Revisiting 3DIC Benefit with Multiple Tiers

Wei-Ting Jonas Chan[†], Andrew B. Kahng[‡†] and Jiajia Li[†]
[‡]CSE and [†]ECE Departments, University of California, San Diego
wechan@ucsd.edu, abk@cs.ucsd.edu, jil150@ucsd.edu

***Abstract*—3DICs with multiple tiers are expected to achieve large benefits (e.g., in terms of power, area) as compared to conventional planar designs. However, few if any previous works study *upper bounds* on power and area benefits from 3DIC integration with multiple tiers. In this work, we use the concept of *implementation with infinite dimension* to estimate the upper bound of power and area benefit from 3DICs. We observe that the maximum power benefit evaluated with infinite dimension is only 18% for particular designs. Such benefits further reduce under the assumption of inter-tier variation. In addition, we study power of designs across various dimensions (e.g., pseudo-1D, 2D, 3D with two, three and four tiers).[1] We observe that design power sensitivity to implementation with different dimensions correlates well with placement-based Rent parameter of the netlist. Therefore, placement-based Rent parameter can possibly be a simple indicator of 3D power benefit. Our study also shows that netlist synthesis and optimization should be aware of the target implementation dimension (e.g., 2D versus 3D).**

I. Introduction

3DIC with multiple tiers is a promising technology in the "More-than-Moore" era to integrate more functionality with greater bandwidth and less power. Many previous works propose 3DIC optimization approaches to achieve better design quality over conventional planar implementations. Further, due to higher integration and reduced wirelengths, 3DICs with more than two tiers are expected to offer larger benefits (e.g., less power). A recent work [35] shows that 3DICs with three tiers achieve 15% more power reduction as compared to corresponding two-tier 3DIC implementations and 36% power reduction versus 2D implementations. A much smaller body of work addresses the fundamental question of predicting 3DIC benefits over conventional 2D implementation, and upper-bounding these benefits. Chan et al. in [4] derive a 67% upper bound of wirelength reduction from a two-tier 3DIC over 2D designs. However, no previous works propose *upper bounds* on the power and total cell area reductions achievable by 3DICs over 2D designs.

SLIP '16, June 04 2016, Austin, TX, USA
 ISBN 978-1-4503-4430-2/16/06...$15.00
DOI: http://dx.doi.org/10.1145/2947357.2947363

[1]Design implementation with high aspect ratio layout.

In this work, we revisit the 3DIC benefit in terms of power and area with multiple tiers. More specifically, we propose the concept of *implementation in infinite dimension* (that is, where all gates can be placed as close as possible – essentially, adjacent – to each other) to derive an upper bound on 3D power and area benefits for given design, technology, and tool/flow. Such implementation in infinite dimension is achieved by synthesis and netlist optimization with zero wireload model (0-WLM). Our studies show that the power benefits can only be 18% for particular designs, even with an infinite-dimensional layout resource. We further evaluate design power across various dimensions (i.e., pseudo-1D, 2D, 3D with multiple tiers). We observe that design power sensitivity to different implementation dimensions correlates with Rent parameters of netlists, especially placement-based Rent parameters. Based on this observation, we suggest that netlist synthesis should be aware of implementation dimension so as to minimize design power.

We summarize our contributions as follows.

- We propose the concept of implementation with *infinite dimension* (i.e., netlist optimization with 0-WLM), based on which, we study the upper bound of power and area benefits of 3DICs.
- We show that upper bounds on 3DIC power and area benefits can be quite small – at most 39% power reduction and 10% area reduction even with infinite dimension.
- We study design power sensitivity to various implementation dimensions (i.e., pseudo-1D, 2D, 3D with different tier numbers and infinite dimension) and show the empirical evidence of a correlation between the power sensitivity and the Rent parameter of the netlist.
- We suggest that there is potential benefit from netlist synthesis optimizations being aware of the implementation dimension.

The remainder of this paper is organized as follows. Section II reviews related works on 3DIC optimization and prediction of 3DIC benefits. Section III describes our implementation flows as well as design power and area estimation flows in different dimensions. In Section IV, we describe experimental setup and results that quantify 3DIC power benefits and power sensitivity to dimensions. Section V concludes and gives directions for ongoing work.

II. Related Works

Various approaches have been proposed for implementation and optimization of 3DICs. Table I summarizes

area, power and wirelength benefits reported in previous works. In the table, "—" indicates "not applicable", i.e., not addressed in the cited work. We note that although we show total cell area reported by previous works in the table, area reduction is typically not the major objective in 3DIC optimization. In addition, most of these works use wirelength reduction as their major design objective [2] [6] [7] [8] [9] [12] [13] [17] [29] [30].

We first review previous works for integration and optimization of 3DICs. Liu et al. [20] propose transistor-level 3D monolithic integration. Bobba et al. [3] propose a cell-mapping and placement flow for monolithic 3D. Thorolfsson et al. [37] propose a 3D placer based on mPL. Lim [19] obtains power benefit by studying the capacitance reduction in TSV-based 3D implementations. Song et al. [35] propose a block-level folding approach and show corresponding power benefits of a three-tier stacking. Chang et al. [5] apply the flow [28] to 7nm technology node. Nayak et al. [24], Athikulwongse et al. [1], and Panth et al. [27] study the power benefits from various 3D integrations (e.g., monolithic 3D, mini-TSV, and TSV-based integration). Song et al. [34] study power reduction with consideration of the power distribution network. Jung et al. [10] [11], Lee et al. [18], and Ok et al. [25] achieve power benefits by applying block-level integration to the *OpenSparcT2* processor, a multicore GPU, and a stereo image processor.

