Design Solutions for a Low-Power SoC Platform Using Near-Threshold Voltages

Mika Kutila, Marko Ylitolva and Jonas Eriksson
University of Turku, Technology Research Center (TRC)
20014 Turun yliopisto, Finland
{mika.kutila, marko.ylitolva, jonas.eriksson}@utu.fi

***Abstract*—In the future the number of small sensor-based devices is going to increase. The Internet-of-Everything (IoE) vision brings smart miniature electronic devices to objects surrounding our everyday lives. These devices are going to monitor our environment or ourselves. They are also going to execute small tasks based on their observations. Operations of this kind can be done with small low-performance devices. In this paper we present building blocks for a low-power system for sensor-based applications. Our work is aiming for reliable operation with supply voltage of 0.5 V instead of the nominal of 1.2 V. We describe low-power solutions used in 130 nm CMOS technology microprocessor, memory and clock generator. Depending on the component, we have managed to achieve reliable operation with voltages of 0.5–0.8 V. Power consumption of components reduce significantly when operating voltage is scaled down to near-threshold region.**

***Index Terms*—Microprocessors, Near-Threshold Computing, RC Oscillator, System-on-Chip, 6T SRAM.**

I. Introduction

In the future the number of small sensor-based devices is going to increase. The Internet-of-Everything (IoE) vision connects objects surrounding our everyday lives into a smart ensemble. The objects are equipped with devices that monitor our environment and ourselves. Furthermore, they are also going to execute small tasks based on their observations. For instance, a device integrated into a shirt can monitor the user's ECG signal for the heart rate and report on a specific event. These kinds of operations can be performed with miniature, low-performance devices.

In this paper, we are creating a basis for a system, which could be used for a low-power and low-performance sensor-based System-on-Chip (SoC). In this kind of systems processors, memories and clock generators are key elements. Optimization of these components has a significant impact on the power consumption of the whole SoC, and we present our low-power solutions for these. Usage of Near-Threshold Computing (NTC) techniques has been the starting point of our designs. In NTC, supply voltage (V_{dd}) is lowered from the nominal value, but it is kept above the threshold voltage (V_{th}) of transistors; this keeps I_{on}/I_{off} ratio of transistors large enough for reliable operation.

Each our component needs a different approach when power savings are pursued. The clock generator is usually always on, so reducing its dynamic power consumption should be the main target. A single memory cell on the other hand is accessed quite rarely; reducing leakage is important there. The same applies for the processor, which spends most of the time in standby mode, waiting for a triggering event. The processor has numerous logic gates and registers. Therefore, designing the whole near-threshold compatible library from scratch involves much work. We have approached this problem by carefully selecting a set of standard library logic gates, that would work with near-V_{th} voltages. The processor is a freely available open-source implementation based on the MSP430 architecture by Texas Instruments.

The used 130 nm CMOS technology makes it possible to manufacture chips with a reasonable price. Low-leakage variants of transistors and gates were used. They are suitable for our low speed requirements–the aim is to have the system operate at 1 MHz clock frequency. To achieve low power consumption, the goal is to make the system operate with V_{dd} of 0.5 V. [1] states that for traditional CMOS technologies the optimal balance between low energy consumption and high performance can be found around 0.5 V.

This paper is constructed as follows. The processor and its implementation solutions are presented in Section II. Memory cell and read circuit designs and simulation results are presented in Section III. The clock generator is presented in Section IV. Plans for the future work have been presented in Section V. Conclusions are made in Section VI.

II. Processor

The processor core of our system had to be synthesizable, customizable, and freely available. Potential candidates were open-source processors written in commonly used hardware description languages. Several of them are available, such as LM32 by Lattice Semiconductors [2] and OpenRISC by the OpenCores initiative [3]. We selected OpenMSP430 for our system; it is an open-source microprocessor based on Texas Instruments' MSP430 architecture [4]. The core is written in Verilog and it is freely available at www.opencores.org. The OpenMSP430 is compatible with the MSP430 instruction set and can execute code generated by any MSP430 tool chain in a near cycle accurate way.

