



WP380 (v1.2) December 11, 2012

Xilinx Stacked Silicon Interconnect Technology Delivers Breakthrough FPGA Capacity, Bandwidth, and Power Efficiency

By: Kirk Saban

The programmable imperative—the critical need to achieve more with less, to reduce risks wherever possible, and to quickly create differentiated products using programmable hardware design platforms—is driving the demand for FPGA-based solutions that provide higher capacity, lower power, and higher bandwidth with which users can create the system-level functionality currently delivered by ASICs and ASSPs.

Xilinx has developed an innovative approach for designing and manufacturing FPGAs that address two key requirements of the programmable imperative. Stacked silicon interconnect (SSI) technology is the foundation of a new generation of FPGAs that break through the limitations of Moore's law and deliver the capabilities to satisfy the most demanding design requirements. It also enables Xilinx to reduce the time required to deliver the largest FPGAs with the highest bandwidth in the quantities needed to satisfy end-customer volume production requirements. This white paper explores the technical and economic challenges that led Xilinx to develop SSI technology and innovations that make it possible.

Introduction

As the role of the FPGA becomes more dominant in system design, the designs grow larger and more complex, demanding higher logic capacity and more on-chip resources. To date, FPGAs have depended predominantly on Moore's Law to respond to this need, delivering nearly twice the logic capacity with each new process generation. However, keeping pace with today's high-end market demands requires more than Moore's Law increases can provide.

The most aggressive adopters of FPGA technology are eager to employ the highest capacity, highest bandwidth devices of each new FPGA generation. However, the challenges for vendors to build such FPGAs early in the product life cycle can limit the ability to supply the volumes of devices their customers require for their production runs. That is because the circuitry overhead that enables reprogrammable technology negatively affects the manufacturability (and therefore, the supply) of the largest FPGAs. At the early stages of a new process node, when defect densities are high, die yield declines dramatically as die size increases. As the fabrication process matures, defect density falls and the manufacturability of large die increases significantly.

Thus, while the largest FPGAs are in short supply at product introduction, over time they eventually become available in quantities that support end-customer volume requirements. In response to the programmable imperative, a few leading-edge customers challenged Xilinx to support their volume production requirements with the largest FPGAs as soon as possible after product introduction.

The telecommunications market, for example, needs FPGAs that incorporate dozens of serial transceivers, providing high signal integrity. Devices also need to provide extensive interconnect logic and block RAM for data processing and traffic management while maintaining current form factors and power footprints. To reap first-mover advantage, the equipment makers want to ramp up manufacturing of their new products as rapidly as possible.

Xilinx has responded to these requirements with an innovative approach for building FPGAs that offer bandwidth and capacity equaling or exceeding that of the largest possible monolithic FPGA die with the manufacturing and time-to-market advantages of smaller die to accelerate volume production. These benefits are enabled by SSI technology, which uses passive silicon interposers with microbumps and through-silicon vias (TSVs) to combine multiple highly manufacturable FPGA die slices, referred to as super logic regions (SLRs), in a single package. The technology also allows die of different types or silicon processes to be interconnected on the interposer. This type of construction is referred to as a heterogeneous FPGA.

The Challenges of Interconnecting Multiple FPGAs

SSI technology solves the challenges that had previously obstructed attempts to combine the interconnect logic of two or more FPGAs to create a larger, "virtual FPGA" for implementing a complex design. These challenges include:

- The amount of available I/O is insufficient for connecting the complex networks of signals that must pass between FPGAs in a partitioned design as well as connecting the FPGAs to the rest of the system.
- The latency of signals passing between FPGAs limits performance.
- Using standard device I/O to create logical connections between multiple FPGAs increases power consumption.

Key Challenge: Limited Connectivity

System-on-chip (SoC) designs comprise millions of gates connected by complex networks of wires in the form of multiple buses, complicated clock distribution networks, and multitudes of control signals. Successfully partitioning an SoC design across multiple FPGAs requires an abundance of I/Os to implement the nets spanning the gap between FPGAs. With SoC designs including buses as wide as 1,024 bits, even when targeting the highest available pin count FPGA packages, engineers must use data buffering and other design optimizations that are less efficient for implementing the thousands of one-to-one connections needed for high-performance buses and other critical paths.

