

Accurate Shielded Interconnect Delay Estimation by Reconfigurable Ring Oscillator

Eyal Sarfati, Binyamin Frankel, Yitzhak Birk, *Senior Member, IEEE*, and Shmuel Wimer^{ID}, *Member, IEEE*

Abstract—Shielding, which is used in VLSI designs to prevent noise interference from the cross-coupling capacitance between adjacent signals can also be used to tune the propagation delay of the clock signals in designs operating at low GHz frequencies. This paper presents a detailed design for a 16-nm ring oscillator with built-in reconfigurable shielding, and a delay estimation methodology. Together these provide a post-silicon measurement methodology that can derive accurate shielding delays without any direct delay measurements. The shielded ring oscillator and the testing methodology are designed to minimize the effects of on-die variations on estimation accuracy. Comparisons of the estimated delays with SPICE simulations show very good fit across process technology corners. The circuit was fabricated in 16-nm technology. The accuracy and robustness of the estimation methodology were verified by cross-validation, obtained from both pre and post-silicon measurements.

Index Terms—Integrated circuit interconnections, parameter estimation, ring oscillator, wire shielding.

I. INTRODUCTION

INTERCONNECT shielding is used in Very Large Scale Integration (VLSI) designs to prevent noise interference between signals. The clock signals spread over the entire silicon die to synchronize the operation of the underlying circuits in digital systems are the noisiest, and hence are shielded. They are a source of *signal integrity* problems, which can be avoided by extensive usage of shielding [1]. Clock signals connected to each sequential element (e.g., latch, flip-flop) are sometimes delayed with respect to each other. This is done by inserting *delay buffers* into the clock distribution network [2]–[5], among others. Often, intentional delay buffers are also inserted into logic signal paths to solve *min-delay* (hold) problems [6]. The internal delay of the buffers is subject to wide, unpredictable changes due to *process variation*, and this has been aggravated by recent progress in reducing VLSI technologies to the nanometer scale [7], [8].

Manuscript received February 27, 2018; revised April 8, 2018; accepted April 9, 2018. Date of publication April 26, 2018; date of current version August 30, 2018. This work was supported in part by the Israel Chief Scientist under the HiPer Consortium of the MAGNET Program and in part by the Marvell Corporate. This paper was recommended by Associate Editor E. Blokhina. (*Corresponding author: Shmuel Wimer.*)

E. Sarfati and Y. Birk are with the Electrical Engineering Department, Technion, Haifa 32000, Israel (e-mail: eyal.sarfati@gmail.com; birk@ee.technion.ac.il).

B. Frankel and S. Wimer are with the Engineering Faculty, Bar-Ilan University, Ramat-Gan 52900, Israel (e-mail: binyamin.frankel@gmail.com; wimers@biu.ac.il).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2018.2825999

Using shields as delay elements requires knowledge of their behavior in real silicon, which may be quite different from the delays obtained in RC layout extraction and SPICE simulations. Due to the lack of observability of internal nodes, direct delay measurements in silicon are very expensive, and in many cases impossible. Hence, there is a real need for indirect measurement of shielding delay effects. To the best of our knowledge, a post-silicon measurement and a methodology that can provide an accurate estimate of the in-circuit shielding impact on delays has not yet been proposed.

An early study on ways to lessen the cross-coupling delay burden incurred by shielding by allowing variable spacing was presented in [9] and [10]. In a recent work, Frankel and Wimer [11] proposed tapering the space of the clock shielding wires to solve the clock tuning problem by useful skew as an alternative to the insertion of expensive delay buffers. It was shown that the optimal space tapering yielding the desired propagation delay with minimum area consumption was proportional to the square root of the distance from the driver to the receiver.

A later work [12] turned shield insertion into a practical clock tuning design flow as a part of the backend clock tree synthesis (CTS). The authors showed that for a memory controller and an ARM[®]-based processor, about 90% of the useful skew insertion prerequisites could be solved by appropriate shielding implementation. This work examined the impact of process variations on intentional required skew propagation delays [13]. It was shown in [12] that delays obtained by shields were 50% less sensitive to variations than those obtained by delay buffers. The other advantages of using shields rather than delay buffers such as the ease of late design changes (ECOs) were also discussed.

The main contributions of this paper are the following:

- a special reconfigurable ring oscillator accompanied by a testing circuit to measure the delay effect of shielding on silicon,
- a verified methodology that enables an accurate post-silicon estimation of the shielding impact on delays,
- a demonstration that delay tuning by shielding is possible over a wide range, and
- post-silicon measurements on 16nm test-chip, which confirmed the validity of the above.

The remainder of this paper is organized as follows. Section II discusses delay tuning by shielding. Section III proposes a special ring oscillator and a test circuit to indirectly measure shield delays on silicon. Section IV presents an

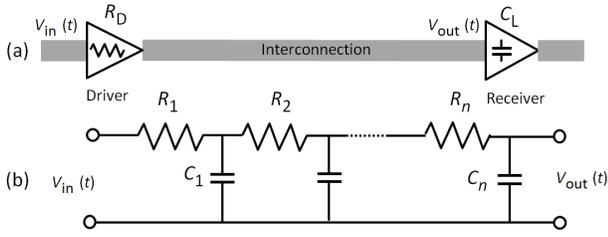


Fig. 1. Driver-to-receiver interconnect (a) and its RC-ladder modeling (b).

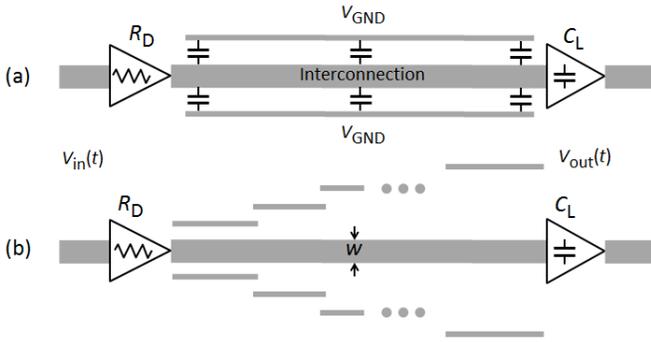


Fig. 2. Shielded interconnect modeling.

estimation methodology to obtain the delays of various shields in the internal, unobservable segments of the ring oscillator. Section V compares the estimated delays to simulated delays, and also demonstrates the robustness and validity of the entire estimation methodology. Section VI presents the post-silicon measurement results and Section VII draws conclusions.

II. DELAY TUNING BY TAPERING OF INTERCONNECTION SHIELDING

The *Elmore delay model* [14] has been widely used in VLSI design since its early days to calculate the interconnect delay [15]. Consider Fig. 1(a), where a driver connected on the near end sends a signal along a wire to a receiver connected on the far end. The driver's resistance R_D characterizes its driving strength. The receiver has an input capacitance C_L . An input unit impulse V_{in} is supplied to the near end at $t = 0$. The interconnection in Fig. 1(a) has distributed resistance and capacitance, and is usually modeled and approximated by the RC-ladder shown in Fig. 1(b), where $R_1 = R_D$ and $C_n = C_L$. The Elmore delay in this case is the limit of the sum over all the resistances multiplied by the downstream capacitance [16], as follows

$$\delta \approx \sum_{i=1}^n R_i \sum_{j=i}^n C_j = \sum_{j=1}^n C_j \sum_{i=1}^j R_i. \quad (1)$$

The connection from the driver to the receiver usually traverses several metal layers. However, we used the simplified but very popular and useful model shown in Fig. 1 where the interconnecting wire resides on a single metal layer.

