Sub-20 nm Design Technology Co-Optimization for Standard Cell Logic

Kaushik Vaidyanathan*, Lars Liebmann**, Andrzej Strojwas*, Larry Pileggi*
* Carnegie Mellon University, Pittsburgh, PA 15213 ** IBM, East Fishkill, NY 12533
kvaidya1@andrew.cmu.edu

***Abstract*—Efficiency and manufacturability of standard cell logic is critical for an IC, as standard cells are at the heart of the nexus between technology definition, circuit design and physical synthesis. Conventional standard cell design techniques are increasingly ineffective as we scale to patterning restricted sub-20 nm CMOS nodes. To meet the constraints and leverage the features of future technology offerings, we propose a holistic design technology co-optimization (DTCO) for standard cell logic. In our holistic DTCO we co-optimize the standard cell architecture to balance manufacturability and efficiency at the cell level while taking into account block level considerations such as pin accessibility and power rail robustness. Our DTCO in a foundry 14 nm CMOS resulted in two standard cell architectures, namely, 10T_BiDir and 10T_UniDir. We evaluated these cell libraries with physically synthesized blocks and ring oscillator test structures in IBM 14SOI process. We observed that 10T_BiDir emerges as the preferred alternative at 14 nm CMOS, with 10T_UniDir promising better scalability to future nodes.**

Keywords—Standard cell logic; 14 nm CMOS; Design Technology Co-Optimization; FinFET device; Multiple patterning.

I. INTRODUCTION

Standard cell logic along with SRAMs and analog components, forms one of the three critical components of a modern IC. Most of the logic blocks in an IC are implemented using standard cell libraries and physical synthesis [1]. Interestingly, with technology scaling over the last decade, the fundamental standard cell layout architecture has not changed significantly. However, as we scale below 20 nm CMOS, three technology elements, namely, FinFET devices, local interconnects and multiple patterning for critical design layers are being extensively used. [2]. Continued use of resolution-limited 193 nm immersion (193i) lithography at the 14 nm and 10 nm nodes requires multiple masks to pattern a single design layer, increasing manufacturing cost and auto-router complexity. These sub-20 nm technology challenges are leading to an increase in cost-per-gate [3]. As conventional standard cell design methods fail to adapt to the changing technology requirements, it is crucial to rethink conventional standard cell design practices as they directly impact manufacturing cost, design efficiency and turnaround time.

Design technology co-optimization (DTCO) has been successfully applied to SRAM bitcells for over a decade to systematically explore different technology options and design styles, to improve design efficiency and manufacturability [4]. We undertake a holistic DTCO for standard cells, where we co-optimize technology, layout, circuit and electronic design automation (EDA) tools to converge on an efficient, manufacturable, productive and scalable standard cell solution. To explore a scalable standard cell architecture, we carry out the proposed DTCO in a 14 nm foundry process while applying additional patterning and process integration constraints from the 10 nm node.

Our holistic standard cell DTCO has four steps (Figure 1(a)), with the first step being technology definition. In the past few years, as discussed by Northrop in [6], there has a significant thrust to move towards a design-aware process technology definition. While the final technology definition is converged upon after several iterations with cell designers, initial standard cell architecture exploration (in step 2) begins with a preliminary technology definition. Apart from considering typical standard cell objectives such as track height, active area efficiency and cell parasitics, additional sub-20 nm-specific objectives such as minimizing manufacturing cost and complexity by minimizing number of patterning exposures are considered. Standard cell architectures that meet cell level design objectives are analyzed next for several block level considerations such as pin accessibility, power rail robustness, color safe boundary conditions etc, in step 3. Contrary to previous process technologies, sub-20 nm CMOS processes require block level objectives to be considered on par with cell level objectives. For instance, narrow sub-40 nm metal widths are very susceptible to electromigration, requiring careful cell level and block level power rail design [2][5]. With increasing auto-router complexity and litho-hotspots seen in standard cell pin connections, pin accessibility of standard cells become important [5]. After analyzing standard cells for such block level considerations, we evaluate them, in step 4, at the library level with a 40 cell library, using physically synthesized logic blocks and with test circuits on IBM 14SOI process.

Key observations made from a foundry 14 nm standard cell DTCO are as follows. First, 10T_BiDir with restricted bidirectional (2D) M1 emerges as the preferred standard cell architecture, balancing manufacturing complexity and design efficiency at the 14 nm node. Second, 10T_UniDir with unidirectional (1D) M1 while being efficient and cost-effective, is still not ready for use in the 14 nm node. As an additional outcome of the holistic DTCO, we also present a list of few critical pattern constructs that a foundry has to support for efficient and manufacturable sub-20 nm standard cell design, detailed in [18]. We conclude the paper with results from patterning experiments that demonstrate the scalability of

these standard cell architectures to future nodes such as 10 nm and 7 nm. Key contributions of this paper are:

- Undertake holistic DTCO for standard cells, to balance design efficiency and manufacturability in sub-20 nm CMOS.
- Demonstrate the effectiveness of holistic DTCO for standard cell logic in foundry 14 nm process.
- Identify critical pattern constructs that are necessary for standard cell design in future nodes.

