Design Considerations for Regular Fabrics

Deepak D. Sherlekar
Distinguished Member of Technical Staff
Virage Logic Corporation
47100 Bayside Parkway
Fremont, CA 94538
1-510-743-8124
d.sherlekar@viragelogic.com

ABSTRACT
Structured ASICs are an emerging new class of ASICs that attempt to bridge the widening gap in per-unit manufacturing costs, non recurring engineering (NRE) costs, power consumption, and performance between zero-mask programmable devices such as FPGAs and devices such as cell based ASICs, which require new custom designed masks for every ASIC. They offer an intermediate trade-off point between the two extremes of the very high per unit cost, but zero non-recurring cost of FPGAs, and the very low per unit cost, but very high non-recurring cost of cell based ASICs. They also offer a similar, intermediate trade-off point between the two extremes for performance and power consumption. A common theme across all structured ASICs is the use of a circuit fabric that has a regular, repeating pattern of elementary building blocks that can be programmed using one or more masks to implement an ASIC device. In this paper, we describe the considerations involved in designing the regular circuit fabrics underlying structured ASIC offerings.

Categories and Subject Descriptors
B.7.1 [**Integrated Circuits: Types and design Styles**]: VLSI; gate arrays. B.7.2 [**Design Aids**]: Placement and routing.

General Terms
Design, Experimentation, Performance, Economics, Reliability.

ISPD'04, April 18-21, 2004, Phoenix, Arizona, USA.

Keywords
Structured ASIC, regular fabric

1. INTRODUCTION
The high levels of integration offered by modern semiconductor process technologies have made complex, System on Chips (SoCs) the most common form of ICs. However, the advantages of new semiconductor processes come with several challenges. Today's chip designers have to contend with often conflicting demands of performance, power, cost, reliability, manufacturability, and time-to-market.

Structured ASICs have been proposed as a candidate solution for many of the challenges mentioned above. However, the two most important factors that have contributed to the recent spurt in interest in the technology are the cost of masks and the power-performance tradeoff they offer relative to FPGAs and cell-based ASICs.

The cost of masks has been nearly doubling with every process generation for the past several years. The cost of a full set of masks has risen to over $600K for 130nm, and to $1.5M at 90nm. The problem is compounded by a rise in the number of design re-spins needed before a chip reaches volume production. Such re-spins are required for several reasons, such as:

- A revision in the functional specification to respond to issues exposed in system-level prototyping, product tests in the field, or changes to standards or protocols used in the ASIC.
- Errors in the functional specification of the ASIC
- Functional errors that were not caught before tape out.
- Deep submicron related problems such as yield, noise, or crosstalk.

As a result, it is not unusual to see as many as 4-6 re spins before a chip reaches volume production. Thus ASIC

designers using cell based ASICs at 90nm, may face as much as $9M is mask costs alone before their chip reaches volume production.

FPGAs provide an alternative to the high mask costs of cell based ASICs. They offer a significant and attractive trade off point at the other end of the spectrum from cell based ASICs, with zero masks and manufacturing turn around time, but an order of magnitude or higher area and power consumptions, and 2X to 4X lower performance.

The larger area limits the capacity of FPGAs. The steep power penalty results in significantly higher cost of packaging, and limits applicability in portable and battery powered applications. The hefty per-unit cost of FPGAs also limits their use in medium to high volume applications.

A spate of new offerings in the Structured ASIC category is indicative of the urgent need to provide solutions that bridge the gap between the two dominant chip design offerings prevalent today – viz. FPGAs and cell based ASICs.

Section 2 introduces the dominant themes behind structured ASIC offerings available today, and briefly covers the structured ASIC offerings from major providers.

Section 3 discusses the considerations involved in designing the optimal design fabric to implement the programmable logic component of a structured ASIC. It underscores the importance of developing a solution though a collaboration between circuit and layout designers, EDA tool developers, and semiconductor process technology experts.

Section 4 describes a regular, fine-grained, metal programmable design fabric underlying Virage Logic's Metal Programmable cell library products, and the trade-off between number of routing layers and density using this fabric.

