Impact of Interconnect Variations on the Clock Skew of a Gigahertz Microprocessor†

Ying Liu, Sani R. Nassif§, Lawrence T. Pileggi and Andrzej J. Strojwas
Department of Electrical and Computer Engineering, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213
§IBM Austin Research Lab, 11400 Burnet Road, Austin, TX 78758
yingl@ece.cmu.edu, nassif@austin.ibm.com, pileggi@ece.cmu.edu, ajs@ece.cmu.edu

Abstract

Due to the large die sizes and tight relative clock skew margins, the impact of interconnect manufacturing variations on the clock skew in today's gigahertz microprocessors can no longer be ignored. Unlike manufacturing variations in the devices, the impact of the interconnect manufacturing variations on IC timing performance cannot be captured by worst/best case corner point methods. Thus it is difficult to estimate the clock skew variability due to interconnect variations. In this paper we analyze the timing impact of several key statistically independent interconnect variations in a context-dependent manner by applying a previously reported interconnect variational order-reduction technique. The results show that the interconnect variations can cause up to 25% clock skew variability in a modern microprocessor design.

1. Introduction

An important feature of today's CMOS technologies is the existence of manufacturing variations (or process variations) in both devices and interconnect networks. Inevitably both types of variations will affect the clock skew of VLSI designs. Due to the large die sizes of microprocessor designs, usually long metal wires are necessary to distribute clock signals to the entire chip. As a consequence, interconnect is dominant in terms of clock signal delays. Thus, the interconnect process variations can have considerable impact on the clock skew variability for modern microprocessor designs, which in turn will affect the parametric yield of the product.

Traditionally the effect of device variations on circuit timing is captured by the worst/best case corner point methods. However, the effect of interconnect variations on timing cannot be captured by the corner point methods since it is context-dependent. When interconnect geometries are subject to variations, the change of delay depends not only on the specific CMOS technology (or design rules), but also on the physical design details of the specific interconnect network we are analyzing. We explain this in terms of the following simple example.

Consider the interconnect structure which consists of two parallel wire segments in the same layer, the cross section of which is shown in Fig. 1. Assume that the two wires are uniform, with the length of *l*, width of *w*, thickness of *t*, the distance to ground plane of *h*, and a spacing of *s*. Also assume that one segment is grounded. By applying a simple resistance model and an empirical capacitance model[6], the per unit length resistance and capacitance of the signal segment are:

$$\frac{R}{\rho} = \frac{1}{w \cdot t}$$

$$\frac{C}{\varepsilon} = 1.15\frac{w}{h} + 2.80\left(\frac{t}{h}\right)^{0.22} + \left(0.03\frac{w}{h} + 0.83\frac{t}{h} - 0.07\left(\frac{t}{h}\right)^{0.22}\right)\cdot\left(\frac{s}{h}\right)^{-1.34}$$

As a rough estimation, we use the Elmore delay to measure the interconnect delay:

$$t_N = \frac{RCl^2}{2} \quad (1)$$

As we can see from (1), the variation of t_N is a function of both resistance variation and capacitance variation. Whether the delay

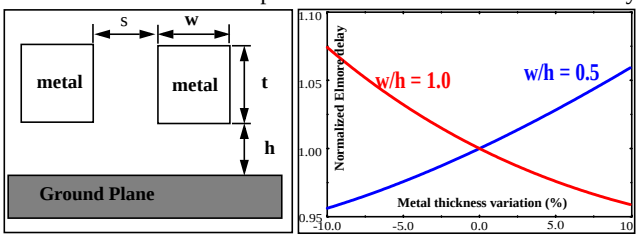

FIGURE 1: Cross section of a simple structure and delay variability

will increase or decrease depends on whether the interconnect is R-dominant or C-dominant. To give a quantitative estimation, we assume that $t = 0.75h$, $w = 2.0s$, we also assume that there is a ±10% variation in the metal thickness *t*. Next we consider the exact same structure in two different situations(contexts), one has the width of $w = 0.5h$, while the other of $w = 1.0h$. The changes of the normalized Elmore delays in these two scenarios are plotted in Fig. 1. From the graph, it is apparent that we have two completely different conclusions for the two "contexts", even though topologically they share the same structure. Because of this context-dependent nature, it is quite difficult to capture the impact of interconnect variations on the clock skew variability. Traditionally this issue was either ignored or incorrectly addressed.

