A Self-Adjustable Clock Generator With Wide Dynamic Range in 28 nm FDSOI

Jaehwa Kwak, *Student Member, IEEE*, and Borivoje Nikolić, *Senior Member, IEEE*

***Abstract*— This work demonstrates a self-adjustable clock generator that closely tracks the voltage dependence of the critical path delay in a microprocessor. A tunable replica circuit (TRC) composed of inverter cells from the standard-cell library performs real-time frequency modulation to reduce the response time to a single cycle, and a digital controller quantizes the clock frequency by selecting appropriate phases of a delay-locked loop (DLL) that generates 16 phases of 2 GHz output. A watchdog unit continuously monitors the clock output to guard against metastability arising from the asynchronous path between the controller and the DLL. The proposed design is implemented in a 28 nm UTBB FDSOI technology and occupies an area of 1120 μm^2 with 2.7 mW power consumption. The generator has a wide tuning range of 550–2260 MHz at 1 V, functions at 35 MHz at 0.4 V, and can continuously synthesize the clock signal for arbitrary voltage waveforms between 0.4 V and 1 V.**

***Index Terms*— Adaptive clocking, dynamic frequency scaling, energy-efficient processor, delay-locked loop (DLL), tunable replica circuit (TRC), Metastability, FDSOI.**

I. Introduction

DYNAMIC voltage and frequency scaling (DVFS) is a widely used technique for increasing the energy efficiency of digital systems [1]. Real-time frequency adaptation in microprocessors also reduces supply noise voltage margins [2]. When applied to a microprocessor system, dynamic voltage scaling trades off performance for quadratic active energy savings. A common DVFS policy is to adjust the supply voltage based on the workload and the energy budget, while keeping the voltage level stable during one fixed configuration. Maintaining a tight relationship between the clock period and the critical path delay in a processor maximizes the energy efficiency. Adding a timing margin to the clock period lowers the performance and increases leakage [3], as the logic is idle for a portion of a clock period.

When operating at a constant supply voltage, a delay margin is added to guard against any voltage disturbances (voltage droop) that may arise during operation. Processor supply voltage is either regulated off-chip, by using a discrete inductor-based DC-DC converter, or can be integrated on-chip, by using a low-dropout or a switched-capacitor regulator [4] [5] [6]. With off-chip regulators, the slow dynamics of the off-chip LC network result in supply voltage droops [7]. Sudden changes in the load current, caused by the wake-up of functional units or other changes in activity, can result in incorrect operation of the chip. Integrated regulators have a much faster response to changes in current demand as they do not have large LC parasitics. Switched-capacitor (SC) DC-DC converters produce an inherent voltage ripple that can affect the digital circuit functionality. To minimize the ripple, many phases of SC DC-DC converters are typically interleaved, which reduces conversion efficiency [8] [9]. The highest-efficiency SC DC-DC converters implement simultaneous switching instead of interleaving, however resulting in a large supply voltage ripple [10].

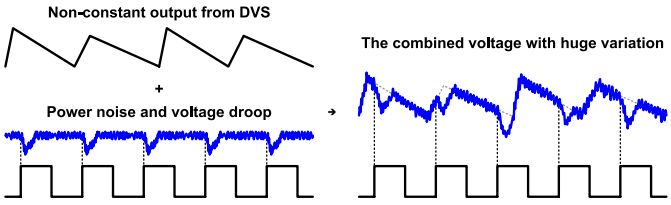

Fig. 1. Worst-case supply noise in the non-constant supply domain.

In the worst case, a voltage droop is added to the system supply ripple as shown in Fig. 1. Given that the droop can be large, various adaptive clocking techniques for mitigating its impact have been introduced. In general, the adaptive clock system consists of the droop detection circuitry and an adjustable clock generator. Current approaches to clock generation in DVFS include the use of integer dividers, per-core analog or digital phase-locked loops (PLLs), prescalers with DPLLs or a main PLL with per-core delay-locked loops (DLLs) [11]–[14]. These adaptive techniques detect the supply droop using an analog droop sensor, determine the operating condition, and adjust the clock frequency accordingly. The entire response time can require several clock cycles because the digital controller needs time to calculate the new target frequency. To reduce the response time, additional resources (e.g., multiple PLLs) may be applied, but this increases complexity, area and power consumption. Another adaptive clocking technique reduces the adaptation time through direct modulation of the PLL's clock frequency by using the supply voltage change [15]. This approach requires simpler logic and shortens the response. However, it

Manuscript received February 15, 2016; revised April 20, 2016; accepted June 8, 2016. Date of publication July 19, 2016; date of current version September 30, 2016. This paper was approved by Guest Editor Noriyuki Miura. This work was supported in part by DARPA PERFECT Award Number HR0011-12-2-0016, SRC/TxACE 1836.136, Intel, a KFAS Fellowship, and fabrication donation by STMicroelectronics.

The authors are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: jhkwak@eecs.berkeley.edu).

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.2016.2582860

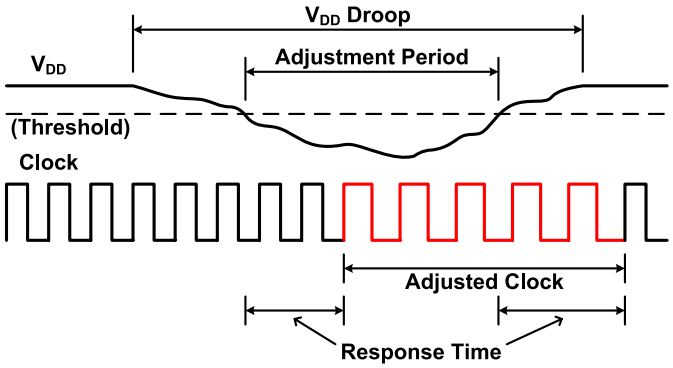

Fig. 2. Typical frequency adjustment scheme.

requires a local PLL for tracking the voltage droop, increasing locking time and area.

This paper presents a small, self-adjustable clock generator with a very fast response time. In contrast to conventional clock generators, the proposed generator constantly adapts to rapidly varying supply voltages. By removing droop detection and tracking the droop, per-cycle frequency adjustment is achieved. The remainder of the paper is organized as follows. Section II describes the operation principle of the clock frequency in the voltage droop (or ripple). The benefits of real-time clock adaptation is also explained here. Section III presents the proposed design details and implementation. Section IV summarizes the results of the experiments with the test chip, and Section V concludes.

II. Principle of Operation

The clock frequency in a DVFS system not only needs to respond to the externally set frequency targets, but also needs to track supply voltage droops caused by load variations. A typical frequency adjustment scheme does not track the continuous droop, and rather applies discrete frequency steps. As shown in Fig. 2, the scheme monitors the supply voltage level in relationship to a pre-set threshold. Once the voltage level is lower than the threshold level, the clock adjustment process starts, and the clock is adjusted after several cycles of the response time. When the supply voltage increases over the threshold, the clock frequency is restored. This scheme is conceptually simple but requires extra timing and voltage margin to guarantee functionality during the transients. In Fig. 2, the original clock frequency must meet the timing constraint until the clock is adjusted, resulting in excessive timing margin during the normal period of the stable supply voltage. Similarly, there is an additional margin during clock restoration.

The conventional frequency adjustment scheme should also cover voltage droops. In this case, the timing margin under nominal conditions (data-to-clock delay, T_{DC}) must be sufficient to guarantee adequate cycle time for worst-case supply noise. A simple logic block is used to simulate these conditions (Fig. 3), and the simulation model shows the timing margin at each clock rising edge in Fig. 4. The clock frequency is set to 1 GHz, and the logic path is adjusted to have 1.2% of timing margin at 0.9 V (12 ps @ 6 ns). At nominal operating conditions of 1 V, design margins account for 17% of the clock period (175 ps @ 1 ns).

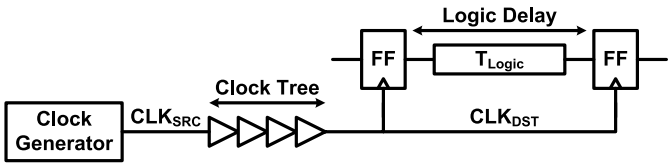

Fig. 3. Block diagram for adaptive clock simulation.