Packaging technology is one of the key factors to this I/O limitation. The most advanced packages currently offer approximately 1,200 I/O pins, far short of the total number of I/Os required.

At the die level, I/O technology presents another limitation because I/O resources do not scale at the same pace as interconnect logic resources with each new process node. When compared to transistors used to build the programmable logic resources in the heart of the FPGA, the transistors comprising device I/O structures must be much larger to deliver the currents and withstand the voltages required for chip-to-chip I/O standards. Thus, increasing the number of standard I/Os on a die is not a viable solution for providing the connections for combining multiple FPGA die.

Key Challenge: Excessive Latency

Increased latency is another challenge with the multiple FPGA approach. Standard device I/Os impose pin-to-pin delays that degrade the overall circuit performance for designs that span multiple FPGAs. Moreover, using time-domain multiplexing (TDM) on standard I/Os to increase the virtual pin count by running multiple signals on each I/O imposes even greater latencies that can slow I/O speeds down by a factor of 4X–32X or more. These reduced speeds are often acceptable for ASIC prototyping and emulation, but are often too slow for end-product application.

Key Challenge: Power Penalty

TDM approaches also result in higher power consumption. When used to drive hundreds of package-to-package connections across PCB traces between multiple FPGAs, standard device I/O pins carry a heavy power penalty compared to connecting logic nets on a monolithic die.

Similarly, multichip module (MCM) technology offers potential form-factor reduction benefits for integrating multiple FPGA die in a single package. The MCM approach, however, still suffers from the same restrictions of limited I/O count as well as undesirable latency and power consumption characteristics.

Key Challenge: Signal Integrity for High-Speed Serial Connectivity


Particularly in communications applications where high speed serial I/O connectivity is common, poor signal integrity can become a major bottleneck in reaching design closure. FPGAs must provide the appropriate transceiver signal fidelity, or countless hours must be spent refining I/O parameters, modifying the PCB design and performing channel optimization to achieve design success. With some applications requiring line rates exceeding 25 Gb/s, delivering adequate signal integrity can be a non-trivial task.

Xilinx SSI Technology

To overcome these limitations, Xilinx has developed a new approach for building production volumes of high-capacity and high-performance FPGAs. The new solution enables high-bandwidth connectivity between multiple die by providing a much greater number of connections. It also imposes much lower latency and consumes dramatically lower power than either the multiple FPGA or MCM approach, while enabling the integration of massive quantities of interconnect logic, transceivers, and on-chip resources within a single package.

Within the density range of an FPGA family, the medium-density devices represent the “sweet spot.” That is, compared to the previous generation, they offer significantly greater capacity and bandwidth—on a die size that can be delivered earlier in the FPGA product life cycle than the largest devices in the same family. Thus, by combining several of these die in a single device, it is feasible to match or exceed the capacity and bandwidth offered by the largest monolithic devices, but with the manufacturing and time-to-volume advantages of smaller die.

Xilinx arrived at such a solution by applying several proven technologies in an innovative way. By combining TSV and microbump technology with its innovative ASMBL™ architecture, Xilinx is building a new class of FPGAs that delivers the capacity, performance, capabilities, and power characteristics required to address the programmable imperative. Via a passive interposer, Xilinx SSI technology combines multiple FPGA SLRs. The interposer provides tens of thousands of die-to-die connections to enable ultra-high interconnect bandwidth with far lower power consumption and one fifth the latency of standard I/Os. **Figure 1** shows the side view of the die stack-up with four FPGA SLRs, silicon interposer, and package substrate.



WP380_01_112812

Figure 1: Virtex@-7 2000T FPGA Enabled by SSI Technology

Originally developed for use in a variety of die-stacking design methodologies, silicon interposers provide modular design flexibility and high-performance integration suitable for a wide range of applications. The silicon interposer acts as an interconnect vehicle based on a silicon manufacturing process (e.g., a 65 nm or 45 nm process) on which multiple die are set side by side and interconnected. SSI technology avoids the power and reliability issues that can result from stacking multiple FPGA dies on top of each other or an MCM. Compared to organic or ceramic substrates, which are often

used in MCMs, silicon interposers offer far finer interconnect geometries (approximately 20X denser wire pitch) to provide device-scale interconnect hierarchy that enables more than 10,000 die-to-die connections.

