



The configurable logic company™



Paradigm Shift in ASIC Technology

In - Standard Metal

Out - Standard Cell

Zvi Or-Bach, eASIC founder and CEO

Abstract

The time has come for Standard Cell to give up its reign over ASIC technology. Standard Cell can no longer achieve the expected cost and performance benefits of process scaling. Deep submicron (DSM) design and manufacturing constraints will drive the need for a new design technology to replace Standard Cell and again allow affordable ASIC designs.

History repeats itself. Full Custom design was the preferred solution until it became too lengthy and expensive for most designs. At that point, the industry chose the Standard Cell methodology, willing to pay the area and performance penalties. Now, metal interconnect dominates DSM circuit timing, feature sizes have become smaller than the lithography wavelength, and Standard Cell fabrication costs have skyrocketed. The time has come for Standard Cell to be replaced as the dominant ASIC technology. Structured ASIC technology, in which the metal layers are also standardized, will be the new ruler in the ASIC Kingdom.

Introduction

About 20 years ago Standard Cell started to displace Full Custom – the dominant design methodology at the time – as the preferred logic design methodology. In those days, the cost of Full Custom design began routinely to exceed \$10 million, and the design community had elected to sacrifice more than a factor of 2 in device performance and density, in favor of a 10X reduction in design cost.

Standard Cell design cost has escalated since then, and currently, it routinely exceeds \$10 million. It is inevitable for the industry to shift again to a methodology that provides a significant reduction in design costs, yet again at the cost of device performance and density. The class of emerging solutions vying for this role has acquired the name Structured ASIC. A subset of this class is the Via Customizable Arrays proposed by few vendors (eASIC, ViASIC, Leopard Logic) and the academia (Carnegie Mellon's VPGA). We generically call them Standard Metal.

Interconnection - Taking Over 'Delay Domination'

Continuous scaling has provided great economic value along with improved design performance. However, not all circuit elements scale equally. To understand the implications of this differentiated scaling, we need to analyze the elements involved. All digital logic is constructed from two elements: transistors, and metal lines for interconnecting them. Shrinking the transistor size reduces the time it takes it to switch and hence its logic delay. Unfortunately, scaling increases the resistance of interconnection lines and, even more so, the capacitance of those lines, and therefore the overall interconnect delay. Hence, while scaling reduces the logic delay, it also increases the interconnection delay as can be seen in Figure 1 below.

Transistors no Longer Dominate – Metal Interconnections Took Over




Figure 1.

At 0.18μ the first discontinuity occurred. After years of path delay domination by transistor delay, continuous scaling had resulted in a crossover to interconnect domination. This major discontinuity has been resisted by the industry through huge investments in process change from Aluminum to Copper and through an effort—which we are still struggling with—to move to low dielectric isolations layers (a.k.a. “Low-K” processes).

Yet, even with all of that effort, interconnect is now dominating advanced logic designs and it will continue doing so in the future. In fact, James Meindl, in a keynote speech at the International Symposium on Physical Design (ISPD '04) warned that “the tyranny of interconnect is threatening the timing, power, and cost of next-generation chips.”¹ As an example, at 100 nm, Meindl said, interconnect switching energy is five times that of MOSFET switching energy. At 35 nm interconnect switching energy becomes 30 times greater. Recently, it was reported that when Agilent Technologies ASIC products division first moved from 130 to 90 nanometer chip design, it got a nasty surprise. “Signal integrity,” said Jay McDougal, microprocessor design methodology manager at Agilent, “was really an order of magnitude worse,”² as is seen from Figure 2.

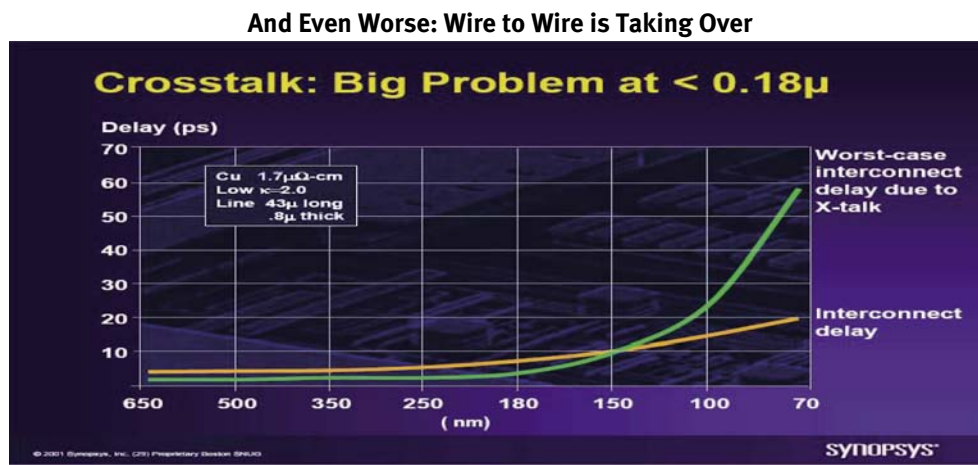


Figure 2.