Estimation of 3D power benefits has also drawn much attention. Priyadarshi et al. [31] propose an architectural framework to estimate 3DIC power for design space exploration. Kim et al. [14] [15] [16] propose models to estimate wirelength and power reductions based on buffer insertion. Chan et al. [4] propose a machine learning-based methodology to estimate power benefits of (3DIC) design flows, and apply their methodology to the flow proposed by Panth et al. [28]. However, no existing work studies potential *upper bounds* on 3DIC power and area benefits. To our knowledge, ours is the first work to address this gap.

In this work, we also examine the correlations between netlist properties and 3D power benefits. Our studies suggest that netlist synthesis could benefit from awareness of implementation dimension. Previous works which study the relationship between netlist structures and placement and routing (P&R) quality of results (QoR) include those of Liu and Marek-Sadowska [21] [22], which propose metrics to predict 2D P&R wirelength from netlist structure. They also propose optimization during the technology mapping stage of logic synthesis to improve 2D P&R results. Rahman et al. [32] propose a low-power gate-sizing scheme using a rich library with complex and large-size cells for logic synthesis, and a library with only simple cells for P&R. Seo et al. [33] argue that the benefit of using complex cells in advanced nodes will diminish due to routing congestion. However, [33] does not consider how 3D integration might mitigate the routing congestion seen in a 2D design implementation. (We consider this aspect of 3D integration below.)

TABLE I
SUMMARY OF 3D BENEFITS STUDIED IN PREVIOUS WORKS.

Work	% Benefit (area)	% Benefit (power)	% Benefit (WL)	# tier
[1]	7%	9%	WL increases	2
[2]	—	—	50%	2
[3]	—	15%	13%	2
[4]	—	39%	—	2
[5]	—	17%	46%	2
[6]	—	—	54%	4
[7]	—	—	50%	2
[8]	—	—	30%	2
[9]	—	—	26%	3
[10]	8%	21%	26%	2
[11]	—	20%	9%	2
[12]	—	—	20%	2
[13]	—	—	4%	2
[14]	—	—	37%	2
[15]	—	28%	41%	4
[16]	—	23%	28%	4
[17]	—	—	32%	2
[18]	—	22%	—	2
[19]	—	3%	25%	2
[20]	—	7%	16%	2
[23]	—	37%	48%	2
[24]	—	37%	—	2
[25]	—	13%	14%	2
[27]	—	35%	33%	2
[29]	—	—	19%	2
[30]	—	—	40%	2
[31]	—	20%	—	2
[34]	—	19%	—	2
[35]	3%	36%	42%	3
[37]	—	13%	21%	2

III. IMPLEMENTATION IN VARIOUS DIMENSIONS

We now describe our implementation and benefit estimation flows across various dimensions – pseudo-1D, 2D, 3D with multiple tiers, and infinite dimension.

A. Pseudo-1D Implementation

To estimate the power penalty of design implementations in a limited dimension, we propose *pseudo-1D implementation*, that is, design implementation with high aspect ratio layout. In this work, we refer to an implementation with layout aspect ratio 0.1 (block height equal to block width / 10) as a pseudo-1D implementation.

B. Optimal 2D Implementation

We pursue an "optimal" 2D implementation so as to quantify the true benefits from 3D integration with multiple tiers and from implementation in infinite dimension. To achieve this, we obtain multiple conventional planar implementations by sweeping several key parameters such as synthesis clock period, placement utilization and BEOL stack options; we then select the best (e.g., minimum power) outcome. Figure 1 shows an example where we vary the synthesis clock period and placement utilization for different 2D implementations. We observe that when the design has tight timing constraints, slightly smaller synthesis clock period eventually leads to smaller power after placement and routing. On the other hand, for a design with loose timing constraints, slightly larger synthesis clock period results in smaller design power after routing.

Furthermore, design power versus placement utilization exhibits a roughly unimodal behavior. In other words, if a placement is too compact, the power increases due to routing

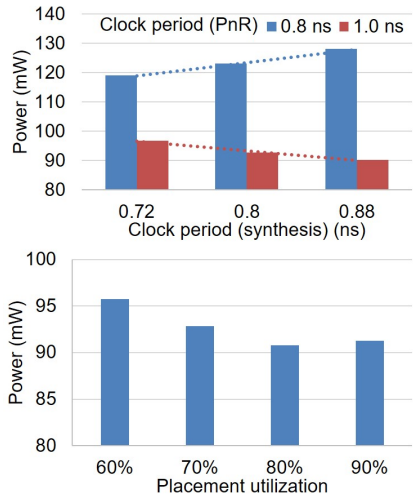

Fig. 1. Design power with (i) varying synthesis clock period, and (ii) placement utilization. Design: JPEG. Technology: 28FDSOI.

congestion (detouring, crosstalk, etc.). On the other hand, if the placement is too sparse, the power again increases due to longer wirelength (larger wire capacitance). Regarding BEOL stack options, we vary the number of layers from six to 11 in our experiments and compare the resultant design power after routing. However, we observe that the power variation is quite small (e.g., less than 3% for design JPEG). This is because the six-layer stack is able to offer enough routing resources for the designs and technology used in our experiments. We therefore implement our design testcases by varying (i) synthesis clock period (e.g., 0.9x, 1.0x and 1.1x of clock period used in P&R) and (ii) placement utilization (e.g., 60%. 70%, 80%, 90%), then selecting the outcome with minimum power consumption.[2]

We also apply such implementation flows to other dimensions to obtain (as far as we are able) fair comparisons.