Creating FPGA Die Slices with Microbumps for Stacked Silicon Integration

The foundation of Xilinx SSI technology is the company's proprietary ASMBL architecture, a modular structure comprising Xilinx FPGA building blocks in the form of tiles that implement key functionality such as configurable logic blocks (CLBs), block RAM, DSP slices, SelectIO™ interfaces, and serial transceivers. These resources are organized into columns and then combined to create an FPGA. By varying the height and arrangement of columns, an assortment of devices can be created to match different market requirements (Figure 2). The FPGA contains additional blocks for generating clock signals and for programming the SRAM cells with the bitstream data that configures the device to implement the end user's desired functionality.




Figure 2: Representation of an FPGA Built with ASMBL Architecture

Starting with the basic ASMBL architectural construct, Xilinx has introduced three key modifications that enable stacked silicon integration (see Figure 3). First, each die slice receives its own clocking and configuration circuitry. Then the routing architecture is modified to enable direct connections through the passivation on the surface of the die to routing resources within the FPGA's logic array, bypassing the traditional parallel and serial I/O circuits. Finally, each SLR undergoes additional processing steps to fabricate microbumps that attach the die to the silicon substrate. It is this innovation that enables connections in far greater numbers, with much lower latency, and much less power consumption than is possible using traditional I/Os (100X the SLR-to-SLR connectivity bandwidth per watt versus standard I/Os).





Figure 3: FPGA SLR (Die) Optimized for SSI Technology

Silicon Interposer with TSV

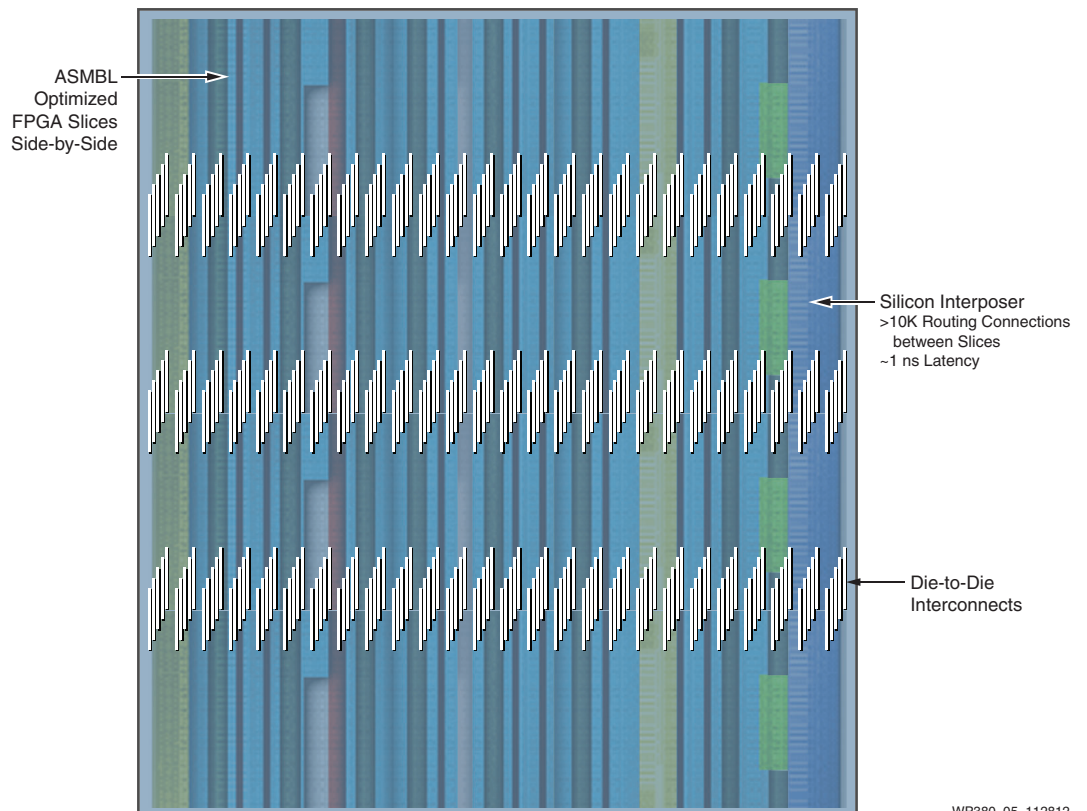
The passive silicon interposer interconnects the multiple FPGA SLRs together. It is built on a low-risk, high-yield 65 nm process and provides four layers of metallization for building the tens of thousands of traces that connect the logic regions of multiple FPGA die (Figure 4).