The key source of efficiency loss in conventional adaptive clocking schemes is the lag through the clock distribution network [21]. The clock signal, although generated with a constant frequency at the source, ends up being modulated as it passes through the distribution network that experiences droop. Operating the chip with a simultaneously-switching SC DC-DC converter [16] can be viewed as the worst-case droop. The simple block in Fig. 3 is re-used for the analysis of this condition. The clock period is matched to the logic delay of 4 FO4 delays, where FO4 is the fanout-of-four inverter delay. Other contributions to timing margin, such as clock-to-Q delay (T_{CQ}) and setup time of the flip-flop (T_{SU}), are ignored for simplicity. Fig. 5 shows the changes of the clock period and the logic delay as the supply voltage (V_{DD}) varies for the two cases: zero clock insertion delay, and 1 FO4 clock insertion delay. For simple illustration, the logic delay is assumed to be linearly dependent to supply changes. In both cases, the timing constraints for the clock period of 4.0 arbitrary units are met when the supply is stable. However, the logic delay in the variable voltage domain is not constant. In the first case of decreasing supply, the clock period is still constant because of the zero clock insertion delay, while the delay is stretched to 4.6, resulting in a setup time violation. In the second case, the clock period tracks the logic, but the late rising edge of the clock due to the clock-tree delay puts the logic gates in the lower supply voltage. Therefore, the logic delay becomes longer (5.0 vs. 4.6) than in the first case. When the supply voltage is increasing, the opposite is observed: the logic delay is compressed due to the higher supply voltage, and thus the timing margin becomes bigger. This does not cause a timing violation, but does result in a loss of energy efficiency.

The proposed adaptive clocking scheme mitigates these issues. It adapts to the continuously varying target voltage, and changes the frequency by directly applying the varying supply to the replica circuit that generates the clock period as shown in Fig. 6. The replica circuit accumulates the supply droop of any amplitude and duration, and determines the clock period after the accumulated delay, which realizes a single-cycle response time. The clock adjustment process is synchronized to the voltage dependency without any preset threshold, allowing continual frequency adjustment with minimal latency. Fig. 7 compares the conventional and proposed approaches as the supply voltage changes and under supply droop.

III. Clock Generator Architecture

A block diagram of the proposed clock generator is shown in Fig. 8. The clock generator has a DLL, a tunable replica circuit (TRC), and a controller. An external global clock supplies the DLL, which generates the equally-spaced reference

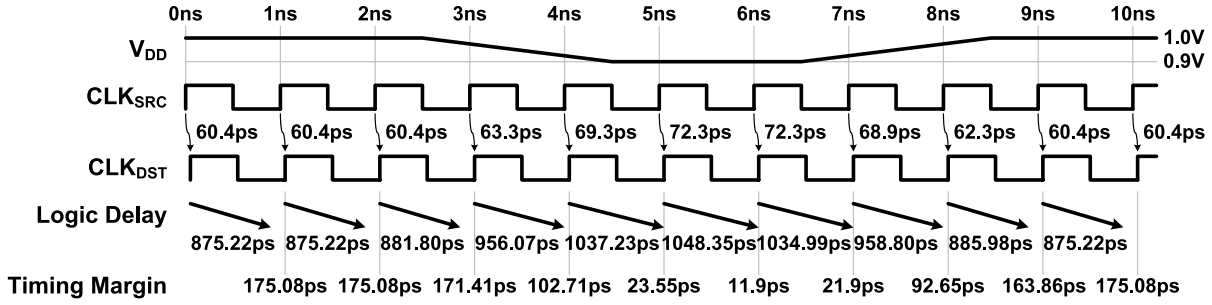

Fig. 4. Timing diagram with delay annotated with supply droop.

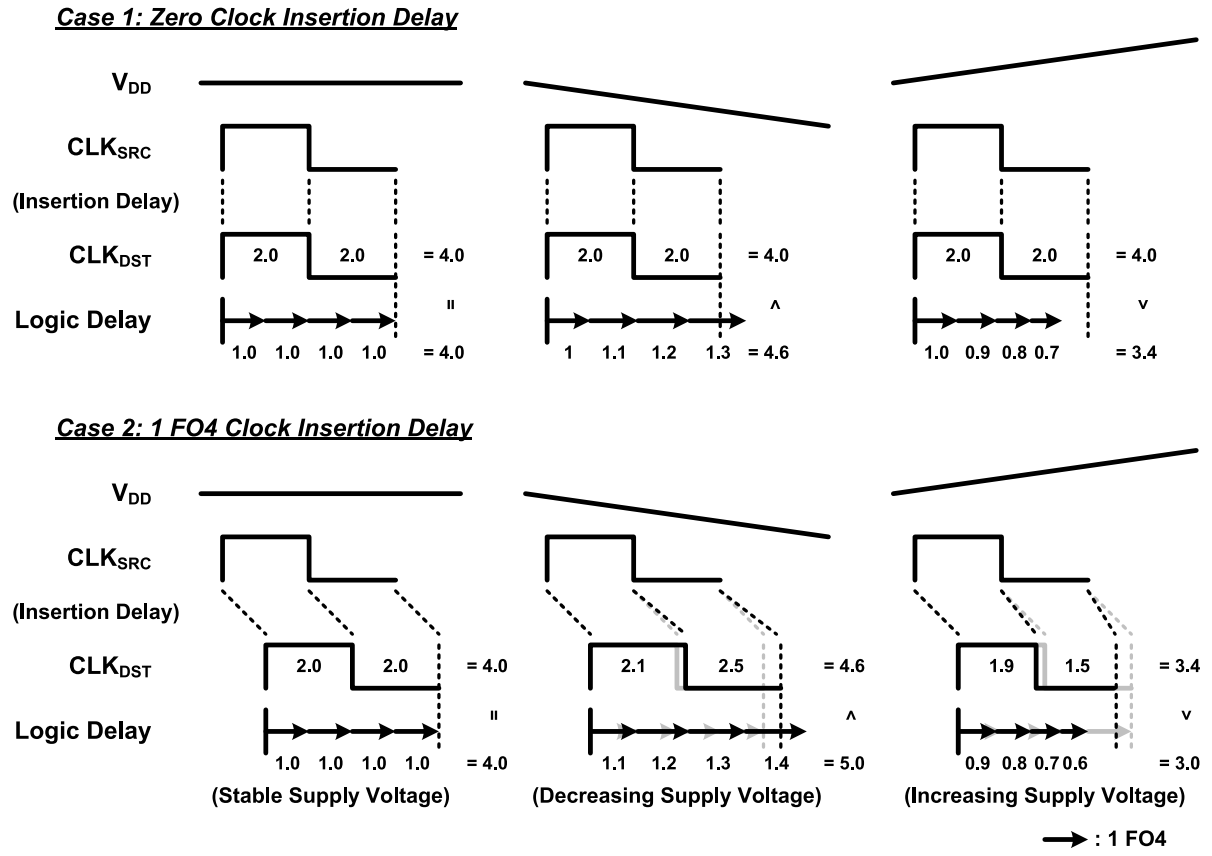

Fig. 5. Changes of clock period and logic delay with varying supply.

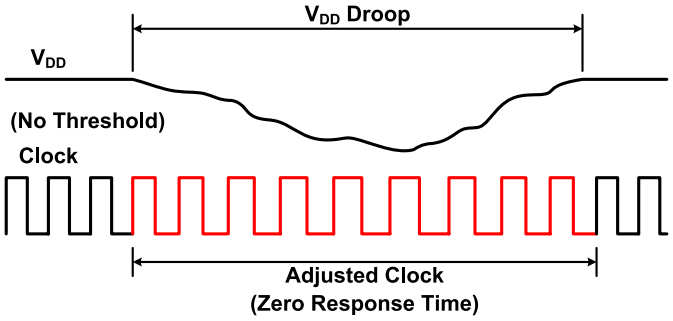

Fig. 6. Continuous frequency adjustment scheme.

clock phases. By using the reference signals from the DLL and the estimated clock delay from TRC, the control logic generates the clock output signal by selecting one of the DLL outputs. A watchdog circuit assures that metastability does not interrupt the clock.

A. Delay-Locked Loop

The schematic of the DLL is shown in Fig. 9. The main blocks are the voltage-controlled delay line (VCDL), the phase detector (PD), and the charge pump (CP) with the loop filter (LF).