978-1-4799-6890-9/14/$31.00

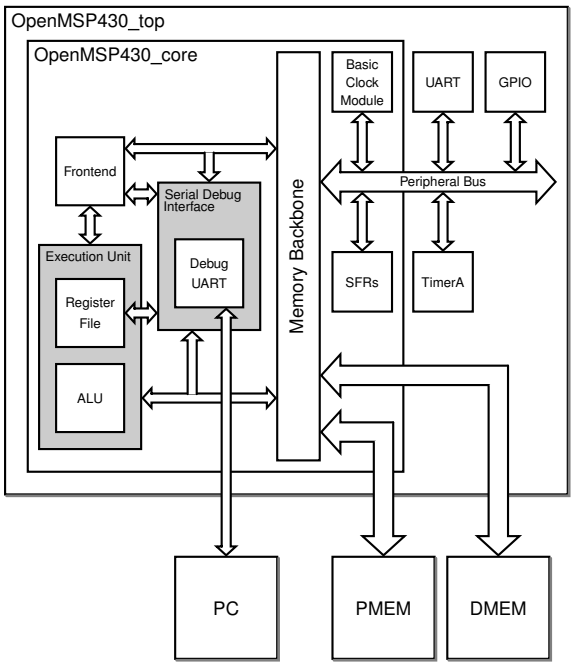

Fig. 1. OpenMSP430 block diagram.

A. Architecture

The processor is targeted for low-power applications. The OpenMSP430 core is simple and small and several general-purpose registers minimize memory access. The architecture supports two clock domains, which make four low-power modes available. Block diagram of the processor is illustrated in Fig. 1.

The OpenMSP430 core is organized of six modules at minimum: **Frontend** performs the instruction fetch and decode tasks and contains the execution state machine. **Execution Unit** contains the ALU and the register file and executes the current decoded instruction according to the execution state. **Serial Debug Interface** contains the required logic for a Nexus class 3 debugging unit. Communication with the host is done with either the UART 8N1 or the I2C protocol. **Memory Backbone** performs a simple arbitration between the Frontend and Execution unit for the program, data and peripheral memory access. **Basic Clock Module** generates clock signals and manages the low power modes. **Special Function Registers** contain various configuration registers, such as Non-Maskable Interrupt and Watchdog Timer.

Peripheral modules were included to enhance the processor capabilities. The TimerA module provides a periodic interrupt capability for e.g. reading data from a sensor. The UART module offers a means for connecting the system to an external device. The 4 bidirectional pins in the GPIO module can be used for e.g. controlling status LEDs or for debugging purposes.

B. Implementation

The 130 nm CMOS technology library has a nominal voltage of 1.2 V. Scaling the voltage down increases the delay

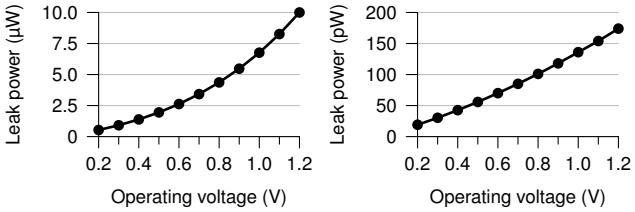

(a) (b)

Fig. 2. Leak powers of OpenMSP430 processor (a) and a single memory cell (b) on various operating voltages, temperature 27 °C.

of logic cells. Thus, the library cells had to be characterized for the lowered operating voltage in order to get new timing characteristics and thus achieve a realistic result in synthesis. Only cells with a fan-in of 2 were used in order to mitigate the performance loss. A reliable voltage was found by running a set of Monte Carlo (MC) simulations on an example circuit similar to the one in [5]. The result of 1 000 MC iterations indicate that the minimum size library cells switch correctly when the operating voltage is reduced to 0.6 V. The simulations were run with different corners at a temperature range of 0 °C to 120 °C. When the voltage is further reduced to 0.5 V, the pass rate is dropped significantly to 38.9 %. Thus, the cells have to be redesigned in order to use the target operating voltage safely.

The cell redesign procedure was omitted in this paper. Instead, we used the hand-selected library cells and characterized them for the safe 0.6 V operating voltage. The processor was designed to run on a 1 MHz clock and it was synthesized to the updated library. Fig. 2 (a) shows that when the voltage is scaled down from the nominal value (1.2 V) to the lower, yet safe value (0.6 V), the leakage power drops from 10 µW to 2.6 µW, yielding a 74 % reduction. When the operating voltage is further dropped to the target value of 0.5 V, the reduction becomes 80 %. Even though sporadic simulations indicate that the processor works as expected on the target operating voltage, the reliability still has to be addressed.