C. Power Benefit Estimation for 3DICs

Given that there is no "golden" 3DIC implementation flow, we propose an implementation flow, based on the conventional 2D implementation tools, to achieve an optimistic estimation of design quality of 3DICs with multiple tiers. With the Shrunk2D flow proposed in [28] as a starting point, we perform the flow described in Algorithm 1 to estimate 3DIC benefits with multiple tiers. To estimate the reduction of wire parasitics from shrunk footprint area, we shrink the cell LEF by a factor of $1/\sqrt{T}$ in both height and width, where T is the number of tiers. We also apply the same scaling ratio to shrink the BEOL LEF to ensure that routing resources remain adequate. We implement designs based on the shrunk LEF (Line 1). With the shrunk LEF implementations, we then divide the die area into M x M grids. For each grid, we perform iterative min-cut partitioning to divide the cells within the grid into T clusters which are assumed to be placed on T tiers (Line 3). Details of our partitioning procedure are described in Algorithm 2. In the procedure, we iteratively apply min-cut partitioning based on Fiduccia-Mattheyses (FM) algorithm to partition the cells into two parts with area ratio of $(T-k)$ where k is the iteration index. The min-cut bipartitioning is performed with MLPart [40]. After each partitioning, the cells from the smaller-area part will be assigned to a new tier. We then collapse the cells in the smaller-area part into one super cell, fix the super cell in the first partition, set its area to zero and continue with the partitioning procedure. The iterative partitioning optimization terminates when all cells are assigned to a tier. Based on the partitioning solution, we annotate parasitics to nets which have cells from different tiers. More specifically, we calculate the maximum delta in tier depth across all cells connected to the net, and for each unit of (delta in tier) depth, we annotate RC corresponding to one TSV and vias across six metal layers (Line 4). With the annotated RC, we perform incremental sizing, VT-swapping and buffer insertion optimization to fix timing violations.

Note that in our estimation flow, we can estimate the benefits from wire parasitic reduction in 3DICs. In addition, we ignore the potential performance and power penalties from placement legalization (when we allocate cells to different tiers), from routing congestion caused by cross-tier power delivery, from difficulties in clock tree synthesis in a 3DIC, and from additional die-to-die variability margin. Our proposed flow therefore provides an optimistic estimation of 3DIC design qualities.[3] In our study, we also attempt to back-annotate placement solution with shrunk LEF to physical synthesis optimization. However, our experimental results show that a such back-annotation loop does not lead to further power benefits in the routed design.[4] We therefore apply a one-pass flow in our experiments.

Algorithm 1 Evaluation of 3DIC benefit with T tiers.

1: Perform 2D implementation with shrunk LEF (scaling ratio = $1/\sqrt{T}$ in cell height and cell width)
2: Divide die area into M x M grids uniformly
3: Apply FM-based min-cut partitioning for T tiers.
4: Annotate RC according to tier assignment based on partitioning solution
5: Perform incremental optimization
6: Report design power

[2]Our separate studies perform implementations with fine-grained choices of synthesis clock period (e.g., 0.8x to 1.2x with a step size of 0.05x of the P&R clock period) and placement utilization (e.g., 50% to 90% with a step size of 5%). Results for the JPEG testcase show that the optimal solution (i.e., solution with minimum power) found with fine-grained parameter sweeping is only 2% better than the solution found with our reported methods. We also attempt to vary the synthesis utilization (i.e., the utilization of floorplan as input to physical synthesis). However, experimental results show that for the technology and design testcases studied, sweeping of synthesis utilization does not lead to any power benefits as compared to a flow without physical synthesis.

[3]Due to lack of production-quality 3D design tool/flow, we are unable to precisely capture the penalties from routing congestion, cross-tier power delivery networks, clock tree synthesis, etc.

[4]Based on our experience, physical synthesis typically improves maximum performance when the clock constraints are tight. However, due to the pessimism of wireload in the physical synthesis, we rarely observe power reduction from physical synthesis optimization.

Algorithm 2 FM-based min-cut partitioning for T tiers.

Input: Netlist N, number of tiers T
Output: Subnetlists $Sol = \{N_1, N_2, ..., N_T\}$

$Sol \leftarrow \emptyset$
for $k = 1 : T - 1$ **do**
  $area_1 \leftarrow N.area/T$
  $area_2 \leftarrow \frac{T-k}{T} \cdot N.area$
  $\{N_k, N_{k+1}\}$ = 2WayMinCutPart(N, $area_1$, $area_2$)
  // tolerance = 5%
  $N \leftarrow$ collapse N_k into one super cell c_k
  $c_k.area \leftarrow 0$
  fix c_k in the first partition
  $Sol \leftarrow Sol \cup \{N_k \setminus \{c_1, c_2, ..., c_k\}\}$
end for
return Sol

D. Implementation in Infinite Dimension

To estimate the upper bound on power and area benefits from 3DICs, we propose the concept of implementation with "infinite dimension", where we ignore wire parasitics during the implementation. To achieve this, we perform netlist optimization with zero wireload model (0-WLM).[5] Given that benefits from 3D integrations mainly come from the reduced wire parasitics in a shrunk footprint area, such implementation with infinite dimension is able to provide an upper bound on 3DIC benefits.