The VCDL, shown in Fig. 10, converts a single-ended global 2 GHz clock signal into 8 differential signals of the same frequency. The single-ended input is converted to a differential pair by using an inverter, followed by cross-coupled differential inverter chains. The cross-coupled differential delay unit is composed of two inverters for the main path, and two small inverters that connect the outputs of the main inverters to minimize the phase mismatch between the output nodes.

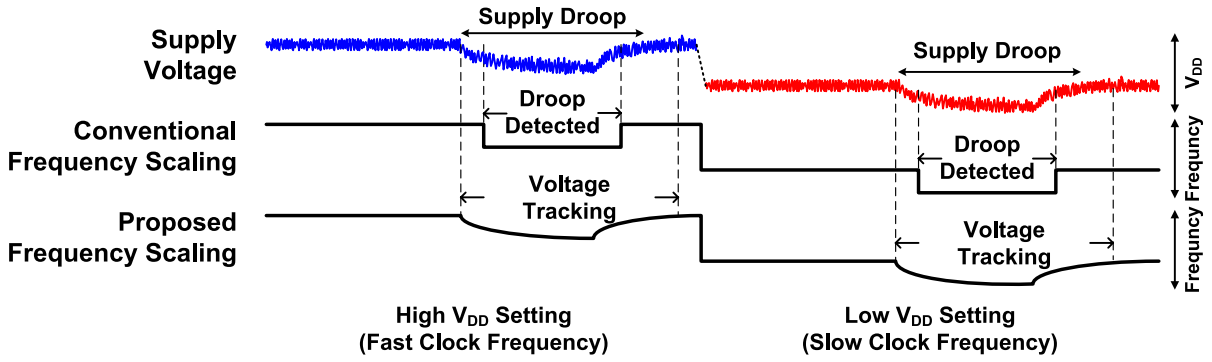

Fig. 7. Illustration of the conventional frequency scaling and the proposed frequency scaling.

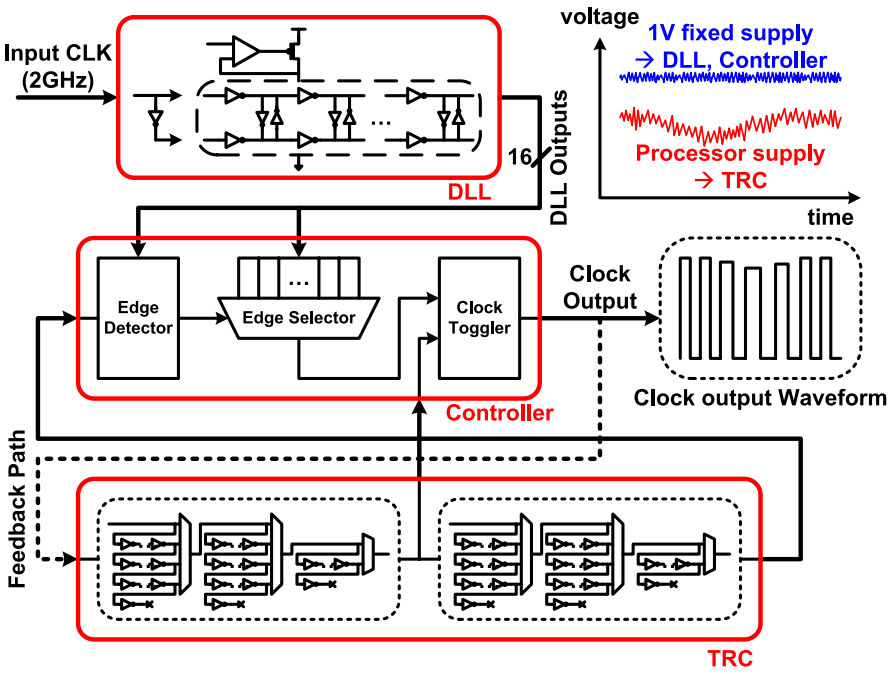

Fig. 8. Block diagram of the proposed clock generator.

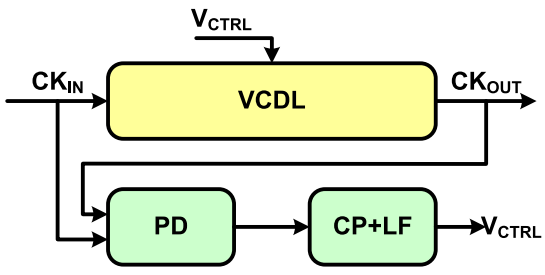

Fig. 9. Block diagram of the DLL.

The small coupling inverters are a quarter of the size of the main path inverters so as not to distort the signal on the path. To achieve better DLL outputs, the same differential units are used before and after the VCDL as a pre-buffer and a dummy load, respectively. The VCDL generates 16 phase references that have the same interval of 31.25 ps (500 ps/16), where the 16th phase matches the 0th phase. The delay line is designed to have the delay between the 0th reference and the 16th reference to be around 375 ps (75% of 500 ps) at the nominal voltage of 1 V, so that the rising edge of the 16th reference can be initially located in the low period (between the falling edge and the next rising edge) of the 0th reference. This initial difference must be guaranteed for the correct operation of the DLL. The VCDL is controlled by a low-dropout (LDO) regulator which is composed of a CML-type differential amplifier and a large PMOS cell as a power gate. A large decoupling capacitor is connected to the power node of the VCDL to alleviate the voltage ripple due to logic switching. Since the output voltage level of the VCDL, which is controlled by the local LDO, is usually lower than a nominal voltage (1 V), a level shifter is added to each output signal. The level shifter is followed by a large buffer to minimize interval mismatch between adjacent outputs of the VCDL due to load mismatches.

The phase detector in Fig. 11 compares the 0th phase and the 16th phase of the reference signals. The detector

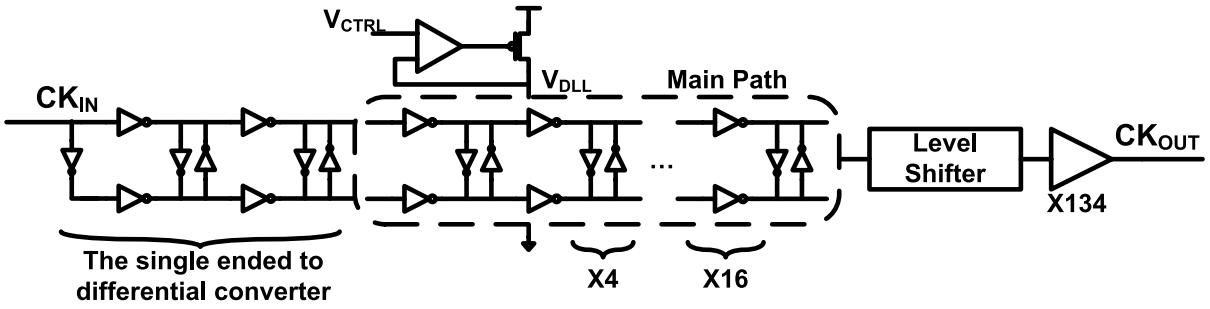

Fig. 10. Schematic of the voltage-controlled delay line.

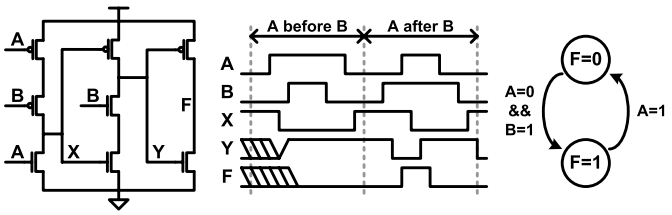

Fig. 11. Phase detector, its waveform, and state diagram of F.

