# Toward Quantifying the IC Design Value of Interconnect Technology Improvements

Tuck-Boon Chan[†], Andrew B. Kahng[‡†] and Jiajia Li[†]
[‡]CSE and [†]ECE Departments, University of California at San Diego
{tbchan, abk, jil150}@ucsd.edu

***Abstract*—As technology scales, wire delay due to interconnect resistance (R) and capacitance (C) is increasing. Thus, improvement of middle-of-line and back-end-of-line (BEOL) materials and process technology (e.g., to achieve reduced barrier material thickness or dielectric permittivity) has always been a key goal in the technology roadmap. However, to date there has not been any systematic quantification of the *value* of BEOL technology improvements on integrated circuit (IC) design metrics. In this work, we create a framework to study the impact of improvements in interconnect technology on IC designs. Using 45nm technology and benchmark designs from public sources, we map reductions of interconnect resistance and/or capacitance to resulting impacts on design power, performance and area – for various types of physical design and operating contexts. By quantifying potential benefits of interconnect technology improvements at a block or core level, our proposed framework complements lower-level (e.g., critical-path) projections. We believe that this type of early assessment can be useful to guide BEOL technology investments and targets, especially as technology improvements require ever-increasing resources and focus in R&D efforts.**

## I. Introduction

With each successive technology node, interconnect resistance (R) and capacitance (C) continue to increase due to geometric scaling of wire dimensions and pitches. The rapid increase of interconnect RC is a major concern to the semiconductor industry, as this leads to increased circuit power and area beyond intrinsic wire performance loss (due to gate sizing and cloning, fanout reduction and buffer insertion), and hence limited circuit performance. Thus, interconnect RC reduction is always a key goal in technology roadmap [4]. For example, improved dielectric permittivity can reduce C; improved CMP uniformity, or novel barrier materials and reduced barrier thicknesses, can reduce R. Beyond first-order effects such as reduced $CV^2f$ dynamic power or smaller loading that permits smaller driver strengths, there are also second-order design benefits such as decreased IR drop margin and reduced P/G mesh area (from R reduction).

To quantify the quality of improved interconnect technology, conventional studies mostly focus on modeling and characterization of R and C values corresponding to a given (new) process technology [1] [8] [9] [11] [12] [13]. However, nominal RC improvements do not match actual block- or chip-level benefits in terms of design metrics such as circuit performance, power and area. For example, circuit performance is affected by interconnect RC as well as active devices. Equally critically, design benefits from reduced interconnect RC are gated by the capabilities of electronic design automation (EDA) tools (e.g., a router may not fully utilize the potential of improved interconnect technology). Overall, since interconnect R and C improvements do not directly correlate to improved design metrics, it has been difficult to assess the actual product benefits of an improved interconnect technology.

Relevant previous works include that of Li et al. [10], which evaluates DRAM performance improvement due to low-k interconnects. Kapur et al. [9] examine the impact of interconnect R and C on signal delay and power consumption. However, the analysis is based on a simple buffered interconnect, and cannot be extrapolated to chip-level consequences. (For example, a larger interconnect RC will increase the delay of setup-critical paths of a chip but will also reduce the number of buffers needed to fix hold-time violations. Such a design optimization cannot be understood from a simple buffered interconnect.) Hoang et al. [3] study the tradeoffs between resistance and capacitance of signal interconnects. They reduce power without introducing performance penalty by thinning down the interconnect height. To quantify the chip-level impact of interconnect variations, Jeong et al. in [7] assess chip-level impact of interconnect variation due to double-patterning lithography. Similarly, Jeong et al. in [5] discuss the impact of guardband reduction (including BEOL) on chip design metrics. In both papers, the discussions focus on interconnect variation but not the ROI from potential *future* interconnect technology improvements. Bamal et al. [2] compare the performance and energy of circuits with different interconnect technologies, such as carbon-nanotube interconnects, 3D interconnects and optical interconnects. However, they use simple benchmark circuits and interactions between interconnect technologies and EDA tools are ignored.

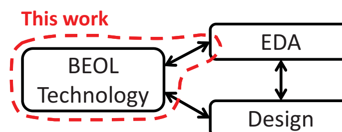


Fig. 1. Key elements that affect IC product metrics in future BEOL-limited technology nodes. We focus on quantification of chip-level benefits of potential future BEOL technology improvements, and on the interactions between BEOL technology improvements and EDA tools. We do not address the evolution of IC designs themselves in response to changing BEOL technology.

