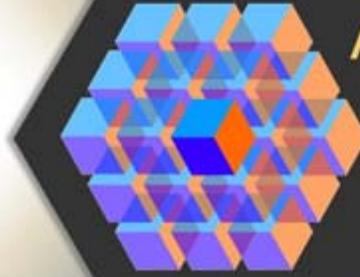


sponsored by



Gabe on EDA • EDAMarket • EDA Confidential



Assembling the Future

A Newsletter About the Design
and Production of Electronics

ISSUE 002 • FEBRUARY 2011

SUBSCRIBE

PDF and Archives

twitter

In this issue:

- [From the Editor](#)
by Gabe Moretti
- [IP Reuse, is it Really That Easy?](#)
by Mike Gianfagna, VP Marketing, Atrenta
- [Why Antifuse OTP NVM Is a Cost-Effective Alternative to Flash, EEPROM and ROM at 65nm and Below](#)
by Bernd Stamme, Kilopass
- [Components, Software, and Glue Logic. Is this all that is needed?](#)
by Gabe Moretti

From the Editor

Gabe Moretti

This issue deals with Intellectual Property. The article by Mike Gianfagna of Atrenta "IP Reuse, is it Really That Easy?" shows that design reuse has been on the mind of leading EDA professionals for some time and how proper documentation plays a critical role in enabling such reuse.

The second contributed article by Bernd Stamme of Kilopass describes a less known form of memory IP. Its title: " Why Antifuse OTP NVM Is a Cost-Effective Alternative to Flash, EEPROM and ROM at 65nm and Below" describes the technology of the device that is growing in popularity as IDMs find more uses for its technology.

Finally, my article looks at IP, its growth, and the obstacles still to be overcome to make its use less costly. The future of the IP market is changing while its growth seems assured. But existing

standards are not used as they should and thus the costs of using IP is not diminishing as it should.

This newsletter, now in its second issue, has been well received by its intended readers, with over 34% of the mailing list actually opening the issue and reading at least one article. I invite all of you to publicize its existence among your colleagues and to contribute to its contents. Every type of article is welcome, and length is not a deciding factor. The newsletter is published monthly. The topics through July are:

- When, not if, will RTL hands-off be the norm?
- From portable to wearable computing: how far is the jump?
- Off-shoring: the Good and the Ugly.
- How is DAC Helping to Shape the Future?
- Impact of Formal Verification: Are We Close to "Correct by Construction"?

If you would like to sponsor the newsletter or advertise in it get in touch with me at gmoretti@gabeoneda.com and I will send you the rate sheet.

IP Reuse, is it Really That Easy?

Mike Gianfagna, VP Marketing, Atrenta

There are some events in your life that somehow stick in your mind. I can remember this one quite vividly. I was sitting in a local bar down the road from where I worked with the manager of a digital design group at the same company. He had just finished the release of a new family of chips, essentially creating a complete data book of functionality that could be “mixed and matched” to build new PC board-based designs. We mused about the tremendous effort involved in building that series of chips. All those SPICE simulations to characterize the process and calibrate the cell library. The countless hours spent in layout design and the seemingly endless DRC runs.

After we finished the first pitcher of beer, an idea came to us, almost at the same time. What if we could somehow reuse all this work to create a new level of product offering? Could we simply strip the bond pads off the newly released designs and, with some supporting software development, create an environment whereby our customers could choose to either use the discrete parts or develop their own level of integration by mixing and matching the core silicon from the parts library onto a structured custom chip design of their own choosing? We coined a name for this idea – Silicon Circuit Board, and we spent the next few months trying to make the idea come to life.

The time was 1983, the place was Somerville, New Jersey and the company was RCA Solid State. The chip series under discussion was a new CMOS implementation of a popular LS-TTL logic family. The idea was, well, a little ahead of its time. We tried for about the next year to make Silicon

Circuit Board a reality. We hit lots of speed bumps along the way and in the end we abandoned the project. The details of what caused us to stop work on this idea is the subject of another discussion, but let's just say that when you start with a library that wasn't built with reuse in mind, it's hard to get the job done. We figured this out almost 30 years ago.