In this paper, we analyze the impact of interconnect process variations on the clock skew of a gigahertz microprocessor design. We apply the variational order-reduction techniques reported in [3] and concentrate on the impact of a few key systematic variations. The results demonstrate that those systematic variations can cause up to 25% variability in the clock skew. Thus interconnect variations can be a potentially important factor in the performance of modern microprocessor designs, and therefore, must be dealt with in the design phase.

† This work was supported, in part, by the Semiconductor Research Corporation under Contract DC-068.067, MARCO/DARPA Gigascale Silicon Research Center and IBM Corporation.

DAC 2000, Los Angeles, California
(c) 2000 ACM 1-58113-188-7/00/0006..$5.00

2. Description of the clock net

The design is a prototype gigahertz microprocessor[2] which is approximately 17×17mm, and is manufactured in a 0.25μm CMOS technology with full copper interconnect. There are several clock distribution networks on the chip, but here we are only interested in the main clock net.

The main clock net consists of two levels. The top level is essentially a balanced H-tree across the entire die, as shown in Fig. 2, which distributes the clock signal from the center of the chip to different sector buffers. Minor modifications are made in its topology when necessary. For the second level, the whole die is partitioned into 4 columns in horizontal direction and 4 rows in vertical direction. There are 16 sectors across the die in total. In each sector, a sector buffer relays the signal further to the tips in the sector through a balanced sub-H tree. To facilitate the analysis, we assign each sector an index number, as indicated in Fig. 2. The sub H-tree is drawn in the figure only for Sector 1. Sub H-trees for the other sectors were not drawn, but appear similar to the one in Sector 1. In total there are approximately 1000 tips in the second-level clock tree.

It should be noted that there is a megacell between Sectors 3 and 7. Because the routing in the megacell is extremely dense, no over-the-cell routing is allowed. As a result, the top level of the H-tree has to detour in those two sectors. Some detours have been also taken in the sector sub H-trees in Sectors 3 and 7. Note also that in Sector 4 there is an extra leg in the top-level. Its purpose is to relay the signal to the driver of a delayed clock net.

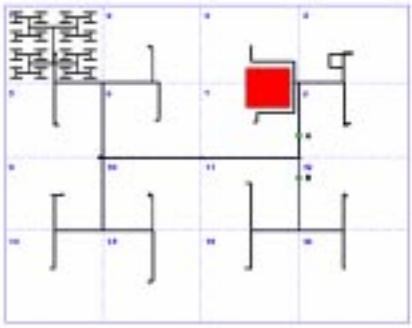

FIGURE 2: Top level of the clock net and sub-H tree in sector 1.

The entire main clock net, both the top level and the 16 sector H-trees, is routed in the top two metal layers, LM and MQ. For both top level and second level H-trees, the width of each segment is optimized by using a proprietary optimization tool. The capacitive loadings at each tip of the second level H-trees are provided based on the distribution of the latches. This load distribution will be used in the skew analysis.

3. Characterization of interconnect variations

Manufacturing variations are random in nature and their true causes are very complex. Generally they manifest themselves in both temporal and spatial manner. According to the geometrical scales of their occurrence, the variations are classified at four different levels[1]: *lot-to-lot*, *wafer-to-wafer* (within-lot), *die-to-die* (within-wafer) and *within-die*. Since our main interest is to determine the effects of variations on circuit timing performance, in this work we focus on the true geometrical variations. More specifically, we are interested in three geometrical variations for each metal layer, namely variations in: metal width, metal layer thickness and inter-layer dielectric (ILD) thickness. We focus our attention on the die level and classify the variations into three components: *global* component, *local* component and random component. At a specific location on the die, if one of the interconnect geometries (width, metal thickness or ILD thickness) can be expressed as:

$$w = w_0(1 + \delta) \tag{2}$$

where w_0 is the nominal design value and w is the true value with the variation, then the variation can be partitioned as:

$$\delta = \delta_{global} + \delta_{local} + \varepsilon \tag{3}$$

In (3), ε is the random residue, which is the component of variation which cannot be accounted for by either the global or local component.