Similarly, EE Times reports that “Crosstalk and power are extremely difficult problems to solve, and will require significant changes to the existing chip design flow,”³ according to Li-Pen Yuan, R&D director for extraction and signal integrity at Synopsys, at a keynote speech at the 2004 Electronic Design Processes methodology conference.

The implication of this trend is that we no longer get the traditionally expected 100% performance improvement by scaling to the next process geometries, as is visible in Figure 3 below. As, Bernard S. Meyerson, the CTO of IBM Microelectronics, said, in Semico 2004 Impact Conference: “there is a paradigm shift here, and it is a very important one...The diminishing returns you get from scaling mean that innovation actually has to happen”⁴.

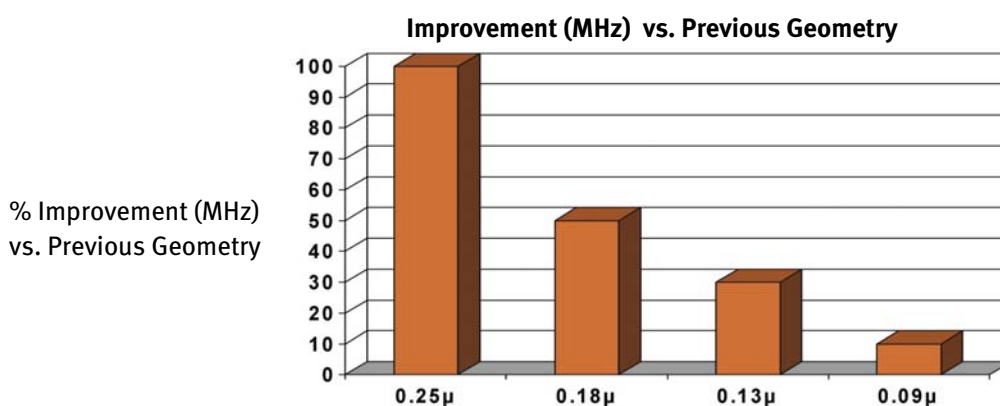


Figure 3

It is very clear by now that the basic logic building block should be changed from fine-grain to coarse-grain. The same driving forces that made coarse-grain Look-Up Table (LUT) the winner in the FPGA domain, where interconnect delay dominates over logic delay, will now drive similar change in the ASIC world. Just as we had transitioned from transistor sizing to gate sizing with the advent of Standard Cell over full custom, it is now time to move to an even coarser building block. It seems quite evident that the natural primitive is the LUT—the winning logic primitive of the FPGA wars. Constructing logic functions using coarse-grain primitives is becoming far more efficient than using multiple fine-grain gates connected by routing wires, since those wires are associated with high delays.

Hence, via-defined logic or, for other considerations, bit-stream defined logic, will become the preferred logic fabric furnishing the first step in the transition from Standard Cell to Standard Metal.

It is quite obvious that the coarser-grain logic of LUT allows for the standardization of the lower metal layers, which are used to construct the underlying logic fabric. As an example, eASIC technology uses Standard Metal 1 through metal 3 for its logic fabric. The real advantage of coarse-grain logic fabric, however, is for the routing layers. Those layers, which naturally belong on top of the logic fabric, are used to connect the logic cells and construct the actual logic circuit. Logic fabric, constructed from repeated coarse logic cells, enables an efficient use of segmented routing, similar to FPGAs. Figure 4 below shows the layout of the basic cell, called eCell, repeated in the logic fabric of eASIC. The eCell is equivalent to 12-15 logic gates and occupies about $200\mu^2$ @ 0.13 process vs. $5\mu^2$ for the smallest NAND gate.

The immediate benefit of coarse grain is in a reduced number of wires which need to be routed by automatic P&R flow, as the wires constructing the coarse grain cell are already part of the custom hand crafted base cell design. This can be clearly seen in the simple analysis in Figure 4 below.