In this paper, we develop a framework to quantify the design-level benefits of potential future improvements in BEOL technology. We measure these benefits according to *design metrics* that include circuit performance, power and area. In our framework, we first create multiple artificial BEOL technology files to model potential improvements in future BEOL technology. Then, we use the artificial BEOL technology files to implement benchmark designs and analyze the effects of the potential R and/or C reductions on design metrics. Such an early assessment of BEOL improvements can be useful to guide BEOL technology investments and targets, especially as technology improvements require ever-greater R&D focus and resources. For example, a framework along the lines of what we propose can provide guidance for setting of R and C reduction targets to maximize ROI. The impact of BEOL technology improvements will be limited by the ability of designs (e.g., SOC architectures and physical implementations) and EDA tools (e.g., performance-driven routers) to exploit them. Thus, depicted in Figure 1, we focus on impacts of BEOL technology improvements, and on assessing whether EDA tools are able to comprehend and exploit such improvements. We do not address potential improvements in future designs (e.g., novel architectures), or in the synergy between BEOL technology and EDA tools.

Our main contributions are as follows.

- We implement a framework to quantify the impact of interconnect resistance and/or capacitance reductions on block- and chip-level performance, power and area. Our framework can capture effects of interconnect R and/or C

reductions within the context of an SP&R (RTL-to-GDS) implementation flow; such effects are not considered in conventional analyses based on analytical delay models or simple buffered interconnects.
- We study the concept of a *flexible BEOL*, in which capacitance and/or resistance of a subset of BEOL layers can be selectively improved. This permits cost-sensitive technology selection, in that product owners designers can choose to pay for only the improved BEOL layers which have largest impact on design metrics. By improving only a subset of the layers, the expected chip manufacturing cost can be reduced compared to using improved and more expensive BEOL on all layers.
- Our experiments quantify the range of potential design benefits of interconnect R and C reductions in advanced technologies. For example, our studies suggest that 45% capacitance reduction (in layers M2, M3, M4 and M5) can lead to 15% chip leakage power and 8% total power reductions.

The remainder of this paper is organized as follows. Section II describes our framework for quantifying the chip-level impact of interconnect R and C reductions. In Section III, we describe experimental setup and results that quantify chip-level impacts of interconnect R and C reductions. Section IV concludes and gives directions for ongoing work.

## II. A Framework to Analyze the Impact of Interconnect RC Reduction

Figure 2 illustrates our proposed framework to quantify the chip-level impact of interconnect improvements. In the framework, we first create a set of modified BEOL technology files to model the potential R and C reductions of future BEOL technology. Then, we implement benchmark designs using the modified BEOL. For comparison, we also implement the designs using the original BEOL. Last, we analyze the design metrics of the implemented circuits and compare the design metrics of circuit implemented with different BEOL files.

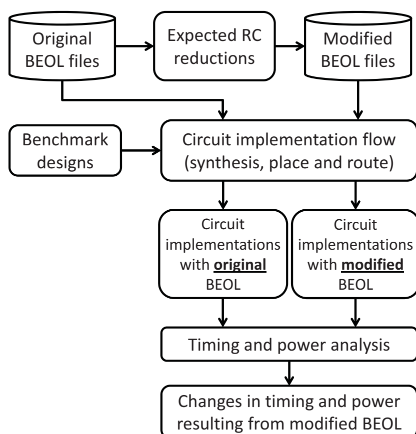


Fig. 2. Proposed framework to quantify chip-level impact of BEOL RC improvements.

### A. Modified BEOL Technology Files

To represent R and/or C reductions in the BEOL, we modify the BEOL technology files (e.g., *.itf* or *.ict*) to generate BEOL files for RC extraction (e.g., *TLUplus* or *captable*) in the circuit implementation flow.

In our experiment, we use *.itf* (instead of *.ict*) files to represent R and C reductions for compatibility with our implementation tool chain, which uses *Synopsys IC Compiler* [16]. The modified *.itf* files are converted to *TLUplus* files using the *grdgenxo* command in *Synopsys StarRC* [17]. In the modified *.itf* files, we model BEOL capacitance and resistance reductions by changing the inter- and intra-layer dielectric constants, and the metal resistivity, respectively. When we report our experimental results in Section III, we denote BEOL with $\alpha$% reduction in dielectric constant as "C = $\alpha$%". Similarly, we denote BEOL with $\beta$% reduction in resistivity as "R = $\beta$%". Note that $\alpha$% (resp. $\beta$%) reduction in dielectric constant (resp. resistivity) is positively correlated, but not equivalent, to the interconnect capacitance (resp. resistance) reduction. This is because actual capacitance (resp. resistance) values are also affected by other factors, e.g., total capacitance reduction is affected by the layout of interconnects, which is orthogonal to the reduction in dielectric constant.