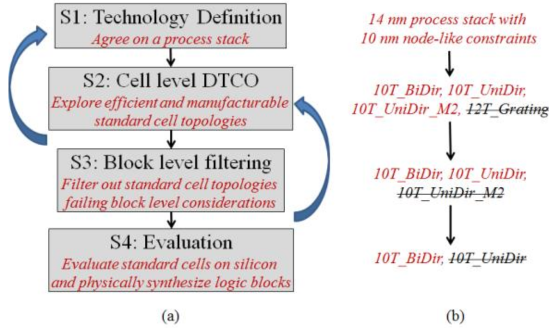

Figure 1. (a) Sub-20 nm holistic DTCO flow for standard cell logic; (b) Example 14 nm holistic standard cell DTCO.

II. Background And Motivation

In this section, we introduce a sub-20 nm CMOS process stack. Following this, we list out standard cell design goals and discuss how conventional standard cell design techniques are inadequate in meeting those objectives in sub-20 nm CMOS.

A. Sub-20 nm CMOS Technology

Three technology elements that will be used extensively in a typical sub-20 nm process are FinFET devices, local interconnect and multiple patterning for critical design layers, increasing process risk, complexity and cost [2][6]. We illustrate these technology features using an example NAND2 gate (Figure 2(a)). FinFET devices use a grating-based (equal line widths and equal spaces) fin layer (FIN), requiring circuit designers to size FinFET devices as discrete number of fins [2]. Active fins are identified using a diffusion layer (ACTIVE), and notches in the diffusion layer are not encouraged to contain process variability. Intersection of poly and active fins forms the transistor. Extremely restricted poly and fin layers are robustly contacted by metal-like rectangular contacts, called local interconnects [6]. Local interconnect contacting the poly is CB and contacting the fin is CA. Owing to their impact on device variability, FEOL and local interconnects are multiple patterned and restricted to be unidirectional and occasionally, gratings. While 1X or minimum width metal layers also run at very fine pitches, their restrictiveness is less agreed upon as they do not directly contribute to device variability. There is a strong incentive to make M1 and other 1X metal layers unidirectional, as it decreases hot-spot risk by limiting the number of unique layout neighborhoods and reduces patterning complexity [7]. On the same token, a more restricted BEOL layer curbs design freedom, potentially resulting in design inefficiencies. Hence, sub-20 nm CMOS technology is significantly different from its predecessors and requires us to investigate the effectiveness of conventional standard cell design techniques in meeting standard cell design objectives.

B. Standard Cell Design Considerations

Due to restrictive patterning, standard cells are at the nexus of technology, circuits and physical synthesis, requiring standard cell design to meet several objectives at the cell level and block level.

1) Cell Level Objectives

Design efficiency of a standard cell is determined by track height, active area efficiency and cell parasitics. Track height is the number of M1 tracks within a standard cell and a moderate track height of 8 to 10 typically balances performance and area efficiency. Active area efficiency is the ratio of active area to the total cell area, and moderate active area efficiency is between 65%-75%. Reducing metal2 (M2) and via1 (V1) usage in the cell reduces cell parasitics, improving efficiency.

With lithography dominating overall manufacturing cost and complexity, curtailing the number of mask levels required to pattern critical design layers is important, while providing reasonable design freedom. The choice of restrictiveness and directionality of 1X metal layers made during standard cell design impacts manufacturing cost and complexity. Cost-effective robust manufacturability is an additional objective for standard cell design in a patterning restrictive regime.

2) Block Level Objectives

As the complexity of technology files for auto-routers and consequently, the associated physical synthesis design turnaround times continue to increase, the choice of appropriate standard cell architecture becomes important [5][8]. Several block level considerations, such as pin accessibility and power rail robustness, have to be co-optimized in conjunction with the BEOL stack and auto-router used in physical synthesis.

Furthermore, to exploit the features and work past the extreme patterning restrictions in sub-20 nm CMOS, a holistic system-on-chip (SoC) view is essential. As a digital SoC contains SRAMs apart from standard cell logic, it is important to be consistent with the restrictiveness and directionality of different layers in the process stack across different components in the SoC. As it may be apparent, many of these objectives arise in response to the challenges posed by sub-20 nm CMOS technology.

C. Conventional Standard Cell Design Techniques

1) Geometric Scaling

One classic technique to design standard cells for current technology nodes is to keep the same layout topology as in the previous technology node, while following design rules in the current node. As this technique has been effective in the previous nodes, we attempted a geometric shrink of a 32 nm NAND gate to meet 14 nm design rules. As shown in Figure 2(a), the resulting NAND gate is area inefficient, has non-connectable inputs pins, requiring at least a triple patterned M1 layer and has severe difficulties in manufacturing minimum M1 area shapes (Table 1). Clearly, geometric

scaling is not an effective in meeting sub-20 nm standard cell design objectives.

2) Extremely Restricted Design

Unidirectional (1D) layout is less complex and more cost-effective to manufacture than 2D layout shapes. Jhaveri et al. proposed to adopt 2D layouts from previous technology nodes and then replace bidirectional layout shapes with two unidirectional layout shapes and a via [7]. To test the effectiveness of this approach in 14 nm CMOS we build a NAND2 gate by keeping the same layout topology as the previous technology nodes, but replacing bidirectional M1 with a vertical M1, horizontal Metal2 (M2) and via1 (V1) (Figure 2(b)). While all layers in the resulting gate are 1D, it is area inefficient (12 M1 tracks tall, 5 poly pitches wide), increasing the number of vias (V1) and containing a number of process sensitive minimum M1 area shapes. Hence, such a localized (patterning-specific) approach only moves the problem from lithography to design and process integration.