And we're still trying to get it right today. The idea of simply integrating semiconductor IP on a silicon substrate sounds seductively simple. The reality is very different though. As we discovered then, and continue to discover today, making semiconductor IP reusable is not an event but a journey. The design needs to be built with that end in mind, right from the very beginning. Consistent formats, extensive documentation, complete simulation data, verification suites, timing constraints, power specifications and solid interface definitions are just the start. Even if the IP is perfect in every way, integrating it onto a silicon substrate poses substantial new challenges when compared with a PC board design. Leakage, on-chip drive requirements, routing delays, lithographic proximity effects and local heating, just to name a few, can get you. If you consider the new die stacking options that are available, thanks to through silicon via technology, things get much more complex.

Also consider that PC board designers are using static, known configurations of each part. Semiconductor IP, on the other hand, can also be delivered in soft, synthesizable form. Exactly what you are getting is often discovered during the physical implementation phase, when it's usually too late to change your mind about anything material. A good scheme to objectively determine the completeness, robustness and integration risks associated with soft semiconductor IP is an important requirement in the current system on chip development marketplace. The good news is that there are solutions to this problem out there, and companies like TSMC are taking the initiative to develop standards for soft IP delivered quality. Atrenta is proud to be part of that effort. Maybe we can learn from the past.

If you spend enough time in semiconductors and EDA, you realize how cyclic the world really is. I'm sure other industries are the same. Old problems don't go away, they just change form a bit and come back to pose another round of challenges. Back to RCA and 1983... The company had recently launched a merchant application specific integrated circuit (ASIC) business. We called it "semicustom" in those days and we had been doing it internally for many years. Now we were ready to take on the likes of LSI Logic with a commercial offering. The closets in my house are littered with memorabilia from 3 decades of trade shows and giveaways. I found this one recently:



We call it “system on chip” today; we called it “systems to silicon” then. Looking at this tote bag reminded me of a song that was written and performed by the late Peter Allen, “Everything Old is New Again”. The premise of the song seems particularly relevant here. To quote, with attribution to Peter W. Allen and Carole Bayer Sager:

Don't throw the past away
You might need it some rainy day
Dreams can come true again
When everything old is new again

What's in your bag of tricks for IP reuse? Always remember to look back before you look forward.

Why Antifuse OTP NVM Is a Cost-Effective Alternative to Flash, EEPROM and ROM at 65nm and Below

Bernd Stamme, Kilopass

As process geometries shrink more functionality gets integrated on-chip. The block diagram of any system-on-chip (SoC) today resembles that of a complex printed circuit board a decade ago. The one function that has resisted this on-chip integration is code storage for on-chip embedded processors. For applications in which the cost of additional mask layers to build embedded flash or

EEPROM onto the standard CMOS process is not cost effective or even available, this function has remained in off-chip flash. For 65nm process node and below, the only alternative to off-chip flash was on-chip masked ROM.

At the 65nm process node and below another solution is beginning to take hold: antifuse one-time-programmable (OTP) non-volatile memory (NVM). Because antifuse OTP NVM is fabricated in standard logic CMOS process and can easily scale to smaller geometries, this technology is starting to offer a more cost effective alternative to masked ROM at smaller process nodes.

What is Antifuse NVM?

Antifuse NVM technology differs from other NVM technologies in that it uses gate oxide breakdown to provide bit storage. During programming an elevated voltage is applied to the gate of the programming transistor in the bitcell, while the source of that transistor is held at a low voltage. This results in a breakdown of the transistor's gate oxide turning what once functioned as a capacitor that blocked current flow into a resistor that enables current flow: a data zero turned into a data one.

