sponsored by

Gabe on EDA • EDAmarket

Assembling the Future

A Newsletter About the Design and Production of Electronics

ISSUE 014 • APRIL 2012

SUBSCRIBE

PDF and Archives

twitter

In this issue:

Standards: Past, Present, and Future

Gabe Moretti

Standards have, are, and will play a major role in EDA. The difficulties inherent in developing semiconductors require the existence of point tools that provide specific solutions to very demanding problems. These tools must be able to work together in an integrated flow that needs the ability to exchange data in a reliable manner. In addition engineers need to be able to exchange tools from various vendors for both technical and commercial reasons. All of these would not be possible without standards.

Past

Our industry recognized the need for standards even before it became Electronic Design Automation. It was still called CAD, for Computer Aided Design, when, shortly after the introduction of proprietary workstations, users found the need to exchange schematic drawings among various workstations. EDIF was the result of this standardization exercise, but we soon discovered that a standard by itself is not very useful. All the workstation companies quickly develop EDIF readers, but you could not find an EDIF writer from any of them. Reality has a way of asserting itself, though.

EDIF 1.0 became a format for storing designs that could be reused and thus, ultimately ported to a new system.

Since then, the EDA industry has develop dozens of standards ranging from design entry languages, to synthesis syntax and semantics, to various physical description formats and languages. One standard that did not start out as an industry standard, but became one by default is the GDS-II format. It was developed by Calma in order to drive the Gerber photo plotting equipment used to prepare masks for photo lithography. It was never meant to last over forty years, but it is still used today. There have indeed been attempts at developing a more modern format with syntax and semantics extensions to deal with today's requirements, but GDS-II is still the format of choice.

Some standard are developed from scratch following a requirements document, VHDL is the most famous of them, while others are developed starting with a popular existing solution, like Verilog for example. It is interesting that the requirement document for VHDL did not ask for a design language, but for a documentation language. Once the standard was complete, a company Vantage Computer Systems decided that the language could be simulated, and thus used to develop an electronic system. it was a good intuition that created a market that is still quite robust today.

In the eighties and early nineties the Design Automation Standards Committee (DASC) of the IEEE was the major body for the development of EDA standards. During the present century the DASC has lost some of its importance as a development body and that role has been taken up by industry consortia. The two most active in standard development are Si2 and Accellera, now renamed Accellera Systems Initiative.

Present

The two organizations have different operating methods and as a result, the Si2 standards are seldom offered to the IEEE for standardization, while all of the standards developed by Accellera are submitted to the IEEE and follow this organization approval and maintenance procedures.

Si2 has five major projects and another two activities that support projects that have terminated their engineering work and are ready to address new industry requirements as the result of evolving technology. These are the Open Modeling project and the LEF/DEF distribution of the specification.

Certainly the most popular Si2 project is the OpenAccess Coalition that this year celebrates its tenth anniversary. Started with a donation by Cadence of an API for its data base, this work is directly or indirectly responsible for many startups in the industry. Startups that targeted specific engineering areas within the EDA flow found a ready made integration path that would support the adoption of their specific product within an existing design flow, as long as it used the Cadence data base. Although it is not clear whether or not Cadence ever benefitted from this strategy, once the offer was made and the technology was part of Si2, it was too late to rethink the strategy.

It can be said, though, that Cadence direct competitors, Magma, Mentor, and Synopsys, have not adopted the OpenAccess data base as their internal data format, nor have they lost significant business as the result of their engineering choice.

Two other projects, Design for Manufacturability and OpenPDK, address back end flow requirements that have attracted significant attention by both foundries and EDA vendors. As

manufacturing technology progresses, the OpenPDK Coalition will see significant obstacles to its work as foundries find it necessary to develop methods that are more and more proprietary due to the manufacturing requirements at 20 nm and below.