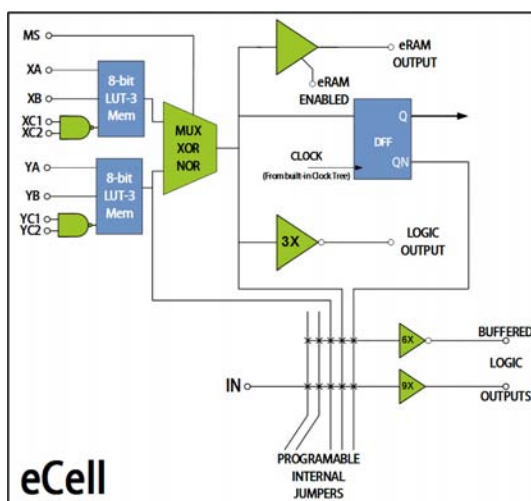


Figure 4.

- Fine grained Standard Cell needs an average of 3 ports per gate to route

$$\blacksquare 3 \text{ ports, 1 gate} \Rightarrow 3 \text{ ports/gate}$$



- Coarse grained eASIC/FPGA need an average of 1 port per gate to route.

$$\blacksquare 11 \text{ ports avg., 15 gates} \Rightarrow \sim 0.75 \text{ port/gate}$$

- ~4:1 reduction in routing needs is gained by using eASIC fabric

The second advantage is that coarse grain cells allow segmented routing. Figure 5 below shows the layout image of a sample coarse grain cell – eASIC's eCell.




Figure 5

An array of logic cells where each cell is 14μ on a side implies that the minimum distance to the next cell is 14μ . Therefore, the interconnection fabric on top of it can be constructed with fixed segments of 14μ length, stitched with vias to form the interconnection fabric.




Figure 6

Figure 6 helps visualizing such interconnection fabric. A Standard Metal segmented connectivity fabric made of four metal layers is illustrated in Figure 7.




Figure 7

This connectivity fabric complements the logic fabric presented previously, to complete the Standard Metal fabric customizable by a via layer. In the case where RAM-based LUT are used, only a single custom layer would be required – via 6 as illustrated in Figure 8.




Figure 8.

Feature-limited - Taking Over the ‘Yield Domination’

The second fundamental discontinuity, which drives to the same conclusion but for completely different reasons, is presented in this section.

A coincidental, but unrelated occurrence, which happened at about the same time, caused another change in silicon fabrication. Figure 9 below shows that until mid 90's the lithography wavelength was greater than the features drawn by it. Since mid 90's, however, the lithography wavelength became longer than the feature size, and while this discontinuity was predicted for many years, all efforts to come up with an alternative to optical lithography have failed. Instead, the industry started using the wave properties of light,




Figure 9

Together with the integration and threshold properties of the photo-resist, in order to achieve sharp sub-wavelength images. Those so-called Reticle Enhancement Techniques (RET) include Optical Proximity Correction (OPC) and Phase Shift (PSM), and are illustrated in Figure 10.





Figure 10

While Figure 10 illustrates the difficulty of the lithographic challenge, the real problem is that RET techniques do not really solve the problem. With dense sub-wavelength layout patterns, we simply cannot get the perfect desired pattern on silicon anymore, as the various required patterns interfere with each other, and the more we scale down, the farther this proximity effect extends. The impact of this issue on device yield is presented in Figure 11. As is clearly visible, the yield once controlled by spot defects and die area, is now controlled by alignment and printability. In fact, the impact of this feature-limited yield is rapidly driving major EDA companies into design-for-manufacturability or design-for-yield.



Courtesy of PDF Solutions Inc.

Figure 11

The natural solution for this yield loss is to use repetitive patterns, just as in SRAM design. SRAM is a very important function occupying significant portion of the die area of most designs, and foundries are spending enormous efforts to optimize the SRAM bit cell. Using special aggressive design rules, the area of the typical foundry-provided bit cell is half of what it would have been, had standard logic design rules been followed. The fact that SRAM bit cells are used in large repetitive arrays enables foundries to overcome the proximity effect by trial and error. They use the aggressively designed bit cell in large arrays, and protect it from the non-repeating surrounding patterns with dummy cells at the array border.

This same technique, of using repetitive patterns in the critical poly and metal layers, is the foundation of the proposed solution to the lithographic yield loss. It is also the essence of the Standard Metal approach.