### B. Flexible BEOL

With advanced manufacturing technologies, improvements can incur a much higher manufacturing cost, e.g., better mask alignment, better-controlled process steps, additional process steps, etc. Thus, applying available improvements on all BEOL layers may increase wafer cost unacceptably. Product engineers may decide to apply improvements on a subset of BEOL layers to control overall manufacturing cost consistent with product goals; this is similar to how the number of Mx, My, Mz etc. layers is optimized today. Even if only a subset of the BEOL layers are improved, the resulting chip-level improvements can still be significant, especially for circuit implementations with non-uniform utilization of different BEOL layers. We use the term *flexible BEOL* to refer to a BEOL stack in which a subset of the layers can be selectively improved. With the availability of flexible BEOL, chip designers might perform "what-if" analyses during circuit implementation and achieve tradeoffs between manufacturing cost and circuit performance. Below, we study the effects of applying R and/or C improvements on different combinations of layers. We also propose guidelines to select BEOL layers to achieve improved design metrics.

## III. Experimental Setup and Results

Our experiments use four RTL designs (obtained from the *OpenCores* website [14]) to quantify the chip-level impact of interconnect R and C reductions. We implement each design with two clock periods (i.e., fast and slow) to evaluate the impact of timing constraints. Table I shows the clock periods and number of instances for each design. We implement designs using TSMC 45nm LVT and HVT libraries, with eight metal layers (M1, M2, ..., M8). Synthesis is performed using *Synopsys Design Compiler vG-2012.06* [15],[1] and placement and routing (P&R) using *Synopsys IC Compiler vG-2012.06-ICC-SP3* [16]. We also use *Synopsys IC Compiler* for timing and power analysis. To reduce the impact of tool noise, we execute each P&R run three times with small perturbations {-2ps, 0ps, +2ps} of clock periods [6]. Among three routed netlists, we choose the one with minimum total power. All implementations have positive hold-time slack and *total negative slack (TNS)* $\geq$ -30ps for setup-time analysis. Further, no implementation has any electrical design rule or electromigration (EM) violation (activity factor = 0.1).[2] Thus, it is fair to compare the power and area among the implementations. The timing analysis configurations are given in Table II.

### A. Impact of R and C Reductions

To evaluate the impact of BEOL R and C reductions, we implement designs with reduced resistance and/or capacitance on particular layers by changing the (*TLUPlus*) BEOL technology files. We compare design metrics between implementations with and without reduced BEOL RC.

[1] In the synthesis flow, we first synthesize and place the design to obtain an initial floorplan, based on which we run a high-quality RTL synthesis flow with topographical feature in Synopsys Design Compiler.

[2] We do not model the effects of modified BEOL on EM limits. Our experiment results are insensitive to EM limits in that all implementations have zero EM violations with a larger activity factor (0.2).

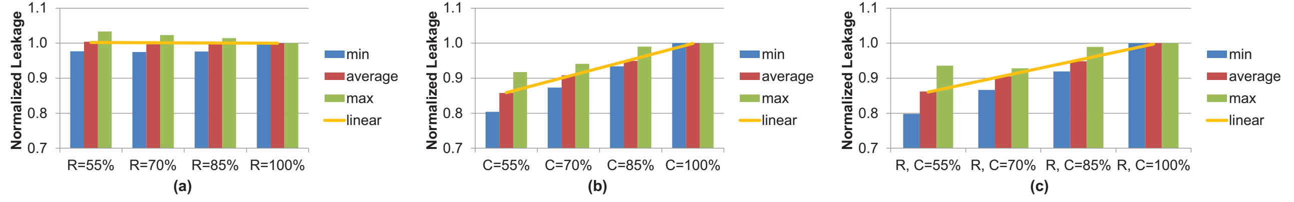


Fig. 3. Leakage power statistics of the circuits implemented with reduced BEOL (on layers M2, M3, M4 and M5). The leakage power is normalized to the R = 100% and C = 100% testcase. (a) C = 100%, only R is reduced; (b) R = 100%, only C is reduced; (c) Both R and C are reduced.

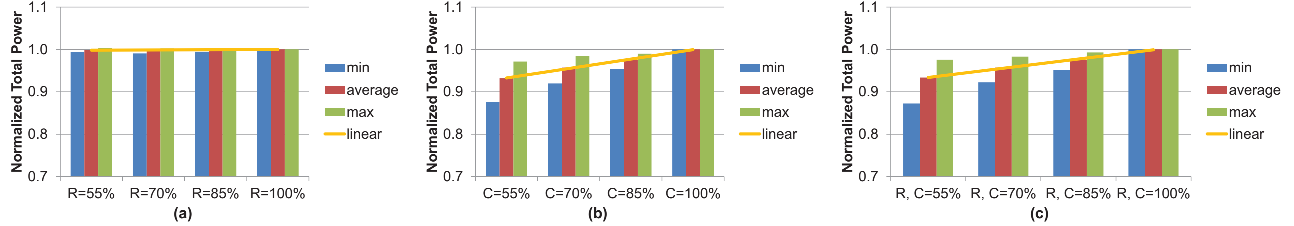


Fig. 4. Total power of circuits implemented with reduced BEOL RC (on layers M2, M3, M4 and M5). The total power is normalized to the implementation with original RC (i.e., R = 100% and C = 100%). (a) C = 100%, only R is reduced; (b) R = 100%, only C is reduced; (c) Both R and C are reduced.

