

Active Mode Sub-Clock Power Gating

Jatin Mistry, James Myers, Bashir M. Al-Hashimi, *Fellow, IEEE*, David Flynn, *member, IEEE*, and John Biggs

Abstract—This paper presents a technique, called sub-clock power gating, for reducing leakage power during the active mode in low performance, energy constrained applications. The proposed technique achieves power reduction through two mechanisms: 1) power gating the combinational logic within the clock period (sub-clock) and 2) reducing the virtual supply to less than V_{th} rather than shutting down completely as is the case in conventional power gating. To achieve this reduced voltage, a pair of NMOS and PMOS transistors are used at the head and foot of the power gated logic for symmetric virtual rail clamping of the power and ground supplies. The sub-clock power gating technique has been validated by incorporating it with an ARM Cortex-M0 microprocessor which was fabricated in a 65nm process. Two sets of experiments are done: the first experimentally validates the functionality of the proposed technique in the fabricated test chip and the second investigates the utility of the proposed technique in example applications. Measured results from the fabricated chip show 27% power saving during the active mode for an example wireless sensor node application when compared to the same microprocessor without sub-clock power gating.

Index Terms—Power Gating, Leakage Control, Active Mode, Low Power, Embedded Microprocessors, Subthreshold

I. INTRODUCTION

Leakage power can be as dominant as dynamic power below 65nm and poses a large source of power consumption in digital circuits during the active mode [1] i.e. when the digital circuit is doing useful work. A number of techniques have been proposed for reducing active leakage power dissipation. These include dual threshold logic [2] which uses high threshold voltage logic gates on non-critical timing paths and adaptive body biasing [3] which raises or lowers the threshold voltage of transistors for active power management. The effectiveness of power gating to reduce leakage power has also been demonstrated during active mode. A finer granularity power gating has been proposed in [4], which involves disabling executional units during active mode. Similarly, a method of power gating part of a multiplier depending on the data width during run-time was proposed in [5]. Recent work has demonstrated the use of power gating on a granularity akin to clock gating. The use of the clock enable signal to power gate the integer execution core was shown in [6]. On the other hand, the use of the clock enable signals to power gate the fan-in logic of the clock-gated registers was shown in [7]. All these methods attempt to reduce power with minimal impact on maximum clock frequency.

In some current and emerging applications such as wireless sensor nodes and biomedical sensor applications, performance is not critical whereas power and energy is the primary design

goal. Active power reduction is primarily targeted through choice of low operating processor voltage and low clock frequency to avoid unnecessary dynamic power consumption. Examples of this principal are reported in [8] which uses a Texas Instruments MSP430 [9] at 1.8V and utilises its 32kHz mode of operation for device control consuming $300\mu\text{W}$. Similarly, an ASIC for wireless sensor nodes reported in [10] is operated at 1V and 10-100kHz. An ASIC for wireless monitoring of an Electrocardiography (ECG) signal that operates at 1.1V and 32kHz is reported in [11]. While the active leakage power reduction techniques described above in [4]-[7] are successful over multiple clock cycles, in applications where low performance is appropriate, leakage power dissipation becomes an increasing problem within the clock period due to idle combinational logic. As the clock frequency is reduced at a fixed V_{dd} , the clock period becomes longer than the evaluation time of the next state thereby increasing idle time of the combinational logic.

In this paper we propose a technique, called sub-clock power gating (SCPG), to capitalise on the increased idle time of combinational logic which results from low performance operation at a fixed V_{dd} . Power reduction is achieved by power gating within the clock period (sub-clock) to reduce leakage power during the active mode. This is unlike previous active mode power gating techniques where the power is shut down over multiple clock cycles [5], [6]. Rather than power gating completely, which we refer to as ‘shut down power gating’, the proposed technique provides a less than V_{th} voltage across the combinational logic (referred to in literature as sub-threshold) to minimise power mode transition energy overheads. To generate this voltage, symmetric virtual rail clamping of both the power (V_{dd}) and ground (V_{ss}) supplies is proposed and can be achieved through the use of both an NMOS and PMOS transistor at the head and foot of the power gated logic. The rest of this paper is organised as follows. Section II introduces the symmetric virtual rail clamping technique and compares the wake-up energy cost with two other power gating techniques. Section III shows how the proposed sub-clock power gating technique can minimise active leakage. Section IV and Section V describe the implementation and experimental validation of the technique. Conclusions are presented in Section VI.

II. WAKE-UP ENERGY COST OF DIFFERENT POWER GATING APPROACHES

Before discussing sub-clock power gating (SCPG), the energy cost of wake-up power mode transitions (moving from the sleep to active mode) is considered for three different power gating techniques: traditional shut down power gating [12], virtual rail clamping [13] and symmetric virtual rail

clamping, which is introduced in this paper. In shut down power gating the energy overhead of moving between power modes is dominated by the recharging of the virtual supply rail and glitching of internal signals resulting from the re-evaluation of logic cones [14]. Recharging of the virtual supply rail in power gating is a discernable energy overhead, but it has recently been reported that logic re-evaluation and glitching accounts for 75.7% of the wake-up energy [14]. These energy overheads stem from the supply rail being fully discharged and subsequent loss of valid logic gate outputs. In traditional applications of shut down power gating during standby mode [12], the wake-up energy associated with moving between power modes (E_{ovhd}) is generally negligible, however as the length of the idle period becomes shorter this overhead becomes comparable to the energy saved (E_{sav}). In applications of active-mode power gating with short idle periods, such as sub-clock power gating which is proposed in Section III, a technique which minimises E_{oh} is preferable to improve energy efficiency.

Virtual rail clamping has been proposed as a way to maintain a voltage across the power gated logic to retain register state [13] but maintaining a voltage across the power gated logic has been proven to reduce the recharge and glitching cost associated with power gating as valid logic outputs are maintained [15]. An inverter employing virtual rail clamping is demonstrated in Fig. 1(a) where a pair of NMOS and PMOS transistors are used at the head of the power gated logic to enable reduction in the supply voltage. In this circuit, the PMOS transistor, marked as *Sleep*, is a conventional power gating transistor and the NMOS transistor, marked as *Ret*, is used to clamp the virtual rail. When *Sleep* and *Ret* are logic 1 the VV_{dd} rail is reduced to $V_{dd} - V_{thn}$ where V_{thn} is the threshold voltage of the NMOS transistor. In this mode, the inverter's PMOS transistor can additionally benefit from reverse body biasing (RBB) by connecting its body to V_{dd} , as shown. The V_{thn} potential across V_{bs} increases the threshold voltage of the PMOS transistor due to the body effect and reduces sub-threshold leakage currents [6]. Using MOS transistors to implement the virtual rail clamping also has the added advantage of being able to achieve shut down power gating by forcing *Sleep* to logic 1 and *Ret* to logic 0.

