must work for different sets of design rules. It may not always be feasible to target designs for all foundries and processes of interest. You may have to take a look at the design rules of each foundry and make a

layout compatible with only a subset of those rules. Supporting too large a set of foundries inevitably increases the size of your layout, resulting in wasted area.

Minimum device geometries and spacing requirements for each foundry can vary widely, so having a custom set of design rules that meets requirements for multiple foundries is a good way to do the physical design. When taping out designs to customers, however, you must ensure that you use the target foundry's design rules rather than your custom set. The most time-consuming form of physical design is manual analog layout, so, for efficient design porting, minimize custom design in favor of digital-synthesized logic that you can place and route. Although some tools can partially automate analog physical design, you should avoid their use for high-speed SERDES applications. IP designers should investigate any opportunity to replace a traditional analog-circuit block with a digital implementation. This extensive use of digital circuits makes the design easier to port to new geometries and allows the design to more easily reap the benefits of the area savings than would an analog implementation. The design must still meet all required specifications, including area, power, performance, and eye-diagram templates (**Figure 1**).

Providing IP solutions in multiple foundries and in multiple process-technology nodes allows IP providers to widen their customer base. To target as many technologies as possible, the IP designer must optimize design, verification, and physical design with porting in mind. This approach means designing analog blocks with a wide tunable range and placing a digital-calibration loop around these blocks. Adding calibration loops around wide-tuning analog-circuit blocks enables the circuit to calibrate process variations and improve yield. Designers should choose digital implementations of analog-circuit blocks wherever feasible. **EDN**

AUTHOR'S BIOGRAPHY

Joseph Pun is an applications engineer at Gennum Corp's Snowbush IP Group. He obtained a bachelor's degree in electrical engineering from the University of Toronto (ON, Canada).