Cost of Test

The RAM based LUT variant of Standard Metal, as proposed by eASIC, has an additional attractive cost advantage – the cost of test. The RAM based LUT allows for large reduction of test costs by taking advantage of the ability to load all the LUTs with an XOR functional pattern at test time. This XOR provides for easy observability and controllability, reducing the number of required ATPG vectors, on average, by 70%, and slashes the test cost accordingly. Figure 12 shows the importance of this reduction. The cost of test does not diminish with scaling – indeed it even grows – while the cost of fabricating the transistors is exponentially reduced over time. Currently, testing is predicted to catch up with the cost of manufacturing, and RAM-based LUT fabrics slow this trend by keeping test cost low. Further, the same principle of XOR-loading can also be used with BIST, creating much higher coverage rates for randomly generated test patterns. Consequently, the increase of die size due to the use of Standard Metal might be more than compensated by the decrease of the associated cost of testing.




Figure 12

Mask Set Cost Becomes Prohibitive

The most apparent impact of the fundamental discontinuities described above is their effect on the cost of a mask set.

As can be seen in Figure 13, mask set cost has been stable at about \$18K for many years. Once scaling hit 0.65μ , however, a new dynamic began; the average cost for a mask set started to double for every new process node. The lithographic difficulties have driven up the average single mask cost. These include; increase in the mask making time due to both OPC and Phase Shift complexity, reduced mask yield, and increase in the cost of mask inspection and repair. Adding fuel to the fire, the domination of interconnects has resulted in an increase in the number of metal layers and accordingly in the number of masks in a mask set. Up to the 0.65μ process most designs were using 2 metal layers. The 0.35μ process deployed 3 metal layers, and at 0.18μ process 6 metal layers became popular. At 90 nm most vendors offer 10 metal layers.




Figure 13

The average mask set cost for 90nm reaches \$1.5M and many analysts expect mask set costs to approach \$10M by the end of the decade.⁵

The economics for justifying a mask set could be viewed from many angles and the same holds true with respect to the impact of the mask set cost on the total cost of design. It is well documented that more than 70% of the design cost is spent on verification. Figure 14 shows the escalating design costs related to scaling. The chart in Figure 13 is quite conservative. Other analysts estimate up to three times higher design cost.




Figure 14

While it is true that the increased design complexity has other impacts in addition to the escalating mask costs, the improvement in EDA and the reuse of IP should have helped to keep the design cost from escalating so rapidly. The well-controlled design cost of FPGAs provides evidence that it is the escalating mask cost, together with the escalating deep sub-micron physical issues, that drives ASIC design costs out of control. Here again we see the need for bringing in Standard Metal to replace Standard Cell.

Figure 15 shows the implication of the design cost increase on the required lifetime revenue to justify such design.

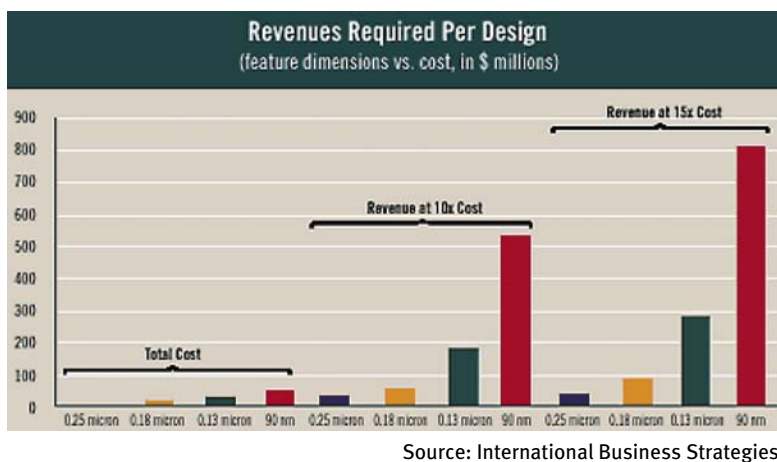


Figure 15

An alternative analysis would evaluate at what revenue the savings in device cost would be worth the increase in design cost. Such analysis requires certain assumptions. The first is the appropriate factor between up-front spending vs. future saving. Such factor depends on the risk assumption for the future revenue and the expected time lag between the design expenditure time and the product revenue. A factor of 5X seems reasonable, given the long lead-time of ASIC production.

The second consideration is what fraction of the cost would be saved. The dominant savings from using Standard Cell would be the die cost, and the die size—which is correlated with it— accounts for about 50% of the final product cost.