Virtual rail clamping enables a single threshold voltage drop reduction across the power gated logic, but to maximise leakage power savings of the power gated logic, it is desirable to reduce the clamped voltage by more than a single threshold voltage, since power is the product of voltage and current [16]. Multiple NMOS transistors placed in series at the head of the power gated logic can enable this, however, we chose to mirror the V_{dd} virtual rail clamping on the V_{ss} supply rail. This proposed symmetric virtual rail clamping technique is shown in Fig. 1(b), where there is now a pair of NMOS and PMOS transistors at the head and foot of the example inverter circuit. When *Sleep* and *Ret* are logic 1 (and thus *nSleep* and *nRet* are logic 0) the VV_{dd} is clamped to $V_{dd} - V_{thn}$ and the VV_{ss} is clamped to $V_{ss} + V_{thp}$. The result is a much more aggressive reduction in voltage across the power gated logic but also has three advantages over single rail clamping [13]. Firstly the charge that is stored in the VV_{dd} supply rail is

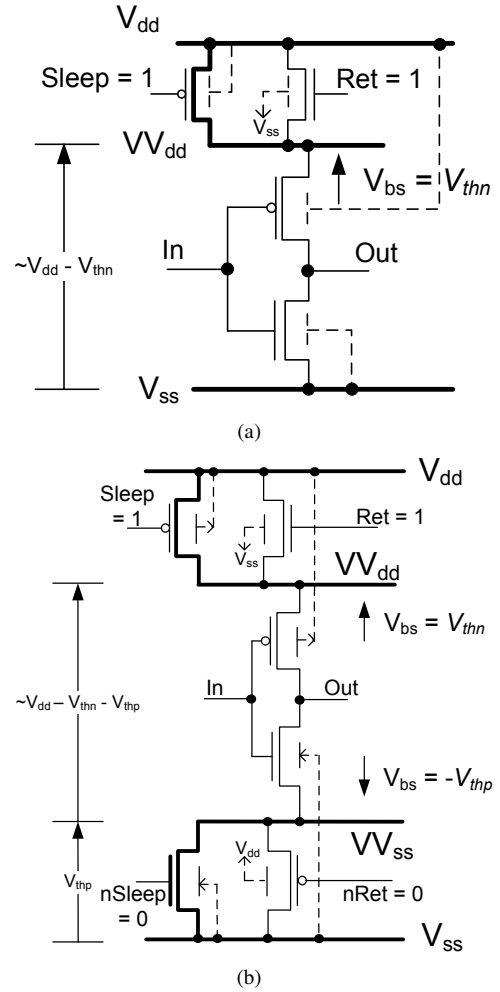


Fig. 1. An inverter with (a) single virtual rail clamping [13] (b) proposed symmetrical virtual rail clamping

recycled to charge up the VV_{ss} supply rail in the sleep mode [17] achieving greater reduction in supply voltage in the same time frame at lower energy cost to single virtual rail clamping. This is shown in Fig. 2(a) from simulation on a 101 stage ring oscillator, where 30% greater reduction in V_{ds} is observed for symmetric virtual rail clamping in the same time frame and is beneficial over short power gated periods. Secondly, both the body of the PMOS and NMOS transistors in the power gated circuit can be connected to the true supplies achieving a V_{th} reverse body bias on all the transistors in the circuit, reducing leakage further. Thirdly, the symmetric RBB ensures better equality of NMOS and PMOS drive strength degradation, as strong RBB on the PMOS transistor from single rail clamping, Fig. 1(a), can result in logic 1 drive. This is because, when a gate's supply voltage is lowered and threshold voltage is increased the I_{on} current of the transistors degrades and can become comparable with the I_{off} current resulting in a battle between the on/off transistors to maintain the correct output voltage [18]. It is known that NOR and NAND gates suffer the greatest effects of this because of large numbers of parallel transistors in the logic gates [18]. Consequently, to observe the benefit of symmetric clamping over single rail clamping, a 4-input NOR gate (NOR4) from a 90nm gate library was




Fig. 2. Symmetric virtual rail clamping effect on (a) V_{ds} reduction speed in 101 stage ring oscillator (b) minimum clamped voltage 4-input NOR gate

simulated with both techniques. Fig. 2(b) shows the percentage deviation of a logic 1 output with respect to clamped V_{ds} in both cases. As can be seen, the reverse body biasing used on only the PMOS transistors in virtual rail clamping causes the NOR4 gate output to sharply deviate at supply voltages below approximately 300mV. Conversely, the symmetric virtual rail clamping enables the NOR4 gate to hold its output for supply voltages down to approximately 200mV. In addition to the clamped voltage, independent control of *Sleep*, *Ret*, *nSleep* and *nRet* in Fig. 1(b) can be used to achieve full shut down power gating.

To quantify the wake-up energy cost of shut down power gating, single rail clamping and our symmetric virtual rail clamping, the three approaches have been implemented on a 101 stage ring oscillator using a 90nm library. In line with the results presented in Fig. 2(b), a 300mV clamped voltage was chosen for virtual rail clamping achieved with a high threshold voltage NMOS clamping transistor and a 200mV clamped voltage was chosen for symmetric virtual rail clamping achieved with low threshold voltage PMOS and NMOS clamping transistors. All three circuits were simulated with 0.6V V_{dd} using HSpice. Table I shows the wake-up energy, sleep mode leakage current saving and wake-up time for the three power gating approaches. As can be seen, the proposed symmetric virtual rail clamping has the lowest wake-up energy and is 3x lower than shut down power gating.

TABLE I
RING OSCILLATOR WAKE-UP ENERGY, STANDBY LEAKAGE AND WAKE-UP TIME

	Wake-up Energy (fJ)	Leakage Saving (%)	Wake-up time (ns)
Shut Down Power Gating [12]	223.0	87.3	12
Single Rail Clamping [13]	76.82	75.4	6.5
Proposed Symmetric Virtual Rail Clamping	74.49	78.4	6.5

This is because the voltage maintained across the power gated logic from equal reduction in both $V_{V_{dd}}$ and $V_{V_{ss}}$ supply rails eliminates signal glitching from the logic re-evaluation present in shut down power gating. Furthermore, despite using a lower clamped voltage than single virtual rail clamping, the charge recycling in symmetric virtual rail clamping results in lower wake-up energy cost. As expected, Table I shows that standby leakage saving is highest in shut down power gating because the power supply is fully disconnected achieving greater reduction in V_{ds} . However, the proposed symmetric virtual rail clamping has a greater standby leakage saving compared to single rail clamping and can be attributed to exploitation of reverse body biasing of both NMOS and PMOS transistors and the lower achievable V_{ds} . Finally, Table I shows that the proposed symmetric virtual rail clamping has a shorter wake-up time compared to shut down power gating, which permits a longer power gated period and is particularly useful over short power gated periods.