IV. Experimental Setup and Results

We perform experiments in a 28nm/12-track FDSOI foundry technology with dual-VT libraries, and 0.85V nominal supply voltage. We perform experiments on an ARM CORTEXM0 core, three design blocks (AES, JPEG, VGA) from OpenCores [41], and LEON3MP from the ISPD-12 benchmark suite [26]. Parameters of these five testcases are shown in Table II. For each design, we determine a range of clock periods starting from a clock period with relative loose timing constraint, up to the clock period which is close to the minimum clock period of the given design and technology. These designs are synthesized using *Synopsys Design Compiler vH-2013.12-SP3* [42] and then placed and routed using *Cadence Innovus Implementation System v15.2* [39]. We further use *Cadence Tempus Timing Signoff Solution v15.2* [43] for timing and power analysis, with wire parasitics (SPEF) obtained from Innovus.

TABLE II
Benchmark parameters.

Design	#Instances	#Flip-flops	Clock period range
CORTEXM0	~9k	840	0.8ns - 1.0ns
AES	~12k	530	0.6ns - 0.9ns
JPEG	~43k	4712	0.8ns - 1.1ns
VGA	~72k	17057	0.7ns - 1.0ns
LEON3MP	~460k	108817	1.1ns - 1.3ns

[5]To ensure a fair comparison to implementations at 2D and 3D, we perform netlist optimization with the same synthesis, placement and clock tree synthesis tool/flow but with 0-WLM and without any routing.

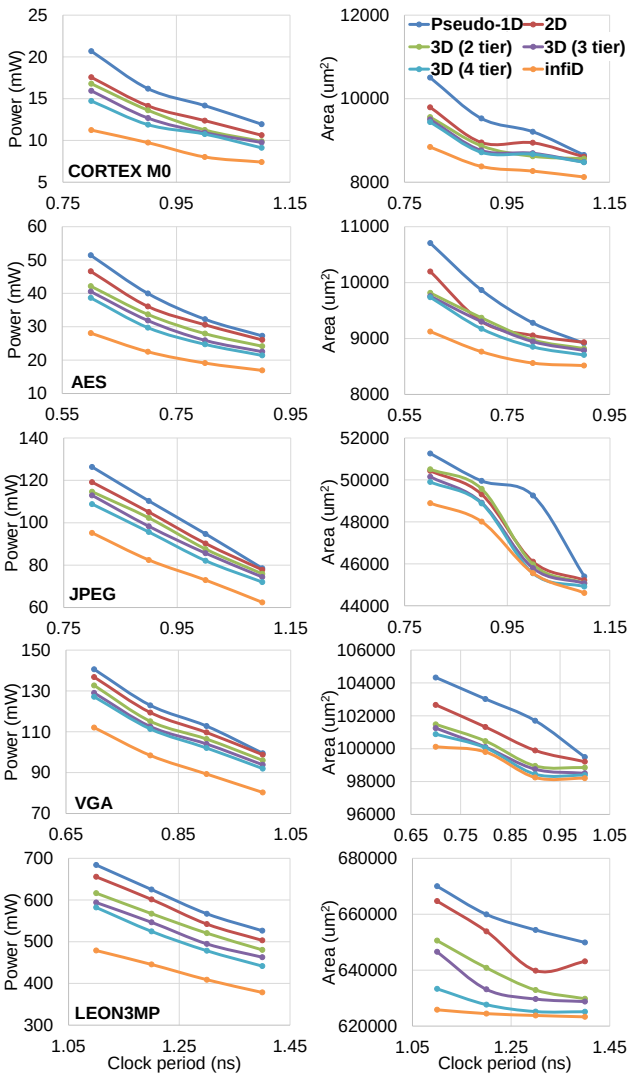

Fig. 2. Design power and total cell area evaluated across various implementation dimensions.

A. Evaluation of 3D benefits

We first compare design power and total cell area across various implementation dimensions across different clock periods. Figure 2 shows the power and area comparison. All the implemented designs have no hold violation and a setup violation less than 10ps. We observe that the maximum power benefits (i.e., the gap between the red curve versus the orange curve) from implementations in infinite dimension are respectively 36%, 39%, 20%, 18% and 26% for CORTEXM0, AES, JPEG, VGA and LEON3MP. The results show a large variation of 3D benefits across different designs. In addition, the power benefits from 3D integration with two, three and four tiers are less than 10% for designs JPEG, VGA and LEON3MP. Furthermore, we observe that the area benefits are small (i.e., $< 10\%$ for all designs, and $< 4\%$ for designs JPEG and VGA).

B. A More Realistic Evaluation

As discussed in Section III, our 3D power estimation ignores potential larger clock skew due to inter-tier process

variation [38]. To achieve a more realistic estimation of 3D benefits, we quantify the impact of clock skew on 3D power reduction. In our experiments, we enable multi-corner optimization by using both slow- and fast-corner libraries during the P&R stage. We further model potential clock skew increase due to difficulties in 3D clock tree synthesis (CTS) as well as inter-tier process variation by applying 0%, 5% and 10% clock uncertainties of the clock periods. The power benefits against the clock uncertainties are shown in Figure 3. The results show that the 3D power benefits diminish when the clock uncertainties increase from 0% to 10% even for two designs which originally have the largest 3D benefits among our benchmarks. More specifically, the power benefit of a two-tier 3D implementation decreases from 11% to 1% for AES and from 5% to -21% for CORTEXM0. Our observation indicates that it is critical for 3D clock tree optimization to minimize the impact of inter-tier variation on clock skew and latency.

