

Circuit Design and Modeling Techniques for Enhancing the Clock-Data Compensation Effect Under Resonant Supply Noise

Dong Jiao, Jie Gu, *Member, IEEE*, and Chris H. Kim, *Member, IEEE*

Abstract—Recent publications have shown that clock jitter can improve timing margin through the compensation effect between the clock cycle and the datapath delay under the influence of resonant supply noise. This paper presents a comprehensive study of this beneficial clock-data compensation effect including an analysis of its dependency on various design parameters and a new phase-shifted clock buffer design that can enhance the effect. Measurement result from a 1.2 V, 65 nm test chip shows an 8–27% increase in the maximum operating frequency while saving 85% of the clock buffer area compared to prior art. An accurate timing model is derived to estimate the beneficial jitter effect.

Index Terms—Clock-data compensation, clock distribution, resonant noise, supply noise.

I. INTRODUCTION

POWER supply noise is considered to be one of the major performance limiting factors in scaled technologies [1]. Traditionally, passive decoupling techniques based on on-chip decoupling capacitors or damping resistors were used to reduce the supply network impedance [2]–[4]. Several circuit techniques have been proposed to alleviate the supply noise impact on circuit performance by minimizing clock jitter [5]–[11]. Supply noise in the 100 MHz to 300 MHz resonant frequency band has been shown as the dominant noise component in a typical high performance microprocessor [12]–[14]. Fig. 1 shows the measured supply network impedance of an Intel Nehalem microprocessor which exhibits a large impedance peak at around 150 MHz [15]. Resonant supply noise is caused by the *LC* tank formed between the package/bonding inductance and the die capacitance and thus affects the entire chip. This so called “first droop noise” can be excited by a sudden current spike caused by a clock edge or a processor wakeup operation [16], [17]. Due to its large magnitude, resonant noise constitutes the worst-case supply noise scenario which has triggered




Fig. 1. Measured supply network impedance response of Intel's Nehalem microprocessor [15].

a flurry of research activities in the circuit design community [3], [14], [18].

Recent papers have revealed an intriguing timing compensation effect between the clock cycle and the datapath delay in the presence of resonant supply noise [16], [17], [19]. This phenomenon, which we will refer to as the “beneficial jitter” effect, is illustrated in Fig. 2. A simple pipeline circuit consisting of a phase-locked loop (PLL), a clock path and a datapath is shown. In traditional analysis, the clock period is assumed to be constant while only the datapath delay is assumed to change under the influence of supply noise. Fig. 2(b) illustrates example waveforms based on the traditional analysis showing several sampling failures during the event of a supply voltage undershoot. In reality, however, the clock path delay is also modulated by the supply noise and therefore stretches the clock period during supply downswings. As a result, the clock path delay and datapath delay compensate for each other which alleviates the timing margin. Fig. 2(c) shows example waveforms for this scenario in which the output is always sampled correctly by the stretched clock cycle.

Adaptive clocking schemes utilizing this principle have been employed in Intel Nehalem processors [15], [20]. There, a PLL-based clock generator is designed to track the supply noise so that the clock period stretching effect is maximized. An alternative way to enhance the beneficial jitter effect is shifting the phase of the supply noise seen by the clock path [16], [17], [19], for example by using an *RC* filtered supply voltage for the entire clock path. A similar approach has been used in Intel Pentium4 processors where the supply noise of the clock buffer is reduced

Manuscript received November 22, 2009; revised May 09, 2010; accepted July 15, 2010. Date of current version September 24, 2010. This paper was approved by Associate Editor Stephan Rusu. This work was supported in part by the Semiconductor Research Corporation under award 2008-HJ-1804.

D. Jiao and C. H. Kim are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: dong@umn.edu).

J. Gu is with the MaxLinear Incorporated, Carlsbad, CA 92011 USA.

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/JSSC.2010.2063931

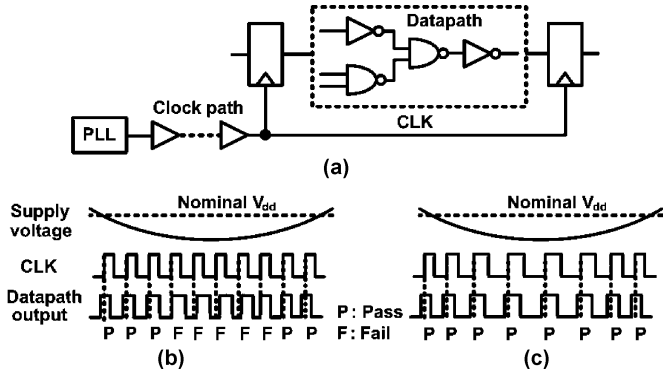


Fig. 2. Illustration of the beneficial jitter effect. (a); (b) constant-period clock; (c) real clock.

by using a local *RC* filter [21]. However, these existing designs require a large area overhead due to the small resistance and large capacitance requirements.

In this work, we propose a novel phase-shifted clock buffer design that can improve the maximum operating frequency by 8–27% while saving 85% of the clock buffer area compared to previous phase-shifted designs for a typical resonant frequency range between 100 MHz and 300 MHz. The proposed design can be used in conjunction with adaptive clocking schemes for further improvement in chip performance. This paper also presents an accurate timing model for estimating the beneficial jitter effect with a 4–10 \times lower prediction error compared to previous models.

II. BENEFICIAL JITTER EFFECT

A. Definition of Timing Slack

We first define the term timing slack in the context of a standard register-based pipeline shown in Fig. 3. For this circuit to operate without errors, a certain amount of timing margin must be ensured so that the final outputs of the logic block are evaluated before the next clock edge. Therefore, we define “slack” as the clock period T_{CLK} minus the actual datapath delay T_{DATA} . Obviously, the slack has to be positive for the circuit to be error free. That is:

$$slack = T_{CLK} - T_{DATA} > 0 \quad (1)$$

Here, the setup time requirement is ignored but it can be easily incorporated later by adding a timing offset.