Signals that are a source of significant noise are shielded, where the shielding wires are connected to a constant voltage V_{GND} or V_{DD} [1], as shown in Fig. 2(a). The cross-coupling

capacitance between the shielding wires and the interconnect signal introduces further driver-to-receiver propagation delays. This was studied in [17] and [18] in the context of optimizing a global interconnection design methodology, in conjunction with wire width and repeater insertion. These studies considered the delay incurred by shielding as an undesirable burden. However, it can be seen as an aid to solving per-signal problems by satisfying the required delay constraints. Both studies assumed shielding of fixed spacing from the signal wire, as shown in Fig. 2(a), whereas optimal per-signal delay tuning requires variable, piecewise-constant spacing, as shown in Fig. 2(b) [10]–[12].

As shown in Fig. 2(b), let us consider a wire of constant width w connecting the driver and the receiver. A two-sided shield extends along the wire, spaced at $s(x)$, $0 \leq x \leq L$. To make the illustration independent of nanometers and microns, the wire-to-wire spacing $s(x)$ is expressed as a multiplicative factor of s_{min} , which is the minimum wire-to-wire spacing allowable by the technology in use. There is typically $s(x) \in \{1 \times s_{min}, 2 \times s_{min}, 3 \times s_{min}\}$. A commonly used approximation for the unit length line-to-line capacitance of two adjacent wires is given by $c_{ll}/s(x)$, where c_{ll} is a technology parameter. The driver-to-receiver delay in Fig. 2(a) can be expressed as

$$\delta = \delta^{wire} + \delta_i^{shield}, \quad i = 1, 2, 3, \quad (2)$$

where δ^{wire} is the contribution of the signal wire and δ_i^{shield} is the contribution of the shield for the $1 \times s_{min}$, $2 \times s_{min}$ and $3 \times s_{min}$ spacings. A back-of-the-envelope calculation of the dynamic range of delay tuning achievable by the shielding which ignores all technology parameters and instead only utilizes geometric parameters was reported in [12]

$$\frac{\delta_1^{shield} - \delta_3^{shield}}{\delta^{wire} + \delta_2^{shield}} \approx \frac{4}{3(2w + 1)}, \quad (3)$$

where w is the multiplicative factor of w_{min} . The derivation of (3) assumed that the line-to-line and the ground capacitance are of the same order. The approximation estimated the delay dynamic range to be 44% ($\pm 22\%$) for wire widths of $1 \times w_{min}$. It was also shown that the shield delay was 50% less sensitive to process variations than the buffer delay. Smaller variations ensured that the useful skew would be sustained across a wide range of operation and silicon conditions. This work employed 16nm technology, but it is important to note that the ability to tune the delay by shielding increases for scaled CMOS due to the ever-increasing ratio between the line-to-line and the ground capacitance [19].

III. MEASURING THE EFFECTS OF SHIELDS IN REAL LAYOUTS

Measuring the delays of interconnects directly from silicon is very difficult and practically impossible. Silicon testing affords very limited probing of delay paths, because there is no visibility of their constituent delay segments or internal nodes. Most works on post-silicon delay measurements have dealt with the degree of fit between the model parameters of the logic cell library and the real silicon parameters.

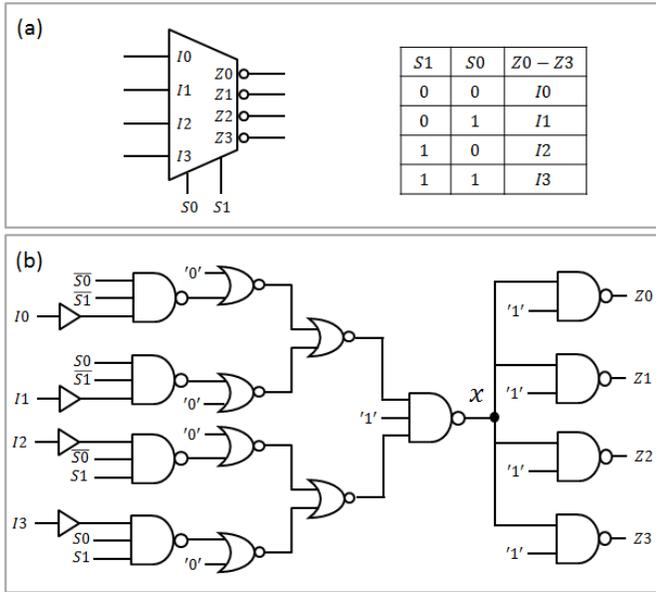


Fig. 3. An inverting MUX.

Chen and Liou [20] used statistical estimation methods for matching purposes. Jang *et al.* [21] derived the delays of the gates comprising the measured paths by an equality-constrained least squares estimation method. They faced the problem of solving the gate delay approximation in an under-determined equation system. The under-determination stems from the small number of measured post-silicon paths compared to the large number of internal unknown segment delays. In a recent work [22] the authors approached the problem of an undetermined system by introducing uncertainties into the simulation, which improved the estimated delays. In contrast to these works, our method provides an overdetermined system which is helpful in achieving accurate estimation.

A CMOS ring oscillator is typically used for the evaluation of the gate delay from silicon, which is an analog phenomenon, by indirect calculation based on counting pulses [23]. A ring oscillator is comprised of an odd number N of inverting gates connected in a closed loop. The oscillation frequency is given by $1/2N\tau$, where τ is the gate delay. To accurately measure the shield delay from silicon we devised a reconfigurable shielded interconnect ring-oscillator circuit, as described below. It is based on a five-stage inverting chain comprised of four inverting MUX stages and an additional inverting stage.

A. Inverting MUX

To support the measurement of different shielded interconnect delays, the inverting stage is implemented by an inverting MUX which makes it possible to select from among four of the input shielding configurations. Since the inverting stages are cascaded in a loop, the output of the inverting MUX is fanned to all four outputs, one of which will be selected by the next stage to be its input. Fig. 3(a) depicts the symbol and functionality of the inverting MUX.

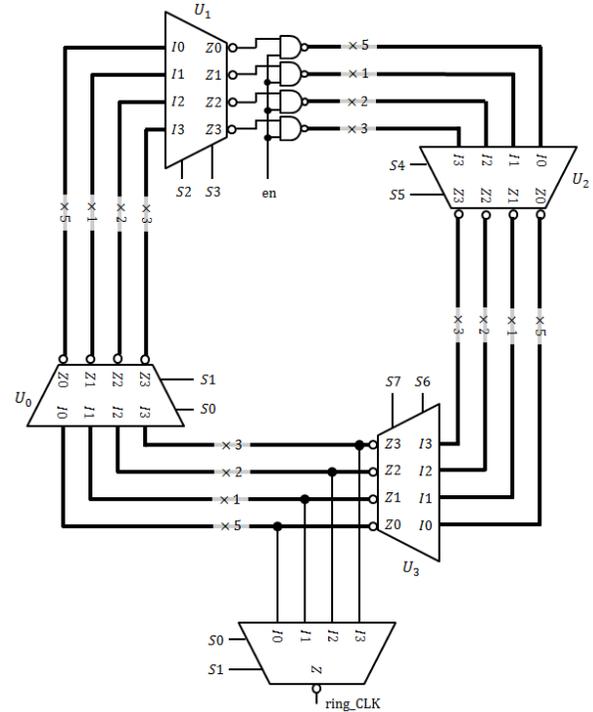


Fig. 4. Five-stage shielded interconnect ring oscillator.