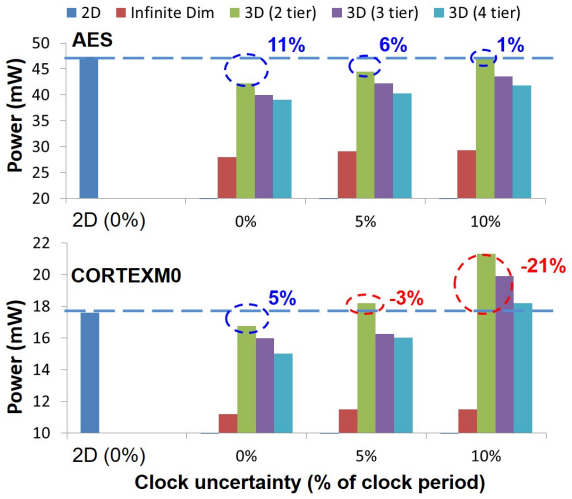

Fig. 3. The impact of larger clock skew (due to complexity of 3D CTS as well as inter-tier variation) on 3D power benefits.

C. Netlist Study

We additionally study the possibility of correlations between (3D) power benefits and various netlist parameters (such as fanout distribution, slack distribution, sequential graph, Rent parameter, etc.) of designs. We observe that the power benefits are well correlated with Rent parameters.

We use the Rentian circuit generator *gnl* by Stroobandt et al. [36] to generate netlists with different Rent parameters, and we evaluate these netlists' power consumption across various implementation dimensions. The inputs to gnl are (i) number of cells, (ii) the target Rent parameter, (iii) the ratio between flip-flops and combinational cells, and (iv) the maximum path delay constraint. The gnl software starts with a set of standard cells and randomly inserts connections among the cells to form logic cones. The gnl software determines number of pins (of standard cells) and input/output terminals according to the user-specified Rent parameter. During the netlist generation, gnl recursively clusters logic cones to form larger ones. The number of terminals on the boundaries of the merged logic cones also follows the specified Rent parameter. The generated netlists thus have desired Rent parameters by construction.[6] Table III summarizes our generated testcases. We generate netlists using cells from the foundry 28nm 12-track FDSOI library and implement them using the flows described in Section III-C. The initial generated netlists (with different Rent parameters) have similar power and area (i.e., within 3% difference) to help establish a fair comparison of power benefits across various Rent parameters. We define timing constraints such that the initial generated netlists have negative slacks, thus inducing non-trivial P&R optimizations.[7] Furthermore, to maintain a similar Rent parameter throughout the P&R flow (i.e., avoiding netlist restructuring), we apply a size-only restriction to all cell instances during the P&R optimization flow.

Figure 4 shows the relationship between post-P&R power and netlist Rent parameter across various implementation dimensions. We observe that the power of the conventional 2D implementation increases with higher Rent parameter, whereas power increase (with Rent parameter) is smaller with 3D implementations. This suggests that implementations in higher dimension can mitigate power penalties due to higher-degree topologies of interconnections, which are indicated by larger Rent parameter values. Accordingly, more 3D power benefit may be expected with netlists having larger Rent parameter. We also observe the existence of *thresholds of Rent parameters* beyond which 3D power benefits seem to increase more rapidly (e.g., 0.69 in Figure 4). Quantitative analysis of the relationship between power benefits and Rent parameter values will be one of our future works.

TABLE III
SUMMARY OF RENTIAN TESTCASES WITH DIFFERENT RENT PARAMETERS.[8]

Rent (input / actual)	Power (mW)	Area (um^2)	Slack (ps)
0.50 / 0.63	46.4 (100%)	39552 (100%)	-72
0.55 / 0.66	46.8 (101%)	40262 (102%)	-74
0.60 / 0.69	46.7 (101%)	40404 (102%)	-68
0.65 / 0.71	47.4 (102%)	40532 (102%)	-110
0.70 / 0.74	46.9 (101%)	40607 (103%)	-73

We also perform similar studies with realistic designs. Figure 5 shows correlation between the maximum 3DIC power benefits estimated in infinite dimension, and Rent parameter values. We extract Rent parameters of the netlists using both partitioning-based and placement-based methods, where we assume that one pin (terminal) is induced by each

[6]We constrain gnl to instantiate equal numbers of DFFX8, INVX8, BUFX8, AND2X8, NAND2X7, OR2X8, NOR2X7, NAND3X12, NOR3X13 and XOR2X8 cells in the generated netlists.

[7]The gnl software constrains the maximum delays of the generated netlists by limiting the depths of the logic cones (i.e., inserting flip-flops at the boundary of the logic cones).

[8]The target clock period is 1ns. We show both input Rent parameters to the gnl software and actual Rent parameters of the generated netlists (placement-based) in the table.

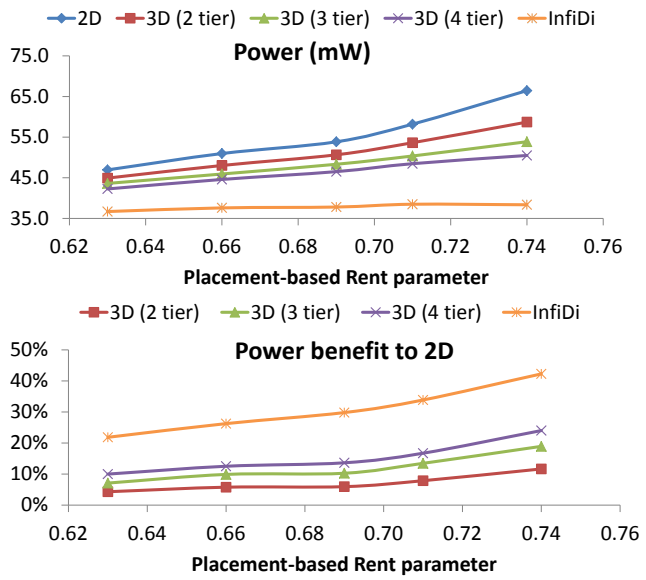

Fig. 4. Power and power benefits versus Rent parameters for the 2D and the 3D implementations with different tiers.