Figure 16 shows the ratio of new logic within future SoCs. It is clearly seen that the die area allocated for logic is decreasing while more and more of the die area is allocated for memory. The net result of this trend is that the die size penalty for the use of Standard Metal vs. Standard Cell is decreasing with scaling.

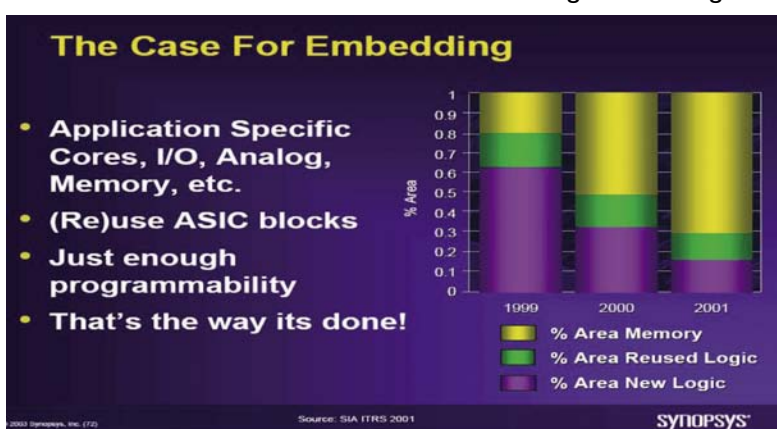


Figure 16

For our analysis we will assume that the Standard Metal is used for logic only, and therefore the average increase in die size by using it vs. Standard Cell would be less than 30%, assuming a density penalty of 2X for Standard Metal.

The net impact is that the use of Standard Metal might increase the manufacturing cost by about 15%. Consequently, in justifying an upfront increase of \$10M in design cost, one needs to save \$50M in future production costs, implying production costs of \$333M and therefore product revenue of more than \$500M. There are very few ASIC designs that represent such revenue! As mask costs and the related design costs escalate with scaling, economic considerations also drive Standard Cell out and Standard Metal in.

It is Time to Talk About Time

Time-to-market is becoming increasingly important as is evident from Figure 17.




Figure 17

In the context of Standard Metal vs. Standard Cell, time plays in two aspects. The first is the time-to-market of an individual design as impacted by the choice of implementation methodology. The second aspect is the rate of adaptation of a new fabrication process (scaling).

Design Time and its impact on Time-to-Market

In fact, Time-to-Market is another dimension of design cost. It takes more than six months and a team of more than 10 designers to tape-out a deep submicron design. Once it is taped out, it will take the fab more than 2 months to process the 40 layers of such design and provide the design team with prototypes.

Figure 18 indicates that there is also a high probability that the first prototypes will reveal significant shortcomings, and that additional iterations will be required, extending the time-to-market by additional 3-9 months




Figure 18

Figure 19 presents the reduction of time-to-market estimated by NEC for a design using their ISSP approach—their Standard Metal — vs. their Standard Cell approach (CBIC- Cell Base IC). Clearly, NEC believes that using Standard Metal would save almost 5 months in Time-to-Market.




Figure 19

In fact, eASIC customers have achieved the remarkable reduction in design time to less than a single day from RTL to tape-out by using the eASICore®:

“the customization of the printer platform for the printer division. We received the final RTL from them in the morning and we shipped the final gds of the chip to Crolles for ebeam customization in less than a day”.

Michele Borgatti, Manager - Front-end Technology and Manufacturing Manager, STMicroelectronics

Using the eASICore, an embedded programmable logic IP, allowed cutting the design time due to multiple factors. First, most of the physical issues, such as power distribution, clock tree, scan chain, antenna rules etc., are already built-in the Standard Metal fabric. Second, the design-for-manufacturing (printability) issues are resolved due to the repetitive fabric. Further, the issue of design convergence is greatly simplified by using coarse grain with its inherent higher intrinsic delay and overdrive of the output signals. And, finally, the large reduction in NRE cost and risk allows for a significant reduction of the verification effort.

Not less important is the significant reduction in device debugging time and effort that is provided by the subclass of Standard Metal with programmable LUT. The re-programmability of such cells greatly enhances the ease of debugging and accordingly reduces the time between the re-spins. The combined advantages of greatly reduced time for physical design, low NRE, and short fabrication time with quick debugging, can reduce the product release by more than a year as presented in the following chart (Figure 20).




Figure 20

Figure 21 is another look at the impact of Time-to-Market on product revenue.