The gate implementation of the MUX is illustrated in Fig. 3(b). Every input-to-output path is comprised of five inverting gates. To maintain perfect symmetry and delay equality among all the paths, appropriate inputs are connected at the gates of each stage. There, all the inputs are isolated via buffers to ensure an identical input load for the driving stages. The remainder of every path is made up of two NAND and two NOR gates. Further symmetry and identity are obtained by alternating between the “low” and “high” pins.

The circuit was implemented in a TSMC 16nmFFC standard cell library. The output was fanned out via a 2-way NAND gate to obtain a similar driving strength as the enabling switch shown in Fig. 4. This switch is another inverting stage of the oscillator loop; it is highly desirable for all the shielded interconnects that are part of the loop to be driven similarly.

As detailed in sections IV and V below, a key part of the shielding delay derivation is the assumption that the inverting MUX in Fig. 3 has a similar propagation delay from any input to any output. This is achieved in our design by appropriate selection of the gate’s inputs and careful layout artwork. To assess the delay sensitivity to the selected input-to-output path, all 16 distinct paths were simulated with SPICE at a (typical P, typical N, 0.8V, 85°C, typical RC wire) corner. An input slope of 100psec was used. Though large in slope, Fig. 4 below shows that the inputs of the MUX are driven through long shielded wires that degrade the slopes significantly. The SPICE model was extracted from the GDSII layout implementation with the StarRC[®] Synopsys tool [24]. Table I lists the delays from the various inputs of the MUX to its various outputs, where all the outputs were connected to same capacitive load. The variabilities in the rise-to-fall and fall-to-rise delays across the entire 16 input-to-output

TABLE I
INPUT-TO-OUTPUT DELAYS OF THE INVERTING MUX

Input	Output	rise→fall [psec]	fall→rise [psec]	Input	Output	rise→fall [psec]	fall→rise [psec]
I ₀	Z ₀	63.6	69.7	I ₂	Z ₀	63.7	71.1
	Z ₁	63.1	69.0		Z ₁	63.2	70.4
	Z ₂	64.1	70.3		Z ₂	64.2	71.7
	Z ₃	64.2	70.5		Z ₃	64.3	71.9
I ₁	Z ₀	63.8	69.9	I ₃	Z ₀	63.7	71.3
	Z ₁	63.3	69.2		Z ₁	63.2	70.6
	Z ₂	64.2	70.4		Z ₂	64.0	71.9
	Z ₃	64.3	70.7		Z ₃	64.1	72.1

TABLE II
INPUT-TO-OUTPUT DELAY VARIABILITIES IN DIFFERENT CORNERS

corner					rise→fall		fall→rise	
P device	N device	V _{dd} [Volt]	temp [°C]	wire RC	average [psec]	max var ± [%]	average [psec]	max var ± [%]
typical	typical	0.8	+85	typical	63.7	0.9	70.6	2.2
typical	typical	1.0	+85	typical	55.5	1.6	61.1	2.3
slow	slow	0.72	-40	worst	79.2	1.4	90.2	2.2
slow	slow	0.72	+125	worst	79.7	1.2	88.6	2.0
fast	fast	1.05	-40	best	45.7	1.8	50.9	1.8
fast	fast	1.05	+125	best	53.6	1.6	57.4	1.8

combinations were ± 0.60 psec and ± 1.55 psec, respectively, which are practically negligible.

To validate that the similarity of propagation delays from any input to any output of the inverting MUX was preserved across process variations and operation conditions, it was simulated in various corners. The results are summarized in Table II. As shown in the grayed columns, the delay variability within a corner was calculated as the difference between the paths yielding the maximum and the minimum propagation delays. Although the propagation delays changed considerably across different corners, which is typical and was expected, the delay variabilities within each corner remained small. As discussed below, the absolute delay values are immaterial to estimating the shielding delay. What matters is the delay similarity of the various input-to-output paths.

B. Shielded Interconnect Ring Oscillator

Fig. 4 illustrates the five-stage inverting loop implementing the shielded interconnect ring oscillator. The loop was comprised of four inverting MUXes, U_0 - U_3 , plus interim enabling NAND gates, yielding altogether an odd number of inversions. The oscillating output signal ring_CLK was selected by a 4-to-1 MUX to isolate the oscillator from the measurement. The oscillator design used the same metal6 $1 \times w_{\min}$ signal width, depicted in Fig. 4 by the bold lines, where w_{\min} is the minimum wire width of the technology. The interconnect connecting a stage to its successor was $200\mu\text{m}$ long.

The layout of the configurable shielded ring oscillator is shown in Fig. 5(a) to illustrate the inverting MUXes and the interconnecting segments. Each stage in the loop is highlighted

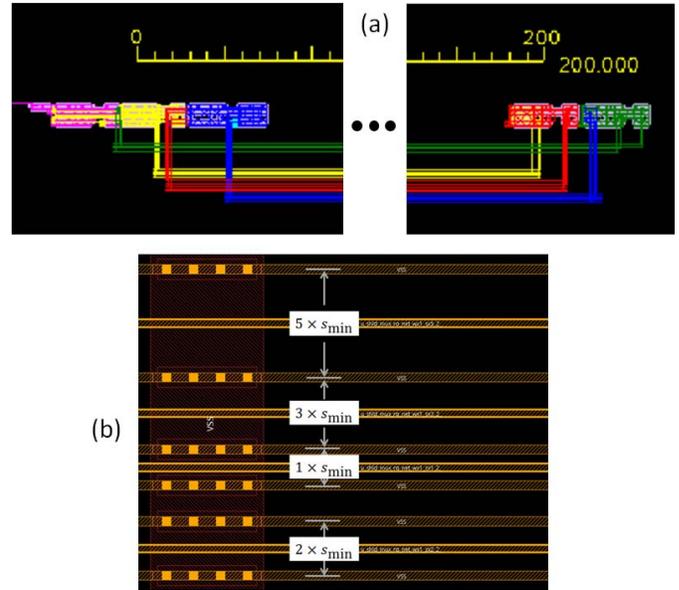


Fig. 5. Layout of ring oscillator and GDSII of shielded wires.

in a different color. To keep the silicon variability effects on the shielded interconnects in the four stages as similar as possible, the ring was placed so that the two MUXes on each side were $200\mu\text{m}$ apart from each other. This preserved the proximity of the shielded wire segments, thus enabling maximum accuracy of the post-silicon estimation. As shown in Fig. 5(b), the wires were shielded with different spaces of $1 \times s_{\min}$, $2 \times s_{\min}$ and $3 \times s_{\min}$. To represent no shielding, we used $5 \times s_{\min}$ spacing instead. The $5 \times s_{\min}$ shielding was a must since otherwise neighboring signals and shields would cause uncontrolled interference and shielding. At each MUX, the chosen input was specified by two selection signals. The oscillator was turned on/off by the enabling signal 'en' of the NAND stage.