WP380_04_101411

Figure 4: **Passive Silicon Interposer**

Figure 5 illustrates the concept of an “X-ray view” of the assembled die stack. It contains a stack-up of four FPGA SLRs mounted side by side on a passive silicon interposer (bottom view). The interposer is shown as transparent to enable a view of the FPGA SLRs connected by traces on the silicon interposer (not to scale).



WP380_05_112812

Figure 5: **“X-ray View” of a Virtex-7 FPGA Using SSI Technology**

The TSVs combined with controlled-collapse chip connection (C4) solder bumps enable Xilinx to mount the FPGA/interposer stack-up on a high-performance package substrate using flip-chip assembly techniques (see [Figure 1](#)). The coarse-pitch TSVs provide the connections between the package and the FPGA for the parallel and serial I/O, power/ground, clocking, configuration signals, etc.

Comprising numerous patent-pending innovations, this SSI technology provides multi-Terabit-per-second die-to-die bandwidth through more than 10,000 device-scale connections—enough for the most complex multi-die designs. Xilinx is using this new technology to enable several members of the Virtex-7 FPGA family.

SSI Technology with Heterogeneous Die

In addition to integrating homogeneous SLRs on a silicon interposer, SSI technology also enables the integration of different types of die. In [Figure 6](#), the Virtex-7 H870T FPGA ties together three SLRs as well as separate 28G capable transceiver circuits via the silicon interposer. Because the SLRs and 28 Gb/s transceiver circuits represent different silicon processes and functions, the Virtex-7 HT FPGA is the world's first heterogeneous architecture—an FPGA consisting of heterogeneous die placed side-by-side to operate as one integrated device.




Figure 6: Heterogeneous 3D FPGA with Integrated 28G Transceivers

One of the key benefits in physically separating the digital FPGA from the transceivers is noise isolation. This ensures the lowest possible jitter and noise to simplify design closure and reduce board cost.

Separating the 28G transceivers from the SLRs is an example of how a heterogeneous architecture can achieve optimal results for specific applications. Because transceivers are complex analog circuits, implementing them on a monolithic device would require a more complex design approach. As a separate slice, the 28G circuit is tailored for the greatest possible capacity and best possible performance and power without compromising the digital logic's features.

The other benefit of a heterogeneous architecture is the ability to provide different ratios of transceivers to traditional FPGA resources. With up to sixteen 28G transceivers, the Virtex-7 HT FPGA enables unprecedented integration and is at the forefront of high bandwidth design.

The Virtex-7 Family

The SSI-enabled devices shown in [Table 1](#) offer unprecedented FPGA capabilities. These devices provide up to: 2,000,000 logic cells; 68 Mb of block RAM; 5,335 GMACs of DSP performance; 1,200 SelectIO pins supporting 1.6 Gb/s LVDS parallel interfaces; and 2,784 Gb/s aggregate bidirectional bandwidth.

Table 1: Virtex-7 FPGAs

FPGAs	Part Numbers					
Virtex-7 T	XC7V585T	XC7V2000T				
Virtex-7 XT	XC7VX330T	XC7VX415T	XC7VX485T	XC7VX550T	XC7VX690T	XC7VX1140T
Virtex-7 HT	XC7VH580T	XC7VH870T				

Notes:

1. The shaded devices are based on SSI technology.

FPGA Design with SSI Technology

With SSI technology, the designer creates and manages a single design project. This is an extremely important advantage because partitioning a large design across multiple FPGAs presents a number of complicated design challenges that do not apply to monolithic implementations.

The typical steps in a monolithic FPGA design flow include:

- Create a high-level description
- Synthesize into an RTL description that matches the hardware resources
- Perform physical place and route
- Estimate timing and adjust design for timing closure
- Generate a bitstream to program FPGAs

When working with multiple FPGAs, the designer (or design team) must partition the netlist across the FPGAs. Working with multiple netlists means opening and managing multiple projects, each with its own design file, IP libraries, constraint files, packaging information, etc.