It should be noted that one of the major causes of the local component is the layout and topology interaction. For example, it has been well-documented that the local ILD thickness is a function of the underlying metal layer density in the CMP process[7]. Additionally, it has been observed that the systematic variations are often of relatively low spatial frequency and smooth across the wafer. Quite often they exhibit symmetrical properties such as a “bull’s eye” (i.e., radial distribution) or a slant plane across the wafer. Fig. 3 shows the metal resistivity measurements from a wafer in this 0.25μm copper technology. The z axis is the normalized metal resistivity (which is a function of width and thickness) of certain metal layer, while the x and y axes are the die locations across the wafer.

Depending on their relative magnitudes, both the global and local components of the variations can have considerable impact on the circuit timing. In today’s state-of-the-art CMOS technologies, there are some methods to minimize the impact of the local components, especially those related to the design-process interaction. For example, the magnitude of ILD thickness variation due to CMP process can be reduced by inserting dummy metal fills[4]. Since we are mostly interested in the impact of some key statistically independent process variations, which have the most significant systematic effects on the timing performance, in this paper we will only analyze the impact of the global components of interconnect variations. However, the impact of local components can also be analyzed, although extra computational cost will be involved.

Due to its relatively low spatial frequency and smoothness, simple models can be used to describe the global component of the variation. Throughout this paper we will apply a simple model, as depicted in Fig. 3. Assume that the global component has a radial distribution. Then across the die, the general expression of the global component is given by:

$$\delta_{global} = \delta_0 + \delta_x \cdot x + \delta_y \cdot y \tag{4}$$

Note that (*x, y*) in the above equation are the normalized coordinates across the die:

$$-1 \le x \le 1, \quad -1 \le y \le 1 \tag{5}$$

For metal layers LM and MQ, three interconnect geometrical variations are considered for each metal layer: the width, thickness and thickness of ILD underneath. The variation of each geometry is modeled as a slant plane as in (4). The total number of variables is $2 \times 3 \times 3 = 18$.

From the process specifications, for both layers LM and MQ, the across-wafer tolerances for each of the width, thickness and ILD is ±15%. Considering the size of the chip and the possible steep ascending/descending in the distribution of the systematic

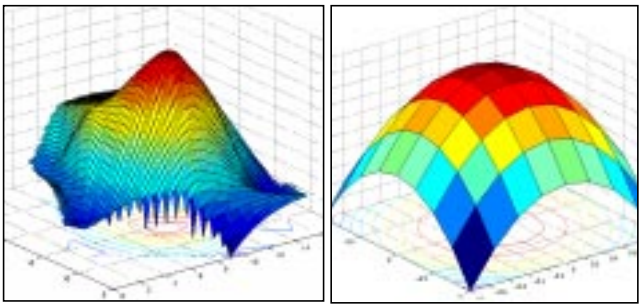

FIGURE 3: A typical radial distribution (LEFT) and decompose it into slant planes for each die (RIGHT).

variations, we choose the across-die variation of each of the geometries to be ±10%. For example, for LM width variation, we assume that:

$$|\delta_0| \le 0.10, \quad |\delta_x| \le 0.10, \quad |\delta_y| \le 0.10, \quad |\delta_0 + \delta_x + \delta_y| \le 0.10 \qquad (6)$$

We make similar assumptions on the metal thickness and ILD thickness of LM layer, as well as the three geometries of layer MQ.

