Fine-Grained Characterization of Process Variation in FPGAs

Haile Yu[1], Qiang Xu[1] and Philip H.W. Leong[2]
[1] *Department of Computer Science and Engineering, The Chinese University of Hong Kong*
{hlyu,qxu}@cse.cuhk.edu.hk
[2] *School of Electrical and Information Engineering, University of Sydney*
philip.leong@sydney.edu.au

***Abstract*—As semiconductor manufacturing continues towards reduced feature sizes, yield loss due to process variation becomes increasingly important. To address this issue on FPGA platforms, several variation aware design (VAD) methodologies have been proposed. In this work we present a practical method of process variation characterization (PVC) to facilitate VAD using only intrinsic FPGA resources. The scheme is based on measuring the difference between ring oscillator (RO) delay at different locations within a die, and can be used to perform process variation characterization for LE delays and interconnect delays including direct connection, double wire and hex wires. The difference in loop delays can also be estimated from equations using parameters extracted from primitives and compared with direct measurements. On a Xilinx Spartan-3e device, it was found that the error between the estimated and measured values was on average less than 10%.**

I. Introduction

As transistor feature sizes continue to be scaled down, increasing process variation becomes a great concern and severely affects delay, power consumption and reliability. Inevitable randomness in manufacturing causes considerable variation in effective channel length L_{eff}, as well as fluctuation in both threshold voltage V_{th} and oxide thickness T_{ox} [1]. Traditional approaches to handle process variations are to increase timing safety margins but doing this in a global manner is wasteful. The reconfigurability available in field programmable gate array (FPGA) devices offers the potential for designers to optimize circuit placement and routing at runtime [2][3], and this feature may be extremely beneficial to tolerate severe process variation and enhance timing yield of FPGA design in the future.

Unfortunately, quantitative measurements of process variation are difficult to extract from an FPGA. Although several previous works have indicated how variation is distributed within a die [4] [5], the granularity of those characterization methods is still not fine enough for practical variation aware design. Both used ring oscillator (RO) based circuits for variation measurement involving several stages of logic elements (LEs). One disadvantage of this approach is that an averaging of the random variations occur, which is undesirable if a single LE characterization is needed.

To address this problem, a fine grained process variation characterization method using a scheme involving differential RO measurements is proposed. It is able to perform process variation characterization for LE delays and interconnect delays including direct connection, double wire and hex wires. This is at a finer granularity than previous on-FPGA approaches.

The contributions of this work can be summarized as below:

- A scheme for fine-grained characterization of FPGA process variation using a RO-based differential measurement method.
- It is shown that the difference in delay of identical ROs at different locations can be accurately estimated from more primitive measurements and used in variation aware design.
- It is shown that the proposed process variation characterization can be implemented entirely with intrinsic FPGA hardware resources.

The remainder of the paper is structured as follows. Related work is surveyed in Section II. The primitives for the PVC scheme are described in Section III. The principle of the proposed methodology is presented in Section IV and the detailed implementation described in Section V. Experimental results and verification of the scheme are given in Section VI. In Section VII, conclusions with possible extensions of the proposed method are stated.

II. Background

Ring oscillators (ROs) have been widely used for delay measurement and diagnosis of process variation on both ASIC and FPGA platforms.

In the field of ASIC design, ROs are widely adopted for delay variation measurement [6][7][8][9][10]. Since ASICs are not normaly tuned in the post-silicon phase, the aforementioned variation characterization technique is usually used for diagnosis of early process development, monitoring mature process in manufacturing, enabling model-to-hardware correlation and tracking product performance [7]. A method for critical path delay measurement using ROs was proposed in [11]. The authors used the target path in a RO loop that also included a reconfigurable delay line with delay equal to one system clock period. The target path delay could then be calculated by subtracting the clock period from the RO loop delay.

On FPGA platforms, Xilinx patented a RO based method to measure delay of an arbitrary path [12][13][14]. Ruffoni et al

978-1-4244-8983-1/10/$26.00 ©2010 IEEE

proposed a method for path delay measurement which compared the delays of two ROs [15]. A reference RO is compared with a RO including the path under test (PUT). Li et al proposed a method using a RO array as a process variation monitor to control and improve yield in a Xilinx Virtex-II pro [4]. A similar technique was used on Altera FPGAs [5]. The latter work experimentally modelled the spatial correlation of process variation and predicted process variation for future technologies. In [16], Zick et al proposed a RO-based online sensing scheme to monitor different information for an FPGA-based processor including delay, leakage, dynamic power and temperature. Moreover, ROs can be used as an IR-drop monitor in processors [17], utilizing the relationship between RO frequency and supply voltage. In [18], Boemo et al utilized relationship between RO frequency and ambient temperature to detect thermal effects in FPGAs.

Apart from RO-based measurement, at-speed transition tests can also be used to measure delay and characterize process variation. Taking a combinational path with flipflops at two ends as the measurement target, transition failure rates can be observed while increasing clock frequency and the path delay deduced. This technique has been realized on FPGA platforms [19][20].

In CAD research, several works on the improvement of integrated circuit performance in the presence of process variation have been published. Lin et al proposed a quantitative timing yield model and process variation aware placement strategy for FPGAs [2]. Process variation aware routing for FPGAs was proposed by Sivaswamy et al [3]. Both methods achieved considerable timing performance improvement. Sedcole et al made a quantitative analysis of FPGA variation, which also showed that statistical static timing analysis could achieve a significant improvement in timing performance compared to the standard worst-case design technique [21]. Process variation information in [2] and [3] were modelled rather than measured.