B. Beneficial Jitter Effect

Conventional analysis only focuses on the increase in datapath delay in the presence of supply noise as shown in Fig. 2(b). However, in reality, the clock path also sees a noisy supply which causes the clock period to gradually stretch during supply downswings (or compression during supply upswings). This clock period modulation effect results in an extra timing margin that compensates for the slowdown in the datapath as shown in Fig. 2(c). Fig. 4 illustrates how the compensation




Fig. 3. Definition of timing slack in a standard pipeline circuit.




Fig. 4. Setup time margin analysis under resonant supply noise.

effect improves the setup time margin. In the presence of supply noise, the maximum datapath delay occurs when the supply voltage is at its lowest point, denoted as “A”. The corresponding clock edge (i.e., the first edge) which triggers the longest datapath delay signal is launched from the clock source at a certain point in time before “A” as it has to traverse through the clock path. The second edge, which will eventually sample the longest delay signal, is launched one clock period after the 1st edge. It experiences a lower average supply voltage due to the supply downswing, and thus takes a longer time to propagate through the clock path. This makes the clock period longer, compensating for the increased datapath delay.

C. Impact of Beneficial Jitter Effect on Hold Time Margin

Now we discuss how hold time margin is affected by the resonant supply noise. Fig. 5 illustrates the setup and hold time margin requirements for a simple register-based (or latch-based) pipeline. Contrary to the setup time margin scenario, the hold time margin is worst when the datapath delay is minimum, denoted as point “B” in Fig. 6. The corresponding clock edge is triggered when the supply voltage is rising. Here, we only need to consider a single clock edge since hold time violations occur due to clock skew for the same clock edge. As the rising supply voltage compresses the clock period, the clock skew becomes smaller, leading to a minor improvement in the hold time margin as depicted in Fig. 6. This improvement may not be noticeable when considering other timing uncertainties as will be shown in later sections. Note that the analysis on setup time and hold time

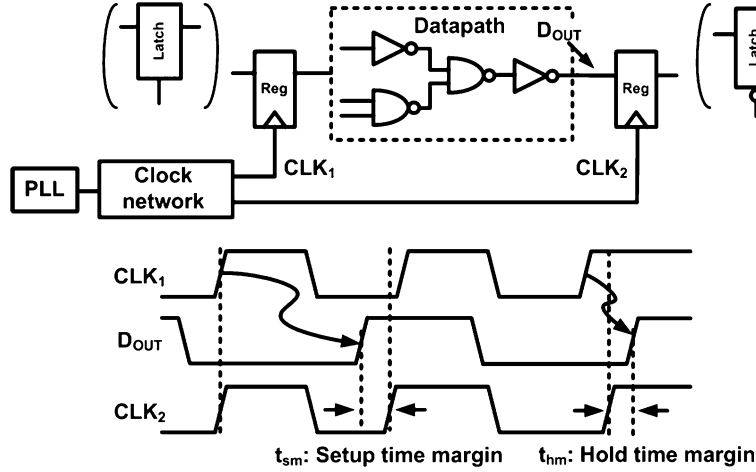


Fig. 5. Illustration of setup and hold time margin in a register-based (or latch-based) pipeline.




Fig. 6. Hold time margin analysis under resonant supply noise.




Fig. 7. Model for clock path delay [17].

margins is applicable to both register-based and latch-based designs.

III. MODELING OF BENEFICIAL JITTER EFFECT

In this section, we propose a numerical model to predict the beneficial jitter effect for different operating scenarios. The proposed model significantly improves the accuracy compared with prior techniques proposed in [17] and [19]. Similar to previous publications, we start our derivation by treating a digital signal in a clock path or a datapath as a travelling wave propagating through a fixed length medium. The velocity of this wave is proportional to the instantaneous supply voltage. Fig. 7 shows the relationship between the traveling distance and traveling time of a propagating signal [17]. In this model, the velocity of the wave is expressed as

$$v(t) = SA_0 + sa \cos(\omega_m t - \theta) \quad (2)$$

where S and s are the sensitivities of $v(t)$ with respect to the supply voltage, A_0 and a are the DC and AC amplitude of the supply voltage, ω_m is the noise frequency and θ is the noise phase. If the velocity is integrated over the total travelling time t_e (the actual delay), the result will be the total distance Y_0 , which is proportional to the nominal delay D_0 . That is,

$$Y_0 = D_0 SA_0 = \int_0^{t_e} [SA_0 + sa \cos(\omega_m t - \theta)] dt. \quad (3)$$

By following the same steps in [17], an analytical expression of t_e can be obtained as

$$t_e = D_0 - \frac{2sa}{SA_0\omega_m} \sin\left(\frac{\omega_m t_e}{2}\right) \cos\left(\frac{\omega_m t_e}{2} - \theta\right). \quad (4)$$

Finally, the worst-case slack is derived from (3) as shown in (5) at the bottom of the page. This equation, which was proposed

$$|slack_{wc}| = \sqrt{4(s_{clk}/S_{clk})(a/A_0)^2 \sin^2(\pi f_m/f_0)(s_{clk}/S_{clk} - s_{data}/S_{data}) + (as_{data}/A_0 S_{data})^2} \quad (5)$$

in [17], predicts a timing improvement only when the delay sensitivity of the clock path is less than that of the datapath. However, detailed HSPICE simulation reveals (see Fig. 12) a significant improvement in the worst-case setup time margin (7% of the clock period) even when the clock path and the datapath have the same delay sensitivities. This fact suggests that an improved model is needed to aid designers in better understanding the impact of the beneficial jitter effect. In this work, we derive a more accurate timing model by using numerical methods to solve (3) rather than resorting to unjustified approximations. Results presented later in this section will show that accurately solving (3) is the most important step in deriving a timing model that is reliable under various operating conditions.

We will use the standard register-based pipeline circuit shown in Fig. 5 to describe the flow for deriving the setup timing margin. Suppose the first clock edge E_1 launched from the clock generation block at time $t = 0$ takes t_{cp1} to arrive at the register. The input data of the first register starts to propagate through the datapath at time $t = t_{cp1}$ and takes t_{d1} to reach the input of the second register. Now assume the second clock edge E_2 is launched at time $t = t_{clk1}$ and takes t_{cp2} to propagate through the clock path. Then, the setup time margin can be calculated as

$$t_{m,s} = t_{clk1} + t_{cp2} - t_{cp1} - t_{d1}. \quad (6)$$

This procedure is repeated numerically by sweeping θ from 0 to 2π and the minimum value becomes the worst-case setup time margin. For modeling the hold time margin, t_{clk1} is set to 0 and two different nominal clock path delays (D_0) are used when solving for t_{cp1} and t_{cp2} , indicating that we are considering the clock skew for the same clock edge propagating through two different clock paths.