III. PROPOSED SCPG TECHNIQUE

In some recent and emerging applications such as wireless sensor networks and biomedical sensors, performance is not of concern whereas energy is a primary concern. In these applications a low clock frequency is utilised with fixed V_{dd} due to low performance demands [8], [10], [11]. As the clock frequency of the digital circuit is reduced the clock period becomes longer than the combined hold time (T_{hold}), evaluation time of the next state (T_{eval}) and the setup time (T_{setup}) resulting in idle time (T_{idle}) within the clock period, Fig. 3. This idle time presents an opportunity to power gate within the clock period to reduce active mode leakage.

The proposed sub-clock power gating technique is shown in Fig. 4 and has three distinct parts. Firstly, the design is split into 2 domains: power-gatable combinational logic, marked as 'Comb. Logic' and separate always-on sequential logic, marked as 'Seq. Logic'. This split is made to avoid the need for state retention registers to store state in sleep mode, which would increase area by 20-50% and time taken to change between the sleep and active modes [12]. Secondly, symmetric virtual rail clamping, described in Section II, is used to power gate the combinational logic because of its low wake-up energy cost. The separation of the sequential logic from the combinational logic has a further benefit as the aggressive voltage reduction attainable with symmetric virtual rail clamping can be used for the combinational logic whilst ensuring that saved state remains intact during active operation, as will be shown in Section V. The third distinguishing feature is the isolation logic between combinational




Fig. 3. Idle time within the clock period resulting from aggressive frequency scaling

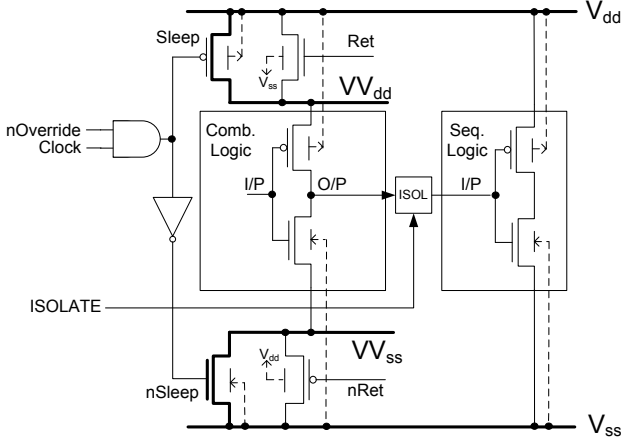


Fig. 4. Proposed sub-clock power gating technique with symmetric virtual rail clamping




Fig. 5. An example of isolation control circuit

and sequential domains, shown as ‘ISOL’ in Fig. 4. This is required to ensure that the output signals of the combinational domain do not cause short-circuit currents in the sequential logic when the combinational domain is powered down.

In traditional power gating schemes, the control to the power gates is usually driven by a power gating controller state machine [12]. However, this is impractical in the proposed SCPG technique since the control needs to be issued within the clock period. The proposed technique instead uses the clock signal, as shown in Fig. 4. The NMOS and PMOS transistors at the head of the combinational logic (Sleep & Ret) use the normal clock signal whilst the NMOS and PMOS transistors at the foot (nSleep & nRet) use the inverse of the clock signal. Therefore, when the clock is high, the combinational logic is clamped to less than V_{th} by the symmetric virtual rail clamping and when the clock is low it is restored to V_{dd} . The energy saved from using the proposed technique is therefore proportional to the length of time the clock is held high, and so, to maximise the power saving achievable with SCPG, it is possible to change the clock duty cycle to extend the high phase of the clock, as will be demonstrated in Section V. The $nOverride$ control signal shown, Fig. 4, provides a method to disable the SCPG technique to achieve normal timing.




Fig. 6. Timing diagram of sub-clock power gating

Also in traditional power gating, a power gating controller would be used to control the output isolation [12]. In the proposed SCPG technique though, the circuit of Fig. 5 is used to drive the ISOLATE signal to the ‘ISOL’ block in Fig. 4. The circuit has two inputs: the clock signal and the value of the combinational logic VV_{dd} which is derived from a TIEHI logic gate. When the clock is logic 1, ISOLATE is driven to a logic 1, thereby isolating the combinational outputs. When the clock is logic 0, ISOLATE is held at logic 1 while the VV_{dd} input remains at logic 0 (clamped). This ensures the combinational outputs remain isolated until the supply rail is charged to an equivalent logic 1, eliminating short-circuit currents during wake-up.

The timing diagram of the combinational logic in the sub-clock power gating mode of operation with symmetric virtual rail clamping is shown in Fig. 6. After the state is captured into the positive edge triggered registers, the voltage to the combinational logic is clamped to less than V_{th} but the amount of time taken for the virtual rails to discharge ensures register hold times (T_{hold}) will be met. At this point, the output isolation is also enforced. The virtual supply rails are held at the clamped voltages for the remainder of the high phase of the clock (T_{pglow}) minimising leakage power dissipation, and the outputs of the combinational domain remain isolated ($T_{isolate}$). Note that by changing the duty cycle of the clock it is possible to extend this off period (high phase of clock), maximising the leakage power savings. The virtual supply rails are restored following the negative edge of the clock but the output isolation is held until the virtual supply rails are fully restored ($T_{pgstart}$). The remainder of the clock period is used for the evaluation of the next state (T_{eval}) and ensuring setup time (T_{setup}) is met before the process repeats in the next clock period.