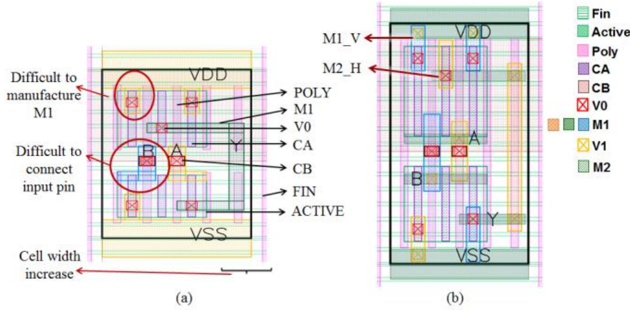

Figure 2. (a) Geometrically scaled NAND2 in 14 nm (3 M1 exposures); (b) 2D to 1D mapped NAND2 in 14 nm CMOS (12 M1 tracks tall).

Table 1. Comparison of conventional standard cell design techniques as applied to sub-20 nm CMOS.

*Attributes**	*Geometric shrink*	*2D to 1D mapping*
Cell area	>25%	>50%
Process integration concerns	Min-M1 area pins	Min-M1 area, excess vias
Patterning concerns	3 M1 exposures	2 or 3 V1 exposures
Router pin access	Poor	Poor
Power rail robustness	Moderate	Moderate

* All attributes compared with 14 nm DTCO'ed cells - 10T_BiDir and 10T_UniDir

III. Sub-20 nm Standard Cell Design Technology Co-Optimization (DTCO)

Design technology co-optimization (DTCO) enables us to systematically explore different technology options and design styles [6]. In this section, we first describe the holistic standard cell DTCO (Figure 1(a)) and then apply holistic DTCO for standard cell logic in a foundry 14 nm process (Figure 1(b)). It is worth noting here that while EDA tools and custom scripts are used extensively in the DTCO process, it is primarily driven by designers with the help of process engineers.

A. Overview of DTCO Process

Step 1: Technology definition. Converging on a technology definition that is both cost-effective and also meaningful to designers is the first step in DTCO [6]. A technology definition specifies the directionality, connectivity, widths, spacings and pitches of different layers in the process stack. While the process of converging on a technology definition is iterative, through DTCO, designers and process engineers can significantly minimize the number of iterations.

Step 2: Cell level DTCO. With a preliminary technology definition, designers explore different standard cell architectures and evaluate them for design efficiency and manufacturability. Keeping cell level considerations in mind (in Section II.B.1), designers create standard cell architectures and design a handful of representative cells, such as, 2 input NAND gate (NAND2), a D flip flop (DFFQ), a 4 input AND-OR-INVERT (AOI22) and a 2 input XOR gate (XOR2). These cells are carefully designed to optimize area efficiency, minimize parasitics, minimize manufacturing complexity and maximize robustness. As good spice and interconnect models are not available early in a process, delay and power areprimarily optimized by minimizing cell parasitics, area, and by choosing appropriate track heights. It is well acknowledged that lower track heights between 8 and 10 are more suited for low power and compact logic blocks, and track heights above 10 are used for high performance logic blocks [5][17].

After a few iterations with different technology definitions, good standard cell architectures are identified and passed over to be evaluated at the block level. While at the outset, this step might seem intractable with infinite possibilities, in reality, with severe patterning restrictions in deeply scaled nodes, the number of standard cell architectures and associated technology definitions dramatically reduces to a small number. It is worth noting that in our holistic approach, SRAM bitcells are also DTCO'ed along with standard cells, and details can be found in [10].

Step 3: Block level filtering. Standard cell architectures that pass cell level analysis are next evaluated at the block level. In sub-20 nm CMOS, the key block level considerations are BEOL stack, pin accessibility, power rail robustness and color-safe boundary conditions (Section II.B.2). We describe the importance of these considerations with examples in Section III.B.2.

Step 4: Evaluation. Standard cell architectures that pass cell level and block level analysis are evaluated in two ways. First, the standard cell performance is characterized on silicon using standard ring oscillator test structures. Silicon characterization is essential to verify if the cells behave as projected by the transistor and interconnect models, especially given that DTCO is done in the pre-production stage of cutting edge CMOS processes. Second, the block level behavior of the standard cell library is evaluated by physically synthesizing logic blocks. Comparison of power, performance and area of physically synthesized blocks would enable us to validate the efficiency of a standard cell library and its compliance to commercial EDA tools and flows.

B. 14 nm Standard Cell DTCO

Step 1: 14 nm technology definition. As the composition of the CMOS process stack is not expected to change significantly until the 7 nm node [2][6], we use the process stack presented in Section II.A. While we could not significantly influence the technology definition as we working in a foundry 14 nm process, we artificially imposed 10 nm technology restrictions (Table 2). Preliminary technology definition is typically converged upon after consideration of different patterning techniques and their associated resolution, misalignment, cost and complexity. The final technology definition is converged upon after a few iterations with leaf cell designers.

Table. 2. Preliminary sub-20 nm technology definition.

Layer	Patterns
Poly, Fin, CA	Pure grating 1D
Active	No Diffusion Notch
CB	1D
M1	1D or Restricted-2D
M2, M3	1D

Step 2: 14 nm cell level DTCO. In this step we explored several standard cell architectures and discuss a few good topologies (10T_BiDir, 10T_UniDir, 10T_UniDir_M2) that progressed to step 3 and also a few not so good topologies (10T_BiDir_CB, 12T_Grating) that were left out.