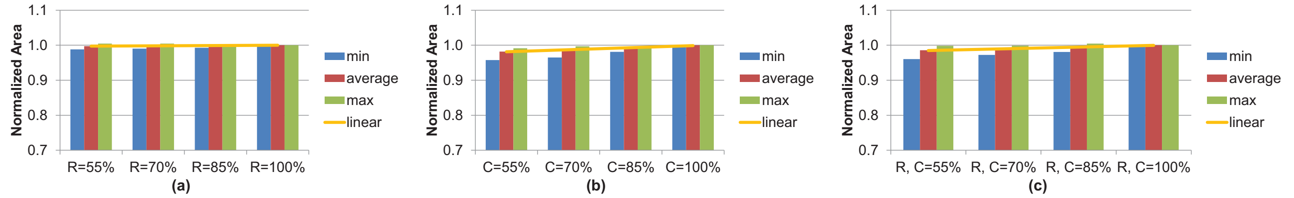


Fig. 5. Total cell area of circuits implemented with reduced BEOL RC (on layers M2, M3, M4 and M5). The total power is normalized to the implementation with original RC. (a) C = 100%, only R is reduced; (b) R = 100%, only C is reduced; (c) Both R and C are reduced.

TABLE I
BENCHMARK DESIGNS

| Design | Clock period (ns) | Total #cells |
|---|---|---|
| aes_cipher | 1.0 | ∼14.0k |
| aes_cipher | 1.2 | ∼13.8k |
| des_perf | 1.0 | ∼16.7k |
| des_perf | 1.2 | ∼16.3k |
| mpeg2 | 1.0 | ∼8.1k |
| mpeg2 | 1.2 | ∼7.8k |
| pci_bridge32 | 1.0 | ∼8.0k |
| pci_bridge32 | 1.2 | ∼7.7k |

TABLE II
TIMING CONSTRAINTS

| Parameter | Value |
|---|---|
| Clock uncertainty | 0.05 × clock period |
| Max transition | 0.08 × clock period |
| Timing derate on net delay (early/late) | 0.90/1.10 |
| Timing derate on cell delay (early/late) | 0.90/1.05 |
| Timing derate on cell check (early/late) | 1.10/1.10 |

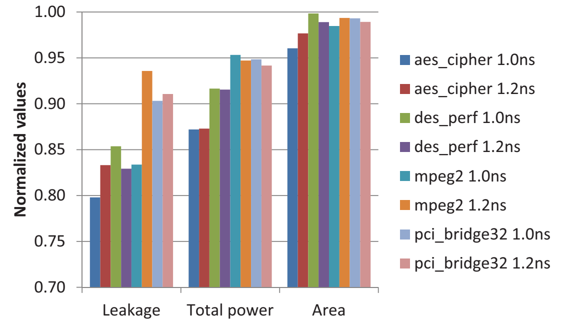


Fig. 8. Normalized circuit performance improvements due to R and C reductions for different designs. R = 55%, C = 55%.

Figures 3, 4 and 5 respectively show leakage power, total power and total cell area of designs with different RC (on layers M2 to M5). We see from Figures 3(b) and 4(b) that 45% reduction in capacitance on BEOL layers leads to 15% and 8% reductions in leakage and total power, respectively.[3] Further, the trend lines (labeled as "Linear" in figures) suggest a linear relationship between capacitance reduction and power reduction.

The reduction in power occurs because the smaller BEOL capacitance not only reduces load capacitance which in turn reduces *net switching power*, but also improves timing which in turn reduces slew, and thus, *internal power*. Moreover, the leakage power reduction indicates that the implementations with reduced BEOL capacitance have fewer cells and/or smaller cell strengths due to smaller wire loads. Note that the leakage power reduction is significantly larger than the area reduction in Figure 5. This is because with improved timing due to reduced interconnect RC, designs tend to have more HVT cells in a multi-Vt implementation. While Vt swapping reduces leakage power, the same footprint of HVT and LVT cells leads to little cell area reduction.

Results in Figures 3, 4, and 5 show that reducing both R and C simultaneously has similar effects as only reducing C. This agrees with the results shown in Figures 3(a), 4(a), and 5(a), where BEOL resistance reductions have little effect on design metric improvements.[4] The effect of resistance reduction hints that in this experiment, the path delays are dominated by wire capacitance.

### B. Impact of R and C Reductions in Advanced Technology

Since wire resistance becomes larger in advanced technology node, the impact of resistance reduction can be critical. To evaluate

[3]We define the total power as including leakage, internal and net switching power.