IV. INTRINSIC BENEFICIAL JITTER EFFECT

In this section, we verify the beneficial jitter effect in an industrial 1.2 V, 65 nm process and analyze its dependency on several design parameters. The test circuit is similar to the one shown in Fig. 3 comprising a 1.9 GHz clock source, an 18-stage inverter chain datapath and an 11-stage clock buffer chain with a nominal delay of 1.0 ns. The delay sensitivities of the clock path and the datapath with respect to supply noise (i.e., s_{clk} and s_{data}) were both set to be 2. Here, we define delay sensitivity as the percentage increase in the path delay normalized to the percentage decrease in the supply voltage at a 10% supply noise condition. That is, a delay sensitivity of 2 means that the delay of a certain path increases by 20% for a 10% decrease in the supply voltage.

A. Intrinsic Beneficial Jitter Effect

Timing slacks for different clock launching times are shown in Fig. 8 for a 200 MHz resonant supply noise. The x axis shows the time when a clock edge leaves the clock source and the y axis shows the corresponding timing slack. The dark line represents the timing slack based on the conventional analysis which only considers the change in the datapath delay while the gray line considers the change in the clock period as well. An 11 ps (or 2.1% of the clock cycle) improvement in the worst-case slack due to the beneficial jitter effect is observed.

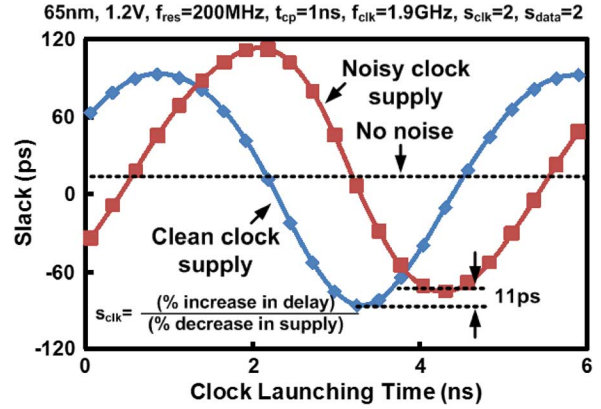


Fig. 8. Slack versus clock launching time under resonant supply noise.




Fig. 9. Dependency of worst-case slack on clock path delay.

B. Factors Affecting the Intrinsic Beneficial Jitter Effect

Clock Path Delay: Fig. 9 shows the dependency of the worst-case slack on the clock path delay simulated by changing the number of clock buffer stages. For extremely long or short clock path delays, the slack considering the beneficial jitter effect (i.e., noisy clock supply) approaches the conventional analysis case (i.e., clean clock supply). This is because a very short clock path makes the clock period modulation effect weaker and conversely, a very long clock path makes each clock edge see a similar average supply voltage.

Delay Sensitivity to Supply Noise: Fig. 10 shows the simulated worst-case slack when the datapath delay sensitivity is fixed at 2 and the clock path delay sensitivity is varied from 0 to 2.4 through the adjustment of the interconnect load, the number of clock buffer stages, and the supply noise amplitude seen by the clock path. The optimal timing compensation effect occurs when the clock path delay sensitivity is around 1.2. A clock path delay sensitivity lower than the optimal point makes the clock period less sensitive to the supply noise making the beneficial jitter effect weaker. On the other hand, a higher sensitivity eventually leads to a worse timing slack due to the excessively compressed clock periods during supply upswings.

Supply Noise Frequency: The worst-case slack for supply noise frequencies from 50 MHz to 1.6 GHz are shown in Fig. 11. At extremely low frequencies, the worst-case slack converges to

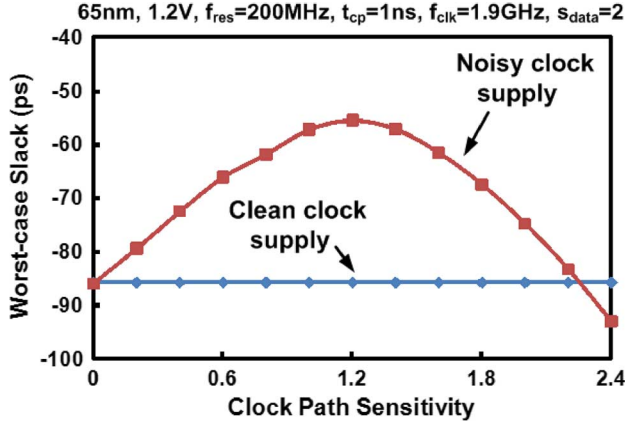


Fig. 10. Dependency of worst-case slack on clock path delay sensitivity.




Fig. 12. Dependency of setup time margin on clock path delay.




Fig. 11. Dependency of worst-case slack on supply noise frequency.




Fig. 13. Dependency of hold time margin on clock path delay.

the clean clock case since two consecutive clock edges see almost the same supply voltage. When the resonant frequency is high, the noisy clock supply case again converges to the clean supply case. This is because of the negligible difference in the supply voltages seen by two consecutive clock edges due to the averaging effect.

C. Modeling of Intrinsic Beneficial Jitter Effect

The methodology described in Section III for modeling the beneficial jitter effect was verified with HSPICE. The clock frequency and the maximum clock skew were assumed to be 1.9 GHz and 20 ps, respectively [22]. A resonant noise with a frequency of 200 MHz and an amplitude of $10\% \cdot V_{dd}$ was used for the simulations.