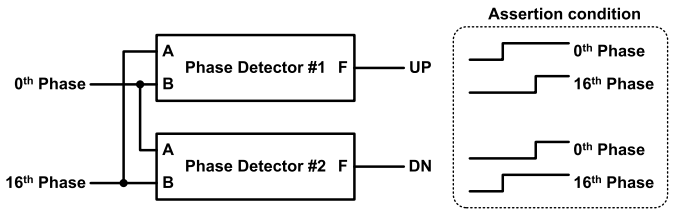

Fig. 12. Two phase detectors configuration and the assertion condition.

is composed of two phase detector units which are used to generate the UP and DN signals respectively, by connecting the appropriate phases to the inputs as shown in Fig. 12. The output F is initially zero. UP is asserted when the VCDL needs to be sped up, and DN is asserted when the VCDL needs to be slowed down. In the detector unit, the output F is set to '1' only when the input B rises faster than the input A, and stays high until the input A is on. The pulse width of the output F is a function of the time difference between the two rising edges. For the DN detector unit, the 16th reference is connected to the input B and the 0th reference is connected to the input A. Similarly, the 0th reference is connected to the input B and the 16th reference is connected to the input A for the UP detector unit. At initialization, the 16th reference has its rising edge when the 0th reference is low, and so the phase detector detects that that the 16th reference has risen faster than the 0th reference and signals the charge pump to slow the VCDL by asserting DN.

The charge pump determines the reference voltage level of the LDO. Using the UP and DN signals from the phase detector, this block charges or discharges the output node. The 500 fF of capacitance is placed at the output of the CP as a loop filter. To match the speed of charging and discharging, the transistor sizes of the PMOS and the NMOS are carefully

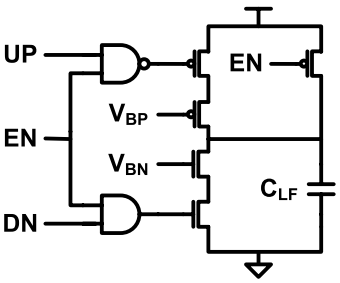

Fig. 13. Charge pump with loop filter.

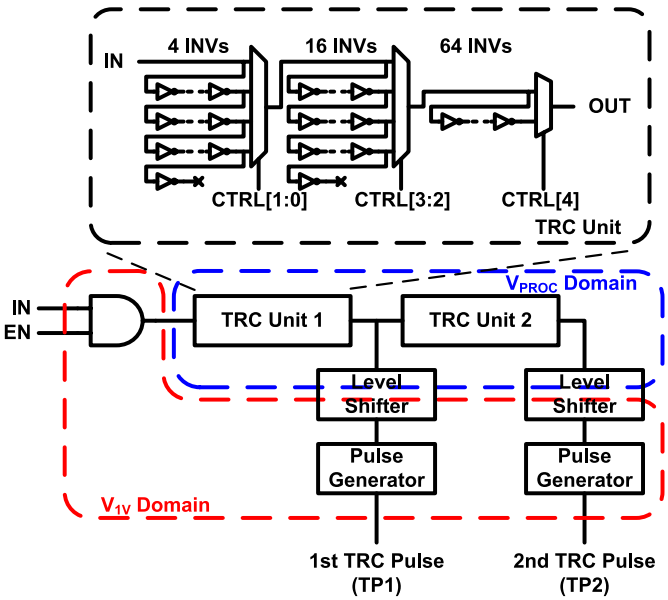

Fig. 14. Block diagram of the proposed clock generator.

selected. It also has a reset signal (EN), which forces the output to 1 V immediately as shown in Fig. 13.

B. Tunable Replica Circuits (TRC)

The tunable replica circuit shown in Fig. 14 is placed in the processor core that is adaptively clocked. TRC is designed as a configurable inverter chain, and has two units that are independently controlled to set the duty cycle of the output clock. Each unit consists of 124 identical inverter cells with a 5-bit control signal. The configuration resolution is 4

FO1 delays, where FO1 is the fanout-of-one inverter delay. Four inverter cells compose the resolution step, and the control signal determines how many resolution steps are used in the delay path. Since the number of inputs increases the mux delay exponentially, the delay path is divided into three stages, with two 4:1 muxes and a 2:1 mux. The first stage mux selects 0, 4, 8, or 12 inverter cells, the second stage mux selects 0, 16, 32, or 48 inverter cells, and the last mux selects 0 or 64 inverter cells. The minimum delay length is 0 inverters with 3 muxes when the control signal is "00000," while the maximum delay length is 124 inverters (12 from the first, 48 from the second, and the 64 from the last stage) with 3 muxes when the control signal is "11111."

The first unit sets the high phase of the clock, which starts at the rising edge and generates the falling edge, while the second unit is for the low phase. With this feature, the period and the duty cycle can be independently adjusted. Two TRC units of 5-bit control each are equivalent of 6-bit control delay chains, resulting in a very wide tuning range of 0 to 248 FO1 inverter delays. The TRC inverters are connected to the processor supply (V_{PROC}) in order to track the critical path under voltage variations, while other logic remains in an isolated 1 V domain (V_{1V}). Therefore, 4 FO1 resolution is guaranteed at any supply voltage level, while the total delay is self-adjusted as a function of the supply. TRC resolution is related to the phase resolution of the DLL. TRC resolution finer than 4 FO1 would eventually be quantized by the DLL with the resolution of 31.25 ps, and will be ignored after passing the quantizer. In the particular implementation, the 4 FO1 is about 25 ps at 1 V.

At the top level, the output clock signal is connected to the TRC input. Two TRC units delay the rising edge of the input successively by the estimated critical path of the processor through appropriate digital calibration of the number of delay stages, and convert the delayed edges to pulses of short duration for the controller by using the pulse generators (PG). The level shifters are located between the TRC and the PG, because the TRC units are on the processor's voltage domain while the PG and the controller operate at 1 V. The TRC goes into sleep mode to reduce the energy consumption when not in use by de-asserting the enable signal (EN). Since the TRC delay directly determines the next clock period without any additional logic, the circuit achieves a response time of a single cycle.

C. Controller

The controller is composed of an edge detector, an edge selector, and a clock flip-flop as shown in Fig. 15. The DLL references (P_0–P_{15}) and the TRC pulses (TP1 and TP2) are the inputs to this block. The clock output signal (CP_{OUT}) is generated from the flip-flop followed by a large clock buffer. The flip-flop data input is connected to VDD, its asynchronous reset input is connected to the first TRC output pulse (TP1), and its clock input (CI) is supplied by the edge selector. The first TRC output pulse asynchronously resets the flip-flop to generate the falling edge of the clock output. The second TRC output pulse (TP2), however, is not directly used to generate the rising edge because the rising edge should be

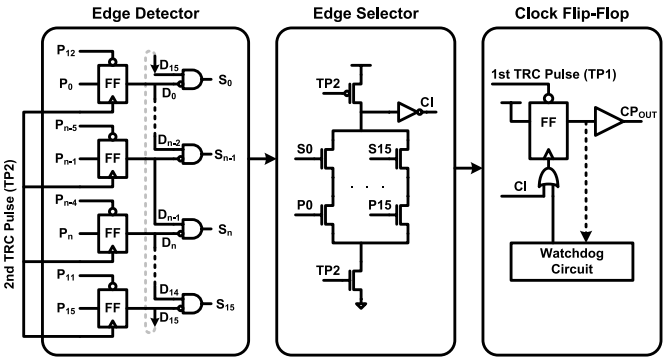

Fig. 15. Controller diagram.

synchronized to the DLL references. In the edge detector, the second TRC output pulse samples the DLL references to detect a reference phase that is about to rise. Since the reference signals have equivalent intervals, the sampled values(D_0–D_{15}) are eight continuing 0's and eight continuing 1's circularly, and the first 0 after the sequence of 1's corresponds to the phase of upcoming rising edge. Each sampled signal is reset to zero by the DLL reference which is 4 phases ahead of the input (P_{n-4}). The edge selector picks the correct reference phase by using the result from the previous stage (S_n). The dynamic mux gate is used to speed up the operation. The gate is in evaluation phase only when the second TRC output pulse is high, and transitions to precharge phase when the pulse from the TRC is low. The mux output, which is a pulse, triggers the flip-flop as a clock signal. The entire waveform is represented in Fig. 16.