As process variations become dominant, variation aware design (VAD) for FPGAs will become increasingly necessary. In the new design framework, the VAD tool would replace traditional CAD tool including placement and routing, and individual FPGAs must be characterized in terms of process variation. Figure 1 illustrates the envisaged high level design methodology.

To fulfil this need for process variation information in variation aware design, a practical fine-grained, on-chip variation characterization technique is required. A general way to observe process variations at the logic element (LE) level was described in [5]. Wong's work [20] can accurately measure path delay, but is limited to situations where the target path has flipflops at both ends, making delay within a single LE difficult. Furthermore, the resolution of the measurement depends on the step-size of the frequency sweep and a Xilinx FPGA was used to provide a variable frequency clock to the Altera FPGA's clock management module. Our characterization method complements these approaches.

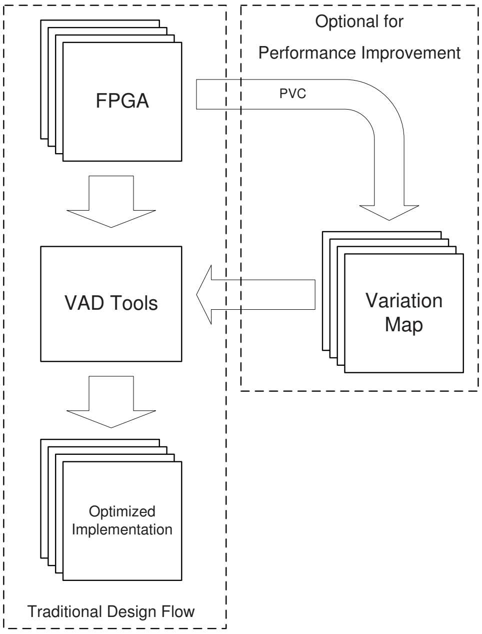

Fig. 1. Variation aware design (VAD) flow.

As the proposed variation characterization technique does not rely on external equipment, the PVC step in figure 1 can be done either by the vendor during testing or after release to end customers. Our proposed method can also aid in speed binning.

III. Characterization Primitives

Although the technique could be applied to any island-style FPGA, Xilinx Spartan-3e FPGAs were used in this work. Reconfigurable logic blocks (CLBs) are arranged in a regular array and connected by wire segments and switch matrices (SM).

Figure 2 illustrates the internals of a CLB. Each CLB is composed from four slices and a SM. Each slice consists of two LEs, each having one 4-input look-up table (LUT) and a flip-flop. As shown in figure 2, a SM is built from wires and programmable interconnection points (PIPs). PIPs can connect pins within a CLB, from CLB to a channel, and vice versa. The PIPs are not fully connected.

Besides connections using SMs, the FPGA interconnect fabric has wire segments of different lengths. There are four types of wire segments – direct, double, hex and long lines. Long lines are not addressed in this research.

Direct connections as shown in figure 3(a) route signals to neighboring blocks in the vertical, horizontal and diagonal directions. The double lines in figure 3(b) route signals to every first or second block away in four directions. Double line signals can be accessed either at the endpoint or at the midpoint and are organized in a staggered pattern. They can be only be driven from their endpoints. The hex lines in figure 3(c) route signals to every third or sixth block in four directions. Hex wire signals can be accessed either at endpoints or at the midpoint. Eight double and eight hex lines are driven by a single CLB. Each combinational output is

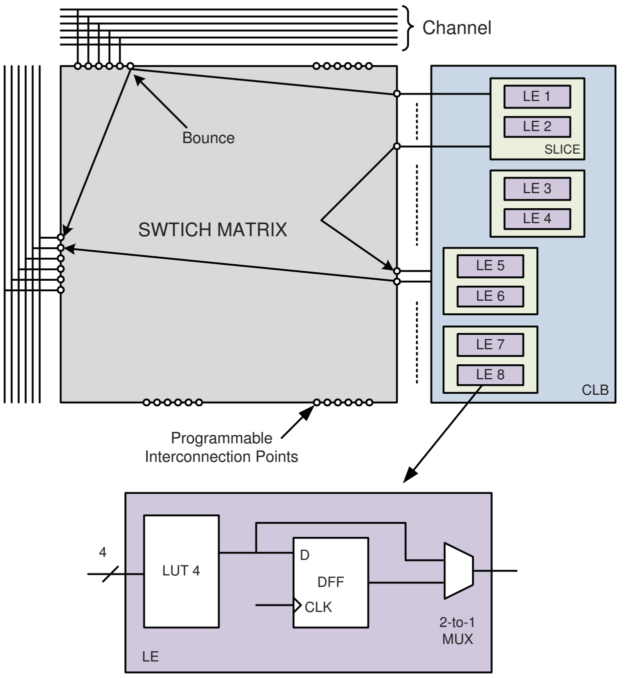

Fig. 2. Block diagram of FPGA "island".