[4]We realize that resistance reductions on Mz or RDL layers may have significant impact on IR drop analysis and hence circuit sizing. This is the subject of ongoing study.

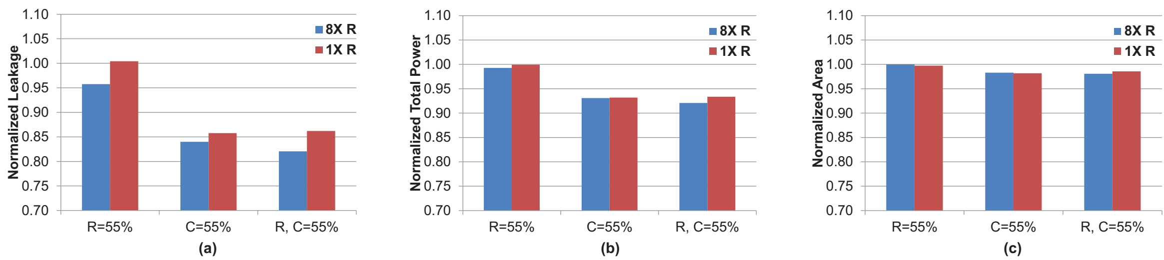


Fig. 6. Comparison of (a) leakage (b) total power and (c) total cell area between testcases with 8× nominal resistance vs. 1× nominal resistance. When not specified, R and C values are 100%.

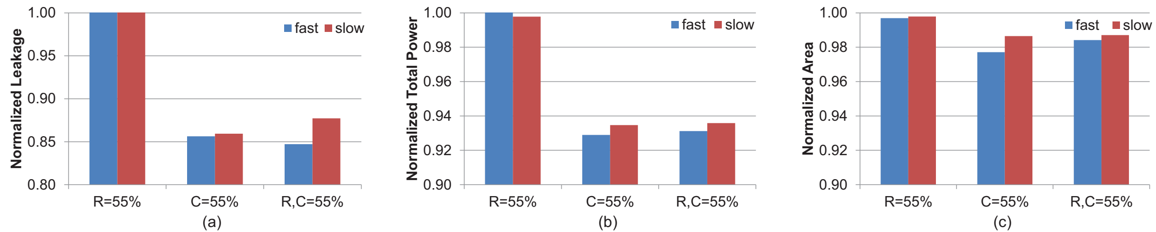


Fig. 7. Effects of R and C reductions on (a) leakage (b) total power and (c) total cell area for fast (clock period = 1.0ns) and slow (clock period = 1.2ns) designs.

the impact of wire resistance in an advanced technology, e.g., 20nm/16nm foundry node BEOL, we increase BEOL resistance by 8× on all layers. We then repeat the experiment described in Section III-A. In other words, the nominal BEOL resistance in this experiment setup is 8× of the resistance in the original TSMC 45nm BEOL, and the R and C reductions are with respect to the 8× resistance.[5]

Figure 6 shows that when the nominal resistance is large (i.e., 8× of nominal resistance at TSMC 45nm), resistance reduction on BEOL (layers M2, M3, M4 and M5) causes approximately 5% leakage reduction, which is larger than the case shown in Figure 3(a). However, the resistance reduction still does not have noticeable effect on total power and cell area. This is because the two major parts of total power, i.e., internal and net switching power, are insensitive to resistance reductions. By contrast, capacitance reduction has significant impact on total power reduction. These results suggest that, for Mx layers, resistance reduction may be less important than capacitance reduction. Thus, interconnect technology research might focus more on low-k dielectric and mask misalignment improvements to benefit design metrics.

Figure 6 shows normalized leakage, total power and area of circuits implemented with TSMC 45nm (1×) and with an artificial advanced technology (8×), with reduced RC. Figure 6(a) shows that when the nominal resistance is larger, leakage reduction resulting from R and C reductions is more significant. This suggests that the improvements in design metrics due to interconnect R and C reductions will increase in advanced technologies.

### C. Design-Dependence of R and C Reduction

We also evaluate the effects of R and C reductions across different designs. Figure 8 shows the leakage power, total power and area of different designs implemented with 45% reduction in both R and C on layers M2 to M5. Table III shows the statistics (such as average drive strength, total cell area, total number of cell instances) of designs shown in Figure 6. We observe that the impact of R and C reductions on total power is generally design-dependent, but the design-dependent trend is not obvious for leakage and cell area. To better understand the impact of R and C reductions across designs, we classify designs based on timing constraints into fast (1.0ns) and slow (1.2ns) clusters and compare design metrics. Figure 7 shows that faster designs are more sensitive to R and C reduction. This is because faster designs have tighter timing constraints, which leads to more critical paths with a larger number of buffers and/or cells with higher drive strength. That is, when R and C are reduced, the relaxed timing allows cell downsizing and/or buffer removal that translate to power and area reductions.