III. Memory

As the memory array has a large number of memory cells, it is reasonable to optimize a single memory cell for low-power operation. Traditionally 8T SRAM cell is considered as more reliable than 6T SRAM cell, but we are trying to gain savings in area and static power by using 6T SRAM cell. However, we must design the 6T SRAM cell to be reliable. Conventionally 8T SRAM cell is read from one side only, and 6T SRAM cell is read simultaneously from both sides. This exposes both of the inner nodes of 6T SRAM cell to precharged bitlines and thus the state of the cell may change in an unwanted way–this is the main reason that makes 6T SRAM cell more vulnerable than 8T SRAM cell. By applying the 8T SRAM read technique to 6T SRAM, we were able to make 6T SRAM work more reliably. Single-sided 6T SRAM read discussed in [6] and [7] is suitable for low-performance and low-power

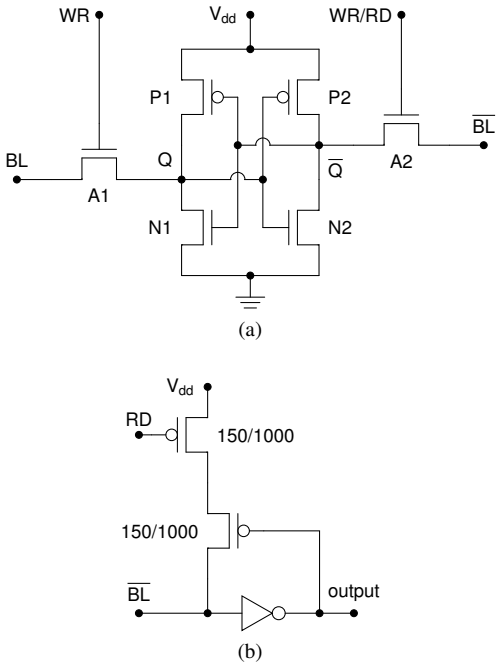

Fig. 3. Schematic diagrams of 6T SRAM cell (a) and read circuit (b), PMOS dimensions: width/length (nm).

systems. Our goal is to use 1 MHz clock in our system, and therefore slower read operation fits our needs. [8] states, that 6T SRAM with single-sided read can be a reasonable choice for near-V_{th} voltages.

Memory cell is illustrated in Fig. 3. Transistor dimensions are determined by considering mainly reliability and static energy consumption. The lengths of all transistors are set to 200 nm, which is larger than the nominal 130 nm; this reduces leakage, and bigger size gives protection against manufacturing inaccuracies and thus reduces device mismatch variations. P1 and P2 have the minimum width allowed by the technology, 150 nm; this minimizes leakage through these PMOS transistors. NMOS transistors N1 and N2 are 350 nm wide. They are wider than the access NMOS transistor A1, which is 250 nm wide, and the same width as the access NMOS transistor A2, which is 350 nm wide. These make the memory cell stable during read operation, which is when 6T SRAM is the most prone to unwanted state change. A2 has the same size as the pull-down transistor N2; this allows precharged $\overline{BL}$ to discharge through them with speed, that makes reliable read operation possible. Leakage through the whole memory cell depends significantly on the widths of N1 and N2, and the minimum leakage can be achieved with widths of 300 nm [9]; the widths of N1 and N2 are close to that.

A. Read Circuit

The read operation starts with precharging $\overline{BL}$. After that, the reading is started by setting WR/RD to V_{dd}. If bit 1 is stored to the cell, Q is at V_{dd} and $\overline{Q}$ is at ground. During read operation, precharged $\overline{BL}$ starts to discharge through A2 and N2. Discharging should be fast enough compared to the rate that $\overline{BL}$ leaks through all the unaccessed SRAM cells connected to it. This way, the read circuit is able to identify the data correctly.

If bit 0 is read from the cell, $\overline{BL}$ should not discharge and it should stay at or near V_{dd}. In this case, the read circuit should identify the value of $\overline{BL}$ before leakage causes it to drop too much. We have placed a weak pull-up connection to the read circuit to help keep the value one in the bitline. This connection is illustrated in the Fig. 3 (b). Also, our bitlines are only 64 bit long, and therefore leakage does not seem to interfere with the read operation. Short bitlines have low capacitance, and they are easier to charge and discharge. This adds up to small transistor sizes, low leakage and better reliability.