IV. IMPLEMENTATION

The proposed SCPG technique has been proven by incorporating it with an ARM Cortex-M0 microprocessor made available to us as a soft IP core from our industrial project partner and fabricated in a 65nm process. The Cortex-M0 is a 32-bit RISC microprocessor with a three-stage pipeline and is chosen because of its relevance to low performance, energy constrained applications. The design flow for augmenting a digital circuit with the sub-clock power gating technique is shown in Fig. 7; three additional steps are added to a traditional

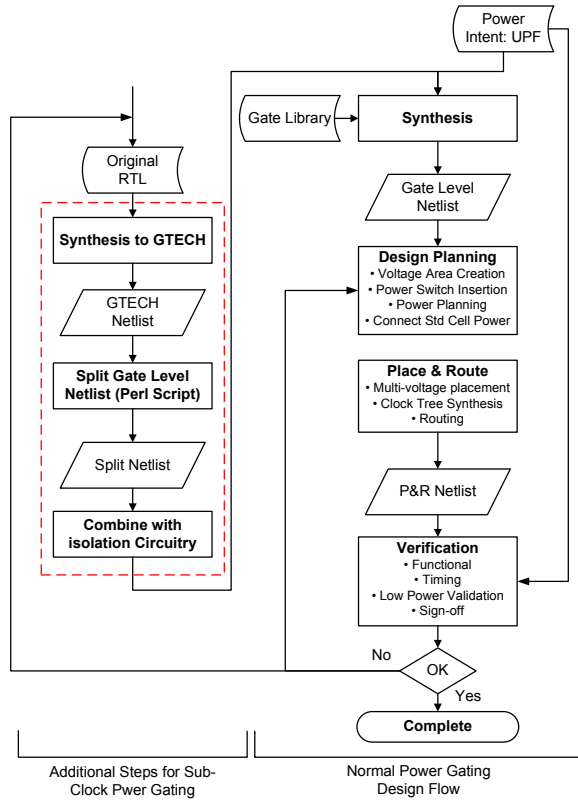


Fig. 7. Design flow of sub-clock power gating

power gating design flow. We assume the use of the IEEE 1801 Unified Power Format (UPF), a leading power intent standard for defining the strategy of a multi-voltage or power gated design and the Synopsys EDA tool suite is used. To achieve the power domain split shown in Fig. 4, the RTL must be written with separate combinational and sequential logic Verilog modules for compatibility with UPF. The first two additional steps in Fig. 7 are used to achieve this split. Firstly the RTL is synthesised to a generic gate library, such as the GTECH library in the Synopsys tools, to give a gate level representation of the circuit, and secondly a Perl script is used to identify and split combinational and sequential logic into two separate Verilog modules. The final additional step, wraps the new Verilog modules together with the isolation control circuit (Fig. 5) and the power gate control statements before being synthesised with a traditional power gating design flow and target gate library.

The Cortex-M0 together with the proposed SCPG technique was part of a 2x2mm system on chip (SoC) which is shown to the right of Fig. 8. The entire Cortex-M0 microprocessor, marked as CM0 in Fig. 8, has its own power supply in the SoC to allow power measurement, and analog pads are included for observation of the virtual supply rails. The rest of the SoC was made up of SRAM for instruction and data, an ASCII Debug Protocol (ADP) to facilitate communication and control of the chip via USB and a clock modulator circuit. As mentioned in Section III, the high phase of the clock can be changed to capitalise on all the idle time of the combinational logic to maximise power savings with the proposed SCPG technique.




Fig. 9. Clock modulation circuit to convert system clock to Cortex-M0 clock with optimised duty cycle

The circuit in Fig. 9 is used to achieve this modulation of clock duty cycle. An external clock is fed into the modulator and is divided down to a period of $(1 + n) \cdot T_{clk}$, where n can be programmed to values up to 2^{32} . The resulting output clock from the modulator is low for T_{clk} and high for $n \cdot T_{clk}$ as shown in Fig. 9.

Shown to the left of Fig. 8 is the layout of the ARM Cortex-M0 with the proposed SCPG technique. The final layout was $200 \times 310 \mu m$, 7% of the total area can be attributed to the inclusion of the proposed SCPG technique. As mentioned in Section III, the proposed technique requires separation of combinational and sequential logic and Fig. 8 shows this separation with the combinational logic located in the middle of the layout and the sequential logic on its periphery. Note that these two areas of the layout directly correspond to the ‘Comb. Logic’ and ‘Seq. Logic’ blocks in Fig. 4. Highlighted on the boundary between the combinational and sequential areas of the layout are the isolation cells which corresponds to the ‘ISOL’ block from Fig. 4. The pairs of NMOS and PMOS transistors in Fig. 4, used for symmetric virtual rail clamping of the combinational logic, were placed in a grid pattern throughout the combinational logic and are labeled in Fig. 8. The power gating transistors were sized using the EDA tools for 5% IR drop. The minimum clamped voltage on the other hand was estimated through HSpice simulation. A logic path of 40 NOR3 and NAND4 gates were used to represent a worst case critical path in the Cortex-M0. It is found that a correct output value of the entire path is held at supply voltages down to 160mV with less than 6% deviation. With the 65nm library used, HSpice simulation showed 8 regular threshold voltage transistors of width $3.6 \mu m$ enabled a clamped voltage of 180mV. At the top of the layout is an always-on region which accommodates the power gating and isolation control circuitry. For experimental comparisons, three modes of operation were implemented in the Cortex-M0: the proposed SCPG with symmetric virtual rail clamping, no power gating enforced using the $nOverride$ signal in Fig. 4, and SCPG with complete shut down [19] achieved by disabling the Ret & nRet transistors shown in Fig. 4.

V. EXPERIMENTAL VALIDATION

Two sets of experiments were carried out to demonstrate the proposed sub-clock power gating and symmetric virtual rail clamping techniques. The first experimentally validates the functionality of the sub-clock power gating and symmetric virtual rail clamping techniques. The second shows the utility of the sub-clock power gating technique in example applications and compares traditional shut-down power gating with symmetric virtual rail clamping. All experiments on the