In the first test, setup and hold time margins were examined for different clock path delays. The results in Fig. 12 show a 45 ps change in the setup time margin and the detailed behavior is precisely captured by the proposed model. When compared with previous models, the maximum estimation error is improved from 26 ps to only 3 ps. Moreover, our proposed model also closely matches the simulation results for hold time margin as shown in Fig. 13. The maximum error is less than 1 ps for all

TABLE I
MAXIMUM MODELING ERROR FOR DIFFERENT CLOCK PATH DELAYS
($f_{clk} = 1.9$ GHz, $f_{res} = 200$ MHz, $s_{clk} = 2$, $s_{data} = 2$)

	Register-based		Latch-based	
	Setup	Hold	Setup	Hold
[17]	41ps	N/A	37ps	N/A
[19]	26ps	N/A	32ps	N/A
This work	3ps	1ps	7ps	1ps

clock path delays used in the simulations. A latch-based pipeline circuit was also simulated and the results are summarized in Table I.

We also tested the accuracy of the model for different supply noise frequencies. As shown in Fig. 14, the setup time margin is improved due to the beneficial jitter effect for a typical resonant frequency range of 100 MHz to 300 MHz. Similar to the previous test, both setup and hold time margins were simulated for register-based and latch-based pipeline circuits and the results are summarized in Table II. A significant improvement in the modeling accuracy is achieved.




Fig. 14. Dependency of setup time margin on supply noise frequency.

TABLE II
MAXIMUM MODELING ERROR FOR DIFFERENT NOISE FREQUENCIES
($f_{clk} = 1.9$ GHz, $t_{cp} = 1$ ns, $s_{clk} = 2$, $s_{data} = 2$)

	Register-based		Latch-based	
	Setup	Hold	Setup	Hold
[17]	111ps	N/A	105ps	N/A
[19]	92ps	N/A	96ps	N/A
This work	10ps	1ps	10ps	1ps




Fig. 15. Concept of the phase-shifted clock buffer design.

V. PHASE-SHIFTED CLOCK DISTRIBUTION

A. Enhancing the Beneficial Jitter Effect Using a Phase-Shifted Clock Distribution

The beneficial jitter effect in its intrinsic form provides modest timing margin relief for pipeline circuits. This is because the point when the clock period is stretched out the most (i.e., point “A” in Fig. 15) does not coincide with the point when the delay is the longest (i.e., point “B” in Fig. 15). It is important to note that the former situation occurs when the supply voltage has a negative slope while the later occurs when the instantaneous supply voltage is the lowest. In order to maximize the timing compensation effect, the phase of the supply noise seen by the clock path should be shifted such that points A and B are aligned.

Fig. 16 (left) shows the schematic of a conventional buffer and various phase-shifted clock buffers for enhancing the beneficial effect [16], [17], [19]. The previous *RC* filtered buffer contains

TABLE III
POWER CONSUMPTION OF DIFFERENT CLOCK BUFFER DESIGNS
($f_{clk} = 1.9$ GHz)

	Conv.	RC Filtered (prior art)	Stacked (this work)
Clean V_{dd}	5.013mW	4.868mW	4.922mW
Noisy V_{dd}	5.116mW	5.493mW	5.024mW

a pMOS pull-up device and an nMOS capacitor to generate a phase-shifted supply. The main drawback of this design is the large area. The resistance of the *RC* filter must be very small to minimize the IR drop (e.g., 50 mV or less) which in turn requires a large capacitance to obtain the desired supply phase shift. As shown in Fig. 16 (right), the layout area of the *RC* filtered buffer is about $10\times$ larger than that of a conventional buffer.

In this work, we propose a phase-shifted clock buffer using stacked devices to significantly reduce the buffer area while achieving a similar timing improvement. Fig. 16 shows the schematic and layout of the new circuit where header and footer devices controlled by separate *RC* filters are used instead of an explicit *RC* filter for generating a phase shifted supply. MOSFETs operating in the linear mode are used for implementing the resistors, enabling a much smaller layout area. The beneficial jitter effect can be further enhanced by using high V_t header/footer devices to make the buffer delay more sensitive to the phase-shifted supply noise. Hence, the proposed stacked buffer design was evaluated for both low V_t (LVT) and high V_t (HVT) header and footer devices. Since the actual switching current no longer flows through the resistor in the new design, small devices with large resistances can be safely used for the *RC* filter which in turn reduces the capacitor area for achieving the desired phase shift. As shown in Fig. 16 (right), the layout area of the proposed buffer is only 10% of the previous *RC* filtered buffer area. Even after considering the fact that the proposed stacked buffer has to be 50% larger than the conventional buffer for the same drive current, an 85% saving in buffer layout area can be achieved.

Power consumption is another major consideration for clock network designs. Table III compares the power consumption of a representative 9-stage clock path using the three different clock buffers. Simulation results show that both phase shifted designs consume slightly less power than the conventional buffer in case of no supply noise (i.e., clean V_{dd}). This is because the header/footer devices reduce the effective supply voltage seen by the buffer which reduces the CV^2 and short circuit power dissipation. Applying a 120 MHz resonant noise to the supply voltage (i.e., noisy V_{dd} case, the noise amplitude is 10% of the nominal supply voltage) led to a 12.8% increase in power consumption for the *RC* filtered buffer due to the power wasted for charging and discharging the large capacitor. In contrast, the proposed stacked buffer design shows only a 2.1% power increase owing to the greatly reduced capacitor size.

B. Modeling of Phase-Shifted Clock Distribution

Our proposed model can be applied to the phase-shifted clock distribution design by introducing a parameter φ which indicates the amount of supply noise phase shift. More specifically,

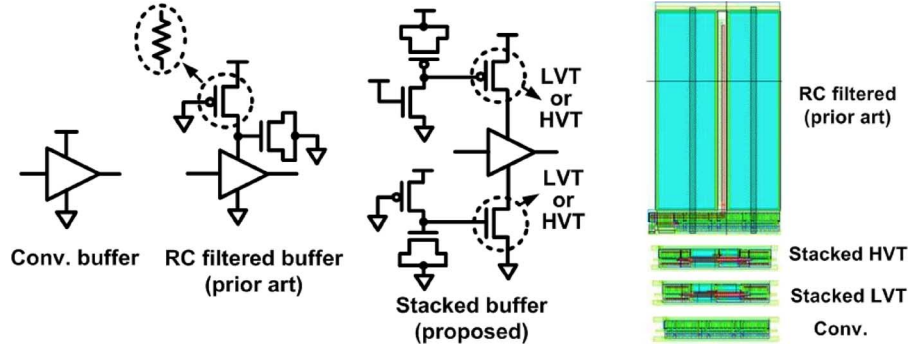


Fig. 16. (left) Schematic of a conventional buffer, an RC filtered buffer, and the proposed stacked high V_t and low V_t buffers. (right) Layout of the different clock buffers.