From the previous analysis we know that using Standard Metal implies about 15% increase in device cost, which is well justified by the over 20% increase in design revenue!

Lost Revenue from Late Designs (percentage)												
Market Characteristics	3 Months			6 Months			9 Months			12 Months		
Speed of Market	Fast	Med.	Slow	Fast	Med.	Slow	Fast	Med.	Slow	Fast	Med.	Slow
Revenue Reduction	14.0	6.6	2.7	21.2	10.1	4.4	27.9	15.5	7.1	35.1	21.0	9.9
Higher Development Costs Due to Design Delays	5.6	5.2	4.7	11.1	10.3	9.4	16.7	15.5	14.1	22.2	20.6	18.8
Opportunity Revenue Loss	12.9	12.5	12.1	26.2	25.0	22.4	39.7	37.5	32.3	55.7	50.0	43.6

Source: International Business Strategies


Figure 21

Hence again: Standard Cell out Standard Metal in.

Rate of New Process Adoption

Figure 22 presents the well-known fact that scaling provides for cost reduction. In general, scaling complies with Moore's Law of density increase by 2x and cost reduction of about 50% each generation. Figure 22, provided by TSMC, shows that 6 months after a new process is available it catches up in costs with the previous process, and it takes about another two years to provide the full 50% cost reduction. While density is the most important factor driving the device cost, yield issues and high wafer price are the reasons for those two years of lag. It takes about a year for a new process yield to catch up, and the foundries typically amortize the cost of a new process over the first 3 years, impacting the wafer cost accordingly.

Advanced technology continues to reduce cost



Source: Genda Hu, TSMC, Jan 7 2004

Figure 22

Figure 23 reveals another interesting aspect of scaling. Logic designs trail memory design by about two years with respect to new process adaptation. As the table shows, in 2006 about two thirds of all memory devices will use 90nm process, while only about one third of logic devices will do the same. The reason for that is the need to support a full set of libraries, design tools and many variations of circuit patterns.

Total Semiconductor Units by Technology (estimated)			
Memory	2004	2005	2006
130 nm	5,243,248	3,971,902	2,712,272
90 nm	781,274	2,912,944	5,102,978
65 nm			189,831
Logic	2004	2005	2006
130 nm	1,709,396	2,008,202	2,463,219
90 nm	219,608	740,847	1,544,176
65 nm		20,583	127,668

Source: Semico Research

Figure 23

In fact, the escalating cost of design, the yield challenge, and the slew of deep sub micron design effects, slow such adoption even further. Figure 24 shows how severe this problem is becoming. It is currently expected that it will take 6 years for 15% of the designs to move into the 65nm process.

Will Developers of 22nm technology watch their children implement volume production?

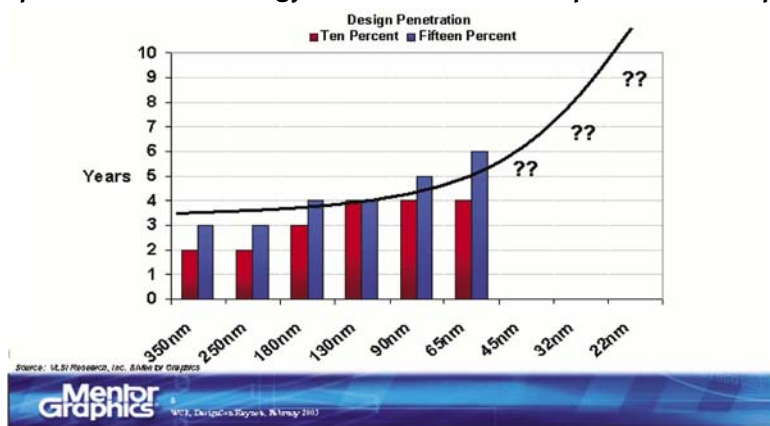



Figure 24.

Actually, only the super high volume designs tend to move early to the new processes, while the adoption for typical ASIC often takes many additional years. Figure 25 shows the severity of the current trend. The number of Standard Cell designs is dropping sharply with each new process node, and so is the rate of adoption.



Source: VLSI Research, Inc.

Figure 25

Standard Metal, on the other hand, and especially its SRAM LUT variety, is actually very much like memory, which enables it to keep adoption pace with memory. Early adoption of new processes for logic designs will enable a 50% reduction in cost, which by far exceeds their 30% logic density penalty. Hence, Standard Cell – out, Standard Metal – in.

