Changing the Approach to High Mask Costs

The ever-rising cost of semiconductor masks is making low-volume production of systems-on-chip (SoCs) economically infeasible. This economic reality limits the production of the prototypes, derivative designs, and low-volume/high-value designs that represent the greatest opportunity to extend the market for SoCs.

What if low-volume SoCs could be manufactured with virtually no mask costs? Is this possible without revolutionary new hardware changes? How would the design-to-manufacturing design chain have to adapt in order to make this possible?

Maskless e-beam direct write (EbDW) technology, enabled by the character or cell projection (CP) capability in today's production equipment, has the potential to eliminate the need for the most costly mask layers. Coupled with design for e-beam (DFEB) software and design technologies, use of EbDW could reduce mask costs dramatically and provide a new, economically feasible path for the production of low-volume designs.

With the most complex and time-consuming mask production steps eliminated from the manufacturing cycle, this new paradigm also promises to deliver faster time-to-market – a critical benefit in the markets for prototypes and derivative designs.

Perhaps most importantly, these benefits can be realized with relatively low investment on the part of the rest of the design-to-manufacturing chain.

E-beam Builds a New Bridge

Traditionally, the mask has served as the bridge between design and manufacturing. If the goal is to reduce mask costs, another bridge is needed.

E-beam lithography equipment, first developed over 20 years ago, uses an electron beam to expose the resist and create semiconductor features on a mask, which is then used to create wafers. E-beam's strength is accuracy. Even at 65nm and 45nm, no optical proximity correction (OPC) or reticle enhancement technology (RET) is required.[1] However, e-beam's historic challenge is its throughput time: the serial nature of e-beam technology means that it is orders-of-magnitude slower than standard, optical lithography.

Over the years, e-beam equipment has evolved, and today, variable-shaped beam (VSB) and CP capabilities help to address these limitations. VSB fractures the complex shapes of design features into multiple rectangles, each of which requires a separate exposure or

[1] In e-beam lithography, proximity effect correction (PEC) is required to correct for both back- and forward-scattering of electrons. Unlike with OPC, the effects are small enough in the case of forward scattering, and large enough in the case of back scattering that the complex interaction of adjacent features that is so troublesome with OPC is avoided. Thus, character projection of a 2-input NAND gate is stamped the same everywhere on the wafer using EbDW.

shot from the e-beam. CP uses stencils to project a larger character in one shot (see Figure 1).

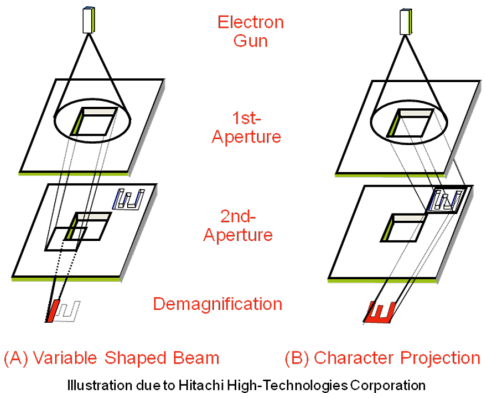

Figure 1: VSB vs. CP

The most advanced CP-capable EbDW machines write directly on wafers, eliminating the need for a mask. And at 65 and 45 nanometers, entire standard cells can fit within the projection area of these EbDW machines. Standard cells and RAM now dominate SoC designs, and these repetitive features can be converted to characters easily. These advances have resulted in a 3-5X improvement in throughput depending on the design and the layer being processed – a major step forward, but not enough to enable EbDW to be used for all critical layers of low-volume designs.

Recently, DFEB technologies and design techniques have been developed that speed up e-beam production a total of 10-25X over the traditional VSB method – to a rate of about one wafer per hour – and make it practical for use on all critical layers of an SoC.

Together, EbDW and DFEB technologies represent a new bridge between design and manufacturing that provides a path to silicon production that is both dramatically less expensive and significantly faster than conventional masks.

Design for E-beam

DFEB is a combination of software and design technologies that optimizes the design process to take maximum advantage of today's most advanced CP EbDW equipment to reduce shot count and in turn, make EbDW feasible for low-volume applications. First, design library elements from IP suppliers, ASIC vendors, and/or foundries must be re-optimized for character projection, and translated into characters laid out in a stencil-mask design that can be projected by EbDW. Then, the design process, implemented through electronic design automation (EDA) tools and services, must be made aware of shot count as a design optimization criterion.

Because they represent over 80% of the cost of a typical SoC mask set, DFEB is aimed at eliminating the complex computer-aided design (CAD) layers of the mask set. These layers contain the design data created by the design team that define the function and performance of the SoC.