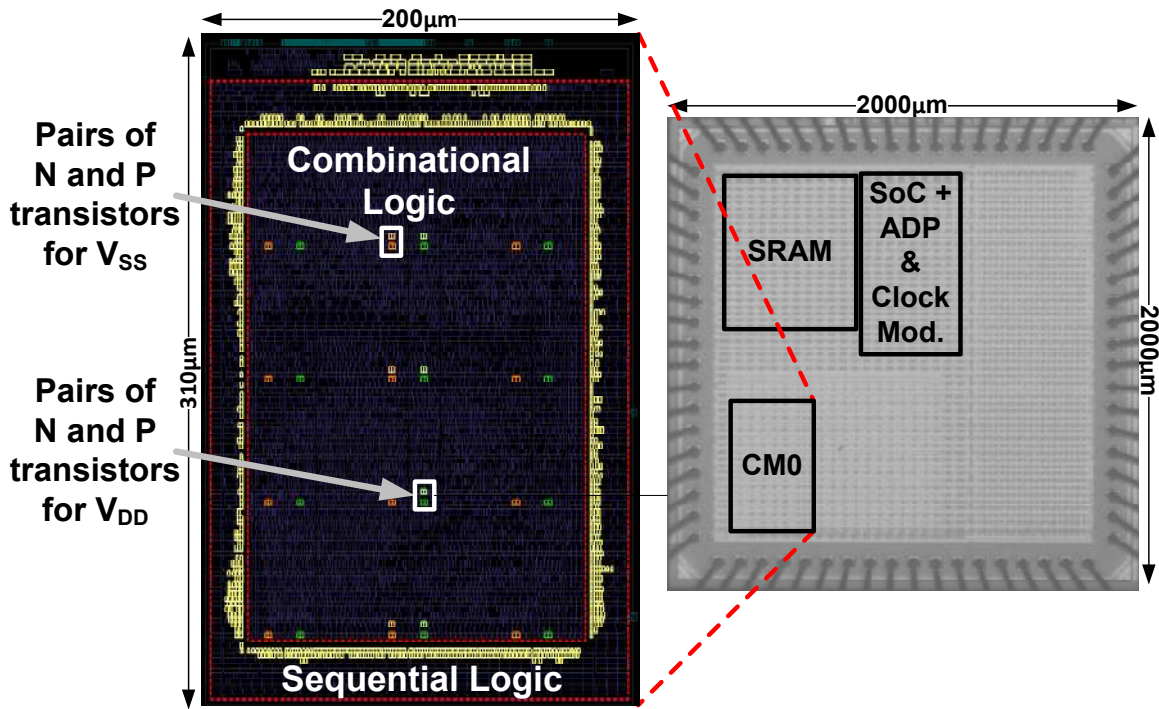


Fig. 8. Complete floorplan of SCPG Cortex-M0 microprocessor (left) and test chip die photograph (right)




Fig. 10. Testboard for experimental measurement

fabricated chip were performed using the testboard shown in Fig. 10. In line with the scaled voltage typically found in processors designed for the target applications [10], a 0.7V external power supply is used for the Cortex-M0's independent V_{dd} . To emphasize the negative impact of leakage on the microprocessor a temperature of 90°C is used. An ammeter with 10nA resolution is connected in series with the power supply to allow current measurement of the microprocessor. A USB interface is used to download the test program and set up the test chip. The analog pads included in the test chip are connected to the board's test points shown in Fig. 10, for observation of the combinational logic virtual rails.

A. Test Chip Validation

Fig. 11 shows an oscilloscope trace of the VV_{dd}/VV_{ss} supply rails when using the proposed sub-clock power gating technique with symmetric virtual rail clamping. The clock used in this trace is 8kHz with 2:1 (high:low) duty cycle. Over the first part of the clock period (T_{clk}) the VV_{dd} and VV_{ss} rails are clamped to 450mV and 270mV respectively, aggressively reducing the combinational voltage to 180mV (T_{pglow}). Over the second part of the clock period, the rails are restored returning the combinational logic to the full 0.7V supply voltage. While the combinational logic supply voltage is clamped down to below the threshold voltage it is not operated at this voltage and so if some logic gate outputs were to flip whilst clamped, then this would not be an issue as they would be rectified when the supply is returned to its nominal value. As discussed in Section IV, the duty cycle of the clock can be modulated to maximise the leakage power saving of the sub-clock power gating technique. The effect of duty cycle on power saving when using sub-clock power gating with symmetric virtual rail clamping is shown in Fig. 12. The clock frequency used in these measurements is 10kHz and the power values are normalised to the Cortex-M0 operating at 0.7V with no power gating. The high phase of the clock period increases from left to right in Fig. 12 and as can be seen, the power goes down (savings increase) as it does so. It is interesting to note that the normalized power reduces but a lower bound is slowly reached. This is because the combinational logic has some finite leakage in the clamped state and the registers and other always-on logic also remain active. The charge up time of the virtual rails in symmetric virtual rail clamping is measured to be 45ns from the oscilloscope trace. Taking into account the critical path length found through static timing analysis

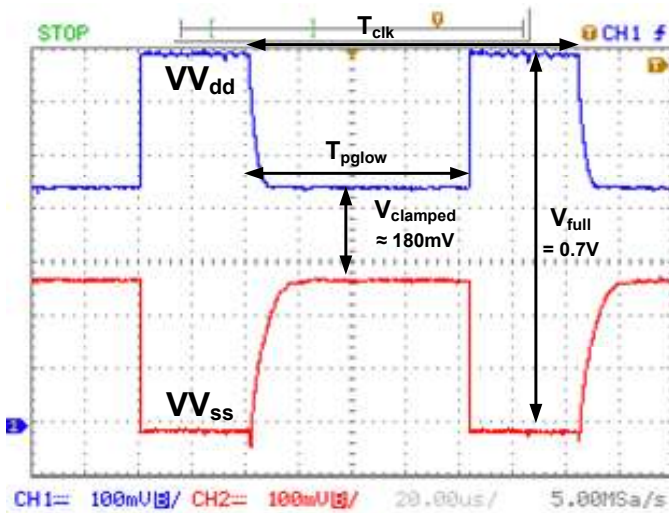


Fig. 11. Measured VV_{dd} and VV_{ss} behaviour with sub-clock power gating operation and symmetric virtual rail clamping

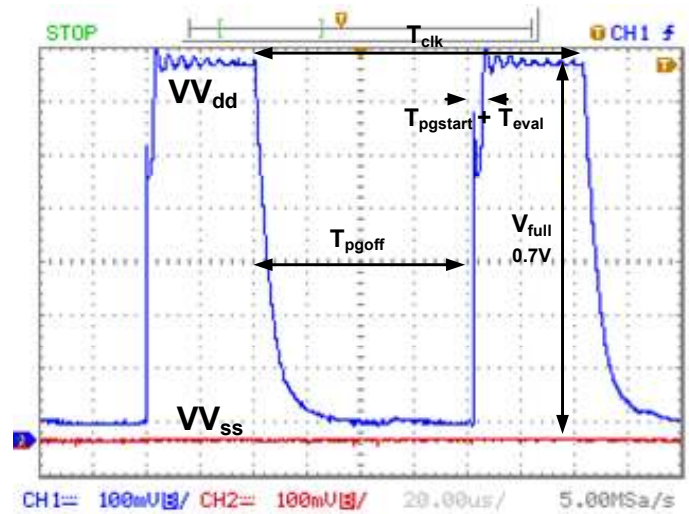


Fig. 13. Measured VV_{dd} and VV_{ss} behaviour in sub-clock power gating with shut down power gating [19]




