

A 16-Bit Low-Power Microcontroller with Monolithic MEMS-LC Clocking

Eric D. Marsman, Robert M. Senger
University of Michigan
Ann Arbor, MI 48109-2122, USA
{emarsman, rsenger}@eecs.umich.edu

Michael S. McCorquodale
Mobius Microsystems, Inc.
Detroit, MI 48226-1686, USA
mccorquodale@mobiusmicro.com

Matthew R. Guthaus, Rajiv A. Ravindran,
Ganesh S. Dasika, Scott A. Mahlke
University of Michigan
Ann Arbor, MI 48109-2122, USA
{mguthaus, rravindr, gdasika, mahlke}@eecs.umich.edu

Richard B. Brown
University of Utah
Salt Lake City, UT 84112-1109, USA
brown@coe.utah.edu

Abstract—Low-power, single-chip integrated systems are prevailing in remote applications due to the increasing power and delay cost of inter-chip communication compared to on-chip computation. This paper describes the design and measured performance of a fully-functional digital core with a low-jitter, monolithic, MEMS-LC clock reference. This chip has been fabricated in TSMC's 0.18 μm MM/RF bulk CMOS process. Maximum power consumption of the complete microsystem is 48.78mW operating at 90MHz with a 1.8V power supply.

I. INTRODUCTION

To satisfy the broad range of workload requirements for microsystems and Systems-on-a-Chip (SoCs), an adaptable microcontroller unit (MCU) must be designed with a wide spectrum of communication capabilities and operating specifications. The processing requirements and allotted power budget for the embedded MCU in PDAs, cell phones, remote environmental sensors, bio-medical devices, etc. varies significantly with the application. By building an MCU that can meet these design requirements and by leveraging an intellectual property based design methodology [1], design effort and design cycle can be greatly reduced without sacrificing significant power or performance.

The Center for Wireless Integrated Microsystems (WIMS) has two demonstration vehicles for which the microsystem has been designed: a remote environmental monitor and a cochlear implant [2]. The low-power, single-chip microsystem presented here is adaptable enough to be used in a variety of embedded applications.

The next section describes the architecture of the microsystem composed of the MCU and micromachined LC tank clock reference. The remaining sections report measured results from the fabricated microsystem, including the MCU, loop cache, dynamic frequency control (DFC) unit, and the monolithic LC clock reference.

II. MICROSYSTEM ARCHITECTURE

Fig. 1 shows the microsystem architecture consisting of the digital core and the CMOS-MEMS LC tank oscillator used as an on-chip clock reference. The digital core includes a 3-stage pipeline, 16-bit datapath, a 24-bit unified address space, 64KB of on-chip SRAM, and an external memory port supporting up to 64KB of additional memory. The interface to the monolithic CMOS-MEMS clock reference is a software controlled memory mapped register to select the clock frequency that supplies the MCU.

A. Microcontroller Architecture

The primarily load-store instruction set architecture (ISA) contains 77 instructions supporting eight addressing modes and single- and multi-word arithmetic, shift, logical, and control-flow operations [3]. Instructions in the custom ISA were carefully chosen to minimize decode complexity and power without sacrificing functionality. One level of interrupt and subroutine support is available in hardware. Nested interrupts and subroutines are enabled through software control of the hardware stack and frame pointer.

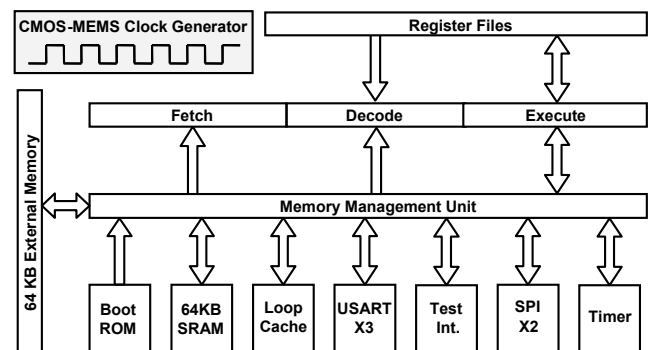


Figure 1. Microsystem architecture

A 3-stage pipeline was chosen to provide adequate performance for remote sensing and bio-medical applications, yet still remain low-power with minimal pipeline hardware overhead. The pipeline utilizes sixteen 16-bit general purpose registers and four 24-bit address registers, divided evenly over two windows. The windowing scheme reduces the register encoding field to enable 16-bit instructions while providing additional registers for temporary storage. Ref. [4] gives a detailed analysis of the WIMS compiler’s efficient utilization of the register windows to achieve up to 19% reduction in power consumption and 30% improvement in performance when compared to a non-windowed architecture. Address register manipulation is enabled through direct memory mapped access or by using any of the several address update modes available in the ISA.