To ensure full symmetry and identity in the driving power of the interconnections, the NAND gates were identical to those in the last stage of the MUX shown in Fig. 3(b). For purposes of measurement accuracy, the NAND gates were located two stages before the measured output. Since the time aperture used to count the number of rising clock edges of the oscillator starts simultaneously with the oscillator enabling signal, we needed to have the first counted edge output present late enough into the oscillation counting aperture. Note that in Fig. 4 different units can have different output loads. Whereas U_0 , U_2 and the NAND stage only drive shielded wires, U_3 drives the shielded wires and the output MUX. To compensate for this, appropriate loads (not shown) were added to ensure that all the outputs of the interconnection drivers had an identical load.

Though ideally we wanted all the $200\mu\text{m}$ interconnections having same shields to behave identically, inevitably there will be differences due to the distributed nature of the layout and the on-die variations. The ring oscillator was made up of 16 unknown delays, denoted by δ_i^j , where $i \in \{1, 2, 3, 5\}$ is the shielding distance of $i \times s_{\min}$, and $j \in \{0, 1, 2, 3\}$ is

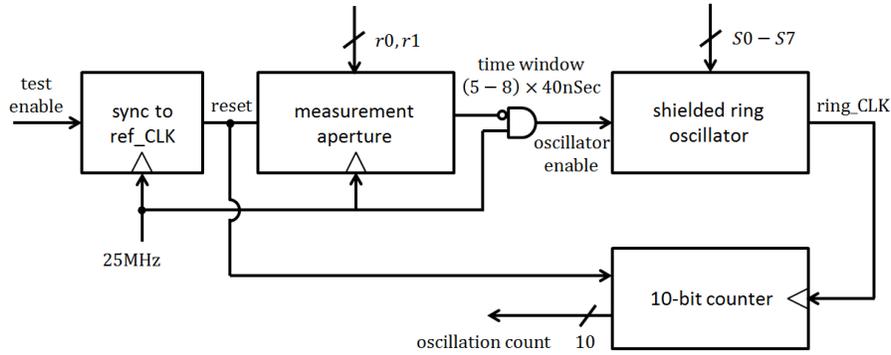


Fig. 6. Testing circuit.

the inverting MUX driving unit U_j . The selections of $S0 - S7$ defined a total of $4^4 = 256$ distinct path compositions, yielding an oscillation frequency determined by the internal delays of the inverting stages in the ring and the particular selection of the shielded interconnections in the ring. Note that $S0$ and $S1$ were also used to output the oscillating signal in accordance with the path composition under test.

C. Testing Setup and Measurement Methodology

Fig. 6 depicts the testing circuit to derive the delays of the 16 shielded wires in Fig. 4. It was comprised of the ring oscillator, a tunable aperture circuit to count the oscillation pulses, a counter and a synchronizer. The circuit used a 25MHz reference clock signal (40nsec cycle). The testing program first decides on one of the 256 possible rings of Fig. 4 by setting $S0$ to $S7$ appropriately. Each test measures the delay of a shielding configuration by counting the number of oscillations within the measurement aperture. The test is launched by a test enable signal synchronized to the rising edge of the 25MHz reference clock. The synchronized reset signal resets a 10-bit counter and starts the measurement aperture which counts the number of ring oscillator pulses. The ring oscillator is enabled via an AND gate. The duration of the measurement aperture is set by $r0$ and $r1$, which in turn set the value of a shift register to 200nsec – 320nsec in steps of 40nsec. Appropriate tuning of the aperture duration ensures that the counter will not overflow during the measurement but rather will extend the aperture duration to maximize the number of counted oscillations, thus increasing the measurement accuracy. Once the aperture is closed, the testing program records the number of oscillations, from which the delay of the ring oscillator can be calculated.

IV. SHIELDED INTERCONNECT DELAY ESTIMATION METHODOLOGY

Measuring delays directly on silicon is complex and expensive, whereas measuring the frequencies of a ring oscillator to any desirable accuracy is relatively simple, as shown in Fig. 6. Once the delays are derived from the oscillator frequencies, the question is how to deduce the effects of various shielding on the delays indirectly.

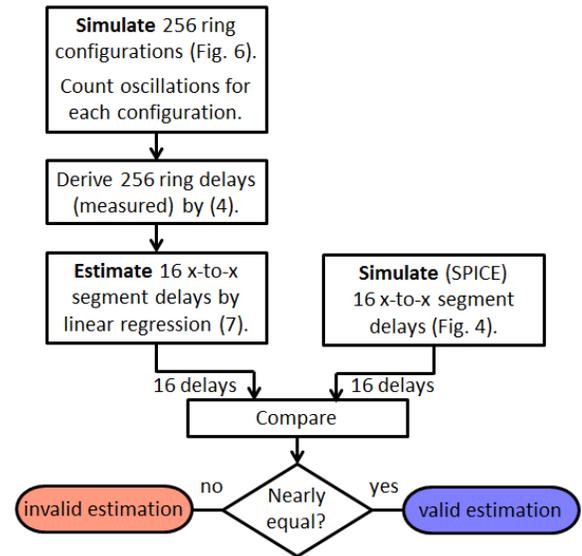


Fig. 7. Pre-silicon delay estimation and validation flow.

Given that this estimation methodology is designed to take place in silicon without any direct delay measurements, it requires validation by comparison to a SPICE simulation. The SPICE results relying on technology parameters and characterizations provide the pre-silicon delays. SPICE makes it possible to validate the delays of the shielded wires which are parametrically estimated [25] from oscillation frequency samples (as in real applications in silicon). This estimation and validation flow is depicted in Fig. 7.

After the matching is confirmed, one can rely on oscillation frequencies and parameter estimation procedures to deliver the real silicon delays of the shielded wires. It is important to note that the parameter estimation is blind to any technology parameter. The accurate delay tuning range can be easily obtained from the estimated delays of the shielded wires by using expressions similar to the left hand side of (3).

Let $0 \leq k \leq 255$ be the index of a ring oscillator configuration obtained by the selection of $S0 - S7$. Let Δ_k be the corresponding delay obtained by dividing the duration t of the measurement aperture by the number n_k of the counted

oscillations; namely,

$$\Delta_k = \frac{t}{n_k}, \quad 0 \leq k \leq 255. \quad (4)$$

The delays obtained in (4) are the measurements. The delay is comprised of four components, each of which starts at an input of a NAND gate denoted by x in Fig. 3(b), and terminates at the output of another NAND gate within the inverting MUX of the successive stage in Fig 4, also denoted by x . There are altogether 16 such segments, whose delays are estimated. Since the SPICE simulations provide these delays directly, the comparison between the estimated and the measured delays can evaluate the accuracy of the estimation methodology as illustrated in Fig. 7. The x -to- x delays also include the MUX delay shown in Table I. Since this work focused on finding the dynamic delay range obtained by shielding, the MUX delay had a negligible effect. If one is interested in the shielded wire delay per-se, the MUX delay in Table I should be subtracted.