Fig. 12. Normalized measured power of ARM Cortex-M0 microprocessor with 10kHz clock at varying duty cycle in sub-clock power gating with symmetric virtual rail clamping, $V_{dd}=0.7V$

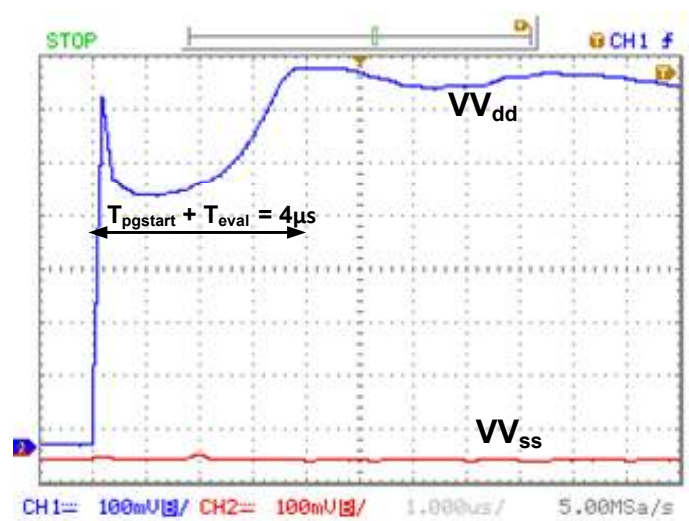


Fig. 14. Measured VV_{dd} charge-up and evaluation time in sub-clock power gating with shut down power gating [19]

(75ns), the minimum allowable low period of the clock is determined to be 200ns to allow enough timing margin for the rails to recharge and the combinational logic to evaluate the next state whilst avoiding timing errors. Using the clock modulator circuit (Fig. 9), a 200ns low period corresponds to an external clock frequency of 5MHz while n can be programmed as necessary to vary the clock frequency.

Fig. 13 shows the behaviour of the virtual supply rails when using SCPG with shut down power gating. In this trace an 8kHz clock with 2:1 duty cycle is used. Unlike symmetric virtual rail clamping, that was shown in Fig. 11, the VV_{ss} rail is unclamped and the VV_{dd} is fully discharged in the first part of the clock period (T_{pgoff}). In the second part of the clock period the VV_{dd} rail is restored to the full 0.7V supply. Notice that the charge-up ($T_{pgstart}$) and evaluation time (T_{eval}) of the combinational logic is clearly visible in Fig. 13 and has been expanded in Fig. 14. The droop seen in the VV_{dd} rail can be attributed to the current demanded by the high volume of signal glitching that occurs as the

combinational logic is brought out of shut down and re-evaluates which opposes the recharging of VV_{dd} [12]. This droop subsequently slows the combinational re-evaluation, exacerbating the length of T_{eval} to $4\mu s$ as shown. This is unlike the symmetric virtual rail clamping shown in Fig. 11 where the voltage maintained across the combinational logic helps to eliminate signal glitching during charge up and avoid a virtual rail droop, allowing the combinational logic to be charged in 45ns. The consequence of this increased wake-up time when using shut down power gating is the need for a clock with a low phase of at least 4000ns to ensure correct operation, as opposed to a 200ns low phase achievable with symmetric virtual rail clamping.

Experimental measurement from the chip shows that when the Cortex-M0 is fully powered but the clocks are stopped, the leakage power dissipation is $7.51\mu W$. On the other hand, when the combinational logic is fully shut down, power dissipation is $1.46\mu W$, representing an 80.6% reduction in

power. Alternatively, when the combinational logic supply is clamped using symmetric virtual rail clamping the power dissipation is $2.44\mu\text{W}$, a 67.5% reduction in leakage power. This is to be expected, since shut down power gating completely disconnects the supply whereas symmetric virtual rail clamping maintains a voltage across the combinational logic and matches with the trends shown in Table I.

B. Applications

Next we compare the power consumption of the Cortex-M0 with and without the proposed sub-clock power gating with symmetric virtual rail clamping technique over a range of clock frequencies. To investigate the utility of SCPG in one of the target applications, we use a program used in an actual wireless sensor node for the ‘Next Generation Energy Harvesting Electronics’ project which tracks the vibration frequency to tune a vibrational energy harvester to maintain resonance (between 42Hz and 55Hz) [20]. The program takes a set of 1000 samples from an accelerometer to calculate the current frequency of vibration which is then used to set a stepper position on an energy harvester. A clock duty cycle with 200ns low phase is used in the sub-clock power gating with symmetric virtual rail clamping and the measured results across five test chips are presented in Fig. 15. As can be seen, the proposed SCPG technique achieves lower power consumption at all frequency points up to a clock frequency of just over 400kHz. At all of these frequency points, the energy saved (E_{sav}) from using the proposed SCPG technique exceeds the energy overhead (E_{oh}) of power gating resulting in the savings seen. However, as clock frequency increases, E_{sav} reduces because of the shorter combinational idle time and eventually becomes comparable to E_{oh} resulting in the convergence point around 400kHz in Fig. 15. At clock frequencies above 400kHz, $E_{oh} > E_{sav}$ and the power consumed by the Cortex-M0 when using SCPG exceeds that of the Cortex-M0 without power gating. In the intended applications of sub-clock power gating, if clock frequencies above and below 400kHz are required, the processor could be switched to no power gating mode by using the $nOverride$ signal (Fig. 4) for clock frequencies above 400kHz.

Five test chips were used for the data shown in Fig. 15 to compare results across multiple dies, and as can be seen, the measurements all follow the same trend. The spread between plotted points can be explained by die to die process variation. The average power and energy per operation of the five test chips is shown in Table II. In the final column the percentage saving achieved when using the proposed technique is stated. As can be seen, the proposed technique saves up to 67% of the energy compared to without power gating and demonstrates sub-clock power gating’s ability to improve energy efficiency for a circuit operating at low clock frequencies. At 455kHz, the processor would need to switch to no power gating with the $nOverride$ signal to remain in the lowest energy mode of operation. In the real wireless sensor node application, the accelerometer is sampled at a frequency of 2kHz to improve accuracy in the frequency calculation. As the program loops around a maximum of 85 instructions, at a sampling rate of