The memory architecture is a banked style with the 64KB of SRAM split into four 16KB banks. This allows instruction and data accesses to occur simultaneously without stalling the machine pipeline as long as they address different banks. This is easily done with minimal organizational control by the software compiler. In addition, this memory configuration allows for unused banks to be shut down on a cycle-by-cycle basis when not being accessed, yielding an overall power savings in the memory. The energy of this configuration is 69.2% less and the area penalty is 15.9% when compared to a single 64KB bank [5].

Considerable power savings in the memory architecture are realized by the addition of a low-power, 512-byte loop cache [6]. Unlike traditional caches, the loop cache is a tag-less bank of low-power memory intelligently managed by the WIMS compiler. The cache is intended to contain the most commonly executed instructions or accessed data, typically found in program loops. Contents of the cache are determined by the compiler and are not under hardware control, as is typical of memory hierarchy caches. It is still possible to change the contents of the loop cache at run time by loading the proper instructions or data into the cache and resuming program execution. The loop cache introduces minimal hardware overhead due to the banked memory structure, but yields significant power savings that will be presented in the results section.

Serial interfaces and timer peripheral components provide the general-purpose functionality required by most embedded systems. A special serial test port facilitates remote, on-site testing of the MCU.

B. Clock Generation Architecture

Clock sources for most SoCs and MCUs consist of a low-jitter, off-chip crystal or ceramic reference with an on-chip phase-locked-loop (PLL) to multiply the external reference frequency. If applicable, a less stable on-chip ring, relaxation, or phase-shift oscillator may be used in lieu of the

external reference. The work described here utilizes a complementary, cross-coupled, negative transconductance (g_m) LC tank shown in Fig. 2. A design overview for this low-jitter, 1.1GHz CMOS compatible reference oscillator is given in [3]. With the proposed LC oscillator, neither a PLL nor off-chip reference are required, thus reducing component cost, form factor, package pin count, and system power. Moreover, the clock is significantly more stable than the aforementioned on-chip clock generation techniques [7].

The quality (Q) factor of the LC tank is determined primarily by the inductor because it is more lossy than integrated fixed metal-insulator-metal (MiM) capacitors, which comprise the C in the tank. To improve the Q -factor, the inductor is fabricated using thick top metal that is released from the surrounding passivation using a wet oxide etch with an ultra-low aluminum etch rate. The remainder of the microsystem is protected from the etch by the silicon nitride passivation, which is impervious to the etch chemistry utilized. The release, which does not require an extra mask step, increases the inductor’s measured Q -factor by up to 13%, from 7.5 to 8.5, as shown in [7]. Frequency trimming is achieved by modulating the current in the cross-coupled pair with v_{trim} and thus modulating the g_m . Absolute frequency deviation due to process variation can be corrected using this scheme. Temperature compensation of the free-running oscillator can be managed similarly [7].

A buffer amplifier is required to isolate the free-running oscillator from the frequency divider (not shown), remove amplitude variation, and present sufficient signal swing to the frequency divider. The frequency divider is a divide-by-two circuit using D flip-flops with complementary feedback. Any type of frequency division may be utilized; the topology implemented here was selected for simplicity.

III. MICROSYSTEM RESULTS

This section provides the measured results for the microsystem fabricated in *Taiwan Semiconductor Manufacturing*




Figure 2. CMOS-MEMS LC tank schematic




Figure 3. Die micrograph of the complete microsystem

Company's (TSMC) 0.18 μ m mixed-mode/RF bulk CMOS process using the design methodology given in [8]. Fig. 3 is a die micrograph of the fabricated system with the major components outlined and labelled. The 128 pin die is 3.54mm per side (12.53mm²) and contains 3.5 million transistors.