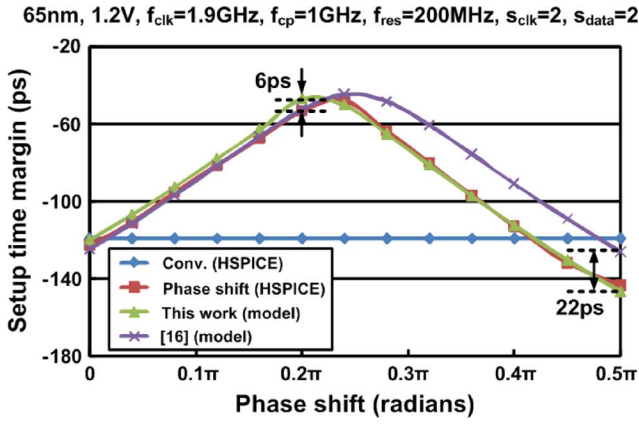


Fig. 17. Dependency of setup time margin on phase shift.

when solving for t_{cp1} and t_{cp2} in (6), we use the following expression for the propagating velocity:

$$v(t) = SA_0 + sa \cos \varphi \cos(\omega_m t - \theta - \varphi). \quad (7)$$

HSPICE simulations were performed for the phase-shifted clock distribution to evaluate the accuracy of the proposed model. The test circuit is similar to the one shown in Fig. 5 with RC filtered buffers used in the clock network. The value of R is chosen to be as large as possible while satisfying the IR drop requirement of less than 50 mV. Fig. 17 shows the setup time margin for different phase shift values. An optimal phase shift value makes the maximum clock period point coincide with the maximum datapath delay point. Simulation results and the estimated values using different models are given in Fig. 17, from which we can see that our proposed model reduces the maximum estimation error from 22 ps to 6 ps. The hold time margin was also simulated for a phase shift value of 0.2π which gives the best setup time margin. The maximum modeling error for this configuration was only 4 ps.

VI. TEST CHIP DESIGN AND MEASUREMENT RESULTS

A. Test Chip Organization

A 65 nm test chip was designed to verify the performance of the proposed phase-shifted clock buffers. Fig. 18 shows the block diagram of the proposed test chip which contains two VCOs, a clock path block, a core logic block, two 13-bit

counters, a noise injection block, a supply noise sensor, and a read-out block. Two starved ring oscillator based VCOs are used to generate the clock signal and the supply noise. By adjusting the external bias voltage V_{BIAS} , the VCO frequency can be raised up to 3.4 GHz. Five clock paths are implemented with different clock buffers: the conventional buffer, the RC filtered buffer, the stacked LVT buffer, the stacked HVT buffer and a “no buffer” design in which the output of the clock VCO is directly connected to the local registers. Each path contains nine buffer stages and long interconnects giving a clock path delay of 1.0 ns. One clock path is selected at a time to test each clock buffer design separately. The datapath circuit consists of two standard D-flip-flops and a ten-stage FO4 inverter chain in between to represent a critical path with a nominal delay of 0.6 ns. Input to the datapath is toggled between 1 and 0 in each cycle. Additional control logic increments the “data counter” only when the sampled output and the corresponding input are identical (during input ‘1’ cycles only). A “reference counter” increments every other cycle, and is used for counting the total number of sampled outputs. By scanning out the number stored in the data counter when the reference counter overflows, the percentage of correct samples can be conveniently measured. The noise injection block has 32 nMOS devices that can be clocked by the noise VCO. By adjusting the noise VCO frequency and activating different number of noise injection devices, the desired noise current can be injected into the supply network. A supply sensor is also designed for on-chip noise measurements. This circuit receives the noisy supply and ground signals as differential inputs, and the output indicates the supply noise frequency and amplitude [18]. The read-out block consists of a 10-bit parallel-to-serial shift register and additional control logic. In COUNT mode, the shift register captures the upper 10 bits of the data counter whenever the reference counter overflows. In READ mode, an external clock is provided to scan out the stored data serially. Fig. 19 shows the read-out waveforms including a mode selection signal, an external clock, and a read-out scan value. The read-out value we record is the average of 512 scan values to eliminate transient noise effects.

Note that a VCO-controlled noise injection block generates supply noise at a specific frequency (plus harmonics) making it easier to characterize the various clock buffers at a given noise frequency. As explained in the introduction section,

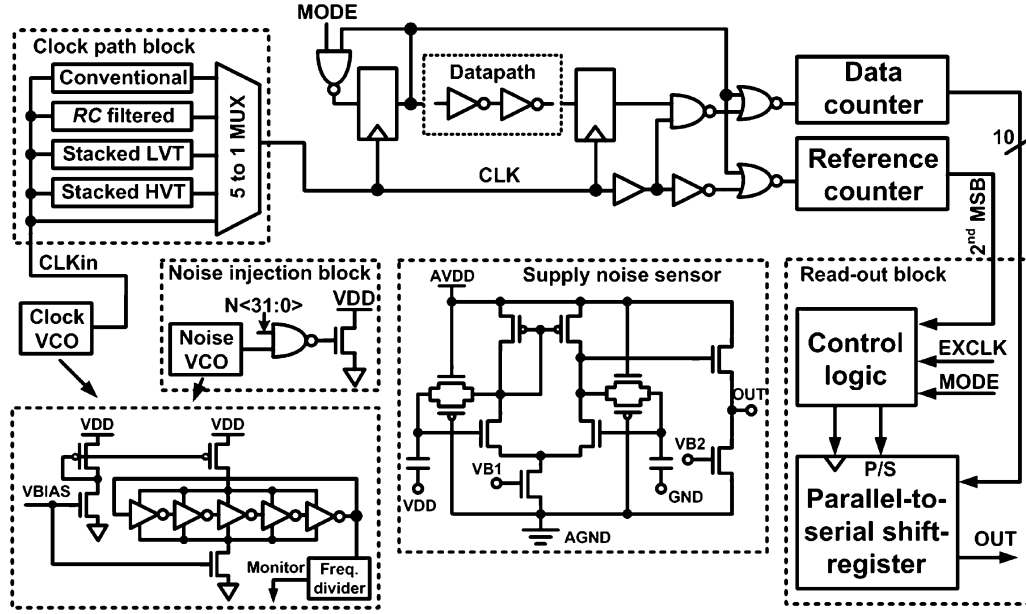


Fig. 18. High level block diagram of the 65 nm test chip.