4. Variational order-reduction of interconnect

In order to analyze the clock skew, the entire clock distribution network is modeled as lumped RC circuits. We then apply the variational interconnect order-reduction technique reported in [3] to construct reduced-order models for the interconnect networks. Unlike regular reduced-order models, the variational reduce-order models constructed here have direct inclusion of the 18 variations. When the magnitudes of the variations are specified, a regular reduced-order model can be constructed at almost no computational cost. Along with the nonlinear models of the drivers, it is then used in HSPICE to calculate accurate timing information. A complete description of the variational order-reduction technique can be found in [3]. We omit the details here due to the limited space.

5. Results

5.1. Nominal skew

The nominal clock skew, which is the clock skew without any manufacturing variations, is computed to be 187ps. The surface plot of the delays across the whole chip is shown in Fig. 4. Note that besides the megacell, there is also a small gap around Sector 10 in which there are no tips. This is due to the congestion in that region.

As we can observe from the plot in Fig. 4, the smallest delays occur mostly in Sector 13, while the biggest delays occur in Sector 7. Therefore, the clock skew is dictated by the delays at the tips in Sectors 7 and 13. Intuitively, the reason for the fast nodes in Sector 13 can be attributed to the relatively light loading in this sector. However, the reason for the slowest nodes in Sector 7 requires some explanation.

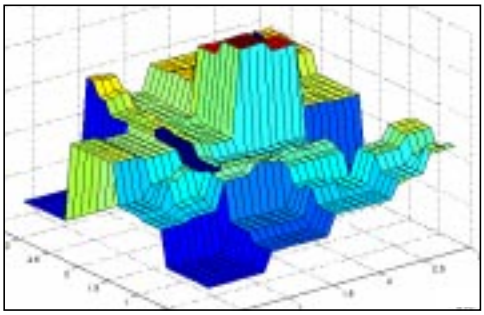

FIGURE 4: Surface plot of the delays of all leave nodes.

As we can see in Fig. 2, because of the existence of the megacell, a portion of the top-level H-tree in Sector 3 is replaced by a Y-shaped fork. The same modification is made for a portion in Sector 7. The extra wiring of the Y-shaped fork and the extra leg to the delayed clock net in Sector 4 cause the down-stream capacitive load seen from node A to be bigger than the down-stream capacitive load seen from the mirror node B. It can be calculated that Elmore delay at node A is bigger than at node B. As a consequence, the top level H-tree branch in sectors 3, 4, 7 and 8 will be slower than the branch in sectors 11, 12, 15 and 16. The situation is further deteriorated by the unbalanced sub H-tree in Sector 7 due to the presence of the megacell.

5.2. Skews under uniform worst-case variations

First we assume that the variation of each of the 6 geometry parameters is modeled as a horizontal plane with the maximal allowable magnitude, which can be viewed as an extension of the traditional corner point methods. For example, for LM width variation, we assume that δ_0 is either +10% or -10%, while δ_x and δ_y are identically zero. Since there are only limited number of possible combinations ($2^6 = 64$), the search is straight-forward. Some results are listed in Table 1. Note that we include the variations with the largest and

Variations						Skews	
LM:w	LM:h	LM:t	MQ:w	MQ:h	MQ:t	ps	chg
+10%	-10%	+10%	-10%	-10%	+10%	186.47	-0%
-10%	-10%	-10%	+10%	-10%	+10%	215.54	+15%
-10%	-10%	-10%	-10%	-10%	+10%	211.03	+12%
+10%	+10%	+10%	-10%	+10%	-10%	169.01	-9%
-10%	-10%	-10%	-10%	-10%	-10%	209.96	+12%
-10%	+10%	+10%	-10%	+10%	-10%	174.51	-6%

TABLE 1: Clock skew under worst-case variations.

smallest skew in the table. The last column indicates the percentage change compare to the nominal skew of 187ps. As we can see in the table, the biggest skew is 15% larger than the nominal case, while the smallest is 9% smaller than the nominal case. Observe that the variation corresponding to the biggest skew lies in the opposite corner to the one with the smallest skew. The results from all 64 runs have a median and mean around 190ps, and a standard deviation of 11ps.