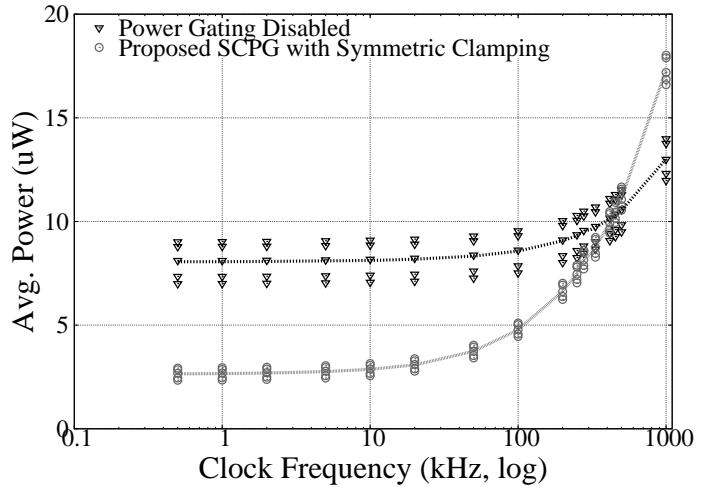


Fig. 15. Measured average power of ARM Cortex-M0 microprocessor at varying clock frequency in tuning program, $V_{dd}=0.7\text{V}$

TABLE II
AVERAGE MEASURED POWER AND ENERGY OVER FIVE TEST CHIPS WITH POWER GATING DISABLED (NO-PG) & SUB-CLOCK POWER GATING (SCPG) ENERGY HARVESTER TUNING

Clock Freq. (kHz)	No Power Gating		Proposed SCPG		
	Power (uW)	Energy (pJ)	Power (uW)	Energy (pJ)	Saving (%)
0.5	8.06	16117	2.65	5300	67.11
1	8.06	8061	2.66	2660	66.99
2	8.07	4035	2.69	1342	66.72
5	8.10	1619	2.76	551.0	65.97
10	8.12	812.4	2.87	286.6	64.72
20	8.18	408.8	3.08	154.2	62.27
50	8.33	166.5	3.74	74.71	55.14
100	8.57	85.74	4.78	47.84	44.20
200	9.07	45.33	6.64	33.20	26.77
250	9.31	37.26	7.48	29.91	19.72
312.5	9.62	30.78	8.45	27.05	12.14
384.6	9.98	25.94	9.52	24.74	4.60
416.6	10.13	24.32	9.97	23.94	1.59
454.5	10.32	22.71	10.50	23.10	-1.74
500	10.54	21.09	11.12	22.24	-5.44
1000	13.01	13.01	17.31	17.31	-33.06

2kHz, the Cortex-M0 can be operated at 200kHz without missing a new sample. At 200kHz, without sub-clock power gating the processor would dissipate $9\mu\text{W}$ consuming 45pJ/operation and with sub-clock power gating the processor dissipates $6.6\mu\text{W}$ consuming 33.20pJ/operation. This represents a 27% reduction in power and 1.4x improvement in energy efficiency.

Table I showed that the wake-up energy associated with shut down power gating was higher than the proposed symmetric virtual rail clamping circuit through ring oscillator simulations. To compare sub-clock power gating with shut down power gating [19] and the proposed symmetric virtual rail clamping, both techniques have been compared across a range of clock frequencies. Fig. 16 compares graphically, the measured average power of sub-clock power gating using shut down power gating, sub-clock power gating using the proposed symmetric virtual rail clamping technique, and power gating disabled. To permit direct comparison of the two power gating techniques, both sub-clock power gating modes used a clock with a $4\mu\text{s}$ low phase achieved with a 250kHz external clock (Fig. 9). As can be seen, at 500Hz and 1kHz in Fig. 16, SCPG with shut down power gating has lower power consumption




Fig. 16. Measured power of Cortex-M0 with power gating disabled, proposed SCPG with symmetric virtual rail clamping and SCPG with shut down power gating [19]

than without power gating but is higher than the proposed symmetric virtual rail clamping. This can be attributed to the high wake-up energy cost associated with the signal glitching that occurs when restoring the virtual rail in shut down power gating. Note also, this high wake-up energy cost causes the energy overhead of shut down power gating to exceed the energy saving at all frequency points above 1kHz resulting in higher power consumption in comparison to no power gating. The increasing power trend of the shut down power gating mode is reversed after 20kHz because the virtual rail does not discharge fully during shut down due to the shorter idle time within the clock period. However, despite the V_{dd} rail remaining partially charged at these frequency points it still dissipates more power than the proposed symmetric virtual rail clamping technique and is a result of the combination of asymmetric reverse body biasing of the logic gates and lack of charge recycling discussed in Section II. The observations seen here provide further validation for the symmetric virtual rail clamping proposed in Section II to reduce wake-up energy cost and improve the energy efficiency of the sub-clock power gating technique.

We envisage the proposed sub-clock power gating with symmetric virtual rail clamping technique being applicable in a range of general purpose, low performance, energy constrained applications. Therefore we have used the Dhrystone benchmark [21], to validate the power saving trends of sub-clock power gating with symmetric virtual rail clamping in a second test program. The Dhrystone benchmark program is chosen since it uses a combination of integer arithmetic functions, logic decisions and memory accesses which is representative of the data acquisition and manipulation of many general purpose applications [21]. The measured power across the five test chips used with the wireless sensor node example when executing the Dhrystone benchmark is shown in Fig. 17. As can be seen a similar trend to the energy harvester tuning program can be observed with sub-clock power gating with symmetric virtual rail clamping showing power saving over no power gating up to a clock frequency of 400kHz.

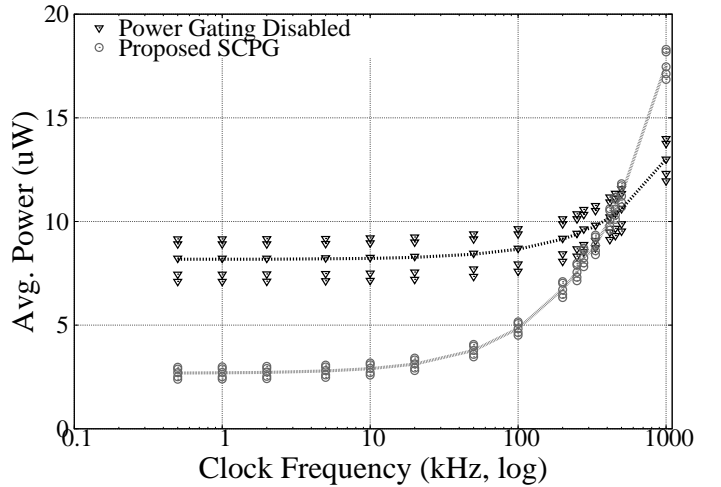


Fig. 17. Measured average power of ARM Cortex-M0 microprocessor at varying clock frequency in Dhrystone, $V_{dd}=0.7V$

VI. CONCLUSION

This paper has proposed a power gating technique that reduces leakage power during the active mode for low performance energy constrained applications by power gating combinational logic within the clock period. Rather than shutting down completely, symmetric virtual rail clamping was proposed to reduce wake-up power mode transition energy cost. The proposed sub-clock power gating with symmetric virtual rail clamping technique has been demonstrated with an ARM Cortex-M0 microprocessor, fabricated in 65nm technology. Measured results up to a clock frequency of 400kHz from the fabricated chip showed that it is possible to reduce average power by up to 67% during the active mode, with only 7% of the layout area accountable to extra circuitry. Using an actual wireless sensor node program example, it was shown that the microprocessor with sub-clock power gating can achieve 27% reduction in power and 1.4x improvement in energy efficiency. Comparison between sub-clock power gating with shut down power gating and symmetric virtual rail clamping provided experimental validation for the need to use symmetric virtual rail clamping to improve energy efficiency.