1) 10T_BiDir – Restricted Bidirectional-M1 Standard Cell

In the 10T_BiDir standard cell layout (Figure 3(a)), we have attempted to retain most of the attributes of a typical bidirectional-M1 standard cell layout, such as power rails shared across the adjacent rows of standard cells and input pins pushed towards the center of the cell. However, to comply with patterning restrictions in sub-20 nm processes, while still retaining its design efficiency, the bidirectional standard cell layout has had to evolve further. Its key distinguishing features are:

- 10T_BiDir contains a robust power rail structure made of local interconnects and metals. The use of this novel and robust power rail structure, instead of the traditional M1 taps, serves as a key feature of 10T_BiDir.
- CA taps are used to connect the source/drain of the transistors to the CB power rail, avoiding difficult to manufacture M1 taps (Figure 3(a)).

These features enable 10T_BiDir to be area efficient while not pushing lithography resolution limits. The 10T_BiDir cells are 10 tracks tall and have an acceptable active area efficiency of 66.67%, with restricted-2D M1 that is compatible with cost-effective patterning techniques such as self-aligned double patterning [11].

2) 10T_UniDir – Unidirectional-M1 Standard Cell

10T_UniDir standard cell is 10 M1 tracks tall with M1 restricted to be unidirectional (1D) structured grating (Figure 3(b)). A structured grating has unequal line widths and equal spaces [16]. While structured 1D gratings are amenable to cost-effective patterning techniques such as self-aligned double patterning (SADP) [11] and directed self-assembly (DSA) [12], they also provide more design freedom than pure gratings (Section II.C.2). Nevertheless, 1D M1-based standard cells proposed earlier, while having favorable manufacturability, have been unsuccessful in meeting design requirements. 10T_UniDir has been carefully designed to be manufacturable without compromising design efficiency. Key distinguishing attributes of 10T_UniDir are:

- M1 in the cell is perpendicular to poly and M2 is parallel to poly. The choice of these orientations comes from the fact that M1 has to be perpendicular to the first extensively used local interconnect level (CA is parallel to poly) to maximize and ease connectivity.
- To minimize M2 usage in the cell, the CA layer is used – beyond its envisioned usage – to connect N type and P type transistors, making it robust with less parasitics.
- Input pins connect to poly at either the center or the bottom edge of the cell layout. Pin locations have been strategically chosen to enable the use of CA to connect N type/P type transistors, as the presence of all CB gate contacts in the center of the cell will disallow any CA connection between N type/P type transistors. Furthermore, this allows the input pins to be spaced further away from each other, improving the pin access. Inputs pins are also widened to avoid process-risk-prone minimum metal area shapes.
- Power rails are 2X wide improving electromigration tolerance and lowering IR drop. With power rails inside the cell 10T_UniDir tiles without mirroring, unlike a typical standard cell that has shared power rails at the edge of the cell.

These features allow 10T_UniDir to be as area efficient (10 tracks tall, 66.67% active area efficiency) as 10T_BiDir while being amenable to cost-effective patterning.

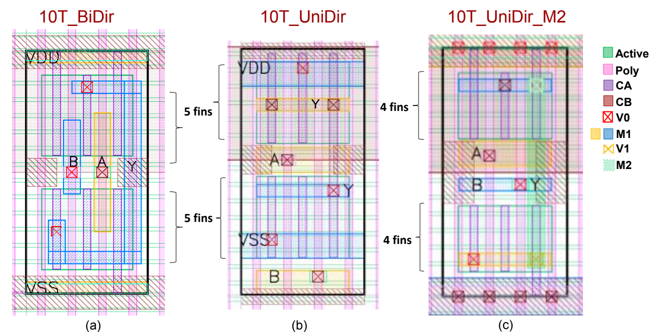

Figure 3. Sub-20 nm standard cell DTCO outcomes. (a) 10T_BiDir, (b) 10T_UniDir, (c) 10T_UniDir_M2.

3) 10T_UniDir_M2 – Unidirectional-M1 Standard Cell with M2

10T_UniDir_M2 (Figure 3(c)) is also 10 tracks tall with M1 and M2 restricted to be unidirectional (1D). 10T_UniDir_M2 and shares similarities with both 10T_UniDir and 10T_BiDir with the following differences:

- Input is moved to center of the cell and the power rails are shared between adjacent rows and pushed to the edge.
- M2, instead of CA, is used to make the output pin connection between NMOS and PMOS transistors.

- Has an active area efficiency of 53.33%.

10T_UniDir_M2 allows for structured 1D grating-based M1 and M2, while using proven standard cell architecture features.

4) 12T_Grating – Grating-based Unidirectional-M1 Standard Cell

Kornachuk and Smayling proposed to build standard cells exclusively out of pure gratings (equal line widths and equal spaces) that are relatively easy to manufacture [9]. To evaluate the effectiveness of this technique in 14 nm CMOS, we restricted all FEOL, local interconnects and BEOL layers to be pure gratings. Our experiment reveals that pure grating-based cells (12T_Grating) are area and power inefficient (Figure 4). Furthermore, these cells i) have minimum M1 area input pins, that raises serious process concerns and ii) minimum width (1X) power rails that are susceptible to increased IR drop and electromigration. Hence, we disregard 12T_Grating as an option going forward to step 3.