Since the DLL references and the TRC output pulse are asynchronous, the flip-flops in the edge detector may occasionally fail to sample the inputs, resulting in metastability. In Fig. 16, TP2 fails to sample P_n if two edges are close, and the output Dn does not settle on time. After passing several logic gates, the metastable signal is usually resolved to the stable logic level. If the D_n is settled to the incorrect value, Sn+1 is asserted instead of S_n, and CP_{OUT} rises late by one DLL delay which is 31.25 ps. In most cases, the metastable outputs are resolved to '0' or '1' in a short time, so the entire operation remains correct. However, the control logic can malfunction if the selection signal, S_n, remains metastable, and thus the edge selector fails to generate the pulse CI. Without enough width of CI, the clock output signal is eventually set to zero, and may remain low. To avoid this, a watchdog block is implemented. This logic continuously monitors the status of the clock output signal, and generates an extra pulse for the clock flip-flop input when the clock output signal remains zero for longer than two clock delays of the previous period as shown in Fig. 17. The clock output (CLK_{OUT}) is synchronized to the global 2 GHz clock (CLK_{REF}) which is also used in the DLL. The sampled clock output is connected to the gate logic which generates the pulse, PU, that is asserted for a cycle when the sampled clock output goes high. The free running counter is reset when PU is asserted, and sends the value to the reference. If the clock output is metastable, the counter value exceeds the previous counter value stored in the reference without the PU assertion. Although two clock signals are asynchronous, it is possible to

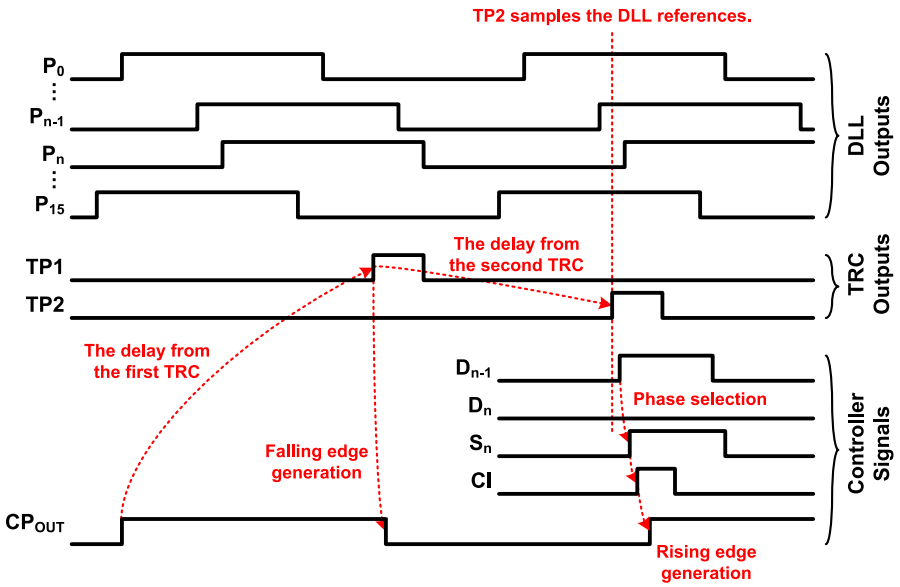

Fig. 16. Clock output generation.

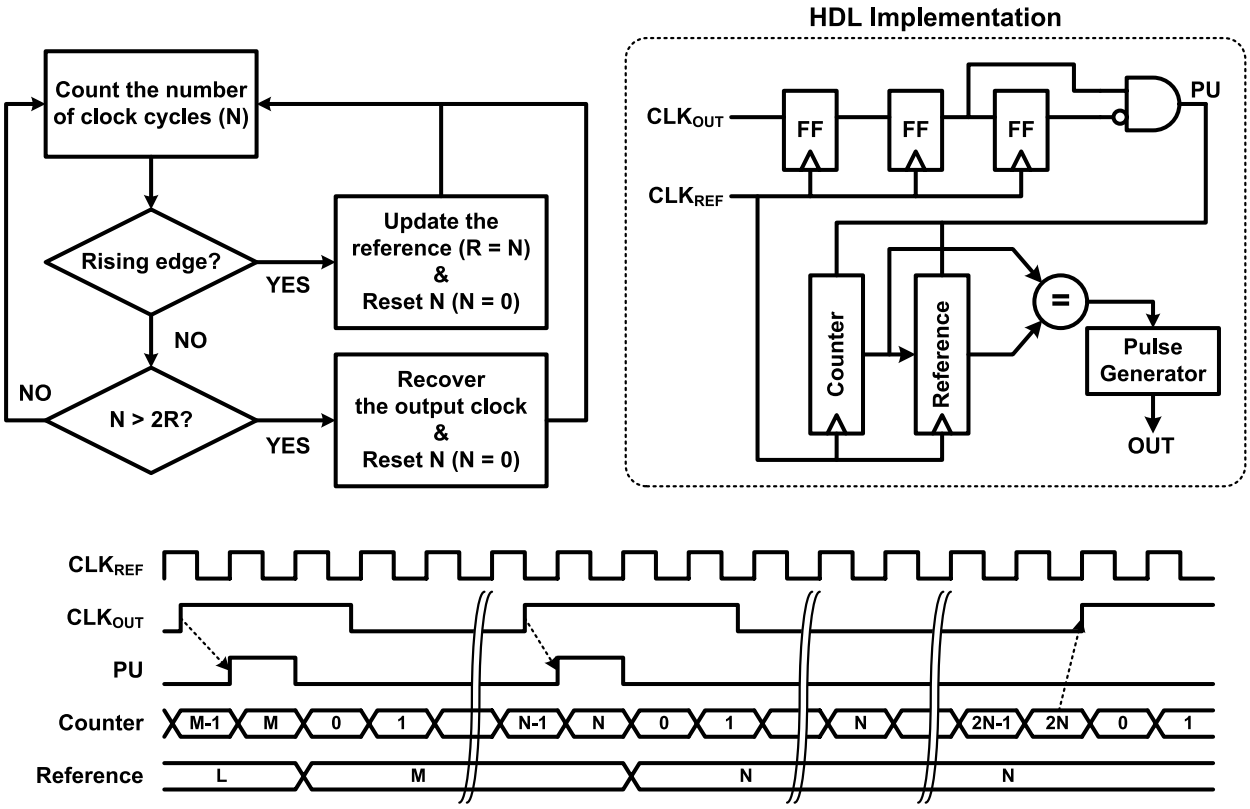

Fig. 17. Watchdog flowchart, schematic for HDL implementation, and the waveform.

estimate the clock skew because the CLK_{OUT} is quantized by the DLL outputs which are originated from the CLK_{REF}. In the design phase, the skew is adjusted to minimize the metastablility.

D. Physical Design

The clock generator is designed to conform to standard cell design rules to enable tight integration with the core logic. By using same design style, and thus by embedding the block directly in the processor, the core and the clock generator can share operating conditions such as voltage and temperature. The signal paths for the DLL references are manually routed, and the wire lengths are carefully matched to minimize any possible discrepancy in the delay intervals. The signal paths in the controller are routed in the same manner. The watchdog is synthesized and automatically placed and routed.

TABLE I
ANALYSIS OF THE CLOCK PERIOD

V_{DD}	Maximum Period (ps)	Minimum Period (ps)	Ratio
	Maximum TRC delay (ps)	Minimum TRC Delay (ps)	
1.0V	1818.2	442.5	4.11
	1613.2	237.5	6.79
0.5V	9708.7	1600.0	6.07
	9503.7	1395.0	6.81

(The controller delay is set to 205 ps in this table.)

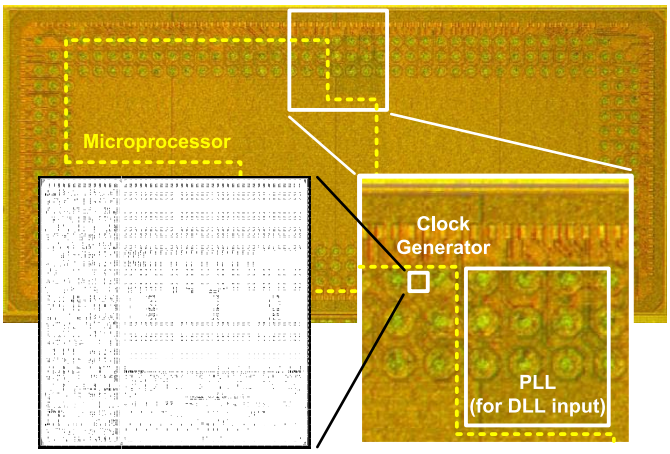

Fig. 18. Chip micrograph.