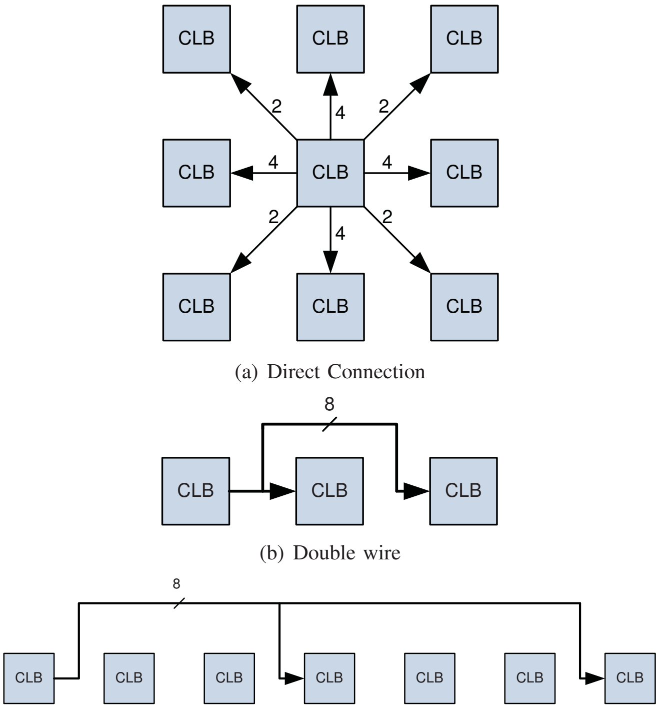

(a) Direct Connection

(b) Double wire

(c) Hex wire

Fig. 3. Direct connection, double and hex wires.

equipped with one double connection and one hex connection in each direction.

A combinational path on the FPGA can be composed from LEs, connections in SM and various wire segments. If the primitive delays can be accurately characterized, optimized variation aware circuit designs become possible.

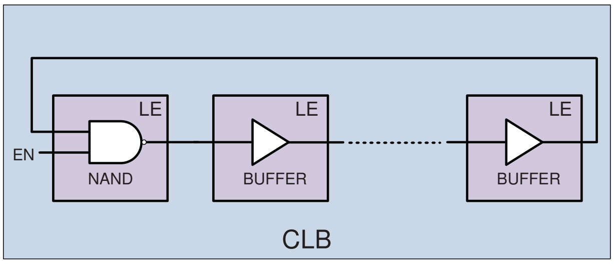

Fig. 4. A ring oscillator.

IV. Methodology

A RO is typically composed of an odd number of inverting stages and each stage can be implemented within a LE. All ROs are implemented with one 2-input NAND and buffer(s), using one of the NAND gate inputs as an enable signal. As the maximum toggle rate of a flipflop in our FPGA is 572 Mhz, the minimum loop delay should be less than 0.874 ns. Over a chosen time interval T, a counter is used to record how many cycles a RO runs. Representing the counter value as C, the RO loop delay D_{loop} can be calculated using equation 1.

$$D_{loop} = \frac{T}{2 \times C} \qquad (1)$$

The process of variation characterization is divided into two phases, namely LE characterization and interconnect characterization. The latter requires information of the former.

A. LE Characterization

LE delay measurement can be realized by implementing ROs in a single CLB, as shown in figure 4. We first create an 8-stage RO utilizing all LEs in a CLB. A 7-stage RO is then built, omitting one LE. In the example shown in figure 2, LE1 is omitted.

D_{loop_8} and D_{loop_7}, the loop delays of the 8 and 7-stage RO, are used to represent delay of intra-CLB connections for these two types of ROs. They are a sum of LE delays and interconnect delay, and are given in equation 2 and 3. Equation 4 gives the difference in loop delay ΔD_{loop} and is composed of two parts, the difference in LE delay ΔD_{LE} and the difference in interconnect delay ΔD_{int}.

$$\begin{aligned}
D_{loop_8} &= \sum_{i=1}^{8} D_{LE_i} + D_{int_8} && (2)\\
D_{loop_7} &= \sum_{i=2}^{8} D_{LE_i} + D_{int_7} && (3)\\
\Delta D_{loop} &= (\sum_{i=1}^{8} D_{LE_i} - \sum_{i=2}^{8} D_{LE_i}) \\
&\quad +(D_{int_8} - D_{int_7}) \\
&= D_{LE_1} + \Delta D_{int} && (4)
\end{aligned}$$

$$
\begin{aligned}
f_{int} &= \frac{\Delta D_{int}}{\Delta D_{loop}} && (5)\\
D_{LE_1} &= \Delta D_{loop} - \Delta D_{int} \\
&= \Delta D_{loop} - f_{int}\Delta D_{loop} \\
&= (1 - f_{int})\Delta D_{loop} && (6)
\end{aligned}
$$

f_{int} is defined as the fraction of ΔD_{int} in ΔD_{loop} (equation 5). Applying equation 5 to equation 4, the delay of LE1 is given in equation 6 and illustrated in figure 5.

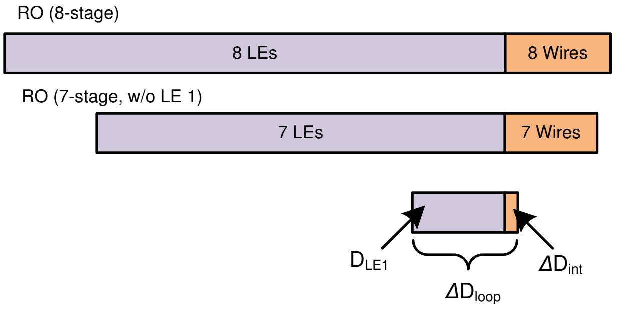

Fig. 5. Delay contribution of a RO.

TABLE I
BOUNCE-FREE INTRA-CLB DELAY.