Section 5 suggests coarser grained variations of the fabric described in Section 4, and ways to explore alternative trade-offs involving the number of routing layers, the cost of different masks, and the potential density/performance achievable by resulting solutions.

Concluding remarks appear in Section 6.

2. What is a Structured ASIC?

A typical Structured ASIC comprises of some or all of the following features:

- A programmable logic fabric
- A collection of pre-placed, configurable memory blocks
- Pre-configured I/O
- Cores such as MPU or DSP
- SerDes cores to interface with various standard high-speed interfaces
- Clocking circuits such as DLLs and PLLs
- Configurable clock distribution networks
- Built in DFT capability
- Pre-configured power distribution grid
- Pre-routed, configurable interconnects

I/O, SerDes, and clocking circuits are likely implemented as custom macros/blocks, and are pre-verified in the context of the product. The same holds for cores such as MPU or DSP, if they are present. Given the nature of these large, complex blocks, structured ASIC products that contain them are usually targeted toward a well-defined class of applications, and are not general purpose.

The programmable design fabric underlying structured ASIC products consists of a two dimensional array of basic cells. These cells are comprised of partially or fully constructed combinational and sequential logic functions. Conventional gate array core cells, which are optimized to implement a NAND2, are the oldest example of such cells.

Figure 1 below, taken from the prior art section of US Patent 5723883, shows a typical conventional gate array cell implementing a NAND2 function using one core cell. The core cell consists of 2 P type and 2 N type transistors, with power and ground terminals VDD and VSS. The inputs of the NAND2 are labeled A and B, while the output is labeled Z.

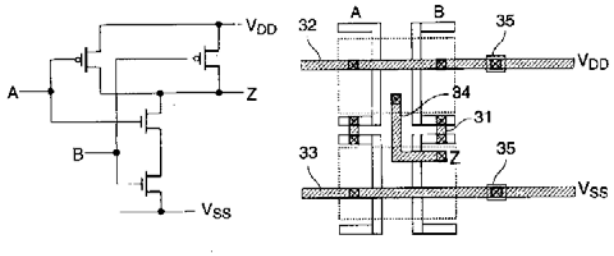

Figure 1

Several vendors have introduced commercial structured ASIC products over the past year. They include

- Altera: Their HardCopy product enables customers to migrate their high volume designs from FPGA to a metal programmable platform in a cost effective and predictable manner.
- AMI Semiconductor: Their XpressArray product uses newer process technology for the base layers, and

older, more mature, fabrication process for metal programmable interconnects.

- ChipExpress: Offers 2 or more layer programmable arrays called Modular Arrays
- Faraday: Their structured array, called Metal Programmable Cell Array, or MPCA, uses 3 or more metal layers for interconnect. It is covered in a separate paper in this conference.
- NEC: Their Instant Silicon Solution Platform (ISSP) also uses a metal programmable structured array with built in support for clock, power and test, and a rich set of embedded IP cores. It is also covered in a separate paper in this conference.
- Fujitsu: Their AccelArray product also has built in support for clock, power, and test, with IP cores tailored for telecommunications and industrial automation markets. The IP cores and base array are implemented with the bottom 3 interconnect layers, and the top three layers are user programmable.
- LSI Logic: Their RapidChip product line offers metal programmable array for user specific logic, with a carefully chosen set of hard IP blocks tailored for specific market segments.

Other vendors offering structured ASIC products include Lightspeed semiconductor, Leopard Logic, and eASIC. A lot of the structured ASIC vendors rely on custom design tools, especially for physical design, to complement commercial EDA tools and offer a smooth and reliable design flow. However, some vendors have tailored their offerings to leverage existing FPGA or ASIC design tools.

3. Programmable Logic Fabric Design

The programmable logic fabric underlying commercial structured ASIC offerings available today share the following traits:

- The fabric consists of a regular pattern formed by tiling one or more basic cells in a two-dimensional array, with holes carved out to accommodate other components such as IP cores mentioned above.
- The basic cell(s) forming the fabric consists of wholly or partially built combinational and/or sequential logic functions.