IV. EXPERIMENTAL RESULTS

The clock generator has been fabricated in a test chip shown in Fig. 18 [17] implemented in 28 nm ultra-thin body and BOX fully-depleted silicon-on-insulator (UTTB FDSOI) technology [18] [19]. It occupies an area of 1,120 μm^2, and consumes 2.7 mW from a 1 V supply.

A. Measurement Setup

The measurement setup is depicted in Fig. 19. The custom flip-chip is embedded in the test board. The test board is controlled from an Opal Kelly FPGA board (Shuttle LX1) which has two interfaces to the test board. The Inter-Integrated Circuit (I^2C) protocol controls the on-board voltage generator which provides the 1 V and 1.8 V supply to the chip, while the Serial Peripheral Interface (SPI) is used to manage the clock generator's control registers. The FPGA board is operated by the host through the Universal Serial Bus (USB). The on-chip PLL, which uses a 50 MHz input clock from the off-chip clock source, generates the 2 GHz clock signal for the DLL input. The varying supply for the TRC of the clock generator is provided by the external sourcemeter, and the oscilloscope is connected to measure the output clock signal.

B. Frequency Measurement

Fig. 20 shows the tuning range of the clock generator. At 1 V, the frequency range of 550–2260 MHz can be achieved, depending on the TRC configuration. The TRC delay, which determines the clock frequency, is longer at lower supply voltages, and thus the frequency range scales down to 103–625 MHz at 0.5 V. The ratio of maximum and minimum frequency may be different as the supply voltage changes because of the delay from the controller. The clock period is the sum of the TRC delay and the controller delay; the TRC delay is scalable when the supply varies, whereas the controller delay is constant (about 200 ps) because the controller is operated in the 1 V domain. Table I shows the decomposition of the clock period at 1 V and 0.5 V. If the controller delay is excluded, the TRC delay tracks the critical path with the supply voltage. Since the clock generator does not contain large wire delay, the controller delay is a proxy to the wire delay in the processor, since the RC wire is largely voltage independent.

Fig. 21 shows the frequency change when the supply decreases from 1 V to 0.4 V. Selected points (0.4 V to 1.0 V, 0.1 V step) illustrate the clock frequency changes, and the frequency of each selected point follows the trajectory. The clock generator has been verified to operate at 0.35 V with the slowest TRC configuration. If the supply is below 0.35 V, the clock generator fails to generate the stable signal, and the watchdog logic resurrects the clock output. Fig. 22 shows that the watchdog logic recovers from the clock signal failure when the supply voltage is 0.34 V. The clock signal remains zero when the failure happens, but is revived after 2 cycles of the previous period when the watchdog is activated. When deactivated, the clock signal is not recovered.

C. Frequency Quantization and Clock Jitter

Since the clock output is generated from the DLL outputs, the rising edge of the clock has a discrete distribution same with the DLL outputs distribution of 31.25 ps interval, verified by the plot in Fig. 23. The non-discrete components of the distribution result from clock jitter.

The clock output jitter is shown in Fig. 24. The clock generator is set to generate 2 GHz clock signal at 1 V to match the frequency with the PLL output clock. Since jitter is meaningless in the context of rapidly varying frequencies, a 1 V fixed supply is used to hold the output clock frequency. The measured long-term absolute jitter of the output clock at 2 GHz is 4.24 ps_{rms}/38.33 ps_{pp}. The jitter from the PLL is the dominant source of the clock output jitter, which is 3.97 ps_{rms}/35.67 ps_{pp} at 2 GHz. Since the I/O pad for the clock output is selected from the standard-cell library and not

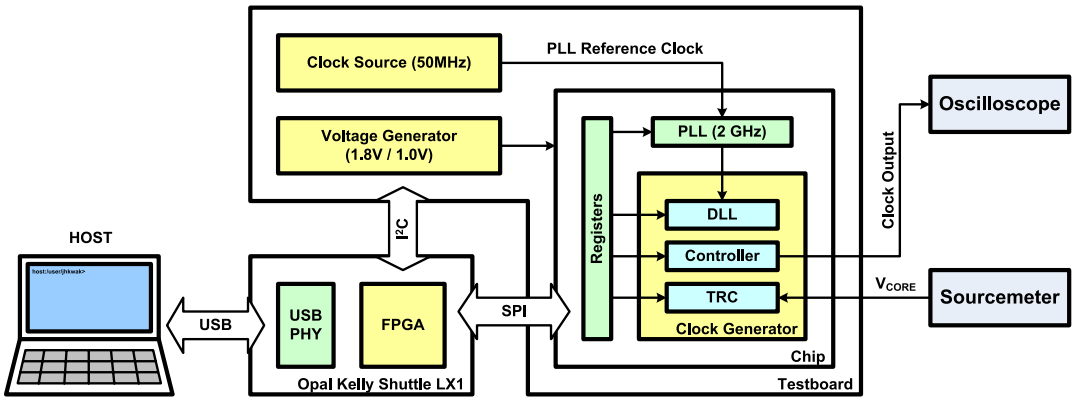

Fig. 19. Block diagram of the measurement setup.

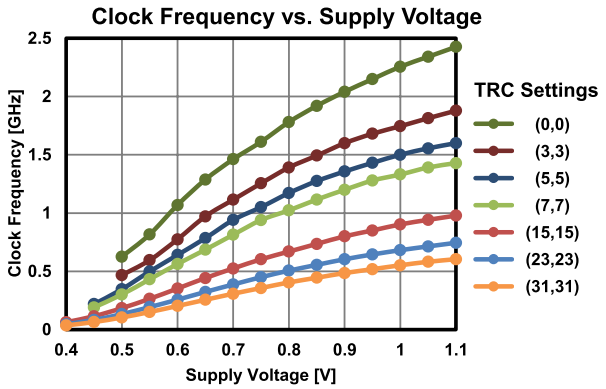

Fig. 20. Tuning range of the clock generator with various TRC settings.

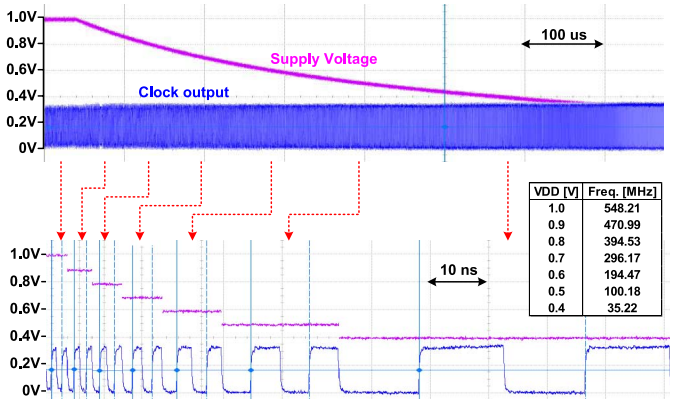

Fig. 21. Frequency change when the supply decreases, and zoomed waveform of selective points.

carefully designed for noise reduction, the PLL jitter may contain the noise from the pad. The jitter from the clock generator itself is 1.49 ps_{rms} from

$$\sigma_{ClockGenerator} = \sqrt{\sigma^2_{Total} - \sigma^2_{PLL}} \quad (1)$$

The contributors to the clock generator's jitter are the phase noise from the DLL and the supply noise from the logic. Because the clock generator shares the 1 V supply voltage domain with the microprocessor, the supply noise from the microprocessor is inserted into the clock generator, and this

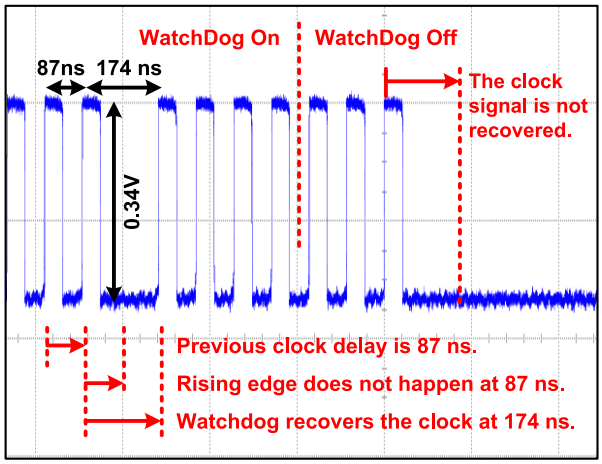

Fig. 22. Waveform of the watchdog operation.