Similar to 12T_Grating, several standard cell architectures were explored and discarded as they were either inefficient and/or had manufacturing concerns. A handful of representative cells were designed for 10T_BiDir, 10T_UniDir and 10T_UniDir_M2 and evaluated for compliance with block level considerations.

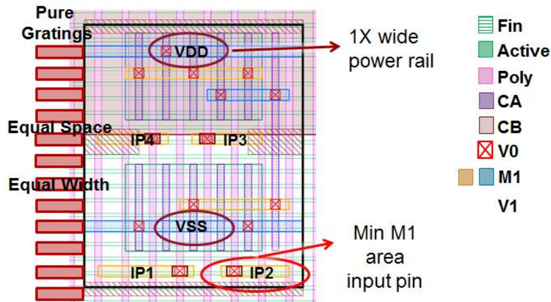

Figure 4 12T_Grating - Pure grating-based cell (AOI22) is inefficient and has several manufacturing concerns.

Step 3: 14 nm CMOS block level filtering. In this sub-section we analyze the 10T_BiDir, 10T_UniDir and 10T_UniDir_M2 for critical block level attributes, namely pin access, power rail, BEOL stack and color safe cell boundaries.

1) BEOL Stack

Backend of line (BEOL) metal stack of a process is also a part of technology definition and impacts the efficiency and manufacturability of a logic block. Mindful of design and manufacturability requirements we converged on two BEOL stacks, one for 10T_BiDir and another for 10T_UniDir and 10T_UniDir_M2 (Table 3). Except for M1, both stacks are equally restrictive and only differ in the directionality/preferred-orientation of a design layer. Unidirectional M2 and M3 layers improve manufacturability by limiting the number of layout patterns while also allowing the use of cost-effective patterning techniques, such as, DSA and SADP. Restricted BEOL stack also allows the use of simplified gridded auto-routers, reducing hotspot risk and design turnaround time. However, unidirectional metal stack inhibits conventional via redundancy techniques, requiring the use of alternative via redundancy schemes such as local loops [13]. For more relaxed pitches seen in M4 and M5 we used a restricted 2D BEOL stack with preferred orientation, to lower manufacturing cost while improving design efficiency.

Table 3. BEOL stack for 10T_BiDir and {10T_UniDir, 10T_UniDir_M2}.

Layer	10T_BiDir BEOL Stack (Direction/Pitch)	{10T_UniDir, 10T_UniDir_M2} BEOL Stack (Direction/Pitch)
Poly	V/y	V/y
M1	H/x & V/y	H/x
M2	H/x	V/y
M3	V/x	H/x
M4	H/1.5x and V/3x	V/1.5x and H/3x
M5	V/1.5x and H/3x	H/1.5x and V/3x

2) Pin access

Connecting to the input and output pins of standard cells is one of the challenging steps in detail routing in the physical synthesis flow, exacerbated more with double patterning [8]. While advances in auto-router algorithms continue to be made, detail routing challenges could be alleviated significantly by improving pin accessibility of standard cells. With extreme patterning restrictions requiring design layers to follow specific grids, the pin access problem becomes more tractable. Pin accessibility for a standard cell can be studied with two parameters, i) the total number of V1 access points that the M2 tracks have to connect to the standard cell M1 input and output pins, and ii) the maximum M2 run length for such a connection.

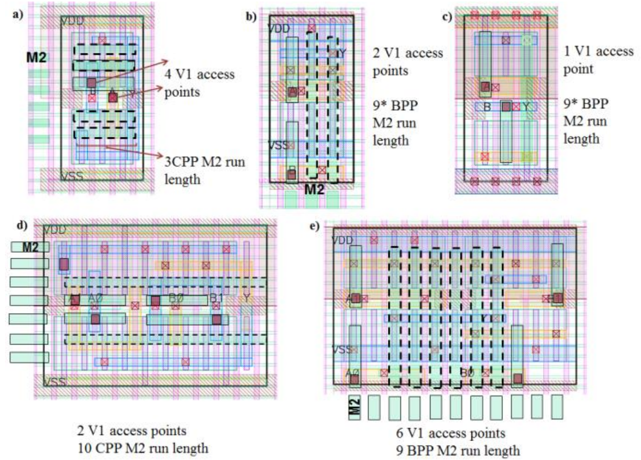

Figure 5. Pin accessibility (a) NAND2_X1 in 10T_BiDir ; (b) NAND2_X1 in 10T_UniDir ; (c) NAND2_X1 in 10T_UniDir_M2 (d) AOI22_X2 in 10T_BiDir ; (e) AOI22_X2 in 10T_UniDir.

We evaluate the pin accessibility of a NAND2_X1 in 10T_UniDir, 10T_UniDir_M2 and 10T_BiDir cell library in Figure 5. In this exercise, we assume all the input pins (A,B) are connected to a M2 track and observe how many V1 or M2 access points would remain for the Y pin and what would be the run length for the M2 connections. We observe 10T_UniDir_M2 has much fewer V1 access points compared to 10T_UniDir and 10T_BiDir. For 10T_UniDir_M2 to have

the same number of V1 access points as 10T_UniDir and 10T_BiDir it would require the cell to grow in height, which would degrade its design efficiency further. Therefore, we discard 10T_UniDir_M2 owing to its poor pin accessibility.