Let $\delta_{i(j)}^j$ be one of the 16 x -to- x delay segments, where $0 \leq j \leq 3$ designates one of the four stages of the ring oscillator in Fig. 4, and $i(j) \in \{1, 2, 3, 5\}$ designates the selection of one out of the four possible shield spacings $i(j) \times s_{\min}$ in the corresponding stage. The following equality holds:

$$\Delta_k = \delta_{i(0)}^0 + \delta_{i(1)}^1 + \delta_{i(2)}^2 + \delta_{i(3)}^3, \quad 0 \leq k \leq 255, \quad (5)$$

where k is spanned over all possible selections of the ring configurations in Fig. 4. The linear system in (5) can be written in the following matrix notation

$$\mathbf{H}\boldsymbol{\delta} = \boldsymbol{\Delta}, \quad (6)$$

where $\boldsymbol{\delta} = [\delta_1^0, \delta_2^0, \delta_3^0, \delta_5^0, \dots, \delta_1^3, \delta_2^3, \delta_3^3, \delta_5^3]^T$ is a 16×1 vector of unknown x -to- x segment delays of the ring oscillator, $\boldsymbol{\Delta} = [\Delta_0, \Delta_1, \dots, \Delta_{255}]^T$ is a 256×1 vector of delay measurements obtained by the testing circuit in Fig. 6. Finally, \mathbf{H} is an 256×16 zero-one matrix, each row of which comprises four ones representing a specific configuration under testing in (5).

The 256 equations involve 16 unknown parameters, yielding an overdetermined linear system. Note that any x -to- x segment is involved in 64 configurations, as dictated by the three other stages. If an x -to- x segment had an identical impact on each of the 64 configurations, one could choose any 16 row-independent equation out of the 256 of (5) to solve the system. In reality, however, the impact of the same x -to- x segment can vary across configurations. This in turn results in some noise in the measured Δ_k , thus making it impossible to obtain an accurate solution. In this case, least square parameter estimation is needed [25]. An ordinary least square would be solved (6) by

$$\hat{\boldsymbol{\delta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \boldsymbol{\Delta}, \quad (7)$$

where $\hat{\boldsymbol{\delta}}$ is the estimated solution. Unlike in silicon, the x -to- x segment delay $\delta_{i(j)}^j$ can be measured by the SPICE simulation, and compared to its estimated value $\hat{\delta}_{i(j)}^j$ as obtained by the approximated solution in (7).

Unfortunately, the rank of the 16×16 $\mathbf{H}^T \mathbf{H}$ matrix is less than 16, and hence not invertible. The rank deficiency follows

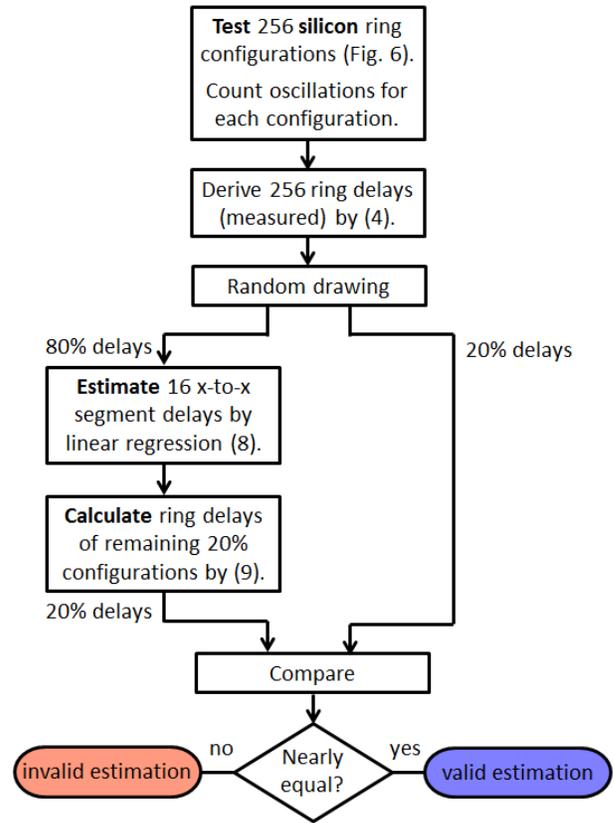


Fig. 8. Testing the quality of the post-silicon delay estimation.

from the fact that not all four variables (the four ones in a row of \mathbf{H}) out of the 16 can be chosen arbitrarily. Rather, these are divided into groups of four variables each, where an equation involves one and only one variable in each group stemming from the selection of a single x -to- x segment within a stage of the ring oscillator. Overall, this dependence yields a matrix of rank 13. Such cases are usually treated with a method called *pseudo inversion* [26], which is another kind of least square approximation.

While this work presents a reconfigurable ring oscillator with identical stages, the rationale can be used to design any ring oscillator, where each stage has different wire lengths. Appropriate configurability can support any combination of wire lengths and shield spacing. The estimation methodology elaborated below only requires solving the appropriate linear equations system with an appropriate choice of variables.

Recall that there is not any direct post-silicon delay measurement, so the 16 estimated post-silicon delay segments cannot be compared to anything. How can we be confident that the linear regression in (7) yields a valid post-silicon estimation? To this end we used Monte Carlo cross-validation. There, the unknown parameters are estimated by a portion of the measurements drawn randomly. The remaining portion is first computed by using the estimated parameters and then compared to the corresponding measurements [27].

This type of flow is illustrated in Fig. 8. Here, 80% of the measurements are drawn randomly from the vector $\boldsymbol{\Delta} = [\Delta_0, \Delta_1, \dots, \Delta_{255}]^T$, denoted by $\boldsymbol{\Delta}_{80\%}$. These measurements

TABLE III
DYNAMIC RANGE OF INTERCONNECT DELAY TUNING BY SHIELDING

corner					Sim. period [psec]		Sim. range [%]	Est. period [psec]	
P device	N device	Vdd [Volt]	Temp [°C]	wire RC	$1 \times S_{\min}$	$5 \times S_{\min}$	$1 \times -5 \times$ $0.5(1 \times +5 \times)$	$1 \times S_{\min}$	$5 \times S_{\min}$
typical	typical	0.8	+85	typical	989	662	39.6	988	662
typical	typical	1.0	+85	typical	909	588	42.9	908	588
slow	slow	0.72	-40	worst	1050	759	32.2	1046	760
slow	slow	0.72	+125	worst	1160	809	35.6	1163	809
fast	fast	1.05	-40	best	767	482	45.6	767	482
fast	fast	1.05	+125	best	963	593	47.5	963	594

with their corresponding rows in matrix \mathbf{H} , denoted by $\mathbf{H}_{80\%}$, are used to estimate the delays of the 16 shielded wires, denoted by $\hat{\delta}_{80\%}$. The following system is solved to yield the delays

$$\hat{\delta}_{80\%} = (\mathbf{H}_{80\%}^T \mathbf{H}_{80\%})^{-1} \mathbf{H}_{80\%}^T \Delta_{80\%}. \quad (8)$$

To verify the accuracy of $\hat{\delta}_{80\%}$, the remaining 20% measurements of the vector $\Delta = [\Delta_0, \Delta_1, \dots, \Delta_{255}]^T$, denoted by $\Delta_{20\%}$, are compared to their corresponding predicted values. A predicted delay $\hat{\Delta}$ is calculated by summing the appropriate estimated delays of $\hat{\delta}_{80\%}$, defined by the ring oscillator configuration corresponding to $\Delta \in \Delta_{20\%}$ as follows

$$\hat{\Delta} = \hat{\delta}_{80\%,i(0)}^0 + \hat{\delta}_{80\%,i(1)}^1 + \hat{\delta}_{80\%,i(2)}^2 + \hat{\delta}_{80\%,i(3)}^3. \quad (9)$$

If $|\Delta - \hat{\Delta}| \approx 0$ for any $\Delta \in \Delta_{20\%}$, we consider the estimation to be reasonably accurate. This is tested in the next section.