	LE1	LE2	LE3	LE4	LE5	LE6	LE7	LE8
LE1	N/A	23	24	23	122	119	122	119
LE2	195	N/A	195	156	86	23	86	23
LE3	55	23	N/A	23	110	72	110	72
LE4	75	101	75	N/A	21	23	21	23
LE5	75	101	75	101	N/A	23	21	23
LE6	55	23	55	23	110	N/A	110	72
LE7	195	156	195	156	86	23	N/A	23
LE8	24	23	24	23	122	119	122	N/A

A connection exists between any two LEs within a CLB. However, some are directly connected, while others require a "bounce" as illustrated in figure 2. The delay of a connection with bounce is considerably larger than a direct one. To reduce interconnect delay, we try to only use direct connections. Table I summarizes the direct connection delays, obtained using Xilinx's timing analysis tool. The rows denote combinational inputs of a CLB, and the columns denote the corresponding outputs (refer to figure 2). For example, the underlined entry with value 23 gives the delay of a connection from LE1 to LE2 in picoseconds.

According to the datasheet, LE delay is nominally 760 ps and connection delay is considerably less. Table II gives the RO composition and interconnect delays estimated using the timing analysis tool. The LE connection sequences in the table ensure a minimum value for interconnect delay, D_{int}, and this is less than 5% of D_{loop} for the device studied, mitigating associated inaccuracies in variation estimation. For ease of expression, LEs are indexed from 1 to 8 according to figure 2. LE delay can be measured using the differential method described earlier.

TABLE II
8-STAGE AND 7-STAGE RO COMPOSITION AND ESTIMATED DELAYS.

Composition	Est. D_{loop} (ns)	Est. D_{int} (ns)	% of D_{int}
1,3,2,6,4,5,7,8	6.262	0.182	2.91%
2,6,4,5,7,8,3	5.478	0.158	2.88%
1,3,6,4,5,7,8	5.528	0.208	3.76%
1,2,6,4,5,7,8	5.478	0.158	2.88%
1,3,2,6,5,7,8	5.568	0.248	4.45%
1,3,2,6,4,7,8	5.481	0.161	2.94%
1,2,5,3,4,7,8	5.595	0.275	4.92%
1,3,2,6,4,5,8	5.481	0.161	2.94%
1,3,2,5,7,6,4	5.595	0.275	4.92%

Differences in LE delay can be derived after LE characterization. For example, ΔD_{LE_1}, the LE1 delay difference between CLBs *j* and *j'* is given by equation 7.

$$
\begin{aligned}
\Delta D_{LE_1} &= D_{LE_1(j)} - D_{LE_1(j')} \\
&= (1 - f_{int})(\Delta D_{loop(j)} - \Delta D_{loop(j')}) && (7)
\end{aligned}
$$

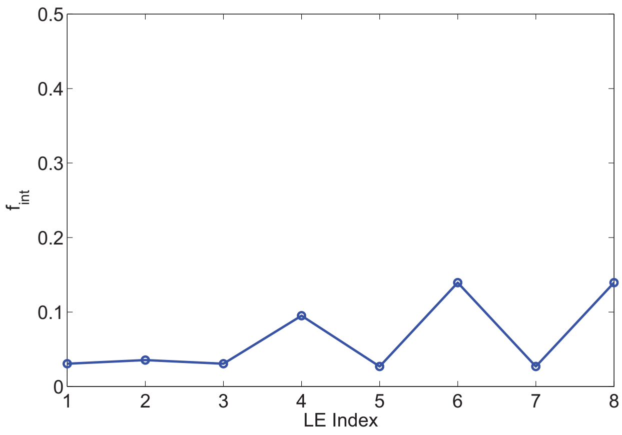

Fig. 6. f_{int} for each LE.

f_{int} can be estimated for each LE using data in table II together with equations 5 and 7. The values range from 0.027 to 0.139 as shown in figure 6.

B. Interconnect Characterization

Due to enhanced connectivity and higher logic capacity, interconnect circuits have become very complicated in modern FPGAs, making interconnect delay characterization difficult.

We create a calibration RO using two LEs and a pair of interconnects as shown in figure 7. The interconnects can be

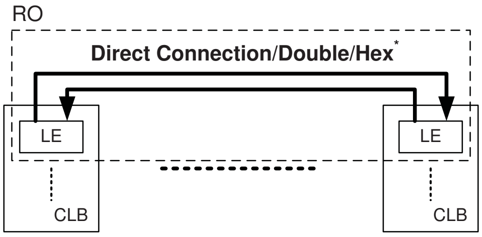

Fig. 7. Illustration of wire delay.