5.3. Skews under tilted variations in NE direction

In the next analysis we assume that the variation is a tilted plane. We also assume that the horizontal component of each geometrical variation, δ_0, is identically zero. We further assume that the x-direction component, δ_x, is equal to the y-direction component, δ_y, for each geometrical variations, and their values are either +5% or -5%. By doing so, we are assuming that each variation is in a tilted plane in the (0,0)-(1,1) or NE direction with maximal allowable slope. Some of the results are tabulated in Table 2. Note that in this situation, the worst skew is 25% bigger than the nominal case. Considering the tight design margin, such a big change may have significant impact on the microprocessor performance.

For an intuitive explanation for this phenomenon, recall that the biggest delays occur in Sector 7, while the smallest delays occur in Sector 13. Therefore, any combinations of the variations which enlarge the difference between the delays in these two sectors will have a big effect the skew. Since Sectors 13 and 7 are in

the diagonal direction on the chip, it is not surprising that we observe the biggest change in the skew for NE direction variations. In all of 64 runs, the clock skew has median of 191.46ps, mean of 193.64ps, and standard deviation of 18ps.

Variations						Skews	
LM:w x y	LM:h x y	LM:t x y	MQ:w x y	MQ:h x y	MQ:t x y	ps	chg
-5%	+5%	+5%	-5%	+5%	+5%	168.64	-9%
-5%	-5%	-5%	+5%	+5%	-5%	221.27	+18%
-5%	-5%	-5%	+5%	-5%	-5%	235.21	+25%
+5%	+5%	+5%	-5%	-5%	+5%	172.69	-7%
+5%	+5%	+5%	-5%	+5%	-5%	159.17	-14%
+5%	-5%	-5%	+5%	+5%	+5%	205.06	+9%

TABLE 2: Selected simulation results under variations in NE direction.

5.4. Skews under tilted variations in NW direction

Following the same argument as in the previous subsection, if the variations are in a plane of NW, or (0,1)-(1,0) direction, then the change in the clock skew should be smaller. We performed another set of analysis based on this argument. This time we still keep the horizontal components zero. We once again pick up the x-directional component to be either +5% or -5%, but we keep the y-directional component the opposite of the x-directional: $\delta_x = -\delta_y$. Some of the results are tabulated in Table 3. We can see that the change in the skew is much smaller compared to the results in NE direction. Among the total of 64 runs, the median is 190.8ps, the mean is 191.3ps, while the standard deviation is only 2.3ps. Therefore, if all the variations are indeed in the NW direction, the clock skew variability will be very much under control.

Variations						Skews	
LM:w x -y	LM:h x -y	LM:t x -y	MQ:w x -y	MQ:h x -y	MQ:t x -y	ps	chg
-5%	-5%	-5%	-5%	-5%	+5%	194.58	+4%
-5%	+5%	-5%	-5%	+5%	+5%	190.61	+1%
+5%	-5%	-5%	-5%	+5%	+5%	191.33	+2%
+5%	-5%	-5%	+5%	-5%	-5%	196.27	+4%
+5%	-5%	-5%	+5%	-5%	+5%	195.75	+4%
+5%	+5%	+5%	-5%	+5%	-5%	188.25	+0%

TABLE 3: Selected simulation results under variations in NW direction.

5.5. Skews under random variations

In actual designs, besides the above three extreme scenarios, we will see more cases in which all three components are present. To demonstrate effect of those realistic variations on the clock skew variation, we performed another analysis by conducting a random search. In the search, we constrain the overall magnitude of each geometric variation to be 10%: $|\delta_0 + \delta_x + \delta_y| = 0.10$. Multiple runs are performed. We tabulate some typical results in Table 4. In the table, for each column of the variations, the first number is the horizontal component of the variation, δ_0. The second number is the x-directional component, δ_x, while the third number is the y-directional component, δ_y. As we can see in the table, although the magnitudes of the variations are not eye-catching, we still observe considerable numbers of cases in which the change of clock skew is over 10%. The worst case is 44ps, or a +23% increase over the nominal skew, which is not too far away from the +25% increase we observed in the extreme cases.