Continuing the pin accessibility comparison of 10T_UniDir and 10T_BiDir by looking at a simple gate like NAND2_X1 and a complex gate like AOI22_X2 reveals:

- The number of V1 access points doesn't scale well with the number of pins in the standard cell for 10T_BiDir, whereas for the 10T_UniDir, the V1 access points scale with the number of pins in the cell.
- The M2 run length varies with cell width for 10T_BiDir, whereas, it remains constant for 10T_UniDir.
- No M1 routing is possible in the block level for 10T_BiDir, some M1 routes can still be drawn at the block level for 10T_UniDir.

From these observations it is highly likely that 10T_UniDir has slightly better pin access, and thereby more efficienct routing than 10T_BiDir.

3) Robust Power Rail Structure

The power distribution network in a modern chip starts from the topmost metal levels and traverses all the way to the M1 power rail owned by the standard cell. The first consideration in designing power rails in a standard cell layout is electromigration. With the rapidly scaled down metal widths, especially M1, the current-carrying-capacity of the power rail drastically reduces, worsening electromigration. The second design consideration is IR drop in the power rail. These can be addressed typically by having wider power rails. However, wider power rails consume routing resources. Good power rail planning for a standard cell architecture, both at the cell level and block level, is critical to achieve a reasonable tradeoff between power rail robustness and routability.

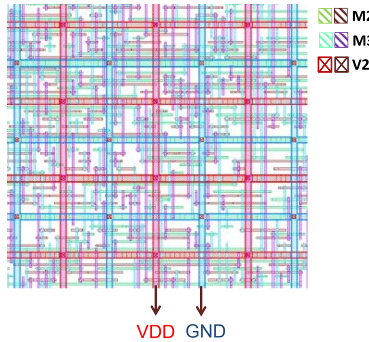

Figure 6. Structured grating-based power rail structure for better EM tolerance and power delivery.

At the cell level, the 10T_BiDir has a robust 2X M1 power rail and a 1X CB power rail, shared between two standard cell rows. Similarly, 10T_UniDir has an even more robust 2X M1 power rail per standard cell row. At the block level, both 10T_UniDir and 10T_BiDir use a structured 1D grating-based M2 and M3 power rail with via-bar and via-large (Figure 6). Structured 1D gratings allow for a wide power rails while also allowing for min-width routes. Via-bar and via-large placed on wide power rails are electromigration tolerant while reducing IR drop. In summary, power rail structures for both 10T_UniDir and 10T_BiDir can be designed to be robust.

4) Other Sub-20 nm Specific Block Level Considerations

Color-safe boundary conditions: In sub-20 nm CMOS nodes, all critical layers use multiple patterning. Decomposing a design level into multiple exposures is analogous to the graph coloring problem. One approach to ensure there are no coloring conflicts in different design layers is to ensure that every leaf cell is colored correctly and also has a color safe boundary [14]. We have incorporated this approach in standard cell design and present it in [5].

Compliance with SRAM bitcell: As SRAM bitcells are integrated on the same IC as standard cells, compliance between them is crucial in a patterning restricted sub-20 nm technology node. For instance, while designing 10T_UniDir standard cells, we ensured that an efficient SRAM bitcell can be designed using the same BEOL stack, i.e., horizontal M1 and vertical M2. Similarly, for 10T_BiDir we created a horizontal M2 based SRAM bitcell. A detailed discussion on compliance between SRAM bitcells and standard cell logic can be found in [10].

IV. EXPERIMENTAL RESULTS

In the last step in holistic DTCO of standard cell logic we evaluate the two competing standard cell architectures, 10T_BiDir and 10T_UniDir, in three ways:

- Library level comparison
- Physically synthesized logic blocks using these cells
- Silicon evaluation in IBM 14SOI process.

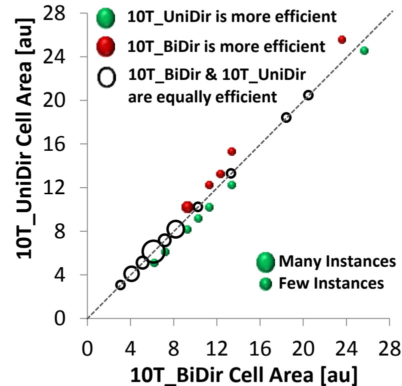

Figure 7. Library level comparison of 10T_BiDir and 10T_UniDir.

A. Library Level Assessment

A library of 40 representative cells was designed using 10T_BiDir and 10T_UniDir standard cell architectures in a 14 nm foundry process. Transistor level simulations using preliminary 14 nm foundry models indicate that both 10T_UniDir and 10T_BiDir exhibit similar power and performance. This is an expected trend as both cell libraries are 10 tracks tall, and a given cell such as NAND2_X1 has the same schematic in both 10T_BiDir and 10T_UniDir.

Furthermore, 10T_UniDir and 10T_BiDir cell libraries have similar area efficiencies as shown in Figure 7.

B. Physical Synthesis Evaluation

An effective method to evaluate standard cell architectures is to compare the design efficiency of physically synthesized design blocks. Using in-house place and route technology files created for 14 nm process assumptions, we physically synthesized a few sample design blocks (ISCAS'89 benchmarks) using 10T_BiDir and 10T_UniDir standard cell libraries. Experimental results indicate that both 10T_UniDir and 10T_BiDir can create design blocks with comparable design efficiencies (Table 4). If these trends hold in silicon, this could be a significant result as 10T_UniDir has unidirectional-M1 that is more restrictive than 10T_BiDir, while still achieving similar design efficiency.

Table 4. Results from physical synthesis of design blocks using 10T_BiDir and 10T_UniDir.