cut hyperedge. Partitioning-based Rent parameter values are extracted based on recursive bipartitioning using the min-cut hypergraph partitioner MLPart [40]. To calculate placement-based parameter values, we perform fast placement with a commercial P&R tool [39] without any sizing, VT-swapping or buffering optimizations. We then perform rectangle sampling based on the placement solutions to estimate Rent parameters.

Even for a larger testcase (LEON3MP with 436K instances), the runtime of the placement used to evaluate the placement-based Rent parameter is only 16 minutes. Our results show that the placement-based Rent parameter can possibly be a simple indicator of 3DIC power benefit for a given netlist.

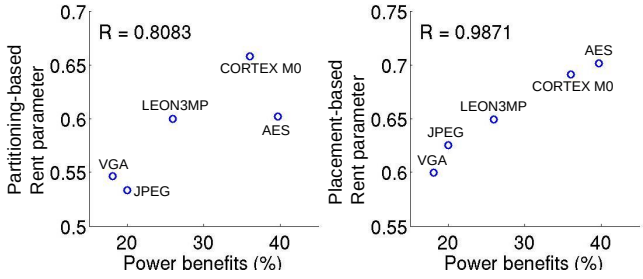

Fig. 5. Power benefits correlate with Rent parameters.

In light of the correlation between power sensitivity to implementation dimensions and the netlist Rent parameter, we propose to modulate the cell usage in synthesis stage to control the Rent parameters and achievable 3D power benefit. We categorize library cells in the 28nm FDSOI design enablement according to their input pin numbers – (1) one-input cells (buffers and inverters) (2) two-input cells (NAND2, NOR2, etc.) (3) three-input cells (NAND3, AOI21, etc.) (4) four-input cells (NAND4, OAI22, etc.), and (5) >four input cells (AOI212, MUX41, etc.). We then scale cell area in Liberty files to modulate the cell usage during synthesis so as to achieve netlists with different Rent parameters. More specifically, we choose design JPEG (which originally has a small Rent parameter) and scale down the area of complex cells; this induces the synthesis tool to use more complex cells and to increase the netlist Rent parameter.[9] We plot the placement-based Rent parameters against the portion of complex cells (cells with more than three input pins) of various synthesized netlists in Figure 6. We observe that the Rent parameters are highly correlated to the incidence (proportion) of three-input cells. This demonstrates that we can modulate Rent parameters of the synthesized netlist. However, more precise control of Rent parameters during synthesis optimization remains as a direction for future work.

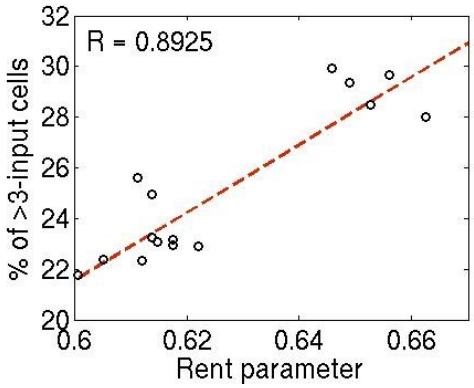

Fig. 6. Correlation between incidence of cells with >3 inputs vs. Rent parameter.

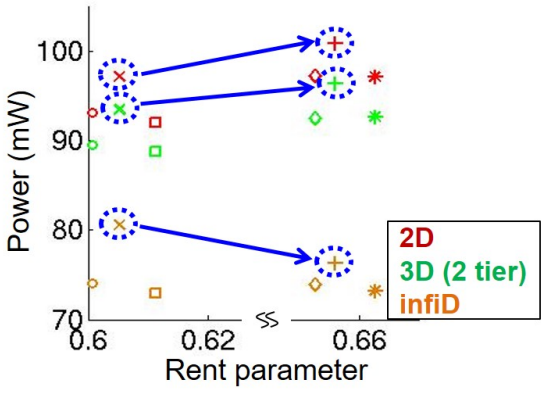

Fig. 7. Power vs. Rent parameter with Rent Modulation.

In Figure 7, we further show power (after routing) of six synthesized netlists of design JPEG which have the

[9]We synthesize JPEG with area scaled by {1x, 2x} for 2-input cells, and by {1x, 0.5x} for 3-input, 4-input and >4-input cells. An alternative way to modulate cell usage and Rent parameters during synthesis is to set a dont_use attribute for certain Liberty cells. However, the dont_use attribute cannot be assigned to NAND2, NOR2 cells in our EDA tooling. In addition, setting the dont_use attribute for a group of cells might degrade synthesis solution quality due to limited available cell types.

TABLE IV
AREA SCALING RATIOS FOR IMPLEMENTATIONS IN FIGURE 7.