Summary

The various aspects of Standard Metal, as discussed in this white paper, will become evident only later in the adoption process. The transition from Standard Cell to Standard Metal requires a learning curve. The early transition will be dominated by low-end designs adopting Standard Metal, to reduce NRE cost, while willing to pay the corresponding higher device cost. Figure 26 presents the disruptive nature of Standard Meta technology. Accordingly, as more designs adapt to Standard Metal, the other factors described above will start to kick in and as Standard Metal becomes adapted as a technology driver, it is clear that the pendulum will tilt completely towards Standard Metal.




Figure 26

Appendix 1

The other Alternatives – Not Really an Alternative

High NRE and Time-to-Market drove the pursuit of many solutions over the years. Many attempts have been made to build advanced computing devices with parallel computing elements to replace the need for Standard Cell. Many tens of millions of dollars of Venture Capital money were spent on dozens of new architectures, yet no one was able to come up with a commercially viable alternative. We believe that the root cause for the failure of those new architectures is the lack of design environment support. The current EDA tools that support Standard Cell were developed in the course of the last 20 years, with multi-billion dollar investments. It makes it very hard for any new architecture to match such investments.

However, two alternatives did manage to achieve commercial success. The FPGA technology, which utilizes coarse grain cells, but otherwise utilizes design flow similar to Standard Cell, and the DSP technology that is a variant of the CPU technology, which received billions of dollars of investment for the development of tools, compilers, and application software.

Both technologies have achieved commercial success and proved themselves in certain market applications. Some argue that as Standard Cell gets harder to use, these technologies might become the alternative. Research done at UC Berkeley investigated these issues, and the results are illuminating. The Berkeley group studied alternatives for implementing computational algorithms such as FFT and Viterbi decoders. The research compared three implementation technologies: Direct-Mapped Hardware (Standard Cell implementation), commercial FPGA, and two types of generic DSP. The evaluation was reported by measuring the energy per operation, and by the number of operations per unit area. Their results, shown below, indicate that the conventional technique of Direct Mapped Hardware (Standard cell and Standard Metal) is three orders of magnitude more efficient! It seems difficult to believe that efficiencies of such magnitude will be left untapped by the marketplace.

	Energy		Area	
	64-point FFT Energy per Transform (nJ)	16-State Viterbi Decoder Energy per Decoded bit (nJ)	64-point FFT Transforms per second per unit area (Trans/ms/mm ²)	16-State Viterbi Decoder Decode rate per unit area (kb/s/mm ²)
Direct-Mapped Hardware	1.78	0.022	2,200	200,000
FPGA	683	5.5	1.8	100
Low-Power DSP	436	19.6	4.3	50
High-Performance DSP	1700	108	10	150

Source: Berkeley Wireless Research Center

(numbers taken from vendor-published benchmarks)

Figure 22. Results in fully parallel solutions

Orders of magnitude lower efficiency even for an optimized processor architecture

Appendix 2

Closing the Productivity Gap – Making the Promise of Reusable IP a Reality

The relentless march of Moore's Law has resulted in a gap between the capability to fabricate ever more complex chips and the ability of design engineering to implement them.



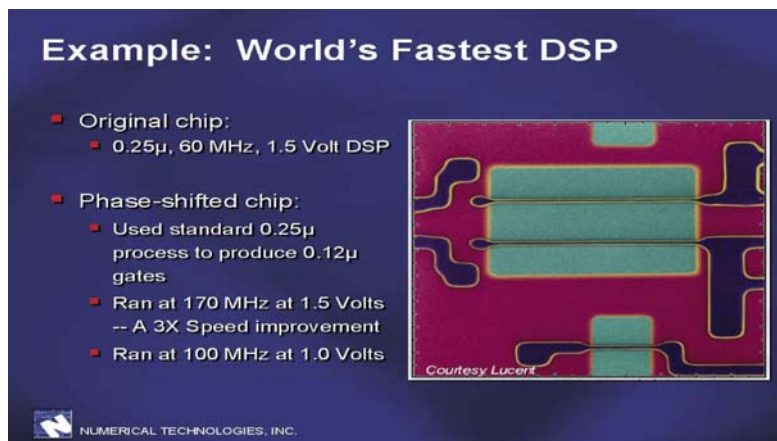
For years, IP reuse has been the industry's mantra for this Productivity Challenge, yet the actual RTL IP reuse has not taken off as predicted. A variety of explanations are offered by industry analysts, and they all boil down to the risk associated with using IP, over which the designer has little control, yet which still requires implementing the physical part of the design and may compromise the multi-million dollar ASIC investment. The relative success of IP reuse in the FPGA domain supports this thesis, as the "tapeout" is a non-event in the FPGA world.