Timing closure for multiple FPGA designs can also be extremely challenging. Calculating and accommodating propagation delays through the board to the other FPGAs poses new and complex problems. Likewise, debugging a design through multiple partial netlists in multiple FPGAs can be extremely complicated and difficult.

In contrast, SSI technology routing is transparent to the user. The user performs a single design bring-up and debug with a standard synthesis and timing closure flow. To accelerate integration and implementation for devices with this capacity (over 2 million logic cells), Xilinx introduced the Vivado™ Design Suite—a development environment designed to support current and future high capacity devices.

Applications

Xilinx Virtex-7 FPGAs with SSI technology break through the limitations of monolithic FPGAs, extending their value in some of the most demanding applications. For example, the Virtex-7 family is ideal for next-generation telecom and networking systems, where dozens of serial transceivers are leveraged to enable a flexible,

single-FPGA solution. These devices are also perfect for use in ASIC prototyping and can serve as a pre-production and /or initial-production ASIC alternative. The Virtex-7 family also enables flexible, scalable, customized high-performance computing solutions for scientific, oil and gas, financial, aerospace and defense, and life science applications. The parallelism inherent in the FPGA architecture is ideal for high-throughput processing and software acceleration. Support for a multitude of high-speed parallel and serial connectivity standards enables the convergence of compute and communications systems. In aerospace and defense, high transceiver count and thousands of DSP processing elements provided by FPGAs with SSI technology enable advanced RADAR implementations.

SSI Technology - Concept to Reality

The development strategy Xilinx employed in the creation of SSI technology began with extensive modeling and the subsequent creation of a series of test devices, or test vehicles, used for design enablement, manufacturability, and reliability validation.

These test vehicles and stress simulation models showed an additional advantage of stacked silicon technology. The silicon interposer functions as a buffer that reduces low-K dielectric stress and improves C4 bump reliability, compared to monolithic solutions.

Extensive simulations and investigations of the thermal impact of the die stack showed that thermal performance of devices with SSI technology is comparable to that of monolithic devices.

The culmination of nearly six years of extensive research and development was realized in September 2011 when Xilinx shipped the world's highest capacity FPGA, the Virtex-7 2000T device, enabled by SSI technology. In May 2012, Xilinx then shipped the world's first heterogeneous device, the Virtex-7 H580T, featuring 28G transceivers targeting Nx100G wired communications applications (see Xilinx Press Release: <http://press.xilinx.com/phoenix.zhtml?c=212763&p=RssLanding&cat=news&id=1700586>). Xilinx has a robust supply chain in place to mass produce FPGAs with SSI technology. TSMC, Amkor, and Ibiden contribute their combined resources and expertise for fabricating 28 nm FPGAs and 65 nm silicon interposers, interconnect layers, microbumps, C4 balls, and package substrates as well as performing wafer thinning, die separation, chip-on-chip (CoC) attach, and package assembly. Xilinx is shipping devices based on SSI technology in volume to customers today.

Summary

As the only FPGA manufacturer to use SSI technology for FPGAs with super high capacity and transceiver bandwidth, Xilinx is breaking important new ground in the system-level integration arena. SSI technology enables Xilinx to deliver the highest logic density, bandwidth, and on-chip resources with the fastest ramp to volume production at every process node.

Designing with FPGAs enabled by SSI technology is significantly easier than the alternative. Flexible tool flows enable design closure automation, yet allow user interaction for achieving even higher performance.

Xilinx is now shipping the world's highest capacity FPGA, the Virtex-7 2000T device, as well as the world's first heterogeneous FPGA, the Virtex-7 H580T—both enabled by SSI technology. For more information, visit www.xilinx.com/virtex7.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
10/27/10	1.0	Initial Xilinx release.
10/21/11	1.1	Updated Introduction . Updated Xilinx SSI Technology , including all figures. Added Table 1 . Updated and interchanged Applications and SSI Technology - Concept to Reality . Updated Summary .
12/11/12	1.2	Added SSI Technology with Heterogeneous Die . Replaced Figure 1 . Updated Creating FPGA Die Slices with Microbumps for Stacked Silicon Integration , Figure 5 , Figure 6 , Table 1 , Applications , and Summary .

Notice of Disclaimer

The information disclosed to you hereunder (the “Materials”) is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available “AS IS” and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.