Because antifuse OTP NVM is implemented in standard CMOS, its bitcell is small and can be programmed on the wafer or after being packaged. Another advantage of antifuse is that its programmed state is not readable when probed with a scanning electron microscope during reverse engineering attempts. And for high security applications, the memory can be altered to obscure its original content.

The 65nm Tipping Point

Antifuse OTP NVM in larger process geometries is less attractive due to the cost of programming memory on expensive automatic test equipment (ATE) during wafer test or final package test. For example, at the 180nm process node the cost of programming a megabit of data is around \$0.10. For mobile SoCs that employ RF testers the cost can be twice to three times greater. Likewise die cost for a megabit of data at the 180nm node can cost from \$0.04 to \$0.08. On a small die this can represent a significant percentage of die area cost.

At 65nm, these problems are eliminated because the bitcell scales with the shrinking of each transistor on chip not only in size, but also in gate oxide thickness. As seen in Figure 1, a bitcell in 65nm process is nearly half the size of the same bitcell in an 180nm process. But it is the reduction in gate oxide thickness at 65nm over 180nm that makes the economics for antifuse OTP NVM compelling. The reduction in gate oxide thickness lowers programming voltage from 8.5V to 5V and cuts programming time an order of magnitude, thus drastically reducing ATE time.

To illustrate the significance of this savings, one chip company building a wireless communications baseband and RF module chip implemented antifuse OTP NVM to replace an external EEPROM for storing 128Kb of boot code. Eliminating the EEPROM decreased system cost by \$0.10 per unit. In production since 2008 with an annual build of tens of millions of units, the company has realized several million dollars in savings each year.

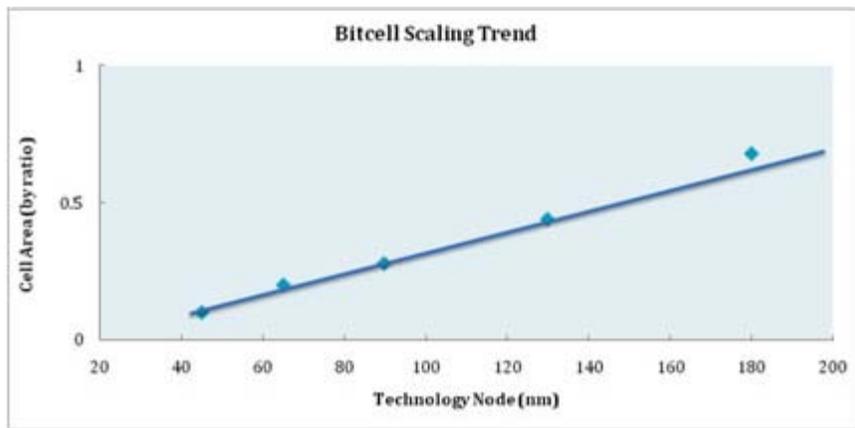


Figure 1: Bitcell Scaling Trend.

On-Chip Program Storage Advantages

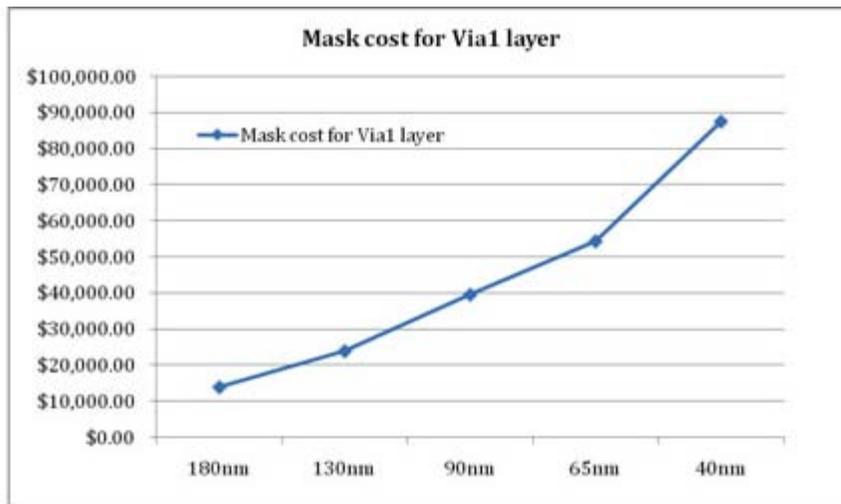
Moving on-chip program code and configuration data previously stored in off-chip flash has the obvious advantage of trimming bill-of-material (BOM) cost. The other advantage of antifuse OTP NVM is its ability to provide execute-in-place (XIP) capability for microcode. This functionality is provided by off-chip flash solutions using multiplexed serial I/O that can achieve data rates to the on-chip CPU on the order of 400Mb/s.