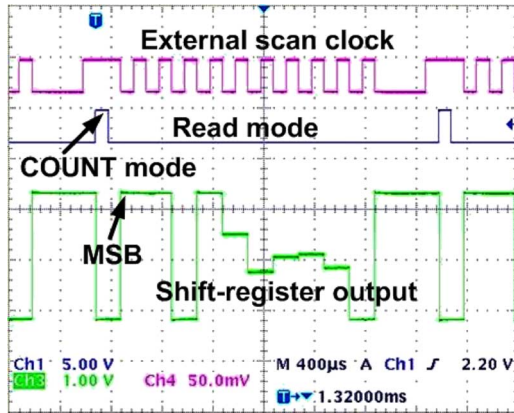


Fig. 19. Example read-out waveforms from the 65 nm test chip.




Fig. 21. Measured bit error rate for different clock buffer designs.




Fig. 20. Chip microphotograph and floor plan.

supply noise at the resonant frequency has been shown to be the dominant component in high performance microprocessors so the global supply noise generated by a VCO-based noise injection block is a simple yet effective way of generating a representative supply noise. Of course, one can consider using more elaborate digital blocks for generating global and local supply noises but the drawback here is that it may be difficult to know the exact supply noise waveform used for the chip testing.

B. Test Chip Measurement Results

The test chip was fabricated in a 1.2 V, 65 nm low power (LP) process and the die photo is shown in Fig. 20. In the first test, eight noise injection devices were turned on and the noise VCO bias was adjusted to generate a 118 MHz noise which corresponds to the resonant frequency of the fabricated test chip. Fig. 21 shows the percentage of correct samples measured from the different clock paths. F_{\max} or the maximum operating frequency is defined as the frequency at which the percentage of correct samples starts to drop. F_{\max} of the conventional buffer design reduced from 1.64 GHz to 1.2 GHz when the supply noise injection circuit was activated. F_{\max} of the RC filtered buffer, the stacked LVT and HVT buffers were 1.33 GHz, 1.31 GHz, and 1.34 GHz, respectively, which translate into roughly a 10% performance improvement compared with a conventional buffer design.

Fig. 22 shows the measured F_{\max} for the different clock buffer designs when increasing the number of noise injection devices. The supply noise frequency is maintained at 118 MHz.




Fig. 22. Measured F_{\max} for different number of noise injection devices.

As expected, the maximum frequency decreases linearly with more number of noise injection devices turned on. The proposed stacked buffer designs improve the F_{\max} by 8–15% when more than eight noise injection devices are turned on. This is similar to what the RC filtered buffer design achieves under the same condition.

The normalized F_{\max} of the different designs are shown in Fig. 23 for a noise frequency range between 10 MHz and 1.2 GHz. The number of noise injection devices is carefully adjusted so that F_{\max} of the conventional buffer design is fixed at 1.2 GHz. The figure shows that F_{\max} of the phase-shifted clock buffer designs is improved by 8–27% for a typical resonant frequency range of 100 MHz to 300 MHz. For noise frequencies higher than 400 MHz or lower than 50 MHz, F_{\max} of the phase-shifted clock buffer designs and the conventional design are similar. This is because the clock cycle modulation effect is very weak in both extreme frequency cases as explained in Section IV-B: when the noise frequency is high, the strong averaging effect makes consecutive clock edges see almost the same average supply voltages; when the noise frequency is low, consecutive clock edges again see almost the same supply voltages since it fluctuates very slowly. At some high frequencies, the phase-shifted buffer designs exhibit some performance degradation but this does not affect the overall performance because the worst-case noise scenario always happens in the resonant band, rather than at higher frequencies [15].

VII. DESIGN CONSIDERATIONS FOR THE PHASE-SHIFTED CLKOK BUFFER DESIGN

A. Comparison With the Adaptive Clocking Scheme

An alternative way of enhancing the beneficial jitter effect is to modulate the clock period at the clock source (e.g., PLL) so that the clock period stretching effect is maximized by the time the clock signal arrives at the flip-flops. Adaptive clocking schemes based on this principle have been recently deployed in Intel Nehalem processors [15]. In this scheme, the clock frequency of the PLL output is carefully designed to track the supply voltage variation with a phase difference as shown in Fig. 24. The proposed phase-shifted clock buffer design can be used in conjunction with existing adaptive clocking schemes to further improve chip performance. The effectiveness of using




Fig. 23. Measured F_{\max} normalized to the conventional buffer case for different noise frequencies.




Fig. 24. The PLL output frequency is modulated by the supply noise in adaptive clocking schemes.

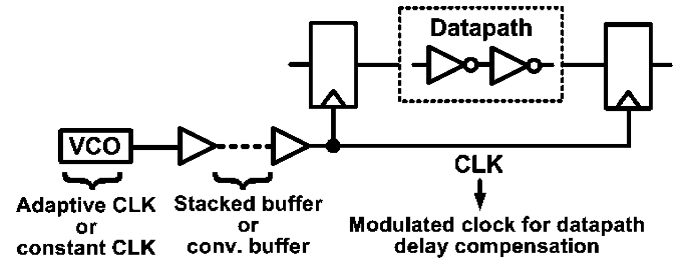


Fig. 25. Clock cycle modulation schemes.

both techniques in tandem for improving chip performance was verified with the test circuit shown in Fig. 25. The VCO output frequency was designed to follow the supply noise with a certain phase shift and a noisy power supply was applied to all blocks. The noise amplitude was set to be 10% of the nominal supply voltage. The simulated timing slack is shown in Fig. 26 for a noise frequency range from 10 MHz to 1.2 GHz. It is shown that the adaptive clocking scheme alone achieves a 17–39 ps worst-case slack improvement for a typical resonant frequency range between 100 MHz and 300 MHz. The phase-shifted buffer scheme provides an additional 30–62 ps improvement in timing slack.