TABLE III
AVERAGE DRIVE STRENGTH, TOTAL CELL AREA, NUMBER OF INSTANCES AND TOTAL WIRELENGTH OF DESIGNS.

| Design | Clk period (ns) | Ave drive strength | Area ($\mu m^2$) | Total #cells | WL ($\mu m$) |
|---|---|---|---|---|---|
| aes_cipher | 1.0 | 1.75 | 15641 | 13877 | 201295 |
| | 1.2 | 1.58 | 15111 | 13558 | 201477 |
| des_perf | 1.0 | 1.62 | 22380 | 17094 | 191623 |
| | 1.2 | 1.43 | 21737 | 16298 | 201735 |
| mpeg2 | 1.0 | 1.92 | 21518 | 7975 | 150039 |
| | 1.2 | 1.80 | 21280 | 7733 | 145254 |
| pci_bridge32 | 1.0 | 1.80 | 21366 | 8137 | 147358 |
| | 1.2 | 1.77 | 20965 | 7726 | 146690 |

### D. Impact of Layer Selection

As discussed above, R and C reductions on all layers require high manufacturing cost. Further, inappropriate selection of BEOL layers for R and C reductions can be suboptimal which leads to less improvement in design metrics. Therefore, we seek to understand the optimum selection of a limited number of BEOL layers for R and C reductions. For instance, if circuit designers can choose only two layers for R and C reductions, which two layers will most improve design metrics? To answer such a question, we select all combinations of two layers among layers M2, M3, M4 and M5. For each combination, we characterize the corresponding BEOL technology files with reduced R and C.

Results in Figure 9 suggest that selecting different layers to reduce R and C can lead to different improvements in design metrics. On average, R and C reductions on layers M2 and M3 have the largest leakage and total power improvements. By contrast, R and C reductions on layers M2 and M5 have the smallest leakage

[5]Transitioning between the foundry 40nm and 20nm nodes, BEOL pitch and cell drive strength have near-traditional scaling, but wire resistance is relatively larger due to material properties. Thus, we scale up the resistivity by 8× while keeping other parameters the same.

Fig. 9. Effect of layer selection on (a) leakage, (b) total power and (c) total cell area.

(a) All nets of all designs

(b) Top 50 setup-critical nets of all designs

(c) Top 50 hold-critical nets of all designs

Fig. 10. Average Δwire distribution (%) over all designs for (a) all nets in the design, (b) top 50 setup-critical nets in the design, and (c) top 50 hold-critical nets in the design. Positive (resp. negative) Δwire distribution is shown as red (resp. blue). Red (resp. blue) dotted ovals indicate when the expected Δwire distribution is positive (resp. negative). When a dotted oval has the same color as the bar in the oval, the Δwire distribution matches the expected distribution. R = 55% and C = 55% for all BEOL layer combinations.

and total power improvements. The area improvement is negligible for all testcases.

We observe that the improvements on design metrics highly depend on the location of layers where R and C are reduced. If the RC-reduced layers are far from each other, more vias are inserted to utilize the reduced R and C on two separate layers, which limits the improvements (e.g., Figure 9 shows that reducing R and C on layers M2 and M5 does not lead to better results than reducing R and C on two adjacent layers). Moreover, we observe that it is more effective to reduce R and C on lower layers, at least in part because lower layers are utilized more, as shown in Table IV.

Last, we observe that reducing R and C on adjacent and highly utilized layers leads to better power reduction (e.g., layers {2,3} or {3,4}). (However, highly utilized layers are not necessarily adjacent. For example, a design with high aspect ratio might have higher utilization on vertical routing layers (e.g., M2 and M4) which are not adjacent.)

TABLE IV
WIRE DISTRIBUTION AMONG LAYERS FOR ALL DESIGNS.

| | min | average | max |
|---|---|---|---|
| M2 | 20.1 | 27.4 | 33.0 |
| M3 | 31.9 | 39.4 | 45.2 |
| M4 | 16.2 | 21.7 | 28.3 |
| M5 | 3.0 | 8.7 | 16.7 |
| M6 | <0.05 | 1.1 | 4.9 |
| M7 | <0.05 | 0.3 | 2.8 |
| M8 | <0.05 | 0.3 | 2.1 |

### *E. Utilization of Improved BEOL*

The benefits realized from R and C reductions depend on the circuit implementation flow and the performance of EDA tools. With this in mind, we assess the capability of a router to utilize BEOL layers with reduced RC. When the BEOL have layers with non-uniform RC, the router that is aware of BEOL RC distribution should route critical paths for setup time constraints on layers with smaller RC; and route critical paths for hold time constraints on layers with larger RC. Based on this hypothesis, we analyze the changes in wire distributions of circuits implemented with modified BEOL (45% reduction in both R and C on subsets of layers) with respect to the circuit implemented with the original BEOL. The changes in wire distributions are denoted as Δ*wire distribution*.