Without compromising power and area to accommodate multiplex I/O, on-chip antifuse OTP NVM can achieve at least 4X performance improvement over the external XIP flash alternative. Being on-chip, the antifuse OTP NVM employs a wider data bus to achieve these significantly higher data rates to the CPU.

Using OTP in Few-Time-Programmable (FTP) Applications

There are many small form-factor, low-power applications that require 512kb (48 KB) of program storage that undergo less than three code revisions. Examples include fixed and wireless broadband (LAN, RF, Wi-Fi, WiMax, LTE, Bluetooth...), MCUs, media processors, among others. For these requirements, antifuse OTP NVM with a storage capacity of 4Mb can easily serve in place of rewriteable flash, EEPROM or masked ROM.

As shown in Figure 2, the cost of three iterations of masked ROM steadily rises with smaller geometries as the cost of a mask rises with each process shrink. This is in addition to the more costly time penalty associated with a silicon spin: weeks of process time and lost productivity.



Source: IC Knowledge

Figure 2. Mask cost continues to rise at progressive smaller process geometry.

Conclusion

At 65nm process geometry and below the advantages of antifuse OTP NVM are opening up applications previously not practical with the technology. Because antifuse OTP NVM is implemented in a standard logic CMOS process, the memory scales with each new process node requiring no additional or special manufacturing steps. One application that antifuse OTP NVM serves well is on-chip program code and configuration data storage previously contained in off-chip flash or EEPROM or in on-chip masked ROM. And applications that require few code revisions are also finding antifuse OTP NVM an attractive option.

Kilopass was the first to pioneer antifuse OTP in a standard CMOS process with no additional processing steps. Its patented 2T bitcells are implemented using standard core devices and follow DFM design rules. Today, Kilopass offers antifuse OTP NVM products from 180nm to 65nm with densities up to 4Mb. In deep submicron, Kilopass continues to build high-density NVM products with its patented antifuse technology. Through design optimization and technology innovation, antifuse NVM technology is expected to expand its usage and penetrate into more applications.

Author Biography: Bernd Stamme is a Director for Marketing and Applications at Kilopass Technology. He has more than 15 years of experience in the IP and semiconductor industry. Prior to Kilopass, he was the Director of IP Technology at SiRF Technology managing the licensing and successful integration of 3rd party IP into SiRF's GPS chip sets. Before SiRF, he held management positions in LSI Logic's CoreWare organization and worked on high speed SerDes IP, communication interfaces and processor cores. Bernd holds a Dipl.-Ing. degree in Electrical Engineering from FH Bielefeld in Germany.

needed?

Gabe Moretti

To listen to some speakers at conferences or to read some articles in the press, one would come to the conclusion that architecting a leading edge SoC is just a matter of integrating IP blocks with application specific glue logic to create the hardware system that will execute the appropriate software. As designers and their managers know, this is hardly the case.