A. Microcontroller Results

The MCU is fully-functional up to 92.5MHz at 1.8V and consumes a maximum of 33.9mW. The target frequency of 100MHz is achieved when restricting memory use to SRAM bank zero and the cache. It was determined that the other memory banks lie on the critical path and limit the speed of the MCU. When put into a 2kHz low-power idle mode, the core consumes only 1.15mW from a 1.8V supply. Further power savings can be realized by lowering the power supply (V_{DD}) when putting the chip into idle mode. Power consumption in idle mode drops to 740 μ W when the power supply is reduced to 1.15V. Digital output pins are available to control

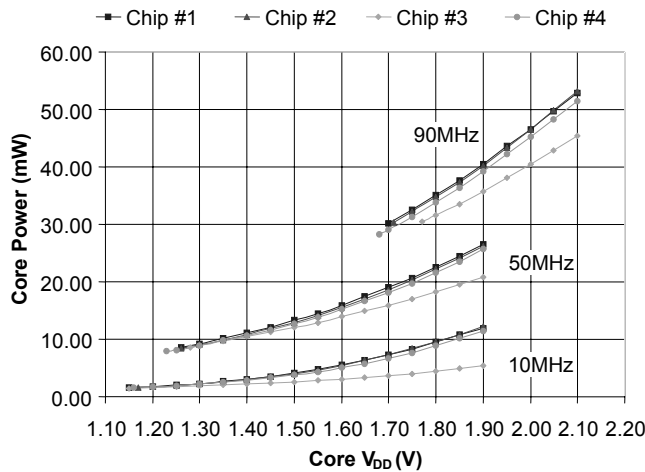


Figure 4. Measured V_{DD} scaling across frequency ranges

an off-chip voltage regulator that can modulate the power supply voltage. Fig. 4 shows the minimum power supply voltage level required for correct operation at key frequencies as measured across four WIMS chips.

B. Loop Cache Results

A single access to the loop cache consumes 45% of the energy that an access to the SRAM consumes. Fig. 5 demonstrates the measured core power consumption for six test programs operating at 100MHz that use only the SRAM or utilize the cache in addition to the available SRAM. The percentage of power savings and the percentage of cache accesses out of total memory accesses are also given.

The custom-built WIMS ISA compiler implemented a dynamic cache filling algorithm that simulated the power efficiency of the loop cache. Across a subset of the embedded benchmarks MiBench [9] and MediaBench [10], an average energy savings of 43% was achieved [11].

C. Dynamic Frequency Control

The ability to dynamically modify the frequency of an MCU with low latency is an important feature for any application that seeks to balance workload with energy consumption. Most systems do not have a constant workload, especially embedded and remote systems which are often interrupt driven and repeatedly change between active and idle modes. The new HDL-synthesizable, glitch-free DFC implementation used here has available frequencies ranging from $f_0=66.7$ MHz to $f_{15}=2$ kHz, where $f_n=f_0/2^n$; $n=1,2,\dots,15$. The latency from the time the frequency is selected until the core operates at the new frequency is $5/2f_0$, or 37.45ns for this design. This is nearly three orders of magnitude faster than the 20 μ s required for a fast-locking PLL [12] to change frequencies. Future implementations will increase f_0 to 125MHz and reduce the latency to $3/2f_0$. Fig. 6 shows oscilloscope traces of the MCU selecting different frequencies without halting the pipeline.

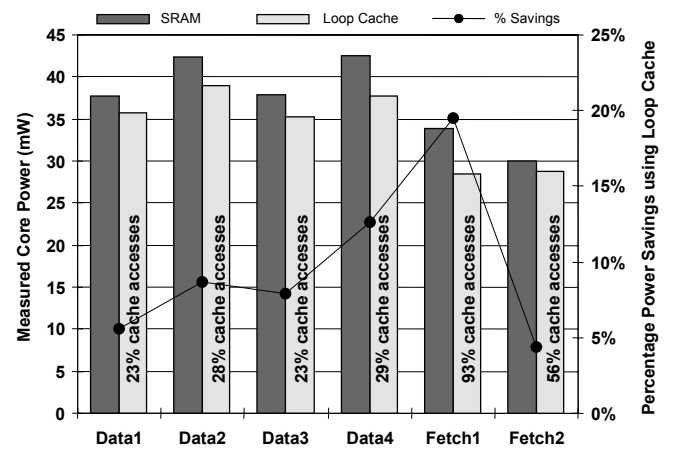


Figure 5. SRAM versus loop cache measured power consumption and percentage savings