V. COMPARISON OF THE ESTIMATION TO THE MODEL SIMULATION

One of the goals of this work was to derive the dynamic range of delay tuning from silicon without any direct delay measurements or any knowledge of the technology parameters or the underlying models. This was first done by simulating the testing circuit of Fig. 6, from which the delays of two rings of Fig. 4 were obtained. The first was comprised of four x -to- x segments with $1 \times s_{\min}$ spacing, and the other was comprised of four x -to- x segments with $5 \times s_{\min}$ spacing. The oscillator was simulated at six corners, and the delays were obtained by dividing the number of oscillations by the measurement aperture. The results are shown in Table III.

The dynamic range of delay was obtained by subtracting these delays and dividing by their average. As shown in the grayed column, the 32% to 47% dynamic ranges of delay tuning were obtained in different corners. The back-of-the-envelope calculation in (3) yielded a 44% dynamic range for wire widths of $1 \times w_{\min}$, which is the width we used in the physical layout in Fig. 5 of the ring in Fig. 4. This is consistent with the dynamic ranges in Table III.

Another main goal of this work was to accurately estimate the delays obtained by using various shield spacings. We used a Monte Carlo cross-validation, for which the system (8) was solved to obtain the estimation of the 16 x -to- x segment

delays in Fig. 4. The measurements were obtained by SPICE simulations of the testing circuit as shown in Fig. 6, and then applying (4).

The comparison simulated delays were obtained by SPICE for every x -to- x segment in Fig. 4. The estimated-simulated comparison was conducted for all six corners. The results are summarized in Table IV. Typical, slow and fast devices are denoted by t, s and f, respectively, whereas typical, worst and best RCs are denoted by t, w and b, respectively. The worst accuracies at each corner are highlighted and the mean error is shown on the right-hand side.

Note that the run-time to simulate the testing circuit in Fig. 6 was extremely long (overnight runs) since it involved hundreds of oscillations per simulation. However, we were not concerned about the length of the run-time since the purpose of these runs was to validate the estimation methodology which was used on-silicon where no simulations are involved.

We can now return to the dynamic range of delay tuning by shielding which was derived from the simulated ring delays for $1 \times s_{\min}$ and $5 \times s_{\min}$ shield spacings, and compare it to the corresponding x -to- x segment estimated delays. The latter were obtained by summing the estimated delays of the segments comprising the $1 \times s_{\min}$ and $5 \times s_{\min}$ rings as in (5). These are listed in the two right-hand columns of Table III, and are shown to be almost identical to the simulated delays of the complete ring.

The estimation quality was validated by using the 80/20 cross-validation methodology elaborated above and illustrated in Fig. 8. Fig. 9 plots the estimated delays $\hat{\Delta}_{20\%}$ obtained by (9) versus the $\Delta_{20\%}$ delays obtained by simulating the ring oscillator at a (typical P, typical N, 0.8V, 85°C, typical RC wire) corner. The results align perfectly with the 45 degree line, with a negligible error of less than 0.1%, thus confirming the validity of the estimation methodology.

VI. POST-SILICON MEASUREMENTS AND VALIDATION

The shielded ring oscillator shown in Fig. 4 and its accompanying testing system shown in Fig. 6 were fabricated in TSMC 16nm technology on a Marvell Corporate test-chip. Testing in the corners adhered to the methodology elaborated in Section III.C. To maximize the accuracy, the longest measurement aperture duration of 320nsec was used. Based on the post-silicon measurements, the delay of the 16 shielded

TABLE IV
ESTIMATED-SIMULATED COMPARISON OF x -TO- x SEGMENT DELAYS

corner	x – to – x	δ_1^0	δ_2^0	δ_3^0	δ_5^0	δ_1^1	δ_2^1	δ_3^1	δ_5^1	δ_1^2	δ_2^2	δ_3^2	δ_5^2	δ_1^3	δ_2^3	δ_3^3	δ_5^3	Mean Err [%]
t P, t N 0.8V +85°C t RC	Sim. [psec]	248	207	180	167	241	199	174	11	254	210	184	167	247	209	182	167	1.56
	Est. [psec]	246	207	180	166	247	206	180	167	250	206	180	163	245	208	181	166	
	Err [%]	0.8	0.2	0.1	0.6	2.4	3.2	3.7	3.5	1.7	1.6	1.9	2.4	0.7	0.6	0.6	1	
t P, t N 1.0V +85°C t RC	Sim. [psec]	228	188	161	149	222	181	156	143	234	190	164	148	226	189	163	148	1.5
	Est. [psec]	226	187	161	148	227	186	161	148	230	187	161	145	225	188	162	147	
	Err [%]	0.7	0.1	0.1	0.6	2.2	3.1	3.6	3.4	1.7	1.7	1.9	2.4	0.6	0.5	0.4	0.8	
s P, s N 0.72V –40°C w RC	Sim. [psec]	262	228	205	192	253	217	195	183	269	230	208	192	264	231	207	193	2.01
	Est. [psec]	260	227	204	191	262	226	204	191	264	226	204	188	260	228	204	190	
	Err [%]	0.9	0.4	0.4	0.8	3.3	4.1	4.7	4.5	1.8	1.7	1.8	2.3	1.3	1.4	1.3	1.7	
s P, s N 0.72V +125°C w RC	Sim. [psec]	292	249	220	205	283	238	211	196	299	251	223	204	292	252	222	205	1.75
	Est. [psec]	289	248	220	203	291	247	219	204	294	247	219	200	289	249	220	202	
	Err [%]	0.8	0.3	0.2	0.6	2.8	3.6	4.1	3.9	1.7	1.6	1.8	2.3	1	1	1	1.3	
f P, f N 1.05V –40°C b RC	Sim. [psec]	192	154	131	122	187	149	127	118	198	157	134	122	190	155	132	121	1.43
	Est. [psec]	191	154	131	121	191	153	131	122	195	153	131	118	190	155	132	121	
	Err [%]	0.5	0	0	0.5	2	3	3.5	3.2	1.8	2	2.4	2.8	0.3	0.1	0.1	0.5	
f P, f N 1.05V +125°C b RC	Sim. [psec]	241	192	162	150	235	186	157	145	249	195	165	150	239	193	163	149	1.37
	Est. [psec]	240	192	162	149	240	191	162	150	245	191	161	146	238	193	163	149	
	Err [%]	0.5	0.1	0.1	0.5	1.8	2.7	3.2	3.1	1.8	2.1	2.5	2.9	0.1	0	0.1	0.2	

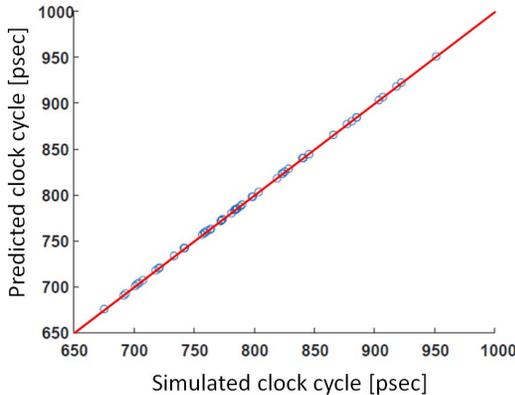


Fig. 9. Cross-validation of the simulated-estimated delay comparison.

interconnects were estimated by solving (8). The estimation quality was validated by applying the cross-validation depicted in Fig. 8.