Functionality and physical characteristics

Any IP is characterized by two types of information: functionality and physical characteristics. In the case of IP implemented in either hardware or software, functionality is the implementation of a standard protocol or function. This is especially true for commercially available IP, for example a USB controller, a multiplier, or a software stack. Some hardware IP, like CPU and MCU cores, also provide the instructions list and description of their execution regarding the utilization of its architectural resources.

Physical characteristics differ considerably between software IP and hardware IP. Software IP physical characteristics are limited to the size of the compiled program and its target execution environment, while hardware IP's needs much more descriptive information. In order to properly choose a hardware IP product, users must know the foundry and process for which the IP has been characterized, the power consumption, the size of the block, the maximum execution speed of all the logical functions supported and more. In fact each process node advancement increases the number of physical characteristics that designers need in order to properly integrate the IP block into their design.

All this is really not new. Those of us that have been around for a while remember data books from semiconductor companies that described in detail the standard parts they were offering. Engineers would use the data books to compare similar parts and choose the provider. Yet the details of how IP blocks are evaluated and chosen are different from that use of standard data books.

A Growing, And Changing, Industry

Although the beginning of the commercial IP industry can be traced to 1985 when EIS Modeling began selling verification models to EDA companies as a service, the third party IP market began about five years later, as Synopsys introduced DesignWare. This was shortly followed, in 1991, by the introduction of the ARM6, an embedded RISC core by ARM. At the same time an American company, HDL Systems, later purchased by Philips, was marketing a Verilog synthesizable model of the MIPS RISC processor under license from MIPS.

Today revenue from the IP market is about one billion dollars and its growth is fueled by the complexity of SoC design as well as the use of standard operating systems, protocols, and hardware resource management systems. As design requirements become more and more stringent with the advance of manufacturing processes, the growth of the IP industry is accelerating

and the profile of IP vendors is changing.

Semiconductor foundries, who in the past have limited themselves to qualifying third party IP, are beginning to offer qualified IP directly to their customers. Although these IP cores are developed by trusted third parties and not by the semiconductor companies themselves, they are qualified by the foundry to its specific process and at times offered with an exclusive license. The practice has gained momentum with the introduction of the 32/28 nm process node and I expect it will become prevalent with the 20 nm node.

The foundries have two primary reasons for the growth of this new market segment: quality control and management of the cost of IP qualification. But other practical reasons also exist. If the cost of second sourcing or porting to another foundry design rules at the 32/28 node is high, the cost of doing so at the 20 nm node will be prohibitive. Thus IDMs must choose the foundry at the beginning of the architectural design and will be committed to that foundry for at least the life of the product.

The new partnership between an IDM and a foundry, which historically was mostly a matter of financial arrangements, is morphing into a technical as well as financial partnership. The choice is not "just" a matter of the capability of a foundry to deliver satisfactory yields and assure consistent fabrication capability, but also will depend on the library of qualified IP the foundry can offer. This does not mean that companies like ARM, Qualcomm, Cadence, and Synopsys will be forced to exit the market. It does mean that the investment they will have to make in order to qualify their own IP on a specific foundry process will required more and better engineering resources and a longer project cycle.

Given the costs of the supporting foundry resources during the qualification project, it will be imperative for each third party supplier to justify the project by showing significant enhanced revenue to the foundry through the availability of their specific IP library.

A lack of standards

The use of third party IP is still, on average, a costly and at times risky design method. The fundamental problem is a lack of standards, and, in the case where a standard exists, the slow adoption of the standard by IP providers. This is true not only for newly developed IP. The effort to document existing IP blocks that existed before a standard was adopted and published is practically non-existent. Evaluating third party IP, and even internally produced IP in a large organization, is a costly and unfortunately at times inexact task. IP developers have historically faced the dilemma of disclosing enough information about their products to make their choice simpler while at the same time protecting proprietary competitive information from disclosure.

The result is that almost each IP block must be evaluated using a unique method. This state of affairs is not just difficult to implement, it is also very costly and prone to mistakes. The EDA industry has tried to address the problem. A company, [ChipEstimate](#) now part of Cadence, was created with the mission of making the evaluation and choice of third party IP simpler and as uniform as possible.