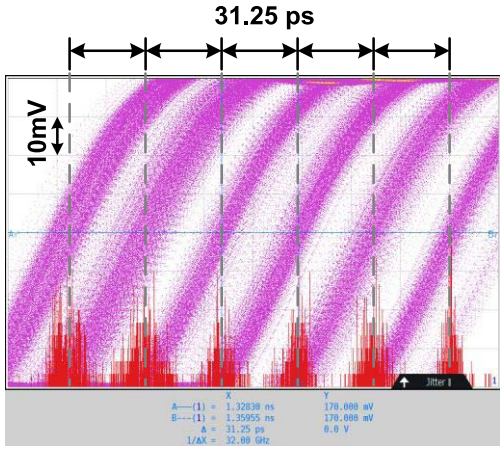

Fig. 23. Distribution of the rising edge of the output clock.

results in the timing uncertainty, and thus clock jitter.

D. Energy Efficiency of the Microprocessor

The clock generator was also embedded in a processor SoC with dynamic voltage scaling, where it was measured to track supply droops that exceed 100 mV and improve overall energy efficiency by more than 10% [20]. The block diagram of the microprocessor is shown in Fig. 25. The clock generator is

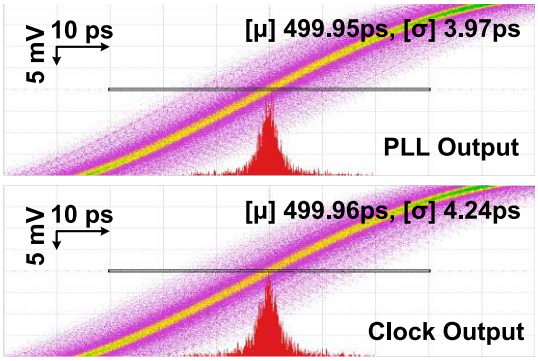

	PLL Output	Clock Output
Mean	499.94771ps	499.95977ps
Std Dev	3.97076 ps	4.24211 ps
p-p	35.67 ps	38.33 ps
Min	483.67 ps	481.22 ps
Max	519.33 ps	519.56 ps
Bin Width	110 fs	110 fs

Fig. 24. Jitters of the PLL output and the clock output, and summary table.

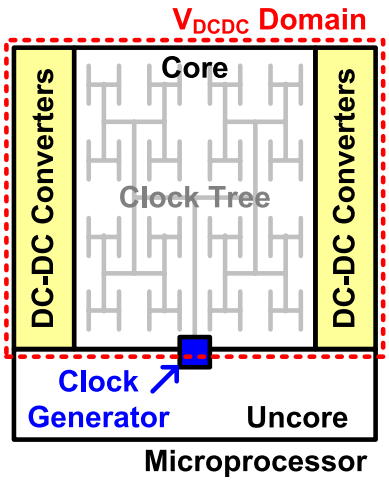

Fig. 25. Block diagram of the test chip.

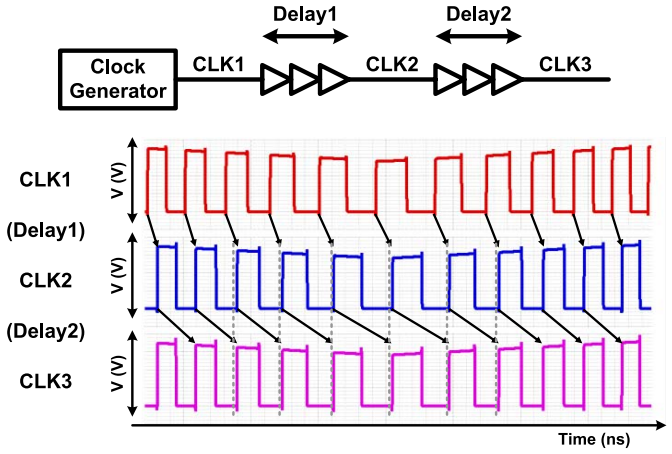

Fig. 26. Simulation of frequency adjustment in the clock tree.

placed adjacent to the core to minimize the clock insertion delay that degrades the frequency adaptation. As explained in Section II, the instantaneous frequency of the clock signal may change as it propagates through the clock tree. To confirm this, the simple clock tree model is simulated in Fig. 26. The clock generator output (CLK1) has adaptive frequency, and its frequency is re-adjusted after passing the buffers

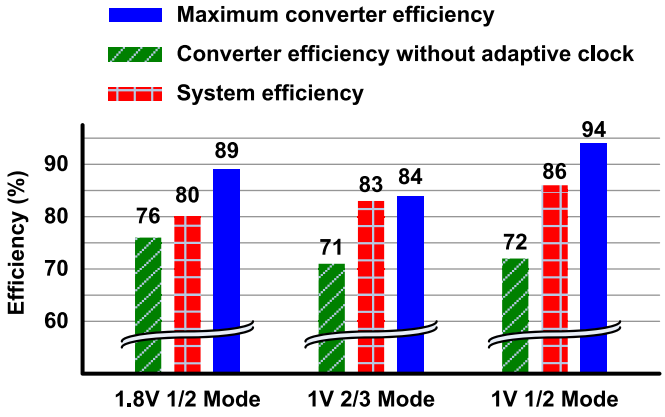

Fig. 27. System efficiency improvements with adaptive clock.

(CLK2, and CLK3). The second clock buffer chain is intentionally set to have the same delay with the clock period at nominal voltage to verify how the clock insertion delay affects the frequencies of CLK2 and CLK3. The simulation result shows that the clock frequency is well-adjusted regardless of the supply voltage and the depth of the clock tree. The SoC features an integrated DC-DC converter to implement DVFS. To minimize the effect of local IR drop, the DC-DC converter unit cells are distributed around the core with a dense power grid. Two layers of thick upper-layer metal were dedicated to a power grid, where the DC-DC converter output and GND each utilize 25% of the chip area in each layer, and power rail analysis estimates a 2 mV voltage drop at 1 V and 100 mA (nominal operating condition). By placing the clock generator next to the core, the voltage droop of the clock generator is matched to the droop in the core. The uncore has an isolated power domain from the core, and uses a separate clock with a fixed frequency. The voltage generated by the DC-DC converter is not used outside the core, so the 1 V voltage, as well as 1.8 V voltage for the DC-DC converter input, use the majority of the power routing resources to connect power coming from the pad frame to the converters.

Fig. 27 shows the energy efficiency improvement for each DC-DC configuration. The DC-DC converter is reconfigurable between three conversion modes with two input supplies: the 1 V input is divided with a 2:1 and 3:2 ratio to generate the 0.5 V and 0.67 V modes (1 V 1/2 mode and 1 V 2/3 mode,

TABLE II
DESIGN SUMMARY AND COMPARISON WITH PREVIOUS WORKS

Parameter	This Work	[2]	[11]	[12]	[13]	[14]
Technology	28nm FDSOI	90nm CMOS	28nm FDSOI	28nm bulk	45nm PDSOI	90nm CMOS
Size [mm^2]	0.000112	N/A	0.013	N/A	0.07	0.75
Minimum Frequency	35MHz at 0.4V	1.26GHz at 1.2V	1.06MHz at 0.33V**	2.0GHz at 0.86V	1.8GHz at V_{NOM}	2.2GHz at 1.0V
Maximum Frequency	2.26GHz at 1.0V	2.5GHz at 1.2V	443MHz at 0.9V	4.0GHz at 1.45V	4.0GHz at V_{NOM}	3.4GHz at 1.4V
Response time	1 cycle	2 cycles	3 cycles	1-3 cycles	N/A	1 cycles
Power	2.7mW at 1V*	N/A	0.67mW at 0.45V	N/A	N/A	N/A

* DLL: 2.1mW, TRC: 0.4mW, and the controller 0.2mW at 1V. The powers of the DLL and the controller are not scalable because the 1V fixed supply is applied. Only the TRC power is scalable.
** 0.8V FBB (Forward body biasing) is applied.

respectively), while the 1.8 V input is divided with a 2:1 ratio to generate the 0.9 V mode (1.8 V 1/2 mode). In each mode, the output voltage is rippling around the target voltage level. The amount of efficiency improvement depends on the DC-DC configuration, but in each mode, the adaptive clock increases the overall system efficiency, which is defined as the total amount of energy required to finish the same workload in the same time. The energy at the maximum frequency with an ideal off-chip voltage source is defined as 100% efficiency. If the DC-DC converter is used instead of the ideal source, the system efficiency becomes lower due to conversion losses from the converters. With adaptive clocking, however, the overall running time decreases due to decreased timing margin, resulting in higher efficiency.