Design	Gates	Area(10T_UniDir/ 10T_BiDir)	Timing(10T_Uni Dir/10T_BiDir)	Power(10T_UniDir/ 10T_BiDir)
S820	192	0.94	1.04	1.07
S1238	315	0.84	1.0	1.0
S1423	425	0.93	1.0	1.0
S5378	801	0.93	0.99	0.95
Mult32	4103	0.98	1.01	1.03

C. Silicon Characterization in IBM 14SOI process

1) Ring Oscillator Test Structures

In an effort to characterize power-performance of 10T_BiDir and 10T_UniDir standard cell architectures on silicon, we designed, fabricated and tested ring oscillator test structures in a preproduction IBM 14SOI process. The UniDir_RO and BiDir_RO use cells from the 10T_UniDir and 10T_BiDir libraries, respectively. The ROs were subjected to tests at different operating conditions, namely, NOM+, NOM- and NOM. The 10T_UniDir based RO was observed to be about 35% slower and 4X leakier than the 10T_BiDir based RO (Figure 8). This trend was not seen in transistor level simulations with extracted parasitics and it could potentially be tracked down to two major differences in cell layouts:

- Gate connection near the edge of the cell in 10T_UniDir, as opposed to the gate connection in the center of the cell in 10T_BiDir. Such a non-conventional gate connection could result in increased gate resistance, slowing down UniDir RO. Additionally, if this gate connection is not manufactured reliably, that could also result in increased gate leakage.

- NMOS and PMOS connection using local interconnect CA layer in 10T_UniDir, as opposed to the M1 based connection seen in 10T_BiDir. CA based NMOS-PMOS connection running close to the gate of the two transistors could increase gate-to-drain capacitance (Cgd), resulting in a slower speed.

Based on these results, 10T_BiDir best trades off performance, area and manufacturability at the 14 nm technology node. With failure analysis and process optimization, *10T_UniDir could emerge as a promising alternative* for future technology nodes.

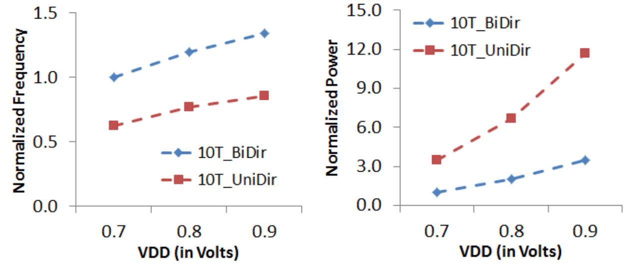

Figure 8. Measurement results (mean) from 10T_UniDir and 10T_BiDir ring oscillators, normalized w.r.t 10T_BiDir measured at 0.7V.

2) Measurements from 10T_BiDir based 32 bit Multiplier

To evaluate the efficacy of 10T_BiDir and the associated BEOL stack and physical synthesis flow, we designed, fabricated and tested a physically synthesized 10T_BiDir based 32-bit multiplier in IBM 14SOI process. Measurements from fully functional blocks working at frequencies in excess of 5GHz demonstrate the efficacy of 10T_BiDir cell library and associated physical synthesis flow (Table 5).

Table 5. Measurements from a physically synthesized 32-bit multiplier.

Attributes	Results
Technology	IBM 14 SOI
Area	2621.44 μm^2
Gates	4469
Metals	5
GOPS per Watt	3000 at 0.6V

V. SCALABILITY TO FUTURE NODES

From our standard cell DTCO exercise, we present a list of critical pattern constructs that are essential to design efficient standard cells in sub-20 nm technology nodes. We conclude with a discussion on the scalability of proposed standard cell architectures to future CMOS processes.

A. Pattern Constructs for Efficient Sub-20 nm Standard Cells

One of the key outcomes of sub-20 nm DTCO was identifying critical pattern constructs which a foundry has to support for efficient standard cell design. While these constructs are described in detail in [18], they are summarized below and shown in Figure 9.

1) 2X wide M1 power rails for electromigration tolerance.
2) 2X wide M1 pins for min-M1 area avoidance and better pin access.
3) Non-grating CB for area efficiency and power delivery.
4) Non-grating poly-cut mask for area efficiency.
5) Via-bar for improved electromigration tolerance and manufacturing yield.
6) Compound 2D grating for M1 is compatible with cost-effective patterning while improving design efficiency.
7) Robust poly contact at the edge of the cell would allow for P and N devices to be contacted via CA, allowing for

patterning friendly and efficient standard cells such as 10T_UniDir.

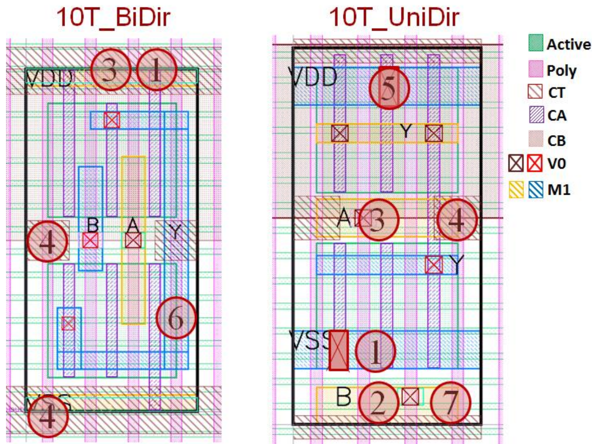

Figure 9. Pattern constructs for efficient sub-20 nm standard cell design.