The newest project within Si2 is the Open3D Technical Advisory Board. One promising area of system integration is the use of 3D structures within a package. For an industry that has thought in two dimensions for practically its entire existence, architecting and developing a system in three dimensions presents very hard problems. Architectural design is critical, since moving blocks from one layer to another will be very expensive. The problem that the industry faces has many dimensions, and the one that Si2 has not been very effective at solving is the business side of the issue.

For reasons that are at the core of Si2, its specifications do not have the power and wide acceptance of international standards, like those of the IEEE. As a result, for example, its Common Power Format (CPF) developed by the Low Power Coalition, generated confusion and additional work in the industry when the Unified Power Format (UPF) developed by Accellera in order to free the industry from the Si2 requirements became the IEEE 1801 Standard.

Accellera became an organization in the first quarter of 2000, the result of the merger of OVI and VI. These two organizations had realized that their missions, centered respectively on Verilog and VHDL had reached the end of their missions and that the industry needed a wider approach to the problems of IC design. As a result of its recent mergerwith OSCI, Accellera has enlarged its mission to cover all aspects of system level design. In addition to its support for six IEEE standards that originated from its technical committees, Accellera has thirteen working groups and committees working on IP issues, Verification proposals that cover analog, IP, coverage measurements and of course various aspects of the SystemC community. May be surprisingly, VHDL support and extension work has remained with the IEEE, as Accellera continues to focus on both SystemVerilog and SystemC.

Both organizations perform a positive role within the industry. But their roles must be understood within the limits of each. Cadence has always played a role within Si2 that is much greater than any other consortium member, while Accellera is "managed" by Synopsys, Mentor, and Cadence who have invested staff resources to direct the consortium since its inception.

Future

Future standards development must address complexity that arises from both ends of the IC development flow. At the front end, the issues that demand attention are reuse and heterogeneity. At the back end the process of fabricating and packaging such large systems as effectively as possible to contain size, power consumption, and costs.

Although IP reuse is now common, it is still more costly and difficult than it should be. Lack of standards is the major contributor to this situation. The industry needs standards for the documentation of IP cores, that will support the design and verification of cores from multiple suppliers as well as portability across manufacturing processes. The latter may become a mute point soon as 20 nm and especially 14 nm processes become more foundry dependent. This will require that every third party IP used in these processes be guaranteed by the foundry up front.

Heterogeneity encompasses digital/analog, hardware/software, and electrical/mechanical issues. Systems will frequently have to handle all three at the same time requiring very careful architecting of the system. Although a few tools are becoming available, much work still needs to take place. Just as an example, after years of discussions and work, we still do not have a strong, reliable standard for analog/mixed signal design or verification language. And this with a problem that is totally within the control of our own industry!

The industry has made some progress in the area of hardware/software integration and verification. The most used standard in this area is TLM, developed jointly by OSCI and Accellera, then still separate entities. But almost nothing is available for hardware/software co-design that would allow engineers to explore at the architectural level the implications of moving functional blocks from hardware to software and vice versa.

Surprisingly not much is talk about in EDA circles about electro/mechanical integration. This is spite of the fact that MEMS have been around for quite a few years, and that aerospace, automotive, and even consumer markets use MEMS extensively. What is available relegates MEMS to the analog domain but very little is done in the area of system level simulation and verification.

In the manufacturing and packaging area, very large SoC present two different, yet related problems. One has to do with size. Is it reliable and cost effective to try to package the entire system on one die? If not how can the system be partitioned and packaged? In addition there are issues with respect to the best process technology to use with different parts of the system. Memories, for example, could be fabricated with a more advanced process than digital logic, ensuring an overall better yield. And analog generally is not well suited to the most advanced processes, yet it should not be the determining factor when picking a process. The area of 3D packaging will certainly see significant development in the next two or three years, but standard development should proceed in parallel, or single source tools will slow progress.