The setup and hold time margins of the adaptive clocking scheme can be mathematically derived through the following steps. Assume that the supply voltage is expressed as

$$V_{dd}(t) = V_{dd0} + v_{dd} \cos(\omega_m t) \quad (8)$$




Fig. 26. Simulated worst-case slack for different clock cycle modulation schemes.




Fig. 27. Setup time margin versus design parameters of clock cycle modulation schemes.

where V_{dd0} and v_{dd} are the DC and AC amplitudes and ω_m is the supply noise frequency. We can expect the clock frequency f_{clk} of this PLL to vary at the same frequency, i.e., f_{clk} can be written as

$$f_{clk}(t) = f_{clk0} + f_{ac} \cos(\omega_m t - \varphi). \quad (9)$$

Here, f_{clk0} and f_{ac} are the DC and AC amplitude and φ denotes the phase shift between the supply noise and the frequency variation. We apply our proposed model to the adaptive clocking scheme by varying t_{clk1} in (6) depending on the time when the first clock edge is triggered, emulating the behavior of the adaptive clock frequency. The detail expression of t_{clk1} is determined by (9). To corroborate the model, we ran simulations using the circuit given in Fig. 2 with a conventional clock path and a supply-tracking PLL. φ in (9) was swept from $-\pi$ to π and f_{ac} was swept from $0.12f_{clk0}$ to $0.32f_{clk0}$. Simulation results in Fig. 27 show that the optimal setup time margin is achieved when φ is 0 and f_{ac} is $0.2f_{clk0}$. The estimation error of the timing model is only 6 ps.




Fig. 28. Partially phase-shifted clock distribution design.




Fig. 29. Slack improvement using a partially phase-shifted clock distribution design.

B. Partially Phase-Shifted Clock Distribution Design

Since the phase-shifted clock buffers are larger than (or have lower drive current than) conventional buffers, a more economical approach would be to limit their use to global clock buffer stages. We refer to this implementation as the “partially phase-shifted design” which is illustrated in Fig. 28. Simulation results of the worst-case slack are shown in Fig. 29 for different numbers of global clock buffer stages using the stacked LVT buffers. Since the number of buffers at each clock hierarchy increases exponentially in an H-tree type topology, the area overhead can be significantly reduced by using conventional buffers in the final stages of the clock network. As shown in Fig. 29, using phase-shifted clock buffers in the first 9 out of 11 stages in the clock network can provide a 52 ps improvement in the worst-case slack (about 71% of the maximal possible improvement) while reducing the clock buffer area overhead by 75%.

C. Impact of Delay Sensitivity, Process, Temperature and Resonant Frequency Variations

Most of the analysis in the previous sections assumes that the clock path and datapath have the same delay sensitivities. In reality, the delay sensitivity may vary depending on the amount of interconnect. For example, a clock path may have a lower sensitivity because of its long interconnect, and a datapath may also have a low sensitivity if it is wire dominated, like in data buses. To verify the performance of the phase-shifted clock distribution technique for different delay sensitivities, we present simulation results of the worst-case slack where the delay sensitivity of the datapath is fixed at 2 while the delay sensitivity of the clock path is swept from 1.6 to 2.4. The figure clearly shows that the worst-case slack is improved using the proposed clock buffer for the entire delay sensitivity range. Also shown are the




Fig. 30. Impact of random process variation on the worst-case slack at 25 °C and 110 °C. Monte Carlo simulations were performed using the following parameters: $V_{t,N} : \sigma/\mu = 3.6\%$, $V_{t,P} : \sigma/\mu = 1.6\%$, $t_{ox,N} : \sigma/\mu = 0.6\%$, $t_{ox,P} : \sigma/\mu = 0.6\%$.

average and 3σ values of the worst-case slack from Monte Carlo simulations with random local t_{ox} and V_t variations. Despite the slight degradation in the timing slack, the proposed stacked clock buffer design provides a consistent timing improvement in the presence of random process variation at 25 °C and 110 °C.


VIII. CONCLUSION

In this paper, we have presented a comprehensive study of the clock data compensation effect occurring in the event of resonant supply noise. A novel phase-shifted clock buffer design was proposed to enhance the timing compensation between the clock path and the datapath. A 1.2 V, 65 nm test chip demonstrates an 8–27% performance improvement in F_{max} for typical resonant noise frequencies from 100 MHz to 300 MHz. Compared with the previous RC filtered phase-shifted buffer, our design saves 85% of the clock buffer area while achieving similar performance improvement. Several practical issues of the proposed clock buffers were discussed including comparison with the adaptive clocking scheme and partially applying the proposed phase-shifted clock buffers. A mathematical framework was also proposed to model the timing margin improvements. The proposed numerical method was verified with extensive HSPICE simulations showing a 4–10× reduction in estimation error compared with prior modeling techniques.

REFERENCES

- [1] M. Saint-Laurent and M. Swaminathan, "Impact of power-supply noise on timing in high-frequency microprocessors," *IEEE Trans. Advanced Packaging*, vol. 27, no. 1, pp. 135–144, Feb. 2004.
- [2] M. D. Pant, P. Pant, and D. S. Wills, "On-chip decoupling capacitor optimization using architectural level prediction," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 10, no. 3, pp. 319–326, Jun. 2002.
- [3] J. Xu, P. Hazucha, and M. Huang *et al.*, "On-die supply-resonance suppression using band-limited active damping," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2007, pp. 286–603.