Implementation	Rent	2-input	3-input	4-input	>4-input
O	0.600	1	0.5	1	1
X	0.605	2	0.5	1	1
□	0.611	1	1	1	1
◊	0.653	2	1	0.5	0.5
+	0.656	2	0.5	0.5	0.5
*	0.663	1	0.5	1	0.5

maximum and minimum Rent parameters (Table IV). As highlighted in blue dotted circles, we observe that although a particular netlist shows small power after synthesis (indicated by infinite dimension), due to its large Rent parameter its power can be larger with a 2D implementation. However, power penalty with a 3D implementation is smaller. This suggests that netlist synthesis should be aware of implementation dimension. For instance, a netlist with small Rent parameter is desirable for a 2D implementation; there are fewer constraints on (or, sensitivities to) Rent parameters for a 3D implementation.

V. CONCLUSIONS

In this paper, we revisit previous assessments of the benefits of 3DIC implementation with respect to area, power and wirelength. Ours is the first work to estimate upper bounds on 3D power and area benefits based on the concept of *implementation with infinite dimension*. We examine several designs with our "infinite dimension" bounding methodology and observe that the available area and power gaps between "best possible" 2D implementation and infinite-dimensional implementation can be small, e.g., power benefits as low as 18% for some designs. Such results indicate that 3D benefits are more likely to be achieved from SoC-level and architectural-level optimizations instead of traditional P&R physical implementation optimizations. We also observe that inter-tier variation causes further significant reduction of available 3D power benefits.

In addition, we study design power across various dimensions and observe a correlation between design power and netlist Rent parameter. Modulation of the netlist Rent parameter during synthesis (that is, by changing the usage and distribution of fanins) suggests that a synthesis optimization that is aware of implementation dimension may be helpful for reduced power in the final physical implementation. We also note that architecture-level improvements enabled by 3D integration (e.g., larger memory bandwidth) are of course still very promising, and these are not addressed in our work.

Directions for our future work include (i) dimension-aware synthesis (i.e., synthesis for multi-tier 3D), (ii) quantitative analysis of the relationship between power benefits and Rent parameters, and (iii) architectural-level benefit exploration.

REFERENCES

[1] K. Athikulwongse, D. H. Kim, M. Jung and S. K. Lim, "Block-Level Designs of Die-to-Wafer Bonded 3D ICs and Their Design Quality Tradeoffs", *Proc. ASP-DAC*, 2013, pp. 687-692.

[2] K. Bernstein, P. Andry, J. Cann and P. Emma, "Interconnects in the Third Dimension: Design Challenges for 3D ICs", *Proc. DAC*, 2007, pp. 562-567.

[3] S. Bobba, A. Chakraborty, O. Thomas, P. Batude, V. F. Pavlidis and G. De Micheli, "CELONCEL: Effective Design Technique for 3-D Monolithic Integration Targeting High Performance Integrated Circuits", *Proc. ASP-DAC*, 2011, pp. 336-343.

[4] W.-T. J. Chan, Y. Du, A. B. Kahng, S. Nath and K. Samadi, "3DIC Benefit Estimation and Implementation Guidance from 2DIC Implementation", *Proc. DAC*, pp. 1-6.

[5] K. Chang, K. Acharya, S. Sinha, B. Cline, G. Yeric and S. K. Lim, "Power Benefit Study of Monolithic 3D IC at the 7nm Technology Node", *Proc. ISLPED*, 2015, pp. 201-206.

[6] J. Cong, C. Liu and G. Luo, "Quantitative Studies of Impact of 3D IC Design on Repeater Usage", *Proc. International VLSI/ULSI Multilevel Interconnection Conference*, 2008, pp. 344-348.

[7] J. Cong, G. Luo, J. Wei and Y. Zhang, "Thermal-Aware 3D IC Placement via Transformation", *Proc. ASP-DAC*, 2007, pp. 780-785.

[8] S. Das, A. Chandrakasan and R. Reif, "Three-Dimensional Integrated Circuits: Performance, Design Methodology, and CAD Tools", *Proc. IEEE Computer Society Annual Symposium on VLSI*, 2003, pp. 13-18.

[9] J. Deguchi, T. Sugimura, Y. Nakatani, T. Fukushima and M. Koyanagi, "Quantitative Derivation and Evaluation of Wire Length Distribution in Three-Dimensional Integrated Circuits Using Simulated Quenching", *Japanese Journal of Applied Physics* 45(4B) (2006), pp. 3260-3265.

[10] M. Jung, T. Song, Y. Wan, Y.-J. Lee, D. Mohapatra, H. Wang, G. Taylor, D. Jariwala, V. Pitchumani, P. Morrow, C. Webb, P. Fischer and S. K. Lim, "How to Reduce Power in 3D IC Designs: A Case Study with OpenSPARC T2 Core", *Proc. CICC*, 2013, pp. 1-4.

[11] M. Jung, T. Song, Y. Wan, Y. Peng and S. K. Lim, "On Enhancing Power Benefits in 3D ICs: Block Folding and Bonding Styles Perspective", *Proc. DAC*, 2014, pp. 1-6.

[12] D. H. Kim, K. Athikulwongse and S. K. Lim, "A Study of Through-Silicon-Via Impact on the 3D Stacked IC Layout", *Proc. ICCAD*, 2009, pp. 674-680.

[13] T.-Y. Kim and T. Kim, "Clock Tree Embedding for 3D ICs", *Proc. ASP-DAC*, 2010, pp. 486-491.

[14] D. H. Kim, S. Mukhopadhyay and S. K. Lim, "Through-Silicon-Via Aware Interconnect Prediction and Optimization for 3D Stacked ICs", *Proc. SLIP*, 2009, pp. 85-92.

[15] D. H. Kim and S. K. Lim, "Through-Silicon-Via-aware Delay and Power Prediction Model for Buffered Interconnects in 3D ICs", *Proc. SLIP*, 2010, pp. 25-31.