B. *Scalability of Proposed Standard Cell Architectures to Future Nodes*

As we considered 10 nm-node-like process stack and patterning constraints, the proposed standard cell architectures scale favorable to 10 nm. SEMs from patterning a random logic block with these standard cells for both 10T_BiDir and 10T_UniDir scaled for 10nm-node-like dimensions demonstrate good patterning fidelity with double/triple patterning techniques (Figure 10). Tsai et al. have also demonstrated that fin patterns (finest pitch in a sub-20 nm technology) in scaled version of our standard cells can be reliably patterned with DSA [15]. Promising results from patterning experiments in 10 nm node-like pitches illustrate the scalability of our DTCO'ed standard cells to future nodes.

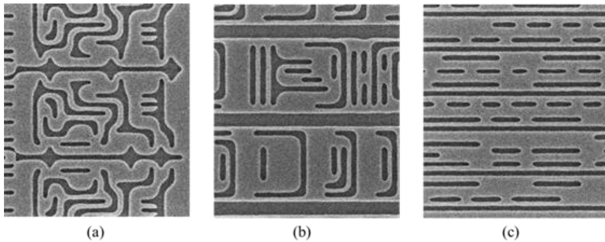

Figure 10. M1-SEMs in 10 nm node-like dimensions. (a) 32-nm node like layout style, (b) BiDir-M1 logic block, (c) UniDir-M1 logic block.

VI. Conclusion

Standard cells created from conventional design techniques fail to meet the varied design and manufacturing requirements in patterning constrained sub-20 nm CMOS processes. To best balance design efficiency for manufacturing cost and complexity we undertake a holistic design technology co-optimization (DTCO) of standard cells. Our holistic DTCO in a foundry 14 nm CMOS process resulted in two competing standard cell architectures, 10T_BiDir and 10T_UniDir. Measurements from test circuits in IBM 14SOI process reveal 10T_BiDir to be a more preferred alternative for the 14 nm node. With process optimization, 10T_UniDir could emerge as a promising alternative for future nodes, such as 10 nm and 7 nm.

Acknowledgment

This work is sponsored by the DARPA GRATE (Gratings of Regular Arrays and Trim Exposures) program under Air Force Research Laboratory (AFRL) contract FA8650-10-C-7038. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Approved for Public Release, Distribution Unlimited. This work was also supported in part by the Intelligence Advanced Research Program Agency and Space and Naval Warfare Systems Center Pacific under Contract No. N66001-12-C-2008, Program Manager Dennis Polla. We thank Greg Northrop, Leon Sigal and Daniel Morris for several insightful discussions.

References

[1] J. Friedrich et al., "Design methodology for the IBM POWER7 microprocessor," IBM Journal of Research and Development, vol.55, no.3, pp.9:1,9:14, 2011.

[2] J. Warnock, "Circuit design challenges at the 14nm technology node," Design Automation Conference, 2011.

[3] S. Iyer, "Keeping Moore's Law Alive – an IDM perspective," GSA Summit, 2012. http://www.gsaglobal.org/events/2012/0426/docs/keepingmooreslawaliv ekeynote-webpdf_000.pdf

[4] L. Liebmann et. al., "Simplify to survive: prescriptive layouts ensure profitable scaling to 32nm and beyond," Proc. SPIE, 2009.

[5] K. Vaidyanathan et. al., "Design and Manufacturability Tradeoffs in Unidirectional & Bidirectional Standard Cell Layouts in 14 nm node," SPIE Advanced Lithography Conference, 2012.

[6] G. Northrop, "Design Technology Co-Optimization in Technology Definition for 22nm and Beyond," Symposium on VLSI Technology, 2011.

[7] T. Jhaveri, "Regular Design Fabrics for Low Cost Scaling of Integrated Circuits," Carnegie Mellon University, PhD Thesis, 2009.

[8] X. Xu et al., "Self-aligned double patterning aware pin access and standard cell layout co-optimization". International symposium on physical design (ISPD), 2014.

[9] S. Kornachuk, M. Smayling, "New strategies for gridded physical design for 32nm technologies and beyond," International symposium on Physical design (ISPD), 2009.

[10] K. Vaidyanathan, "Exploiting challenges of sub-20 nm CMOS for affordable technology scaling," PhD thesis, 2014.

[11] Y. Chen et. al., "Self-aligned triple patterning for continuous IC scaling to half-pitch 15nm," Proc. SPIE, 2011.

[12] W. Hinsberg et al., "Self-Assembling Materials for Lithographic Patterning: Overview, Status, and Moving Forward," Proc. SPIE, 2010.

[13] W. Huang, et al., "Local loops for robust inter-layer routing at sub-20 nm nodes," Proc. SPIE 8327, 2012.

[14] L. Liebmann, et al., "Decomposition-aware standard cell design flows to enable double-patterning technology," Proc. SPIE, 2011.

[15] H. Tsai et al., "Directed self-assembly for ever-smaller printed circuits," SPIE, 2013. https://spie.org/x93535.xml

[16] K. Vaidyanathan et al., "Rethinking ASIC design with next generation lithography and process integration," SPIE Advanced Lithography Conference, 2013.

[17] R. Aitken et al., "Physical design and FinFETs," ISPD, 2014.

[18] K. Vaidyanathan et al. "Design implications of extremely restricted patterning", Proc. of SPIE, 2014.