- [4] J. Gu, R. Harjani, and C. Kim, "Distributed active decoupling capacitors for on-chip supply noise cancellation in digital VLSI circuits," in *Symp. VLSI Circuits Dig.*, Jun. 2006, pp. 216–217.
- [5] M. Mansuri and C. K. Yang, "A low-power adaptive bandwidth PLL and clock buffer with supply-noise compensation," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1804–1812, Nov. 2003.
- [6] L. H. Chen, M. Marek-Sadowska, and F. Brewer, "Coping with buffer delay change due to power and ground noise," in *Proc. Des. Autom. Conf. (DAC)*, Jul. 2002, pp. 860–865.
- [7] T. Fischer, J. Desai, B. Doyle, S. Naffziger, and B. Patella, "A 90-nm variable frequency clock system for a power-managed titanium architecture processor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 218–228, Jan. 2006.
- [8] S. Yasuda and S. Fujita, "Compact fault recovering flip-flop with adjusting clock timing triggered by error detection," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, Sep. 2007, pp. 721–724.
- [9] N. Agarwal and S. S. Rath, "Low-jitter clock distribution circuit," US Patent 6,842,136 B1, Jan. 11, 2005.
- [10] M. Saint-Laurent, "Clock distribution network using feedback for skew compensation and jitter filtering," US Patent 7,317,342 B2, Jan. 8, 2008.
- [11] V. Gutnik and A. Chandrakasan, "Clock distribution circuits and methods of operating same that use multiple clock circuits connected by phase detector circuits to generate and synchronize local clock signals," US Patent 7,571,359 B2, Aug. 4, 2009.
- [12] T. Rahal-Arabi, G. Taylor, M. Ma, and C. Webb, "Design and validation of the Pentium III and Pentium 4 processors power delivery," in *Symp. VLSI Circuits Dig.*, Jun. 2002, pp. 220–223.
- [13] S. Pant and E. Chiprout, "Power grid physics and implications for CAD," in *Proc. Design Automation Conf. (DAC)*, Jul. 2006, pp. 199–204.
- [14] E. Hailu, D. Boerstler, K. Miki, J. Qi, M. Wang, and M. Riley, "A circuit for reducing large transient current effects on processor power grids," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2006, pp. 2238–2245.
- [15] N. Kurd, P. Mosalikanti, M. Neidengard, J. Douglas, and R. Kumar, "Next generation Intel core micro-architecture (Nehalem) clocking," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1121–1129, Apr. 2009.
- [16] T. Rahal-Arabi, G. Taylor, J. Barkatullah, K. L. Wong, and M. Ma, "Enhancing microprocessor immunity to power supply noise with clock/data compensation," in *Symp. VLSI Circuits Dig.*, Jun. 2005, pp. 16–19.
- [17] K. L. Wong, T. Rahal-Arabi, M. Ma, and G. Taylor, "Enhancing microprocessor immunity to power supply noise with clock-data compensation," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 749–758, Apr. 2006.
- [18] J. Gu, H. Eom, and C. H. Kim, "A switched decoupling capacitor circuit for on-chip supply resonance damping," in *Symp. VLSI Circuits Dig.*, 2007, pp. 126–127.
- [19] D. Jiao, J. Gu, P. Jain, and C. Kim, "Enhancing beneficial jitter using phase-shifted clock distribution," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, Aug. 2008, pp. 21–26.
- [20] N. Kurd, J. Barkatullah, and P. Madland, "Adaptive frequency clock generation system," US Patent 7,042,259 B2, May 9, 2006.
- [21] N. A. Kurd, J. S. Barkatullah, R. O. Dizon, T. D. Fletcher, and P. D. Madland, "A multigigahertz clocking scheme for the Pentium 4 microprocessor," *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1647–1653, Nov. 2001.
- [22] J. M. Hart, K. T. Lee, and D. Chen *et al.*, "Implementation of a fourth-generation 1.8-GHz dual-core SPARC V9 microprocessor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 210–217, Jan. 2006.



Dong Jiao received the B.S. degree in electrical engineering from Tsinghua University, Beijing, China, in 2006, and the M.S. degree in electrical and computer engineering from the University of Minnesota, Minneapolis, in 2009, where he is currently pursuing the Ph.D. degree.

He spent the summer of 2008 working as an intern with Seagate involved in MRAM design and simulation. His research interests include power integrity for mixed-signal ICs and low-power digital and SRAM design.



Jie Gu (S'03–M'10) received the B.S. degree from Tsinghua University, China, in 2001, the M.S. degree from Texas A&M University in 2003, and the Ph.D. degree from the University of Minnesota, Minneapolis, in February 2008.

He worked as an intern in Texas Instruments, Dallas, TX, from January 2007 to July 2007. From January 2008 to March 2010, he worked an IC design engineer focusing on power management for wireless ICs in Texas Instruments. He is currently working on mixed signal and RF IC design for highly integrated silicon receiver in Maxlinear, Inc. He has published more than 20 technical papers, holds several patents, and serves as a number of conference and journal committees.



Chris H. Kim (M'04) received the B.S. and M.S. degrees from Seoul National University, Seoul, Korea, and the Ph.D. degree from Purdue University, Lafayette, IN.

He spent a year at Intel Corporation where he performed research on variation-tolerant circuits, on-die leakage sensor design and crosstalk noise analysis. He joined the electrical and computer engineering faculty at the University of Minnesota, Minneapolis, in 2004 where he is currently an Associate Professor.

His research interests include digital, mixed-signal, and memory circuit design for silicon and non-silicon technologies.

Prof. Kim is the recipient of an NSF CAREER Award, a McKnight Foundation Land-Grant Professorship, a 3M Non-Tenured Faculty Award, DAC/ISSCC Student Design Contest Awards, IBM Faculty Partnership Awards, an IEEE Circuits and Systems Society Outstanding Young Author Award, ISLPED Low Power Design Contest Awards, an Intel Ph.D. Fellowship, and the Magoon's Award for Excellence in Teaching. He is an author/coauthor of 70+ journal and conference papers and has served as a technical program committee member for numerous circuit design conferences. He is the technical program committee chair for the 2010 International Symposium on Low Power Electronics and Design (ISLPED) and the guest editor for a special issue of the *IEEE Design and Test Magazine*.