$$\Delta\text{wire distribution for a layer } x = \underbrace{\frac{\text{total wirelength of layer } x}{\text{total wirelength on all layers}}}_{\text{with modified BEOL}} - \underbrace{\frac{\text{total wirelength of layer } x}{\text{total wirelength on all layers}}}_{\text{with original BEOL}} \quad (1)$$

Figure 10(a) shows that 88% of the instances (circled bar) match the expected Δwire distributions. This shows that the router is likely to route more wires on the layers with R and C reduction. We also analyze the wiring layer distributions for setup- and hold-critical nets. Figure 10(b) shows that 75% of the instances (circled bar) for setup-critical nets match the expected Δwire distributions. However, Figure 10(c) shows that only 25% of the instances for hold-critical nets match the expected Δwire distributions. The mixed results show that the router is not fully responsive to BEOL RC characteristics when attempting to satisfy chip timing.

At a higher level, circuit implementation flows built from commercial EDA tools must be able to leverage the potential benefits offered by an improved BEOL technology. To understand whether present SP&R tools have this capability, we implement designs with both original and modified (with improved R and C) BEOL technology files. Then, we report design metrics of both implementations with modified (with improved R and C) BEOL technology files. Ideally, circuits implemented with improved BEOL should have better design metrics than those implemented with the original BEOL, given that the latter was oblivious to the reduced BEOL RC during circuit implementation. Figure 11 shows

the total power and maximum frequency (i.e., 1/(clock period − worst negative slack)) of circuits implemented with and without improved BEOL. In both cases, we report design metrics using the modified BEOL. In the modified BEOL, we reduce RC on layers {M2, M3}, {M4, M5} and {M2, M3, M4, M5}. Further, we study reductions in R and/or C of 15%, 30%, 45%. (This results in a total of 21 BEOL configurations, each of which has one red and one blue dot in the figure.)

In Figure 11, pairs of circuit implementations (with original BEOL, in red, and with modified BEOL, in blue) that use the same modified BEOL setup for timing and power analyses are connected by a directed edge. When the directed edge points downward (resp. leftward), the circuit implemented with improved BEOL has lower power (resp. frequency) than the circuit implemented with original BEOL. Most of the edges in Figure 11 point in a downward-left direction, suggesting that circuits implemented with improved BEOL have lower frequency as well as power compared to the circuit implemented with the original BEOL. This is a likely consequence of the performance target being 1.0GHz for all circuit implementations; we see that the timing slack due to the improved BEOL is traded for power reduction in the implementation tool.

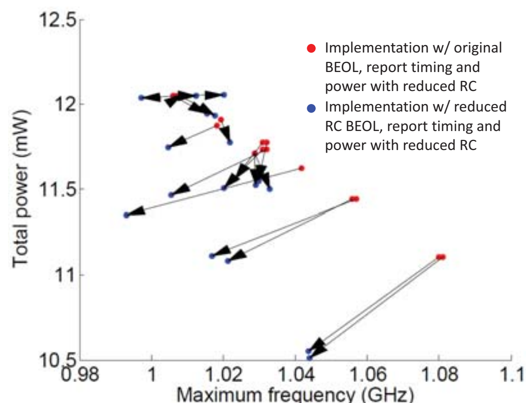


Fig. 11. Power and frequency for pairs of circuits (aes_cipher, 1.0ns) implemented with modified BEOL, and implemented with original BEOL, with power and timing analyses performed according to the modified BEOL.

## IV. Conclusions

In this work, we have developed a framework to quantify the impact of interconnect resistance and/or capacitance reductions on chip-level metrics such as performance, power and area. We observe that reduction in capacitance gives more noticeable benefits than reduction in resistance for the design metrics that we study. Our experiments also suggest that in more advanced technology nodes (e.g., 20nm/16nm foundry node BEOL), there is potential for larger relative improvements in design metrics from interconnect R and C reductions. We also evaluate the benefits of R and C reductions on different layers, within a *flexible BEOL* context. We find that the improvements in design metrics tend to be larger when reduction of R, C parasitics is made on highly utilized and adjacent layers. We also assess the capability of current EDA tools to exploit the availability of reduced interconnect RC on selected routing layers. These experiments suggest that present-day routers may have room for improvement with respect to setup- and hold-aware use of layers with varying parasitics (cf. the last of our directions for future work, below).

Our ongoing and future studies include the following.