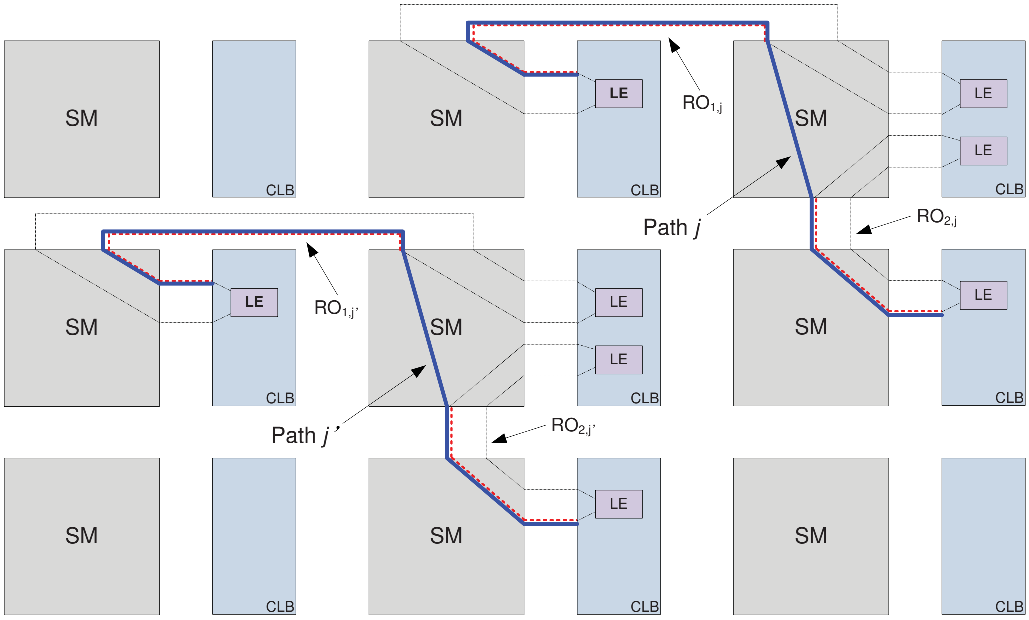

Fig. 8. Bold solid lines denote the target path. Dotted lines highlight the fraction of calibration RO contributed to the target path.

direct connections, double lines or hex lines. RO interconnect delay D_{int} is calculated by subtracting the LE delay from the loop delay as mentioned before. Unfortunately, a pair of hex lines cannot be created in this manner so a further differential method is applied to isolate them. For example, the interconnect pair could be composed of a mix of direct connection and hex lines. Once the delay of the direct connection is known, the hex line delay can be correspondingly derived.

To facilitate VAD tools, the delay difference between otherwise identical delay components rather than their absolute value is required. In figure 8, the two bold solid lines are target paths whose delays we wish to compare. Two types of calibration ROs are used for delay comparison of the target paths. Only the overlapped part in the calibration RO contributes to delay comparison, and this is illustrated by the dotted lines. As it is not always possible to isolate the delay of an interconnect segment, a "contribution factor" (denoted as F_C in equation 8) is introduced, where D_O and D_{int} are respectively the delays of the overlapped part and total interconnect of the calibration ROs.

$$F_C = \frac{D_O}{D_{int}} \tag{8}$$

In figure 8, the target paths (solid lines) are not fully covered by the calibration ROs. A "coverage rate", R_C, given in equation 9 is used to describe the proportion covered, where D_{O_i} denotes the delay of the overlapped part for RO i, and D_{path} is delay of target path.

$$R_C = \frac{\sum_{i=1}^{n} D_{O_i}}{D_{path}} \tag{9}$$

If the delays of two identical interconnect paths in different locations (path *j* and *j'*, as shown in figure 8) are compared, each path is covered by n calibration ROs (in figure 8, $n = 2$). Applying all equations above, the delay difference between two paths (ΔD_{path}) can be calculated as below.

$$\begin{aligned}\Delta D_{path} &= [D_{path}]_j - [D_{path}]_{j'} \\ &= \frac{1}{R_C}([\sum_{i=1}^{n} D_{O_i}]_j - [\sum_{i=1}^{n} D_{O_i}]_{j'}) \\ &= \frac{1}{R_C}\sum_{i=1}^{n}([D_{O_i}]_j - [D_{O_i}]_{j'}) \\ &= \frac{1}{R_C}\sum_{i=1}^{n} F_{C_i}([D_{int_i}]_j - [D_{int_i}]_{j'}) \end{aligned} \tag{10}$$

To calculate ΔD_{path}, it is necessary to know R_C and F_{C_i}. Unfortunately, the Xilinx timing tool only reports pin-to-pin delay (from combinational input to combinational output). Therefore, the delay of the overlapped part can not be explicitly specified and F_C and F_{C_i} can not be explicitly derived. However, we empirically estimate that $R_C = 1$ and $F_{C_i} = 0.5$.

Since the proposed method does not explicitly isolate the delay of the overlapped part (dotted line in figure 8) from the calibration RO, inaccuracies may arise. Taking RO$_{1,j}$ and RO$_{1,j'}$ in figure 8 as an example, if the overall interconnect delay of RO$_{1,j}$ is larger than that of RO$_{1,j'}$ ($[D_{int_1}]_j > [D_{int_1}]_{j'}$), applying a common "contribution factor" F_{C_1} to $[D_{int_1}]_j$ and $[D_{int_1}]_{j'}$, it is estimated $[D_{O_1}]_j$ is larger than $[D_{O_1}]_{j'}$. However, the delay of the overlapped part for RO$_{1,j}$ is actually smaller than that for RO$_{1,j'}$. Fortunately, spatial correlation effects usually mean that if delay of a segment of interconnect is fast, the neighboring ones tend to be fast as well. This property mitigates inaccuracies in delay estimation and errors of this type do not frequently occur.

V. Implementation

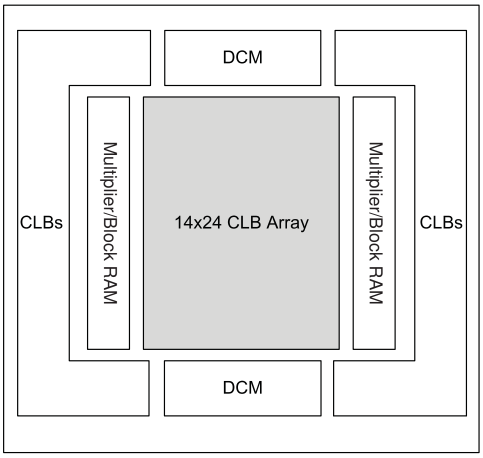

Fig. 9. FPGA architecture and characterization region.