Clock timing margin translates directly to system efficiency loss. Because the delay from the controller is constant, the portion of the adjustable TRC delay is lower at the high supply voltage, and this results in slowdown of frequency adaptation, and timing margin reduction. Since the voltage ripple from the DC-DC converter is largest in 1.8 V 1/2 mode, the accuracy of voltage tracking is the worst in this mode and affords a smaller efficiency improvement. For the other voltage modes, however, the system efficiency is improved by more than 10%, because the clock generator better tracks the voltage ripple.

Table II summarizes the performance of the clock generator in comparison with previous work. It shows that the proposed clock generator has outstanding performance in tuning range and area with comparable power consumption. The main difference of the proposed work from most prior work is the approach to supply droop detection. Prior work assumes that the desired supply voltage is fixed and detect the voltage droop systematically, which increases the size of the detection logic as well as the response time and limits the tuning range. To achieve a single-cycle response time, multiple PLLs are used in [14], which increases the area. The proposed design, however, does not have a target supply voltage, and directly uses the drooping voltage as a supply. No additional logic for droop detection is required with this approach, and any supply voltage can be used as long as the technology can support it.

V. CONCLUSION

The rapidly adjustable clock generator in this paper has wide voltage range of 0.4–1 V and synthesizes a clock signal at frequencies from 35 to 2260 MHz. The design, including DLL, TRC, and controller, was implemented in a 28 nm FDSOI process, and has small size of 1,120 μm^2 with 2.7 mW of power consumption. Direct coupling of the supply voltage to the frequency generation realizes a single-cycle response time, and the frequency quantization reduces the synchronization time to 1–2 cycles. Measurement results show that the clock generator adjusts the clock frequency with changes in the supply voltage and handles clock failure due to metastability. The results also show that the energy efficiency in the microprocessor is improved with the adaptive clock by tracking the supply voltage change. The presented clock generator provides a low-cost and high-performance solution for adaptive clocking for a wide variety of applications.

ACKNOWLEDGMENT

The authors would like to thank BWRC, ASPIRE, Brian Richards, Brian Zimmer, Yunsup Lee, and Ben Keller.

REFERENCES

[1] T. D. Burd, T. A. Pering, A. J. Stratakos, and R. W. Brodersen, "A dynamic voltage scaled microprocessor system," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1571–1580, Nov. 2000.
[2] T. Fischer, F. Anderson, B. Patella, and S. Naffziger, "A 90 nm variable-frequency clock system for a power-managed Itanium-family processor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 294–295.

[3] D. Markovic, V. Stojanovic, B. Nikolić, M. A. Horowitz, and R. W. Brodersen, "Methods for true energy-performance optimization," *IEEE J. Solid-State Circuits*, vol. 39, no. 8, pp. 1282–1293, Aug. 2004.
[4] Z. Toprak-Deniz *et al.*, "Distributed system of digitally controlled microregulators enabling per-core DVFS for the POWER8 microprocessor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 98–99.
[5] N. Kurd *et al.*, "Haswell: A family of IA 22 nm processors," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 112–113.
[6] A. Nalamalpu *et al.*, "Broadwell: A family of IA 14 nm processors," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C314–C315.
[7] H. K. Krishnamurthy *et al.*, "A 500 MHz, 68% efficient, fully on-die digitally controlled buck voltage regulator on 22 nm tri-gate CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2014, pp. 1–2.
[8] M. Turnquist *et al.*, "Fully integrated DC-DC converter and a 0.4 V 32-bit CPU with timing-error prevention supplied from a prototype 1.55 V Li-ion battery," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C320–C321.
[9] S. T. Kim *et al.*, "Enabling wide autonomous DVFS in a 22 nm graphics execution core using a digitally controlled hybrid LDO/switched-capacitor VR with fast droop mitigation," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 1–3.
[10] H.-P. Le, S. R. Sanders, and E. Alon, "Design techniques for fully integrated switched-capacitor DC-DC converters," *IEEE J. Solid-State Circuits*, vol. 46, no. 9, pp. 2120–2131, Sep. 2011.
[11] S. Clerc *et al.*, "A 0.33 V/–40 C process/temperature closed-loop compensation SoC embedding all-digital clock multiplier and DC-DC converter exploiting FDSOI 28 nm back-gate biasing," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 1–3.
[12] A. Grenat, S. Pant, R. Rachala, and S. Naffziger, "Adaptive clocking system for improved power efficiency in a 28 nm x86-64 microprocessor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 106–107.
[13] J. Tierno *et al.*, "A DPLL-based per core variable frequency clock generator for an eight-core POWER7 microprocessor," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2010, pp. 85–86.
[14] J. Tschanz *et al.*, "Adaptive frequency and biasing techniques for tolerance to dynamic temperature-voltage variations and aging," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 292–293.
[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] B. Zimmer *et al.*, "A RISC-V vector processor with tightly-integrated switched-capacitor DC-DC converters in 28 nm FDSOI," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C316–C317.
[17] J. Kwak and B. Nikolić, "A 550–2260 MHz self-adjustable clock generator in 28 nm FDSOI," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2015, pp. 1–4.
[18] P. Flatresse *et al.*, "Ultra-wide body-bias range LDPC decoder in 28 nm UTBB FDSOI technology," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 424–425.
[19] N. Planes *et al.*, "28 nm FDSOI technology platform for high-speed low-voltage digital applications," in *Proc. IEEE Symp. VLSI Technol.*, Jun. 2012, pp. 133–134.
[20] B. Zimmer *et al.*, "A RISC-V vector processor with simultaneous-switching switched-capacitor DC–DC converters in 28 nm FDSOI," *IEEE J. Solid-State Circuits*, vol. 51, no. 4, pp. 930–942, Apr. 2016.
[21] 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.

Jaehwa Kwak (S'10) received the B.S. and M.S. degrees in electrical engineering from Seoul National University, Seoul, Korea, in 2004 and 2006, respectively. He is currently working towards the Ph.D. degree in electrical engineering at the University of California, Berkeley, CA, USA. During his Ph.D. work, he was a design intern at Intel Corporation, Hillsboro, OR, USA, in summer 2011 and at Advanced Micro Devices, Sunnyvale, CA, USA, in summer 2013.

From 2006 to 2009, he was with GCT Research, Seoul, Korea, where he was a staff researcher and worked on designing digital controllers for wireless communication systems. His current research interests are in the area of energy-efficient microprocessor design, including self-adjustable clock generators and advanced synchronization circuits.

Borivoje Nikolić (S'93–M'99–SM'05) received the Dipl.Ing. and M.Sc. degrees in electrical engineering from the University of Belgrade, Serbia, in 1992 and 1994, respectively, and the Ph.D. degree from the University of California at Davis, CA, USA, in 1999.

In 1999, he joined the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, where he is now a National Semiconductor Distinguished Professor of Engineering. His research activities include digital, analog and RF integrated circuit design and VLSI implementation of communications and signal processing systems. He is coauthor of the book *Digital Integrated Circuits: A Design Perspective* (2nd ed., Prentice-Hall, 2003).

Dr. Nikolić received the NSF CAREER award in 2003, College of Engineering Best Doctoral Dissertation Prize and Anil K. Jain Prize for the Best Doctoral Dissertation in Electrical and Computer Engineering at University of California at Davis in 1999, and the City of Belgrade Award for the Best Diploma Thesis in 1992. For work with his students and colleagues, he has received best paper awards at the IEEE International Solid-State Circuits Conference, Symposium on VLSI Circuits, IEEE International SOI Conference, European Solid-State Device Research Conference, European Solid-State Circuits Conference, S3S Conference and the ACM/IEEE International Symposium of Low-Power Electronics. In 2014–2015, he was Distinguished Lecturer of the IEEE Solid-State Circuits Society.