The work proposed in this paper can be considered as an orthogonal approach to the recently proposed subthreshold technique for maximising energy efficiency when operating at low performance. The subthreshold technique enables realization of minimum energy computation by scaling the supply voltage below V_{th} until a minimum energy point is found where dynamic energy equals leakage energy per operation [18]. Due to the aggressively scaled supply voltage, the technique comes at a cost of performance making it suitable for low performance, energy constrained applications. Sub-clock power gating, on the other hand, provides a power/performance trade-off allowing the digital circuit to toggle between low power, low performance and high power, high performance states unlike sub-threshold which is optimised for low performance only. Additionally, since sub-clock power gating is used with limited voltage scaling, it avoids increased design complexity making it fully compatible with standard EDA tools and gate

libraries, and avoids sensitivity to supply voltage and threshold voltage variation associated with operating below V_{th} [18].

REFERENCES

- [1] A. Agarwal, S. Mukhopadhyay, A. Raychowdhury, K. Roy, and C. Kim, "Leakage Power Analysis and Reduction for Nanoscale Circuits," *IEEE Micro*, vol. 26, 2006.
- [2] L. Wei, Z. Chen, K. Roy, M. Johnson, Y. Ye, and V. De, "Design and Optimization of Dual-Threshold Circuits for Low-Voltage Low-Power Applications," *IEEE Transactions On Very Large Scale Integration (VLSI) Systems*, vol. 7, pp. 16–24, 1999.
- [3] N. Mehta and B. Amrutur, "Dynamic supply and threshold voltage scaling for cmos digital circuits using in-situ power monitor," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 20, no. 5, pp. 892–901, 2012.
- [4] Z. Hu and A. Buyuktosunoglu and V. Srinivasan and V. Zyuban and H. Jacobson and P. Bose, "Microarchitectural Techniques for Power Gating of Execution Units," in *International Symposium on Low Power Electronics and Design*, 2004.
- [5] K. Usami and M. Nakata and T. Shirai and S. Takeda and N. Seki and H. Amano and H. Nakamura, "Implementation and Evaluation of Fine-Grain Run-Time Power Gating For A Multiplier," in *International Conference on IC Design and Technology*, 2009.
- [6] J. W. Tschanz, S. G. Narendra, Y. Ye, B. A. Bloechel, S. Borkar, and V. De, "Dynamic Sleep Transistor and Body Bias for Active Leakage Power Control of Microprocessors," *IEEE Journal Of Solid-State Circuits*, vol. 38, pp. 1838–1845, 2003.
- [7] J. Seomun, I. Shin, and Y. Shin, "Synthesis of Active-Mode Power-Gating Circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, 2012.
- [8] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi, "Hardware Design Experiences in ZebraNet," in *International Conference on Embedded Networked Sensor Systems*, 2004.
- [9] Texas Instruments, *MSP430 User's Guide*. TI, 2009.
- [10] B.A. Warneke and K.S.J. Pister, "An Ultra-Low Energy Microcontroller for Smart Dust Wireless Sensor Networks," in *IEEE International Solid-State Circuits Conference*, 2004.
- [11] X. Liu, Y. Zheng, M. W. Phyu, F. N. Endru, V. Navaneethan, and B. Zhao, "An Ultra-Low Power ECG Acquisition and Monitoring ASIC System for WBAN Applications," *IEEE Journal On Emerging and Selected Topics in Circuits and Systems*, vol. 2, 2012.
- [12] M. Keating, D. Flynn, R. Aitken, A. Gibbons, and K. Shi, *Low Power Methodology Manual*. Springer, 2007.
- [13] S. Kim, S. V. Kosonocky, D. R. Knebel, K. Stawiasz, and M. C. Papaefthymiou, "A Multi-Mode Power Gating Structure for Low-Voltage Deep-Submicron CMOS ICs," *IEEE Transactions On Circuits and Systems-II:Express Briefs*, vol. 54, pp. 586–590, 2007.
- [14] D. Juan and Y. Chen and M. Lee and S. Chang, "An Efficient Wake-Up Strategy Considering Spurious Glitches Phenomenon for Power Gating Designs," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 18, 2010.
- [15] H. Singh, K. Agarwal, D. Sylvester, and K. Nowka, "Enhanced leakage reduction techniques using intermediate strength power gating," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 15, no. 11, pp. 1215–1224, 2007.
- [16] N. Weste and D. Harris, *CMOS VLSI Design: A Circuits and Systems Perspective*. Pearson Education, 2005.
- [17] E. Pakbaznia, F. Fallah, and M. Pedram, "Charge Recycling in Power-Gated CMOS Circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 27, 2008.
- [18] S. Hanson, B. Zhai, K. Bernstein, D. Blaauw, A. Bryant, L. Chang, K. Das, W. Haensch, E. Nowak, and D. Sylvester, "Ultralow-Voltage Minimum-Energy CMOS," *IBM Journal of Research and Development*, vol. 50, 2006.
- [19] J. N. Mistry and B. M. Al-Hashimi and D. Flynn and S. Hill, "Sub-Clock Power-Gating Technique for Minimising Leakage Power During Active Mode," in *Design Automation and Test in Europe (D.A.T.E)*, 2011.
- [20] A. Weddell, D. Zhu, G. V. Merrett, S. P. Beeby, and B. M. Al-Hashimi, "A Practical Self-Powered Sensor System with a Tunable Vibration Energy Harvester," in *International Workshop on Micro- and Nano-Technology for Power Generation and Energy Conversion Applications 2012*, 2012.
- [21] R. York, "Benchmarking in Context: Dhrystone," 2002.