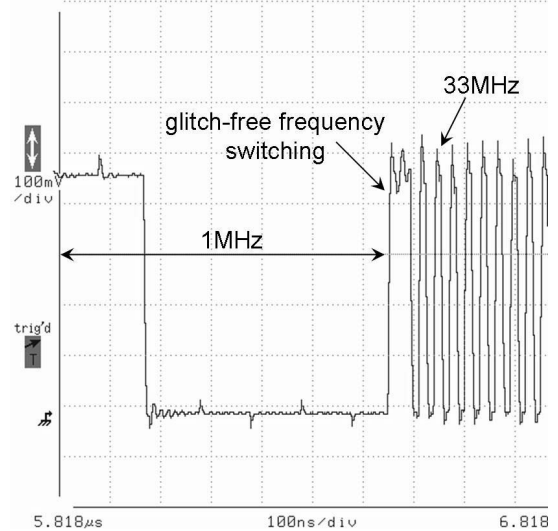



Figure 6. Low latency dynamic clock frequency selection

D. LC Tank Results

Ref. [7] describes the fabricated *LC* reference which oscillates at 1.056GHz with a $\pm 2\%$ precision before electrical trimming. The oscillator achieves a worst case 48/50 duty-cycle with less than 300ppm RMS period jitter while occupying an area of 0.3mm^2 , or 2.4% of the die. The measured single sideband phase noise power spectral density of the 33MHz clock output was -95.4dBc/Hz at 10kHz offset. It consumes 17.28mW from a 1.8V supply.

IV. CONCLUSION

This work reports the single-chip integration of a flexible, low-power MCU that has a custom-built ISA and C-compiler with a CMOS compatible MEMS-*LC* reference oscillator. Maximum active power consumption of the MCU is 33.9mW at 92.5MHz and 1.8V with an idle-mode drawing only $740\mu\text{W}$ at 1.15V. The on-chip, 0.3mm^2 MEMS-*LC* reference supplies a highly accurate, low-jitter clock source while consuming 17.28mW at 1.8V.

An HDL-synthesized DFC unit allows for glitch-free, low-latency modification of the clock frequency to satisfy the changing workload requirements of low-power embedded systems. The loop cache enables efficient execution of repetitive program code in a variety of low-power applications. Simple static filling of the 512-byte loop cache achieved a measured 5 to 20% power reduction compared to the on-chip SRAM. A dynamic, compiler managed cache filling algorithm demonstrated a 43% energy savings when simulated with the loop cache.

ACKNOWLEDGEMENTS

Fabrication of this work at TSMC was funded by the MOSIS Educational Program. The authors wish to thank Artisan Components for supplying the digital standard cell

libraries and memory generators. Work was supported by the Engineering Research Centers program of the NSF under award number EEC-9986866.

REFERENCES

- [1] M. McCorquodale, *et al.*, "Microsystem and SoC design with UMIPS," *IFIP International Conference on VLSI SOC*, Dec. 2003.
- [2] Center for Wireless Integrated Microsystems (WIMS), An NSF Engineering Research Center. <http://www.wimserc.org/>, 2004.
- [3] R. Senger, *et al.*, "A 16-bit mixed-signal microsystem with integrated CMOS-MEMS clock reference," *DAC*, pp. 520-525, Jun. 2003.
- [4] R. Ravindran, *et al.*, "Increasing the number of effective registers in a low power processor using a windowed register file," *CASES*, Oct. 2003.
- [5] S. Martin, *et al.*, "A low-power microinstrument for chemical analysis of remote environments," *11th NASA Symp. on VLSI Design*, Coeur d'Alene, ID, pp. 1-4, May 2003.
- [6] N. Bellas, I. Hajj, C. Polychronopoulos, and G. Stamoulis, "Energy and performance improvements in microprocessor design using a loop cache," *Intl. Conf. on Computer Design*, pp. 378-383, 1999.
- [7] M. McCorquodale, *Monolithic and Top-Down Clock Synthesis with Micromachined RF Reference*, Ph.D. Dissertation, U. of Michigan, 2004.
- [8] M. McCorquodale, *et al.*, "A top-down microsystems design methodology and associated challenges," 2003 *DATE Designers' Forum*, pp. 292-296, Mar. 2003.
- [9] M. Guthaus, *et al.*, "MiBench: A free, commercially representative embedded benchmark suite," *In IEEE 4th Annual Workshop on Workload Characterization*, Dec. 2001.
- [10] C. Lee, M. Potkonjak and H. Mangione-Smith, "MediaBench: A tool for evaluating and synthesizing multimedia and communications systems," *Micro-30*, Nov. 1997.
- [11] R. Ravindran, *et al.*, "Compiler managed dynamic instruction placement in a low-power code cache," *Conference on Code-Generation and Optimization*, Mar. 2005.
- [12] K. Nowka, *et al.*, "A 32-bit PowerPC system-on-a-chip with support for dynamic voltage scaling and dynamic frequency scaling," *J. Solid State Circuits*, vol. 37, pp. 1441-1447, Nov. 2002.