Their method allows designers to evaluate different IP blocks that implement the same function, but does not insure a standard method of integration. Their system helps users explore a wide range of chip architecture options in a short time, including selection of IP, technology nodes and manufacturing technologies, low power strategies, packaging and much more. Using the InCyte Chip

Estimator to develop and optimize chip specifications gives information about chip size, power, leakage, performance and cost that is as accurate as possible. The results generally correlate well with the physical characteristics of the implemented silicon, but there is no guarantee that such will be the case every time. So although the service is valuable, it falls short, not for ChipEstimate's fault, but for the lack of suitable uniform information.

Standard developers are responding

The same lack of precise information is also present in many software libraries of functions. Yet, third party software IP is in better shape since documentation for industry standard RTOS nucleuses from companies like Apple, Google, and Microsoft is both sufficient and precise. In addition, standards like the Transaction Level Modeling (TLM) developed by OSCI and the Unified Verification Methodology (UVM) from Accellera help in the integration and verification process of both software modules and hardware IP blocks.

In 2009, the IEEE approved standard 1685. The title of the standard is: ***Standard Structure For Packaging, Integrating, And Reusing IP Within Tool Flows***. IEEE 1685 is the first standard description for Intellectual Property (IP) blocks for highly automated electronic design environments. It provides the electronics industry with an approach to create and use IP blocks in a standard way so that IP use and re-use is easier and more economic.

IEEE 1685 describes an XML (Extensible Markup Language) Schema for meta-data documenting of IP used in the development, implementation and verification of electronic systems with an API (Application Programming Interface) to provide tool access to the meta-data.

The IEEE 1685 standard offers the capability of having all information available in a machine readable way so that an IP component can be integrated into a design correctly and tested in its environment. The standard can be used for virtually any type of hardware IP (digital, analog and mixed signal) including models.

Thanks to the financial sponsorship of Accellera, the standard can be obtained for free through an IEEE program called "Get IEEE Standards". The Get IEEE 1685™ provides access to the current, active version of the standard in a PDF format at no cost to end-users. IEEE retains the copyright on the standard. For more information or to download the standard, please visit [this IEEE site](#).

"Our Get IEEE Standards programs accelerate and facilitate the dissemination of standards and speed up their implementation and adoption," said Judith Gorman, Managing Director of the IEEE Standards Association. "The sustained and growing interest in Get standards like IEEE 1685 standard benefits everyone in the worldwide electronics industry and shows the popularity of this program."

"IP is a valuable part of System on a Chip designs," said Shishpal Rawat, Accellera Chair. "Our collaboration with the IEEE-SA gives the worldwide design community access to our first IP-creation focused standard at no charge."

Conclusion

The IEEE 1685 standard is almost two years old, and while IEEE and Accellera announced in

November of 2010 that in the previous six months over 1200 copies of the standard had been downloaded (see <http://www.gabeoneda.com/news/accelera-and-ieee-standards-association-report-popularity-intellectual-property-ip-standard>) I am not aware of one single IDM requiring that the IP offered to them be IEEE 1685 compliant. At the same time, no third party IP seller has advertised its IP as being IEEE 1685 compliant.

If the opinion of the industry is that third party IP is costly to qualify and use, why has not this standard been adopted in a timely and industry wide manner? Is it just inertia? Is it lack of information? Maybe a bit of both.

I would like to see Accellera provide better support for its invaluable work by providing adoption tools for the standards it so brilliantly develops, especially those that are not in response to an established industry practice, but those that are truly enabling technologies that either improve on existing methods, or introduce new, more productive, methods.

It is clear to me that IEEE 1685 is the new data book for hardware IP blocks. Will semiconductor foundries embrace it, improve it, and use it? I hope so, since what we have today is costly, and it does not really enable design reuse and portability.