Variations						Skews	
LM:w	LM:h	LM:t	MQ:w	MQ:h	MQ:t	ps	chg
-3.1%	+3.9%	+3.7%	-0.7%	+4.9%	+3.2%	170.21	-8%
-3.4%	+3.6%	+4.1%	-4.1%	+0.9%	+3.9%		
-3.5%	+2.5%	+2.2%	-5.2%	+4.3%	+2.9%		
-2.2%	-6.0%	-1.1%	-1.3%	+3.0%	+2.8%	206.23	+10%
-6.7%	-0.7%	-3.3%	-8.7%	+3.4%	+4.8%		
-1.0%	-3.3%	-5.6%	-0.0%	+3.7%	+2.5%		
+4.4%	+4.0%	+3.0%	-2.3%	+3.2%	-3.6%	162.47	-13%
+5.4%	+0.2%	+3.2%	-3.1%	+2.0%	-3.6%		
+0.2%	+5.8%	+3.8%	-4.6%	+4.8%	-2.7%		
-3.0%	+2.6%	-3.5%	-5.7%	-2.9%	+5.5%	204.08	+9%
-5.0%	+0.1%	-0.4%	-3.6%	-2.5%	+0.4%		
-2.0%	+7.4%	-6.1%	-0.7%	-4.6%	+4.2%		
-0.7%	-0.5%	-0.4%	+2.3%	-0.6%	-4.6%	231.42	+23%
-0.6%	-7.0%	-3.4%	+4.2%	-4.8%	-2.7%		
-8.7%	-2.5%	-6.1%	+3.5%	-4.6%	-2.7%		

TABLE 4: Selected simulation results under random variations.

6. Conclusions

Clock net design for modern microprocessors is still a challenging problem and the parametric yield can be very strongly affected by the sensitivity of the design to the fabrication process fluctuations. In this paper, we have analyzed the impact of interconnect manufacturing variations on the clock skew in a gigahertz microprocessor. From the result, we observe that the clock skew can vary up to 25% due to the systematic interconnect variations. This is a very significant par of the clock skew budget and would indicate a need for a clock net redesign to increase its immunity to such fluctuations.

References

[1]Boning, D.S. and J.E. Chung, "Statistical metrology - measurement and modeling of variation for advanced process development and design rule generation", *Proc. 1998 Int. Conf. on Characterization. and Metrology for ULSI Technology*, March 1998.

[2]Hofstee, P. *el al*, "A 1 GHz single-issue 64b PowerPC processor", *Technical digest of IEEE Intl. Solid-State Circuit Conference(ISSCC)*, 2000.

[3]Liu, Y., L.T. Pileggi and A.J. Strojwas, "Model order-reduction of RC(L) interconnect including variational analysis", *Proc. IEEE/ACM Design Automation Conference*, 1999.

[4]Mehrotra, V., S. Nassif, D. Boning and J. Chung, "Modeling the effects of manufacturing variation on high-speed microprocessor interconnect performance", *Proc. IEDM*, 1998.

[5]Pullela, S., N. Menezes and L.T. Pillage, "Reliable non-zero skew clock trees using wire width optimization", *Proc. IEEE/ACM Design Automation Conference*, 1993.

[6]Sakurai, T. and K. Tamaru, "Simple formulas for two- and three-dimensional capacitances", *IEEE Trans. Electron Devices*, vol. ED-30, February 1983.

[7]Stine, B.E. *et al*, "The physical and electrical effects of metal fill patterning practices for oxide chemical mechanical polishing processes", *IEEE Tran. Electron Devices*, vol. 45, No. 3, March 1998.