Figure 9 shows a block diagram of the Xilinx Spartan-3e FPGA used in this work. Apart from CLBs, dedicated embedded blocks such as multipliers, block RAMs (BRAM) and digital clock managers (DCM) are present and can increase the delay of connection between neighboring CLBs compared with a homogeneous array. As a proof of concept, a 14×24 CLB array (totally 2688 LEs) in the center of die is characterized. This is shown as a shaded area in figure 9. Different types of ROs are built as hard macros using Xilinx FPGA Editor. Placement constraints are specified to control the region to be characterized. The auxiliary circuits are implemented using logic resources outside of the characterized region.

According to the method of LE characterization described in subsection IV-A, nine configurations are needed for a full characterization of LE delay (one for the 8-stage RO and eight for the 7-stage RO). For interconnect characterization, the work associated with switching configurations could be much larger. To completely characterize the interconnect primitives, at least 56 configurations need to be tested. This study is limited to full characterization of a single direct connection, double line and hex line. Others are partially characterized. Currently, a manual approach is used to test different configurations but it is believed a dynamic scheme would greatly speed up the characterization process. Moreover, enhanced architectural support in the FPGA could greatly improve efficiency.

It is well known that transistor delay is very sensitive to temperature and supply voltage [22]. As much as possible, supply voltage and temperature are held constant during measurement. In the future we may study ways to investigate how fluctuation patterns of supply voltage and temperature affects on-chip characterization and develop new ways to reduce their effect.

VI. Experimental Results

A. Scaling Factor

The RO loop delay is estimated before actual measurement. We found that the measured RO loop delay is always smaller than the value stated in timing analysis tool, as would be expected as it is a conservative value over a range of operating conditions and devices. We define "scaling factor" F_S in equation 11, where D_{spec} denotes delay specified by timing tool, and D_{real} denotes real delay by measurement. For five ROs with different numbers of stages, the scaling factor is 0.55 on average. Details of the comparison are summarized in table III.

$$F_S = \frac{D_{real}}{D_{spec}} \tag{11}$$

TABLE III
RO delay comparison.

RO Types	D_{spec} (ns)	D_{real} (ns)	Scaling Factor F_S
4 stages	3.162	1.767	0.559
5 stages	3.915	2.148	0.549
6 stages	4.697	2.574	0.548
7 stages	5.481	3.028	0.552
8 stages	6.262	3.424	0.546

B. Characterization Results

1) LE Characterization Results: Taking one CLB as an example, the delay of each LE in nanoseconds is listed in table IV. Systematic LE delay mismatch can be observed. LE1 to LE4 are all faster than LE5 to LE8, although they are conceptually identical. The differences may be caused by differences in the physical design. From the design tool we know that LE5 to LE8 can serve as distributed RAM, while LUT of LE1 to 4 does not have this functionality. To confirm correctness of our LE characterization, we build two 5-stage ROs, which are respectively composed of LE1 and LE8. By placing two ROs in different locations within the die, it was found that an RO using only LE1 is always faster than one using only LE8. The within-die spatial delay distribution is illustrated in figure 10.

TABLE IV
Statistical analysis of LE delay.

LE#	mean	% of 3-sigma	LE#	mean	% of 3-sigma
1	0.383	11.5%	2	0.403	14.7%
3	0.421	12.3%	4	0.407	11.3%
5	0.486	15.6%	6	0.454	14.2%
7	0.458	12.6%	8	0.465	14.1%

2) Interconnect Characterization Results: As mentioned in section IV-B, we characterize delay of a pair of connections (bold line in figure 7) by subtracting the LE delay from the total RO loop delay. The delay of a single connection is estimated as half the interconnect delay of a calibration RO D_{int}, as given in equation 12.

$$D_{wire} = \frac{D_{int}}{2} = \frac{D_{loop} - 2 \times D_{LE}}{2} \tag{12}$$

We characterize one type of direct connection, double line and hex line in the horizontal direction. Their mean values were respectively 233.1 ps, 271.3 ps and 358.6 ps. A 3-sigma

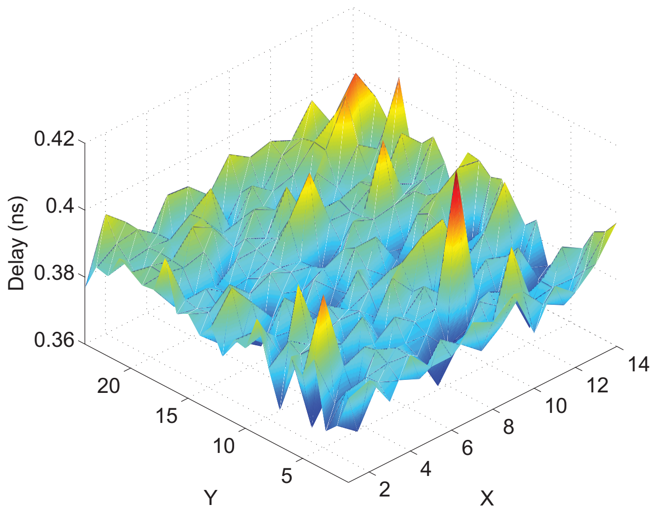

Fig. 10. Spatial distribution of LE1 delay.

variance of approximately 10% of the mean was observed for all three types of wire segments.