Basic cells that implement whole functions often resemble basic logic blocks in FPGAs, but without the overhead required for field programmability. Design tools for such fabrics can leverage ideas from FPGA design tools, but still require sufficient customization to warrant a dedicated tool suite for a streamlined flow with good optimization.

Cells comprised of partially built functions may consist of patterns of transistors that are partially wired to enable efficient creation of frequently used functions. A design fabric may be composed of more than one type of basic cell. Designated terminals in the basic cells may be connected using programmable interconnects to create different logic functions. The interconnection patterns for different functions are represented in the form of a cell library similar to a gate-array or a standard cell library. The locations and types of the underlying basic cells of the fabric must match that of the library cell for legal placement. Thus the job of creating a legal placement that adheres to the constraints of the fabric may be more complex than that of conventional standard cell placement.

Coarse-grained programmable design fabrics, which are implemented with large base cells having many functions, may pose additional challenges in achieving a trade-off between efficient use of all logic inside a base cell, and minimizing the lengths of interconnects between different base cells. This makes it extremely difficult to design a design fabric without simultaneously exploring its amenability to optimization with corresponding development of EDA tools. As a result, many commercial programmable design fabrics for structured ASICs either use fabrics similar to gate arrays, or, like FPGA vendors, simultaneously develop their own design tools. Likewise, significant academic research has taken place on the design of programmable fabrics as well as related design tools, such as via-programmable gate arrays (VPGAs).

4. Virage Logic's ASAP MP Design Fabric

Virage Logic's ASAP Logic product line offers custom standard cell libraries for many of the world's leading foundries such as TSMC and UMC as well as for many of the world's top integrated device manufacturers. These cell libraries employ a patented routing architecture, called NXT (US patents 5,723,883; 5,898,194; 5,923,059; 5,923,060; 6,091,090).

An innovative metal programmable logic design fabric employing the same routing architecture has also been developed. The fabric consists of a single type of base cell that has 4 transistors (two P and two N). A single base cell can implement functions such as a NAND2, a NOR2, a buffer, or an inverter. All cells in the accompanying cell ASAP MP cell library are made up of one or more horizontally adjacent base cells.

The fabric is created by tiling the base cell in a 2 dimensional pattern arranged in rows that are alternately flipped (mirrored) about the X axis. Thus the placement terrain for the fabric looks much like that of a standard cell design, with alternate rows flipped and abutted. The only difference is that unlike standard cells, the MP cells of the

fabric can only be placed at the corresponding locations of the underlying base cells, which is at a granularity of an integral multiple of the width of the base cell from the edge of the row. Further, unlike standard cells, an MP cell may not be flipped about its y axis within a row.

The unit base cell underlying the design fabric for TSMC 90nm, including connections on Metal1 that are made to implement a 2 input NAND function is illustrated in the figure below.

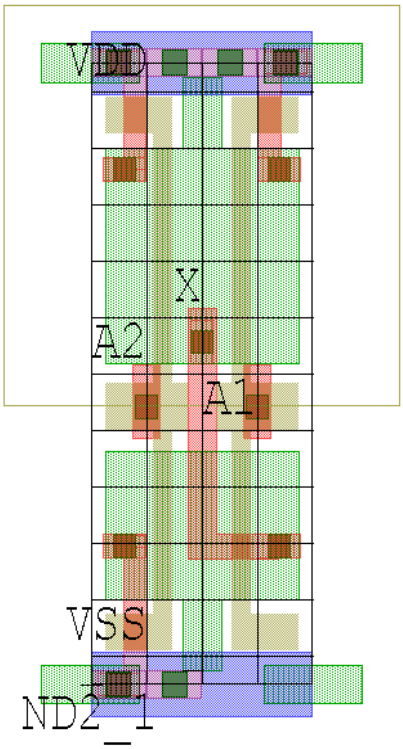

Figure 2.