Read operations are executed only from $\overline{BL}$. The read circuit is illustrated in Fig. 3 (b). The circuit takes less area than sense amplifier circuit and it does not rely on precise timing of control signals.

B. Simulations

The memory cell was simulated by creating a test bench, which models one column of cells in a 128 B memory page. To get estimations about the reliability of the memory cell, simulations were executed with FS and SF mismatch corner parameters with 6σ variation. Temperatures of 15 °C and 40 °C were used to simulate the behavior close to room temperature. With 100 MC simulations at V_{dd} of 0.5 V there were no logical errors.

Leak power is another important measurement, while the most of the memory cells of the memory array are only standing by at any given time. Leak simulations were executed by storing 0 or 1 to the cell, and then waiting for the leak currents to settle. After that, the leakage from V_{dd} was measured. The results are presented in the Fig. 2 (b). The leakage reduces 59 % when the V_{dd} is lowered from the nominal 1.2 V to 0.6 V; this is the safe value for the processor. However, if V_{dd} of 0.5 V can be used, the leakage drops 68 % compared to the nominal.

IV. Clock Generator

Recently, on-chip RC oscillators have been used as sleep timers for their small area, low cost and good stability under voltage and temperature variations. We needed a clock generator with 1 MHz and 2 MHz outputs. We scaled up the frequency of a kHz range on-chip RC oscillator proposed in [10] to MHz range in order to retain the favorable properties of RC oscillator sleep timers while increasing the computational power of the system. The architecture ensures a stable total period time by canceling offset between high and low states of oscillation period. The schematic diagram of our oscillator is illustrated in the Fig. 4. A flip-flop was connected to the output of the oscillator to get a 50/50 duty cycle for the clock signal. The 1 MHz flip-flop output CLK is used as the system clock. Ø is used as the 2 MHz signal–SRAM memory uses it for internal timing.

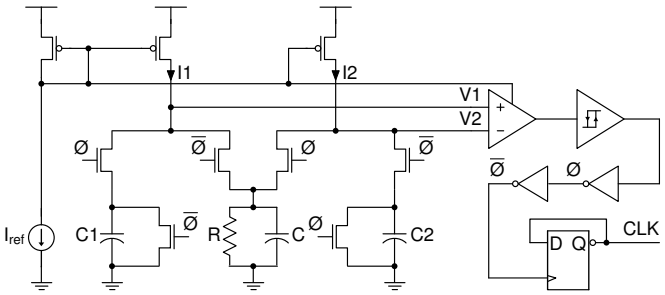

Fig. 4. Schematic diagram of oscillator circuit.

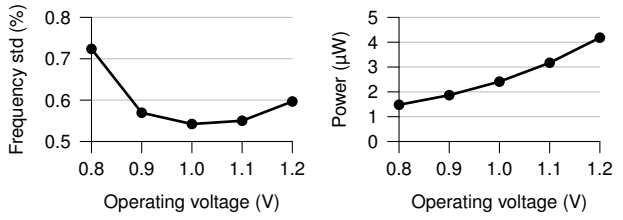

(a) (b)

Fig. 5. Standard deviation of the oscillator frequency (a) and total power consumption of the oscillator (b) on various operating voltages, temperature 27 °C.

When Ø is 1, capacitor C1 is charged by current I1 causing a rising voltage to V1. Current I2 and resistor R generate a reference voltage at V2. When V1 crosses V2, the comparator, Schmitt trigger and following inverters cause Ø to change. When Ø is 0, V1 and V2 change their roles; V1 becomes the reference and V2 acts as the rising voltage. The oscillation period can be represented as $2RC + 2t_{delay}$ (t_{delay} is the combined delay of the comparator and the Schmitt trigger). R and C sizes are adjusted in such way, that node Ø oscillates at 2 MHz frequency; this produces 1 MHz signal to the output CLK.

The oscillator operates within V_{dd} of 0.8–1.2 V and consumes 1.48 µW power at 0.8 V. The oscillator occupies area of 0.01 mm^2. Simulation results of V_{dd} scaling are shown in Fig. 5. The oscillator power consumption is reduced 65 % when V_{dd} is downscaled from 1.2 V to 0.8 V. At 0.8 V and 27 °C the oscillator has frequency standard deviation of 0.72 % and RMS jitter of 3.5 ns. The frequency variation is ±0.49 % over 15–40 °C. Lowering the supply voltage reduces the accuracy of the oscillator. If supply voltage is reduced from 0.8 V closer to our goal 0.5 V, the oscillator needs to be optimized further to maintain the accuracy of the clock.