An analysis of the VSB shot counts for a CAD layer of a typical SoC design reveals that shot density is highest in SRAM features, followed by standard cells, with analog, RF, I/O and other special cells requiring the fewest shots, and the fewest shots per area. Because SoC designs are dominated by RAM and standard cells, with other special cells making up a very small percentage of the total features, improving the shot count for these two feature categories has a decisive impact on overall shot count.

D2S, Inc., a DFEB service provider and managing sponsor of the eBeam Initiative (see page 5), will work with IP suppliers, ASIC vendors, foundries, and design teams with proprietary library elements to study and re-optimize RAMs and standard cells to co-design with the characters that are placed on the stencil mask. The stencil acts, and is sometimes referred to, as a "mini-reticle." Each cell, and each orientation of that cell, becomes a character. EbDW machines project the characters through the stencil directly onto wafers. Depending on the machine, more than 100 characters typically can fit on each stencil, covering the majority of features used in complex SoCs.

Essentially, DFEB changes the process of converting design features to e-beam characters from an inefficient and less-effective process of "search and find" to a streamlined "take and optimize" operation. Currently, when design features are converted to characters, an available character set is searched to find the character(s) which best match that feature. Using DFEB, the entire design library is optimized for shot-count efficiency and translated to characters for EbDW production.

Designers then use these DFEB libraries in conjunction with traditional design libraries during the design process. The DFEB design methodology uses shot count as an optimization criterion along with area, timing, power, and yield during the synthesis process, employing currently available commercial EDA synthesis products without modification.

Analog, RF, I/O and other custom cells are produced using standard VSB techniques. However, because these cells not only require fewer shots than RAMs or standard cells, but also represent a small percentage of the total features of an SoC design, this has little impact on overall shot-count.

DFEB design methodologies impact only the implementation phase of the design process. The register-transfer level (RTL) design is unchanged. Therefore, a systems designer operating above the RTL level is unaware of any differences in designing for EbDW.

Because the DFEB cell library is co-designed with the characters on the stencil, the same set of stencils are re-used for all designs using the same cell library. So, for example, the Fujitsu 65nmLP DFEB library, and the stencil for that library can be applied to all

designs using that library. This saves additional manufacturing time, because stencils are created per library, rather than per design.

The DFEB methodology includes careful steps to make DFEB designs "downward compatible" for the less-accurate processing with masks and standard optical lithography, so that DFEB designs can be used later in larger-volume production using masks, if that is desired. Only the last steps of design that involve DFM/RET/OPC steps must be performed on the DFEB design to complete it for mask production.

What Designs Benefit Most?

The designs that benefit most from DFEB are those that are most sensitive to reticle cost and manufacturing turnaround time: designs such as derivatives, prototypes and low-volume/high-value designs. Significantly, these three types of low-volume designs represent a core of innovation and potential market growth for the semiconductor industry.

Derivative designs – designs that leverage existing IP to create multiple, differentiated versions of a single design platform – represent one of the most promising areas for the expansion of the SoC market as a whole. With minimal additional design investment, derivative designs can create a "long tail" of low-volume products that increase greatly the overall market for a given SoC platform. However, derivative designs can only provide this long tail if the production costs are also minimal. Reticle costs at advanced nodes have kept the production of derivative designs in check, as shown in Figure 2.