3) Die-to-Die Variation: We also compare two different FPGAs of the same model, respectively named chip #1 and chip #2. Die-to-die variation is shown in figure 11. Taking the LE1 delay over all CLBs as the comparison target, chip #1 is 7.6% faster than chip #2 on average. It can also be seen that the 3-sigma variance distribution of chip #1 is larger than chip #2, and that chip #1 is faster than chip #2 by this percentage for all comparisons. This technique is also well suited for FPGA speed binning. Table V summarizes statistical features of the two chips measured.

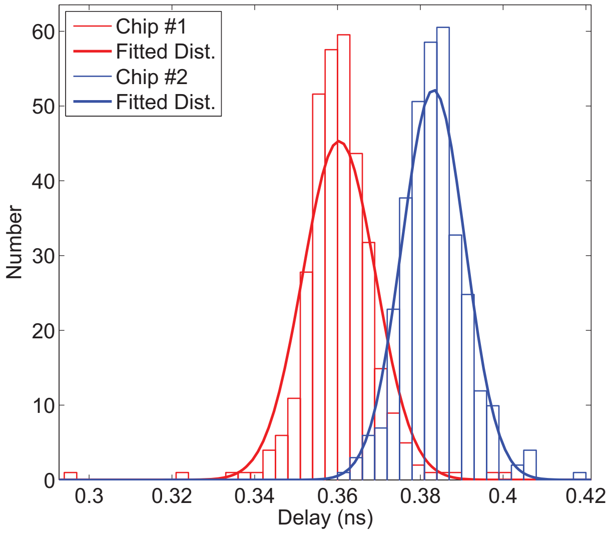

Fig. 11. Delay distribution of LE1 for two different chips.

TABLE V
STATISTICAL LE1 DELAY ACROSS CLBS.

	Chip #2	Chip #1
Mean Delay (ns)	0.383	0.360
% of 3-sigma	11.47%	14.14%

C. Verification

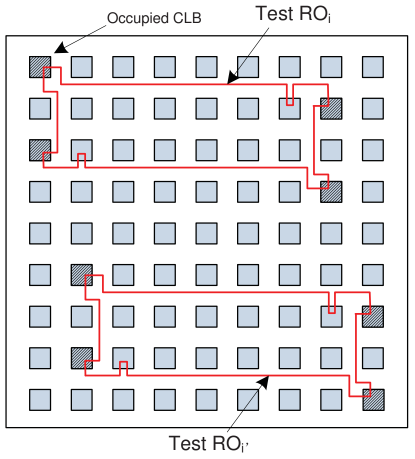

Fig. 12. Two ROs of identical design in different locations within die.

To validate PVC results, we place two ROs with identical physical design in different locations within the die as shown in figure 12. The difference between their loop delays, which is defined in equation 13, can be measured (denoted as $\Delta D_{loop,meas}$) and estimated by characterization results (denoted as $\Delta D_{loop,est}$) respectively.

$$\Delta D_{loop} \quad = \quad D_{loop,RO_i} - D_{loop,RO_{i'}} \tag{13}$$

By allowing a RO to clock a counter over a time interval, the number of rising edges can be recorded. The RO loop delay is calculated by equation 1, and $\Delta D_{loop,meas}$ can be obtained using equation 13 to characterize the fine-grained delay variation. The loop delay can be also calculated from existing information, as a sum of multiple LE delays and interconnect delays. By applying equation 13, the difference in loop delays $\Delta D_{loop,est}$ can be estimated.

The error of the delay estimation R_{err} is given by equation 14.

$$R_{err} \quad = \quad \frac{\Delta D_{loop,est} - \Delta D_{loop,meas}}{\Delta D_{loop,meas}} \tag{14}$$

We build five ROs which are composed of different delay primitives. Proportions of interconnect and LE delays are varied for each RO. Since two tested ROs are placed within the FPGA arbitrarily, their delay difference is not very significant (about 2∼3% of the total delay on average). The RO route goes through different delay primitives, which may have different variation patterns. From a statistical view, long paths could average the process variation effect if the route is chosen without optimization. Process variation aware placement and

TABLE VI
CHARACTERIZATION RESULT VERIFICATION.

Case #	% of D_{LE}	% of D_{int}	$D_{diff,meas}$ (ps)	Estimated $D_{diff,est}$ (ps)	R_{err} %
1	68.9%	31.1%	55.8	62.0	11.1%
2	55.7%	44.3%	76.0	79.1	4.08%
3	49.5%	50.5%	86.0	93.8	9.20%
4	43.0%	57.0%	41.5	38.9	6.27%
5	35.7%	64.3%	52.6	57.2	8.75%

routing [2] [3] could help, however, the problem of finding the fastest path given variation information is beyond the scope of this work. Table VI summarizes the comparison of RO loop delay between real measurement result and estimated value by characterization results. We achieve an error rate less than 10% on average, and the delay differentiation capability is safely within 10 ps. It could be observed that in most cases, the estimated difference is larger than measured value. This is because the "contribution factor" F_C over-estimates delay contribution from the overlapped part of calibration RO.

VII. CONCLUSION

Variation aware design potentially take leverage of FPGA's programmability to counter the effects of process variation and maintain performance. We presented a method to characterize FPGA process variation of logic elements and interconnects at fine granularity. Experiments show that our method can be used to effectively estimate path delays and results show that the delay mismatch estimation error of our variation characterization results is less than 10% on average.