Marvell provided us with typical silicon material, for which we tested 12 corners, covering the following temperatures $\{25^\circ\text{C}, 50^\circ\text{C}, 85^\circ\text{C}, 105^\circ\text{C}\}$ and supply voltages $\{0.8\text{V}, 0.9\text{V}, 1.0\text{V}\}$. Recall that for each corner 256 ring configurations needed to be tested to yield the 16 shielded wire segment delays. To filter random noise that can occur in measurements (e.g. power-supply fluctuations, thermal noise, etc.), each test of the 256 configurations was repeated 50 times, and the number of oscillations within the 320nsec aperture was averaged.

The pre-silicon delay estimation methodology comprised two flows. The first, shown in Fig. 7, compared the estimated delays to delays obtained directly from the SPICE simulation. The second, shown in Fig. 8, validated the estimation quality. While both were essential to proving the correctness of the estimation method, post-silicon can only use the second flow.

Fig. 10 shows the 80/20 cross-validation of the post-silicon measured-estimated clock cycles in 3 out of the 12 corners. Each corner shows the range of the post-silicon clock cycle measurements extending along the 45 degree line. The 20% validating clock cycles are scattered around the 45 degree line. For each corner the relative delay error $|\Delta - \hat{\Delta}|/\Delta$ was calculated for every $\Delta \in \Delta_{20\%}$, where Δ is the measured delay and $\hat{\Delta}$ was obtained by (9). The maximum cross-validation error in each corner is shown in the corresponding plot. The maximum error of all the 12 corners is summarized in Table V and all were below 0.2%, thus confirming the validity of the post-silicon estimation methodology. Note that the cross-validation shown in Fig. 9 was obtained for pre-silicon, which is essential to prove the estimation methodology, whereas these of Fig. 10 and Table V were obtained for post-silicon, thus proving the accuracy of the concrete estimations.

A goal of this work was to demonstrate that the timing of the clock signals could be tuned by utilizing the dynamic delay range obtained by shielding. Table VI shows the post-silicon delay ranges measured for the 12 corners, where the range was defined as in Table III: $(1 \times -5 \times)/0.5 (1 \times +5 \times)$. These are compared to the corresponding ranges obtained by SPICE simulations. Though slightly smaller, the silicon was able to deliver sufficiently large dynamic delay range.

It should be emphasized that the similarity between the pre-silicon and post-silicon delay is not the main point of this paper. In reality they can be quite different. The difference depends on the extent to which the technology model parameters used for the design fit the fabricated silicon. To examine the “typicality” claimed by the silicon manufacturer, we ran SPICE delay simulations of the ring oscillator in Fig. 4 in the 12 corners shown in Tables V to VII. In each corner the 16 post-silicon x -to- x delays were compared to their

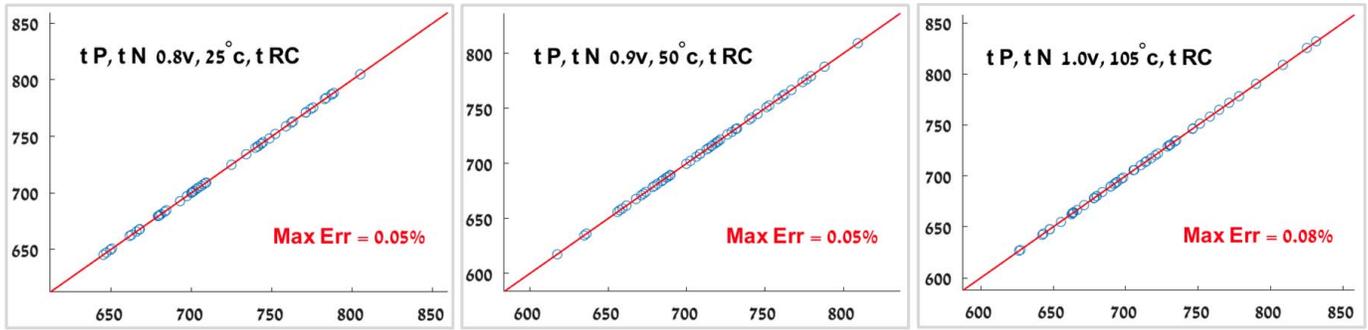


Fig. 10. Post-silicon cross-validation of the simulated-estimated delay comparison for 3 out of 12 corners.

TABLE V
POST-SILICON MAXIMUM VALIDATION ERROR OF 12 CORNERS

	0.8V				0.9V				1.0V			
	25°C	50°C	85°C	105°C	25°C	50°C	85°C	105°C	25°C	50°C	85°C	105°C
Max Validation Err [%]	0.05	0.06	0.07	0.07	0.12	0.05	0.08	0.08	0.17	0.16	0.07	0.08

TABLE VI
POST-SILICON TUNABLE DELAY RANGE

		0.8V				0.9V				1.0V			
		25°C	50°C	85°C	105°C	25°C	50°C	85°C	105°C	25°C	50°C	85°C	105°C
Delay range [%]	Post-silicon	33.33	33.72	34.18	34.51	34.9	35.3	35.71	35.94	36.11	36.53	37.02	37.12
	SPICE	38.32	38.87	39.50	39.83	40.37	40.88	41.33	41.62	41.86	42.25	42.78	42.97

TABLE VII
COMPARISON OF POST-SILICON SHIELDED SEGMENT DELAYS TO PRE-SILICON SPICE

	0.8V				0.9V				1.0V			
	25°C	50°C	85°C	105°C	25°C	50°C	85°C	105°C	25°C	50°C	85°C	105°C
Max SPICE Err [%]	9.17	9.09	8.90	9.08	8.98	9.10	9.20	9.19	9.32	9.28	9.43	9.51
Mean SPICE Err [%]	3.95	3.83	3.65	3.73	3.62	3.66	3.68	3.67	3.60	3.64	3.70	3.76

corresponding SPICE delays; the average and maximum difference are shown in Table VII. The actual differences were small, which has ramifications for practical design.

VII. CONCLUSION

A detailed design of a 16nm ring oscillator with built-in reconfigurable shielding accompanied by a delay estimation methodology was presented, whose accuracy and robustness were validated and demonstrated. The circuit and the methodology enable accurate post-silicon extraction of shielding delays without any direct delay measurements. The same methodology can be adapted to measure the delays of other interconnection structures composed of various metal layers, wire widths, and wire and shield tapering. The accuracy and robustness of the methodology was proved through post-silicon measurements.

ACKNOWLEDGMENT

The authors acknowledge the useful comments by the anonymous reviewers that helped us improve the manuscript.