One factor that complicates the development of standard is that the industry is no longer US centric. Developing countries, like China and India, will play a significant role in the development of standard. if not, we will no longer have world-wide standards, but the industry will find itself facing local standards regulating large, significant markets. this will increase the cost of developing, marketing, and supporting electronic products.

The Move to IP Standards

Josh Lee, President and CEO, Uniquify Inc.

In last month's issue of Assembling the Future, I wrote about the business of IP, a business that's moved impressively from the seemingly Wild West approach of its formative years to become a real business with viable business models. IP vendors have taken a leap forward in business acumen by

supplying a solutions-based approach to IP and become successful to wit.

In this issue of Assembling the Future, I'll cover IP standards, something that may be far more critical than our community realizes. We may soon be forced to modify our business models as Systems on Chip (SoC) morph into Super Systems on Chip (SSoC) with an increasing amount of IP.

Before going any further, let's compare SoCs and Super SoCs, and attempt to differentiate them. A typical SoC includes an on-board processor, memory, a special-purpose processing engine or engines, peripheral interfaces and I/O. In our definition of Super SoC devices, they have multiple processors, special purpose processing engines, on-chip networks, multiple memories and many different peripheral and I/O functions. These super-sized devices are being fabricated at 40nm process technologies, with the latest generation being implemented in 28nm and moving to 20nm.

Likely, we will soon see a new paradigm in SSoC design driven by the automated selection and configuration of IP. This new approach, based on specifications supplied by the design team, would assemble a prototype system by selecting and combining IP with specialized algorithmic blocks developed the team.

A recent study concluded that the IP market is growing faster than the semiconductor market at large — twice as fast, in fact, a strong indication that IP is in demand for the creation of SSoCs. Yet another report suggests that close to 60 to 70% of a super chip is composed of IP. What's more, the responsibilities of SSoC designers have shifted from an expertise in logic design to an understanding about IP selection, qualification, implementation and verification.

And, why not? One super device can contain upwards of billions of transistors. By using IP for common blocks and functions, a design team can implement a cost-effective building-block strategy to save time, while allowing them to focus on value-added hardware features to differentiate the chip.

With this shift to so much more IP on chip, our ecosystem may well be thrown into unfamiliar terrain. We will need to concern ourselves with the overall design, verification and implementation of an SSoC, added nuance and complexity few of us are prepared to manage.

It's almost an understatement to say that IP providers will need to pay more attention to IP standards as we move into the era of SSoC design. Indeed, we may need a standards organization to sort out a well-ordered and structured standardization effort, similar to what EDA has put in place with Accellera, a feeder channel into IEEE.

Of course, there are two other standards groups worth mentioning. OCP-IP (for open core protocol-intellectual property), for example, is dedicated to proliferating a common standard for IP core interfaces that facilitate "plug and play" SoC design. It has a healthy roster of corporate support, including a set of standards for the OMAP 2 platform from Texas Instruments.

And lest we forget VSIA (for Virtual Socket Interface Alliance) an early attempt to develop IP standards founded in 1996 and dissolved in 2008. Its mission to enhance SoC designer productivity through an international standard for measuring quality and examining the practices used to design, integrate and support the IP was noble but a bit early to be successful.

Standardization will play a big role, especially with interfaces between IP blocks interconnected in the SSoC, an important consideration when IP comes from multiple vendors. For example, standards could be developed to coordinate the interface definition, physical interface and timing requirements.

Like so many other industries or market segments, IP is moving into a new phase that makes it more important and predominant, but it also creates new challenges. Standards and standardization is all part of growing into a new phase. How we manage it is up to us.

About Josh Lee

Josh Lee is president and CEO of Uniquify, a leading semiconductor IP and ASIC/ SoC and SSoC design and manufacturing services provider. He holds a Bachelor of Science degree in Electrical Engineering and Computer Science from the University of California at Berkeley.