- We study the potential impacts of interconnect R and C reductions on embedded SRAM scaling and performance, as well as on the scalability of on-chip power delivery networks. For example, less resistance on Mz or RDL may reduce IR drop, which leads to smaller gate sizes. These additional studies will provide a more holistic, comprehensive analysis of chip-level impacts of interconnect technology improvements.
- We extend our BEOL technology improvement analysis to M1 and middle-of-line layers, which affects the standard cells and therefore the chip-level timing and power performance.
- We study the impacts of interconnect R and C reductions across wide supply voltage ranges. In particular, circuit performance is more wire-dominated at high-performance (high supply voltage, overdrive) modes, and in such contexts we expect to see more significant impacts of BEOL technology improvements.
- Last, we seek to develop design optimization methodologies to exploit BEOL R, C reductions (e.g., assignments of nets to BEOL layers) for improved performance and/or power. For example, one can optimize the assignments of clock or timing-critical nets to the subset of BEOL layers with improved R and C.

## Acknowledgments

Research support from Sandia National Labs, Qualcomm, Samsung, NSF, SRC and the IMPACT (UC Discovery) and IMPACT+ centers is gratefully acknowledged.

## References

[1] N. D. Arora, "Modeling and Characterization of Copper Interconnects for SoC Design", *Proc. Intl. Conf. on Simulation of Semiconductor Processes and Devices*, 2003, pp. 1-6.
[2] M. Bamal, S. List, M. Stucchi, A. S. Verhulst, M. V. Hove, R. Cartuyvels, G. Beyer and K. Maex, "Performance Comparison of Interconnect Technology and Architecture Options for Deep Submicron Technology Nodes", *Proc. Intl. Interconnect Technology Conf.*, 2006, pp. 202-204.
[3] V. N. Hoang, G. Doornbos, J. Michelon, A. Kumar, A. Nackaerts and P. Christie, "Balancing Resistance and Capacitance of Signal Interconnects for Power Saving", *Proc. Intl. Interconnect Technology Conf.*, 2007, pp. 126-128.
[4] "International Technology Roadmap for Semiconductors, 2011 Edition, Interconnect Chapter." *http://www.itrs.net/Links/2011ITRS/2011Chapters/2011Interconnect.pdf*
[5] K. Jeong, A. B. Kahng and K. Samadi, "Impacts of Guardband Reduction on Design Process Outcomes: A Quantitative Approach", *IEEE Trans. on Semiconductor Manufacturing* 22(4) (2009), pp. 16:1-16:26.
[6] K. Jeong and A. B. Kahng, "Methodology From Chaos in IC Implementation", *Proc. ISQED*, 2010, pp. 885-892.
[7] K. Jeong, A. B. Kahng and R. O. Topaloglu, "Assessing Chip-Level Impact of Double-Patterning Lithography", *Proc. Intl. Symposium on Quality Electronic Design*, 2010, pp. 122-130.
[8] D. Josell, S. H. Brongersma and Z. Tökei, "Size-Dependent Resistivity in Nanoscale Interconnects", *Annual Review of Materials Research* 39 (2009), pp. 231-254.
[9] P. Kapur, J. McVittie and K. C. Saraswat, "Technology and Reliability constrained Future Copper Interconnects – Part II: Performance Implications", *IEEE Trans. Electron Devices* 49(4) (2002), pp. 598-604.
[10] W. Li, G. S. Sandhu, C.-C. Cho and G. D. Blalock, "DRAM Performance Improvement by Lowering Interconnect Capacitance", *Proc. Intl. Interconnect Technology Conf.*, 2001, pp. 78-80.
[11] A. Naeemi, R. Sarvari and J. D. Meindl, "Performance Modeling and Optimization for Single- and Multi-Wall Carbon Nanotube Interconnects", *Proc. Design Automation Conf.*, 2007, pp. 568-573.
[12] N. Shigyo, "Tradeoff Between Interconnect Capacitance and RC delay Variations Induced by Process Fluctuations", *IEEE Trans. Electron Devices* 47(9) (2000), pp. 1740-1744.
[13] Y. Travaly, B. Mandeep, L. Carbonell, Z. Tökei, J. V. Olmen, F. Iacopi, M. V. Hove, M. Stucchi and K. Maex, "On a More Accurate Assessment of Scaled Copper/Low-k Interconnects Performance", *IEEE Trans. on Semiconductor Manufacturing* 20(3) (2007), pp. 333-340.
[14] "OpenCores." *http://opencores.org*
[15] "Synopsys Design Compiler User Guide". *www.synopsys.com/Tools/Implementation/RTLSynthesis/DCUltra/Pages/default.aspx*
[16] "Synopsys IC Compiler User Guide". *www.synopsys.com/Tools/Implementation/PhysicalImplementation/Pages/ ICCompiler.aspx*
[17] "Synopsys StarRC User Guide". *http://www.synopsys.com/Tools/Implementation/SignOff/Pages/StarRC-ds.aspx*