Nevertheless, there are some limitations in this work. Due to architectural constraints, the delay of a single wire segment can not be explicitly characterized. Instead, we introduce "contribution factor" F_{C_i} and "coverage rate" R_C to handle such delays, which are derived empirically from observation in experiments. Improved methods can be used to estimate these these parameters and will be the target of future studies. Furthermore, since FPGA interconnect circuits have a much larger number of potential configurations, dynamically reconfiguration could be used to speed up the characterization process. We plan to study this problem in an FPGA which supports dynamic reconfiguration. Even using dynamic reconfiguration, a full interconnect characterization may not be possible and a study of architectural modifications to facilitate on-device characterization would be an interesting topic for future research.

REFERENCES

[1] M. Nourani and A. Radhakrishnan, "Testing on-die process variation in nanometer VLSI," *Design & Test of Computers, IEEE*, vol. 23, no. 6, pp. 438–451, June 2006.

[2] Y. Lin, M. Hutton, and L. He, "Placement and timing for FPGAs considering variations," in *Field Programmable Logic and Applications, 2006. FPL '06. International Conference on*, Aug. 2006, pp. 1–7.

[3] S. Sivaswamy and K. Bazargan, "Variation-aware routing for FPGAs," in *FPGA '07: Proceedings of the 2007 ACM/SIGDA 15th international symposium on Field programmable gate arrays*. New York, NY, USA: ACM, 2007, pp. 71–79.

[4] X.-Y. Li, F. Wang, T. La, and Z.-M. Ling, "FPGA as process monitor-an effective method to characterize poly gate CD variation and its impact on product performance and yield," *Semiconductor Manufacturing, IEEE Transactions on*, vol. 17, no. 3, pp. 267–272, Aug. 2004.

[5] P. Sedcole and P. Y. K. Cheung, "Within-die delay variability in 90nm FPGAs and beyond," in *Field Programmable Technology, 2006. FPT 2006. IEEE International Conference on*, Dec. 2006, pp. 97–104.

[6] M. Bhushan, A. Gattiker, M. Ketchen, and K. Das, "Ring oscillators for CMOS process tuning and variability control," *Semiconductor Manufacturing, IEEE Transactions on*, vol. 19, no. 1, pp. 10 – 18, feb. 2006.

[7] M. B. Ketchen and M. Bhushan, "Product-representative "at speed" test structures for CMOS characterization," *IBM Journal of Research and Development*, vol. 50, no. 4.5, pp. 451 –468, jul. 2006.

[8] H. Masuda, S. Ohkawa, A. Kurokawa, and M. Aoki, "Challenge: variability characterization and modeling for 65- to 90-nm processes," sep. 2005, pp. 593 – 599.

[9] S. Ohkawa, M. Aoki, and H. Masuda, "Analysis and characterization of device variations in an lsi chip using an integrated device matrix array," mar. 2003, pp. 3 – 75.

[10] B. Das, B. Amrutur, H. Jamadagni, N. Arvind, and V. Visvanathan, "Within-die gate delay variability measurement using reconfigurable ring oscillator," *Semiconductor Manufacturing, IEEE Transactions on*, vol. 22, no. 2, pp. 256 –267, may. 2009.

[11] X. Wang, M. Tehranipoor, and R. Datta, "Path-RO: a novel on-chip critical path delay measurement under process variations," in *ICCAD '08: Proceedings of the 2008 IEEE/ACM International Conference on Computer-Aided Design*. Piscataway, NJ, USA: IEEE Press, 2008, pp. 640–646.

[12] "Method for characterizing interconnect timing characteristics using reference ring oscillator circuit," *U.S. Patent*, no. 5790479, August 1998.

[13] "Method and system for measuring signal propagation delays using the duty cycle of a ring oscillator," *U.S. Patent*, no. 6069849, May 2000.

[14] "Method and system for measuring signal propagation delays using ring oscillators," *U.S. Patent*, no. 6219305, April 2001.

[15] M. Ruffoni and A. Bogliolo, "Direct measures of path delays on commercial FPGA chips," in *Signal Propagation on Interconnects, 6th IEEE Workshop on. Proceedings*, May 2002, pp. 157–159.

[16] K. M. Zick and J. P. Hayes, "On-line sensing for healthier FPGA systems," in *FPGA '10: Proceeding of the ACM/SIGDA international symposium on Field programmable gate arrays*. New York, NY, USA: ACM, 2010.

[17] Z. Abuhamdeh, B. Hannagan, A. Crouch, and J. Remmers, "A production IR-drop screen on a chip," *Design & Test of Computers, IEEE*, vol. 24, no. 3, pp. 216–224, May-June 2007.

[18] E. I. Boemo and S. López-Buedo, "Thermal monitoring on FPGAs using ring-oscillators," in *FPL '97: Proceedings of the 7th International Workshop on Field-Programmable Logic and Applications*. London, UK: Springer-Verlag, 1997, pp. 69–78.

[19] J. Li and J. Lach, "Negative-skewed shadow registers for at-speed delay variation characterization," oct. 2007, pp. 354 –359.

[20] J. S. J. Wong, P. Sedcole, and P. Y. K. Cheung, "Self-measurement of combinatorial circuit delays in FPGAs," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 2, no. 2, pp. 1–22, 2009.

[21] P. Sedcole and P. Y. K. Cheung, "Parametric yield in FPGAs due to within-die delay variations: a quantitative analysis," in *FPGA '07: Proceedings of the 2007 ACM/SIGDA 15th international symposium on Field programmable gate arrays*. ACM, 2007, pp. 178–187.

[22] G. Quenot, N. Paris, and B. Zavidovique, "A temperature and voltage measurement cell for VLSI circuits," in *Euro ASIC '91*, 27-31 1991, pp. 334 –338.