New Protocols Shaping the Future of Mobile SoCs

Hezi Saar, Marketing manager, Synopsys

The Mobile Industry Processor Interface (MIPI) Alliance establishes specifications for hardware interfaces in mobile devices. Its specifications focus on eliminating proprietary, legacy, often point-to-point or parallel interfaces and promote standardization for system integrators. The MIPI Alliance defines a wide range of chip-to-chip interfaces that are “inside the mobile device,” as opposed to external interfaces such as HDMI and USB. MIPI protocols such as Low-latency interface (LLI), CSI-3, DigRF v4, and DSI-2 protocols as well as USB SSIC and JEDEC UFS are all designed to improve interoperability and reduce power consumption, pin count, and integration costs.

The MIPI Alliance has over 200 members (and growing), including key vendors to the mobile industry, who participate in many working groups and co-author some of the specifications.

The members define and promote interfaces that drive consistency in processor and peripheral interfaces, promote reuse and compatibility in mobile devices, and facilitate product innovation.

MIPI specifications enable multiple chip-to-chip interfaces serving different purposes. On the application processor side, interfaces for display, camera, storage and high-speed communication are found. On the baseband processor side, there are high-speed connectivity and baseband-to-RFIC interfaces.

Market momentum for MIPI interfaces is strong. Analysts from InStat and iSupply project strong adoption of MIPI camera (CSI-2) and display (DSI) interfaces in mobile electronics including mobile computing applications such as tablets. IPNest indicates that almost all worldwide IC shipments in handsets will be using MIPI interfaces by 2013.

There are many applications utilising MIPI-enabled standard connectivity. Mobile multimedia and wireless connectivity devices are the natural fit for mobile interfaces that enable standard chip-to-chip connectivity and higher scalability while minimising power.

It is likely that there will also be penetration of MIPI-enabled interfaces in home, computing, industrial and even telematics applications utilizing the low power and low system cost profile of the MIPI specifications and their wide adoption.

Popular MIPI Protocols

The mobile market demands high performance, low power and scalable interfaces to enable manufacturers to easily upgrade their mobile device platforms and avoid costly re-design.
There are a number of MIPI display and camera specifications that enable standard connectivity to the application processor:

- Display serial interface (DSI): Allows standard, efficient and low power connectivity between the application processor and display that is embedded in the mobile device.

- Camera serial interface (CSI-2): An efficient low power, low pin-count interface that connects the image sensor to the application processor. The CSI-2 protocol provides the functional means to transfer the data between the devices, including detecting and correcting errors that may occur during transmission. A separate I2C-compliant interface is used for camera control interface (CCI) functions. The CSI-2 interface supports a variety of data and color space formats, transfer in packets, lane management, error detection and correction.

- D-PHY: Serves as the electrical and physical connection allowing the implementation of the CSI-2 or DSI protocol between the application processor to the display or image sensor. It provides a high-performance serial differential interface offering up to four data lanes, plus a common differential clock lane meeting the need to connect to multi-megapixel cameras and high resolution displays to the application processor.

Next-Generation MIPI Protocols

The next-generation MIPI protocols will be used as the foundation of several mobile applications to support the needs for higher resolution and increased bandwidth, utilizing the M-PHY physical layer. The most common use of the M-PHY physical layer today is with the DigRFv4 controller, which is used to communicate with the standard baseband RFIC. Design teams can also use the M-PHY physical layer to implement a number of other application-specific and application-agnostic protocols.

M-PHY: The Foundation for Mobile Protocols

The M-PHY protocol uses differential signaling and an embedded clock in a dual simplex lane configuration, and provides a feature set that's optimized for a spectrum of mobile applications.

Achieving low electromagnetic interference (EMI) is extremely important in mobile electronics – the devices are small and their components closely packed, creating an environment that is very sensitive to RF signals. The M-PHY features help to lower EMI and reduce interference through:

- Low amplitude option
- Programmable slew rate
- Limited common mode noise
- Dual bandwidth rates A and B for each high-speed gear

Because the M-PHY interface supports a wide data bandwidth range, design teams can operate it in high-speed mode to avoid interference and reduce noise in the device, while achieving high performance and optimal low power system operation.