NAND2 created with ASAP Logic Base MP Cell

The unit cell is 4 grids wide and 11 grids tall, and was optimized for high performance. The metal pitch is the minimum possible in both directions. The nwell and substrate contacts straddle all four corners, and are shared with adjacent cells. Robust 2 tracks wide power and ground rails on Metal2 are provided.

All placement constraints for the ASAP MP cell library can be easily expressed in the description of the placement terrain for all major EDA tools during the library preparation phase. This enables ASIC designers to use the library in place of a standard cell library on a drop-in basis, with the same flow that is used for standard cell libraries.

4.1 ASAP MP Design Fabric Features

The NXT routing architecture underlying the ASAP MP design fabric and cell library can provide significant advantages in raw as well as routed density. The use of horizontal M2 for routing frees up significant M1 routing resource for cell layout. We have realized significantly lower use of M2 in cell layout outside of the power and ground rails. This enables creation of cells as short as 9 tracks in height. The short cell results in significant reduction in raw cell area.

The use of horizontal M2 at the cell level extends naturally to routing at the block and chip level. Thus the primary direction of M2 is horizontal, with M3 and subsequent layers alternating in direction.

The use of horizontal M2 dramatically increases the number of M2 tracks that can access the pins on a cell. This increase in a pin's capture cross section reduces the number of jogs in routing in the vicinity of a cell, thus minimizing local congestion around it.

The use of vertical M3 enables cells to be placed under M3 power straps without causing blocked pins.

The higher EM limits on M2 make the M2 power and ground rails significantly more robust than M1. Further, the cell layouts are constructed such that all pins are two or more tracks away from the power and ground rails. As a result, for very high performance applications, one can easily widen the default 2 tracks wide rails to either 3 or 4 tracks without blocking access to pins. This programmability of the cell power rail can also be used to alleviate routing bottlenecks. For example, designs that are limited by vertical routing resources can use wider rails with less frequent vertical power straps, while designs limited by horizontal routing resources can use thinner rails with more frequent power straps.

The above advantages in routing, coupled with improvements in raw cell area, result in dramatic improvement in routed density. As a result, the routed area achieved is only 5% to 30% larger than that possible with conventional, 9 tracks tall standard cell libraries. The performance achieved is within 5% of that possible with standard cell libraries. The density can easily be improved further by reducing the cell height fro 11 to 10 tracks.

4.2 ASAP MP Design Fabric Applications for Structured ASICs

The ASAP MP design fabrics and cell libraries available on standard foundries can be readily used to create Structured ASIC products. In fact, two of the Structured ASIC offerings mentioned in Section 2 already employ a process optimized, custom ASAP MP design fabric and cell library. An important feature of these products is that they can use mainstream ASIC design tools, and do not require features

that are not available in mainstream EDA tools for ASIC design. Further, the density and performance is comparable to that of standard cell ASICs.

The biggest uncertainty in putting together a structured ASIC is in the specification of the hard IP blocks that go in it. However, this is a manageable problem for companies trying to develop a closely related set of ASICs using the structured ASIC paradigm. The availability of the ASAP MP design fabric and cell libraries on standard foundry processes enables COT ASIC designers to define their own custom structured ASIC product that can be used to derive multiple, closely related products.

4.3 Density vs. Masks Trade Off

The ASAP MP design fabric and cell library offer a meaningful trade-off between number of routing layers and routed density.

Figure 3 below plots a chart of density against number of masks.

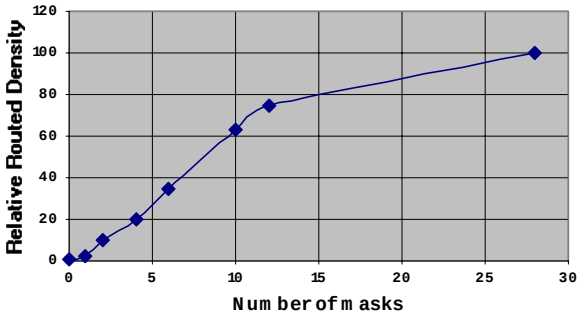

The first point, for zero masks, is for FPGAs. Subsequent points for 1 and 2 masks are for various design fabrics that employ either 1 routing layer, or 1 routing and 1 via layer.