[16] D. H. Kim, S. Mukhopadhyay and S. K. Lim, "TSV-Aware Interconnect Distribution Models for Prediction of Delay and Power Consumption of 3-D Stacked ICs", *IEEE TCAD* 33(9) (2014), pp. 1384-1395.

[17] D. H. Kim, R. O. Topaloglu and S. K. Lim, "Block-Level 3D IC Design with Through-Silicon-Via Planning", *Proc. ASP-DAC*, 2012, pp. 335-340.

[18] Y. J. Lee and S. K. Lim, "On GPU Bus Power Reduction with 3D IC Technologies", *Proc. DATE*, 2014, pp. 1-6.

[19] S. K. Lim, *Design for High Performance, Low Power, and Reliable 3D Integrated Circuits*, New York, Springer, 2012.

[20] C. Liu and S. K. Lim, "A Design Tradeoff Study with Monolithic 3D Integration", *Proc. ISQED*, 2012, pp. 529-536.

[21] Q. Liu and M. Marek-Sadowska, "Wire Length Prediction-Based Technology Mapping and Fanout Optimization", *Proc. ISPD*, 2005, pp. 145-151.

[22] Q. Liu and M. Marek-Sadowska, "Semi-Individual Wire-Length Prediction with Application to Logic Synthesis", *IEEE TCAD* 25(4) (2006), pp. 611-624.

[23] Y. J. Lee, D. Limbrick and S. K. Lim, "Power Benefit Study for Ultra-High Density Transistor-Level Monolithic 3D ICs", *Proc. DAC*, 2013, pp. 1-10.

[24] D. K. Nayak, S. Banna, S. K. Samal and S. K. Lim, "Power, Performance, and Cost Comparisons of Monolithic 3D ICs and TSV-Based 3D ICs", *Proc. SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S)*, 2015, pp. 1-2.

[25] S. H. Ok, K. R. Bae, S. K. Lim and B. Moon, "Design and Analysis of 3D IC-Based Low Power Stereo Matching Processors", *Proc. ISLPED*, 2013, pp. 15-20.

[26] M. M. Ozdal, C. Amin, A. Ayupov, S. M. Burns, G. R. Wilke and C. Zhuo, "ISPD-2012 Discrete Cell Sizing Contest and Benchmark Suite", *Proc. ISPD*, 2012, pp. 161-164.

[27] S. Panth, K. Samadi, Y. Du and S. K. Lim, "High-Density Integration of Functional Modules using Monolithic 3D-IC Technology", *Proc. ASP-DAC*, 2013, pp. 681-686.

[28] S. Panth, K. Samadi, Y. Du and S. K. Lim, "Design and CAD Methodologies for Low Power Gate-Level Monolithic 3D ICs", *Proc. ISLPED*, 2014, pp. 171-176.

[29] S. Panth, K. Samadi, Y. Du and S. K. Lim, "Placement-Driven Partitioning for Congestion Mitigation in Monolithic 3D IC Designs", *IEEE TCAD* 34(4) (2014), pp. 540-553.

[30] M. Pathak, Y. J. Lee, T. Moon and S. K. Lim, "Through-Silicon-Via Management during 3D Physical Design: When to Add and How Many?", *Proc. ICCAD*, 2010, pp. 387-394.

[31] S. Priyadarshi, W. R. Davis and P. D. Franzon, "Pathfinder3D: A Framework for Exploring Early Thermal Tradeoffs in 3DIC", *Proc. IC Design & Technology (ICICDT)*, 2014, pp. 1-6.

[32] M. Rahman, R. Afonso, H. Tennakoon and C. Sechen, "Power Reduction via Separate Synthesis and Physical Libraries", *Proc. DAC*, 2011, pp. 627-632.

[33] J.-S. Seo, I. L. Markov, D. Sylvester and D. Blaauw, "On the Decreasing Significance of Large Standard Cells in Technology Mapping", *Proc. ICCAD*, 2008, pp. 116-121.

[34] T. Song, M. Jung, Y. Wan, Y. Peng and S. K. Lim, "3D IC Power Benefit Study Under Practical Design Considerations", *Proc. International Interconnect Technology Conference and Materials for Advanced Metallization Conference (IITC/MAM)*, 2015, pp. 335-338.

[35] T. Song, S. Panth, Y. J. Chae and S. K. Lim, "Three-Tier 3D ICs for More Power Reduction: Strategies in CAD, Design, and Bonding Selection", *Proc. ICCAD*, 2015, pp. 752-757.

[36] D. Stroobandt, P. Verplaetse and J. van Campenhout, "Generating Synthetic Benchmark Circuits for Evaluating CAD Tools", *IEEE TCAD* 19(9) (2000), pp. 1011-1022.

[37] T. Thorolfsson, G. Luo, J. Cong and P. D. Franzon, "Logic-on-Logic 3D Integration and Placement", *Proc. International 3D Systems Integration Conference (3DIC)*, 2010, pp. 1-4.

[38] H. Xu, V. F. Pavlidis and G. De Micheli, "Effect of Process Variations in 3D Global Clock Distribution Networks" *ACM JETC* 8(3) (2012), pp. 20:1-20:25.

[39] *Cadence Innovus User Guide.*

[40] *MLPart. http://vlsicad.ucsd.edu/GSRC/bookshelf/Slots/Partitioning/MLPart*

[41] *OpenCores. http://opencores.org*

[42] *Synopsys Design Compiler User's Manual.*

[43] *Cadence Tempus User Guide.*