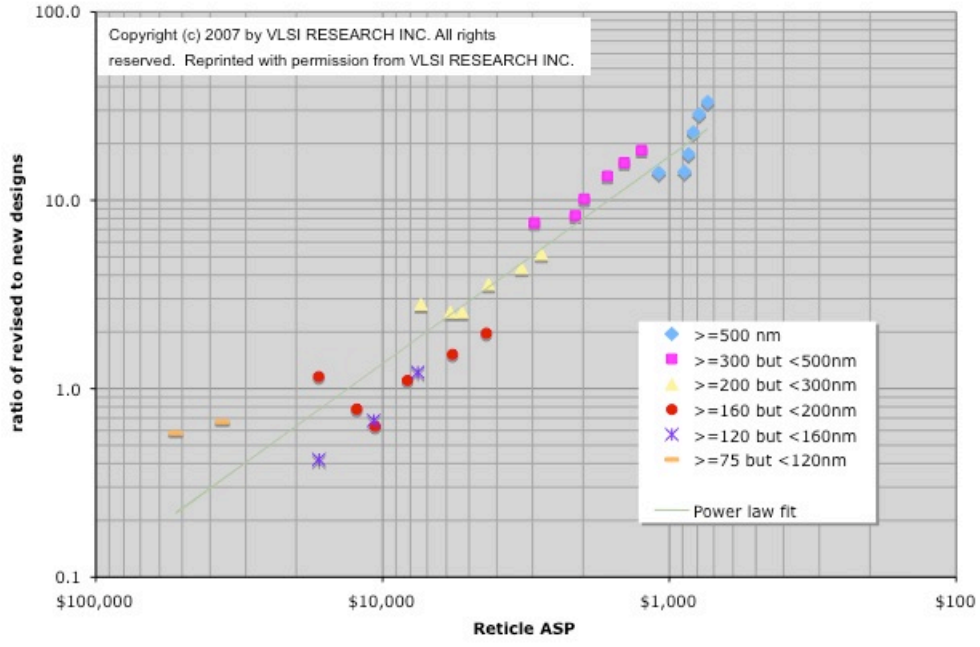

1

Figure 2: Correlation of derived design activity to reticle cost

By minimizing manufacturing costs – and time – DFEB could fuel a surge in derivative design production at advanced nodes.

Prototypes also stand to benefit from DFEB. ASIC and SoC complexity have made prototypes that can plug into a target system and run at speed a necessity to make the monumental debug process manageable. But few design teams can afford two mask sets – one for prototyping and one for volume production – so the dramatically lower cost of EbDW is a key benefit for these designs. And, because the faster a design team can get a prototype into their system, the faster they'll get to market, the accelerated manufacturing cycle offered by EbDW and DFEB is especially attractive for this application.

Finally, low-volume/high-value designs such as those for supercomputing applications or very specialized equipment – designs that will never reach high-volume production – will benefit from the drastically lower manufacturing costs associated with DFEB. By eliminating the majority of the mask layers required to produce these designs, the total cost of these traditionally very expensive chips can be cut significantly.

Collaboration of Design-to-Manufacturing Chain

No one company can enable DFEB alone – it's an industry-wide solution that requires collaboration and coordination throughout the design chain to address the questions about how DFEB will impact the semiconductor ecosystem. These questions include:

- What is the manufacturing accuracy and throughput?
- Is the equipment ready?
- Is the design flow ready?
- Is there a path from e-beam to mask-based production?
- Are there qualified design teams?
- Are design kits and stencils available?
- What is the manufacturing capacity for DFEB?
- Is DFEB scalable?

To answer these questions and to support the development of DFEB solutions for the benefit of the entire ecosystem, a number of leading companies from throughout the design chain have formed an ecosystem educational forum called the eBeam Initiative.

With members representing the entire value chain, from intellectual property (IP) and electronic design automation (EDA) companies to semiconductor manufacturers and equipment makers, system design companies and research entities, service companies and mask makers, the eBeam Initiative is expected to accelerate greatly production-oriented EbDW technology using DFEB. The initiative also includes representatives from the design community who will serve in an advisory capacity (see Figure 3).

Figure 3: Charter members and design team advisors of the eBeam Initiative

Roadmap to DFEB

Driven by DFEB's promise of more design starts and faster time-to-market, the members of the eBeam Initiative have already started working together to create a production-proven path to DFEB-based designs. Various eBeam Initiative members have collaborated to validate maskless manufacturing with successful test wafers for the 45- and 32-nm nodes. The initiative roadmap (see Figure 4) calls for proven manufacturing in 2009, proven design in 2010, and multiple chip suppliers and DFEB certification in 2011.

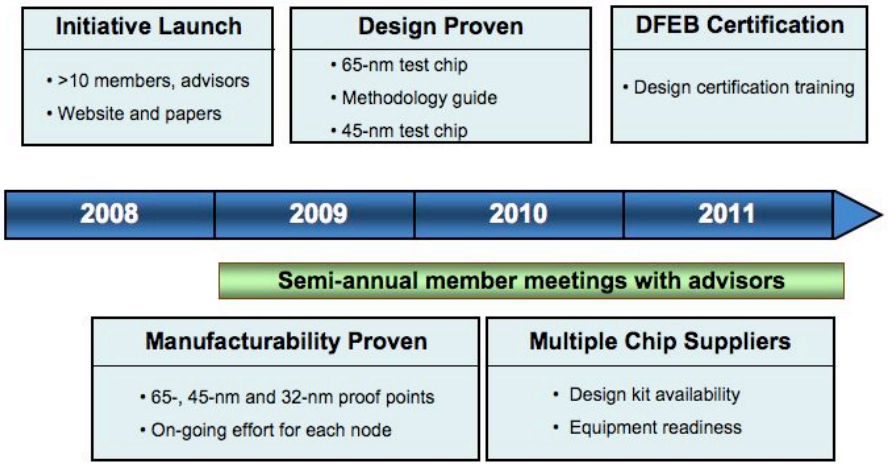

Figure 4: eBeam Initiative Roadmap