To see M-PHY silicon in action, view the video below which showcases large and small amplitude measurements taken from Synopsys' silicon proven DesignWare M-PHY.

Let's review the application-specific M-PHY-based protocols:

Low-latency interface (LLI)

LLI is primarily used for cache refill of a companion chip using the application processor attached memory. This means the interface based on the M-PHY, with multiple lanes in each direction, needs to be simple and scalable to achieve the latency-sensitive chip-to-chip implementation. There are other interesting use cases for this chip-to-chip capability such as connecting companion chips to an application processor or SoC, or transmitting data over non-proprietary but standard interfaces.

LLI-optimized implementation requires transferring data as quickly as possible to achieve the lowest system power, so a high-speed gear like HS-Gear3 is recommended.

UniPro enabling UFS, CSI-3 and DSI-2

UniPro, an application-agnostic protocol that is also based on M-PHY for the physical layer, supports the following applications:

- Universal flash storage (UFS): JEDEC signed a memorandum of understanding with the MIPI Alliance to develop this storage interface based on MIPI M-PHY and MIPI UniPro specifications. JEDEC aims for its UFS standard to be the storage interface in mobile and consumer electronics to enable effective storage and consumption of rich multimedia. The UFS standard is currently available and published by JEDEC.
- CSI-3, the next-generation camera interface, and DSI-2, the next-generation display interface (specifications are under definition). Display and imaging technologies like 3D and high-resolution continue to push the boundaries of bandwidth for image sensors and displays. CSI-3 and DSI-2 are targeting to meet these requirements and add to the capabilities of today's CSI-2 and DSI protocols that are based on D-PHY

DigRFv4

As higher air traffic demand continues to increase, we see more 3G and 4G/LTE networks being deployed and supported by wireless carriers. As a result, more manufacturers are introducing 4G RFICs that support high bandwidth transfers. The typical non-standard implementation could use many pins to implement a 4G data connection from the baseband processor and RFIC, and does not provide the scalability needed for the breadth of mobile electronics flavors.

The MIPI DigRF specification defines a high-speed serial interface between RF transceiver ICs and baseband processors, meeting the increased data throughput requirements for mobile terminals. DigRFv4 is the latest specification targeting 4G standard air interfaces such as LTE and mobile WiMAX, as well as supporting existing 3G standards. It is a high-bandwidth standard interface that makes use of the scalable MIPI M-PHY to support existing and next generation mobile broadband requirements, simplifying system integration as well as reducing power, pin-count, EMI and overall cost.

SuperSpeed inter-chip (SSIC)

The USB 3.0 promoters group signed a memorandum of understanding with the MIPI Alliance allowing the USB working group to develop this USB chip-to-chip protocol, a standard still under definition. This collaboration supports low power and high-speed implementations that allow reuse of the USB drivers to support the widely-used USB 3.0 standard interface. The generic SSIC chip-to-chip interface combines the Super Speed high speeds and the low power of M-PHY to create a flexible and reusable approach allows chip-to-chip connectivity between host and peripherals.

Summary

These new interfaces require high-speed and low-power SERDES implementation. Finding the balance between high-performance and low-power modes allows you to optimize your total system power consumption. Future proofing your design by using high-speed MIPI protocols ensures that your design is ready for tomorrow's mobile and consumer market needs.

MIPI Alliance: www.mipi.org

DesignWare MIPI IP Solutions: www.synopsys.com/mipi

Synopsys videos: http://www.synopsys.com/ip/pages/videos.aspx

A version of this article first appeared in ElectronicsWeekly.com on 05 March 2012: http://www.electronicsweekly.com/Articles/05/03/2012/53130/synopsys-describes-mipi-digrf-protocol-for-4g-mobile.htm