V. Future Work

In the future, a modified cell library should be created and characterized for the processor. This supposedly optimizes the processor further and enables reliable operation at the target supply voltage of 0.5 V. The processor shall face improvements in form of basic signal processing capabilities and a high-speed serial communication interface. Furthermore, power and clock gating will be investigated.

For the memory, the optimal size of memory page should be found. Larger page needs larger peripheral circuits in form of larger address decoders and stronger drivers, but the area cost per byte could be lower. Smaller page on the other hand, can manage with smaller peripheral components, but to achieve the same byte count more pages are needed.

In the oscillator design, it would be interesting to investigate if usage of high speed devices inside the comparator and the Schmitt trigger would be beneficial for the whole oscillator power consumption. These components are now slow and there is significant variation in their delays; their faster operation would give greater weight for the RC delay, which varies less. Also, changing the amplitude of the voltage swing of RC part might be worth of investigation.

VI. Conclusion

In this paper we have presented low-power solutions for a low-performance processor system. Our goal was to use V_{dd} of 0.5 V instead of the nominal 1.2 V. We managed to get the processor design to operate reliably at 0.6 V, and our memory at 0.5 V. Our 1 MHz clock generator takes only 0.01 mm^2 area and its power consumption reduces 65 % when V_{dd} is scaled from 1.2 V to 0.8 V. There is still work to be done to get all the components to operate reliably at 0.5 V. However, our simulations indicate, that static power savings are significant in near-threshold voltages. With the safe V_{dd} of 0.6 V our processor leaks 74 % and our memory cell 59 % less than with the nominal V_{dd}.

References

[1] L. Chang, D. Frank, R. Montoye, S. Koester, B. Ji, P. Coteus, R. Dennard, and W. Haensch, "Practical strategies for power-efficient computing technologies," *Proc. IEEE*, vol. 98, no. 2, pp. 215–236, Feb. 2010.

[2] Lattice Semiconductor. (2012, Dec) LatticeMico32 processor reference manual.

[3] OpenCores. (2012, Dec. 11) OR1200 OpenRISC Processor. [Online]. Available: http://opencores.org/or1k/OR1200_OpenRISC_Processor

[4] O. Girard. (2013, Dec.) OpenMSP430 Specification (Rev. 1.1.4). [Online]. Available: http://opencores.org/websvn,filedetails?repname=openmsp430&path=%2Fopenmsp430%2Ftrunk%2Fdoc%2FopenMSP430.pdf

[5] J. Kwong and A. Chandrakasan, "Variation-driven device sizing for minimum energy sub-threshold circuits," in *Proc. ISLPED*, Tegernsee, Germany, Oct. 2006, pp. 8–13.

[6] G. Chen, M. Wieckowski, K. Daeyeon, D. Blaauw, and D. Sylvester, "A dense 45nm half-differential SRAM with lower minimum operating voltage," in *Proc. ISCAS*, Rio de Janeiro, Brazil, May 2011, pp. 57–60.

[7] C. Piguet, J.-M. Masgonty, S. Cserveny, C. Arm, and P. D. Pfister, "Low-power low-voltage library cells and memories," in *Proc. ICECS*, vol. 3, Malta, Sep. 2001, pp. 1521–1524.

[8] M. Kutila, A. Paasio, and T. Lehtonen, "Comparison of 130 nm technology 6T and 8T SRAM cell designs for near-threshold operation," in *Proc. MWSCAS*, College Station TX, USA, Aug. 2014, pp. 925–928.

[9] M. Kutila, A. Paasio, and T. Lehtonen, "Simulations on 130 nm technology 6T SRAM cell for near-threshold operation," in *Proc. ISCAS*, Melbourne VIC, Australia, June 2014, pp. 1211–1214.

[10] A. Paidimarri, D. Griffith, A. Wang, A. Chandrakasan, and G. Burra, "A 120nW 18.5kHz RC oscillator with comparator offset cancellation for ±0.25% temperature stability," in *Proc. ISSCC*, San Francisco CA, USA, Feb. 2013, pp. 184–185.