The points from 4 to 12 routing layers are with the ASAP MP design fabric, while the last point, with 28 masks, is for standard cells.

The routed density with the ASAP MP design fabric can be improved by minimizing the use of the programmable routing layers for power and clock routing. Thus one can use the top routing layers, which are usually thick layers with low resistance, for power and clock routing at the chip level, and maintain only a skeleton power redistribution grid in the programmable logic region, with connections to the global power grid routed on the top layers.

5. MP Design Fabric Variations

The MP design fabric of Section 4 can be improved by flipping alternate base cells in the same row about the y-axis. This arrangement enables the number of Nwell and substrate contacts to be cut in half, which can help reduce the width of the core cell by over 12%, or increase the total transistor width by over 15%. However, this change in the fabric adds a small complication to the legal locations check for placement, which is not handled by most EDA tool vendors today.

Another area of improvement is eliminating the use of contact and M1 masks, which are somewhat more expensive than other routing layer masks. In this case, the base cell will include features on the contact and M1 layers as well. Macro cells can be created by routing connections on M2 and M3. In this case, the power and ground rails can be buried in M1, thus freeing up more resource on M2, and minimizing the impact of losing a routing layer. One can carry this a step further by moving programmable layers to M3 and above. However, in order to make the implementation efficient, it is necessary to minimize the number of terminals in the base cell that must be moved to the top. This may be done by using a larger base cell with more transistors. However, larger cells are inefficient unless one can pack multiple small functions such as buffers, inverters, NAND2 or NOR2 into a single base cell. Exploration of such architectures should only be done in conjunction with a placement program that can effectively achieve such packing.

The regular, repeating pattern of shapes in Metal Programmable design fabrics minimizes the impact on many DSM manufacturing issues. Thus, spacing rules such as those for poly and diffusion can be tightened significantly since they need not account for worst-case patterns possible in standard cell or custom designs. In one situation dealing with a company with a captive semiconductor fab, we were able to tighten several design rules sufficiently to fit the same 4-transistor base cell of the MP fabric in 10x3 grids, which is not only smaller by over 30% than our current 11x4 base cell, but yields smaller NAND2 and NOR2 gates than possible with even standard cells.

6. Conclusions

The use of metal programmable regular design fabrics in structured ASICs provides significant benefit in mask costs over conventional cell based ASICs with a small loss in density. While this feature provides significant benefit, there are other considerations that are worth examining. For example, in order to reduce manufacturing TAT, it is necessary to limit programming to only the top few layers, rather than all or only the bottom layers. However, the exploration of efficient structures for programming with just the top few layers will require simultaneous exploration of design automation algorithms with the design of such structures.

The regularity in structure also provides better correlation between Silicon and simulated results with SPICE. The tighter correlation can be used to tighten the tolerance between best, worst, and typical case delay models.

The composition of a structured ASIC offering that will be beneficial to a wide array of applications remains a challenging problem. Unless we can develop platforms with universal applicability over a reasonably large set of ASICs, it will be difficult to realize the full potential of this technology.

REFERENCES

[1] Gheewala, Tushar R. Gate Array Cell Architecture and Routing Scheme. *U.S. Patent 5923059, July1999*.

[2] Gheewala, Tushar R. Integrated Circuit Cell Architecture and Routing Scheme. *U.S. Patent 5898194, April1999*.

[3] Gheewala, Tushar R. Integrated Circuit Cell Architecture and Routing Scheme. *U.S. Patent 5923059, July1999*.

[4] Gheewala, Tushar R. Reduced Area Gate Array Cell Design Based on Shifted Placement of alternate Rows of Cells. *U.S. Patent 5923060, July1999*.

[5] Gheewala, Tushar R. Power and Signal Routing Technique for Gate Array Design. *U.S. Patent 6091090, July2000*.

[6] Patel, Chetan, et al. An Architectural Exploration of Via Patterned Gate Arrays. ISPD Proceedings, April 2003.