Standard Metal, and in particular its flavors which support maskless, NRE-free silicon, revive the opportunity for wide spread RTL IP reuse to close the productivity gap.

Appendix 3

Aggressive Phase Shift Masks – Extending the Life of Moore’s Law

Aggressive phase shift masks (PSM) have been around for some time and carried with them the promise of leapfrogging process generations for selected devices. Yet, despite substantial performance gains, as demonstrated in the figure below, they have never become a significant factor in the industry under the Standard Cell regime—they were too expensive, and the delivered results were limited, preventing them from becoming real players.



The results were too limited because aggressive PSM affects only the device characteristics—its switching speed—and has no significant effect on the ability of the device to drive the interconnect. Yet, in the deep submicron, interconnect became the dominant contributor to path delays, and hence the gains in the example above would typically have been smaller applying this technique at 0.13 microns instead of 0.25 as reported above. Couple this with the fact that in Standard Cell regime it is often the cell output buffer that dictates the actual size of the cell, and the reasons for the limited appeal of aggressive PSM become obvious.

This picture is changed, however, with the advent of coarse granularity cells in Standard Metal fabrics. In these cells, multiple levels of logic are consolidated within close proximity, allowing the full benefits of aggressive device sizing with PSM. Intra-cell devices and drives can be customized, allowing the intrinsic cell delays to enjoy the benefits of devices at one or two process generations ahead, while still buffering the outputs, which drive the interconnect, with hefty buffers. Furthermore, as the optimization of a single coarse cell benefits the performance of all designs implemented on the Standard metal fabric, the cost of this approach is amortized over countless designs and becomes negligible.

In summary, the application of aggressive Phase Shift Masks works in tandem with other benefits of Standard Metal and allows it not only to be a technology process driver, but even leapfrog process generations.



References:

1. EE Design, April 19, 2004, “Speaker offers solutions to ‘tyranny of interconnects’”, By Richard Goering
<http://www.eedesign.com/showArticle.jhtml?articleID=18902137>
2. Planet Analog, June 24, 2004, - “Bad signals interfere with 90-nm designs”, By Richard Goering
<http://www.planetanalog.com/showArticle.jhtml?articleID=19202088>
3. Planet Analog, Apr 28, 2004, – “Synopsys keynoter cites crosstalk, power challenges”, By Richard Goering
<http://www.planetanalog.com/printableArticle.jhtml?articleID=19202081>
4. EE Times, January 18, 2004 – “Innovation to drive chip performance curve”, By Mike Clendenin
http://www.eetimes.com/article/printableArticle.jhtml?articleID=17500148&url_prefix=story&sub_taxonomyID=2251
5. Tech Online, June 3, 2004 – “The “Missing Link” of SoC Design—Platform and Structured ASICs”, By Jim Lipman
http://www.techonline.com/community/ed_resource/36652
6. Berkeley Wireless Research
bwrc.eecs.berkeley.edu

Covered by the following patents:

US Patents: US 6,194,912; US 6,236,229; US 6,245,634; US 6,331,733; US 6,331,789; US 6,331,790; US 6,476,493; US 6,642,744; US 6,686,253; US 6,756,811; US 6,819,136.

European Patents: EP 1 161 797 B1

TRADEMARK INFORMATION

FlexASIC is a trademark of Flextronics.

eASIC, eASICore and “The Configurable Logic Company” are registered trademarks of eASIC Corporation in the United States. eCore is an abbreviation of eASICore. Structured eASIC, eI/O, eCell, eUnit, eRAM, eTool, eGenPLD, eDebug, eMacroWare, eGenSim, eGenMem, eGenModule, eMμ are trademarks of eASIC Corporation in the United States. Other company and product names may be trademarks or registered trademarks of the respective owners with which they are associated.

U.S. GOVERNMENT RESTRICTED RIGHTS

Use, duplication, or disclosure by the United States Government is subject to the restrictions set forth in DFARS 252.227-7013 (c)(1)(ii) and FAR 52.227-19.

eASIC Corp.

2242 Camden Ave,
Suite 203
San Jose, CA 95124

Tel: 408-879-9400

Fax: 408-879-9430

info@eASIC.com

www.eASIC.com