REFERENCES

- [1] E. Salman and E. Friedman, *High Performance Integrated Circuit Design*. New York, NY, USA: McGraw-Hill, 2012.
- [2] S. Tam, S. Rusu, U. N. Desai, R. Kim, J. Zhang, and I. Young, "Clock generation and distribution for the first IA-64 microprocessor," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1545–1552, Nov. 2000.
- [3] S. Lee, S. Paik, and Y. Shin, "Retiming and time borrowing: Optimizing high-performance pulsed-latch-based circuits," in *IEEE/ACM Int. Conf. Comput.-Aided Design-Dig. Tech. Papers (ICCAD)*, Nov. 2009, pp. 375–380.
- [4] J. Kim, D. Joo, and T. Kim, "An optimal algorithm of adjustable delay buffer insertion for solving clock skew variation problem," in *Proc. 50th Annu. Design Autom. Conf.*, Jun. 2013, pp. 1–6.
- [5] Y. Kaplan and S. Wimer, "Mixing drivers in clock-tree for power supply noise reduction," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 5, pp. 1382–1391, May 2015.

- [6] N. H. Weste and K. Eshraghian, *Principles of CMOS VLSI Design*, vol. 2. Reading, MA, USA: Addison-Wesley, 1993.
- [7] A. Agarwal, D. Blaauw, and V. Zolotov, "Statistical timing analysis for intra-die process variations with spatial correlations," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design*, Nov. 2003, pp. 900–907.
- [8] C. Constantinescu, "Trends and challenges in VLSI circuit reliability," *IEEE Micro*, vol. 4, no. 4, pp. 14–19, Jul. 2003.
- [9] M. A. El-Moursy and E. G. Friedman, "Exponentially tapered H-tree clock distribution networks," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 13, no. 8, pp. 971–975, Aug. 2005.
- [10] M. A. Karami and A. Afzali-Kusha, "Exponentially tapering ground wires for Elmore delay reduction in on chip interconnects," in *Proc. Int. Conf. Microelectron.*, 2006, pp. 99–102.
- [11] B. Frankel and S. Wimer, "Optimal VLSI delay tuning by wire shielding," *J. Optim. Theory Appl.*, vol. 170, no. 3, pp. 1060–1067, 2016.
- [12] E. Sarfati, B. Frankel, Y. Birk, and S. Wimer, "Optimal VLSI delay tuning by space tapering with clock-tree application," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 64, no. 8, pp. 2160–2170, Aug. 2017.
- [13] M. Alioto, G. Palumbo, and M. Pennisi, "Understanding the effect of process variations on the delay of static and domino logic," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 18, no. 5, pp. 697–710, May 2010.
- [14] W. C. Elmore, "The transient response of damped linear networks with particular regard to wideband amplifiers," *J. Appl. Phys.*, vol. 19, no. 1, pp. 55–63, 1948.
- [15] H. B. Bakoglu and J. D. Meindl, "Optimal interconnection circuits for VLSI," *IEEE Trans. Electron Devices*, vol. ED-32, no. 5, pp. 903–909, May 1985.
- [16] J. Rubinstein, P. Penfield, and M. A. Horowitz, "Signal delay in RC tree networks," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. CAD-2, no. 3, pp. 202–211, Jul. 1983.
- [17] A. B. Kahng, S. Muddu, and E. Sarto, "Tuning strategies for global interconnects in high-performance deep-submicron ICs," *VLSI Design*, vol. 10, no. 1, pp. 21–34, 1999.
- [18] R. Jakushokas and E. G. Friedman, "Resource based optimization for simultaneous shield and repeater insertion," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 18, no. 5, pp. 742–749, May 2010.
- [19] *Interconnect Summary*, Int. Technol. Roadmap Semicond., Malaga, Spain, 2013.
- [20] Y.-Y. Chen and J.-J. Liou, "A non-intrusive and accurate inspection method for segment delay variabilities," in *Proc. Asian Test Symp.*, Nov. 2009, pp. 343–348.
- [21] E. J. Jang, J. Chung, A. Gattiker, S. Nassif, and J. A. Abraham, "Post-silicon timing validation method using path delay measurements," in *Proc. 20th Asian Test Symp. (ATS)*, 2011, pp. 232–237.
- [22] J. Chung and J. Kim, "Segment delay learning from quantized path delay measurements," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 34, no. 6, pp. 1038–1042, Jun. 2015.
- [23] Y. A. Eken and J. P. Uyemura, "A 5.9-GHz voltage-controlled ring oscillator in 0.18- μm CMOS," *IEEE J. Solid-State Circuits*, vol. 39, no. 1, pp. 230–233, Jan. 2004.
- [24] Synopsys. (2017). *StarRC Parasitic Extraction*. [Online]. Available: <https://www.synopsys.com/content/dam/synopsys/implementation&signoff/datasheets/starrc-ds.pdf>
- [25] B. Lei et al., *Classification, Parameter Estimation and State Estimation: An Engineering Approach Using MATLAB*. Hoboken, NJ, USA: Wiley, 2017.
- [26] A. Ben-Israel and T. N. E. Greville, *Generalized Inverses: Theory and Applications*, vol. 15. Springer, 2003.
- [27] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, 1995, pp. 1137–1143.



Eyal Sarfati received the B.Sc. degree in electrical engineering from the Technion-Israel Institute of Technology, Haifa, Israel, in 2009, where he is currently pursuing the M.Sc. degree. Since graduation, he has been with Marvell as a VLSI Backend Design Engineer.



Binyamin Frankel received the B.Sc. and M.Sc. degrees in electrical engineering from Bar-Ilan University in 2014 and 2016, respectively, where he is currently pursuing the Ph.D. degree in computer engineering. His research interests include VLSI circuits and systems design optimization.



Yitzhak Birk (M'82–SM'02) received the B.Sc. (*cum laude*) and M.Sc. degrees from the Technion-Israel Institute of Technology, Haifa, Israel, in 1975 and 1982, respectively, and the Ph.D. degree from Stanford University, Stanford, CA, USA, in 1987, all in electrical engineering. He has been with the faculty of the Electrical Engineering Department, Technion, since 1991, where he heads the Parallel Systems Laboratory. Previously, he was a Research Staff Member with the IBM's Almaden Research Center.

His research interests include computer and communication systems, and in particular storage subsystems and the interplay between storage and communication. The true application requirements are considered in each case. The judicious exploitation of redundancy, coding and randomization for performance enhancement, as well as cross-disciplinary approaches are recurring themes in much of his work.



Shmuel Wimer received the B.Sc. and M.Sc. degrees in mathematics from Tel-Aviv University, Israel, in 1978 and 1981, respectively, and the D.Sc. degree in electrical engineering from the Technion-Israel Institute of Technology, Israel, in 1988.

From 1978 to 2009, he held research and development, engineering, and managerial positions in industry. From 1999 to 2009, he was with Intel, and prior to that with IBM, National Semiconductor, and the IAI-Israel Aerospace Industry. He is an Associate Professor with the Engineering Faculty, Bar-Ilan University, Israel, and an Associate Visiting Professor at the Electrical Engineering Faculty, Technion. His research interests include VLSI circuits and systems design optimization and combinatorial optimization.