

Evaluating Adaptive Clocking for Supply-Noise Resilience in Battery-Powered Aerial Microrobotic System-on-Chip

Xuan Zhang, *Student Member, IEEE*, Tao Tong, *Student Member, IEEE*, David Brooks, *Member, IEEE*, and Gu-Yeon Wei, *Member, IEEE*

Abstract—A battery-powered aerial microrobotic System-on-Chip (SoC) has stringent weight and power budgets, which requires fully integrated solutions for both clock generation and voltage regulation. Supply-noise resilience is important yet challenging for such SoC systems due to a non-constant battery discharge profile and load current variability. This paper proposes an adaptive-frequency clocking scheme that can tolerate supply noise and improve performance when implemented with an integrated voltage regulator (IVR). Measurements from a ‘brain’ SoC, implemented in 40 nm CMOS, demonstrate $2\times$ performance improvement with adaptive-frequency clocking over conventional fixed-frequency clocking. Combining adaptive-frequency clocking with open-loop IVR extends error-free operation to a wider battery voltage range (2.8 to 3.8 V) with higher average performance.

Index Terms—Clock generation, supply noise, System-on-Chip (SoC).

I. INTRODUCTION

ROBOTICS have grabbed much public imagination these days with their promises for a range of versatile applications. A special branch of robotics is the recent development of microrobot which builds miniature robotics with characteristic dimensions less than 1 mm. Despite its small size, the microrobot embodies many of the essential components that can be found in regular-sized robots, such as power source and conversion, actuation, sensing, and autonomous control. Accommodating all these functionalities within the size limit of the microrobot provides fertile ground for the design of a highly integrated system-on-chip (SoC) targeted for robotic applications.

Among the numerous design challenges presented by microrobotic SoC, one critical problem to be addressed is the reliability and performance of the system in the presence of supply noise. Similar to many integrated computing systems, the microrobotic SoC employs a digital processor as its central control unit and thus is susceptible to disturbance on the supply voltage. However, the crucial weight and form factor constraints set the microrobotic SoC apart from conventional systems. Given the extremely stringent weight budget, extra external components must be avoided at all cost, which leads to the integration of

on-chip dc-dc converter and the absence of external frequency reference. With such integrated voltage regulators (IVR) powered directly off of a discharging battery, the microrobotic SoC experiences supply noise characteristics different from conventional digital systems, where existing supply noise mitigation techniques cannot be easily applied [1]–[4].

In this paper, we propose an adaptive-frequency clocking scheme to exploit the synergy between integrated voltage regulation and clock generation in a microrobotic SoC. The resulting supply-noise resilience and performance improvement has been demonstrated by a prototype SoC developed for an aerial microrobot known as RoboBee [5]. Our proposed adaptive clocking scheme not only delivers better reliability and performance, but also extends the error-free operation to a wider battery voltage range, which is beneficial to the microrobotic system.

The paper is organized as follows: a brief background on existing supply noise mitigation techniques is provided in Section II along with an introduction of the unique design considerations in the microrobotic SoC. More specifically, we describe the system architecture and the circuit implementation of the adaptive clocking scheme in our prototype chip in Section III and Section IV. Finally, Section V summarizes the measurements results from the fabricated chip in a 40 nm CMOS process.

II. BACKGROUND

A. Supply Noise

Digital computing systems based on synchronous logic circuits typically employs a fixed frequency clock. To guarantee correct operation, final outputs from the datapath must arrive at the next flip-flop stage before the next clock edge by some time margin known as the “setup time”. Since the datapath delay is a function of the supply voltage, it is susceptible to noise on the supply line.

Supply noise is the result of non-ideal power delivery system and load current fluctuation under varying computation workload. It can come from the parasitic resistance, inductance, and capacitance in the power delivery network, and manifests itself as static IR-drop, which is the static voltage drop due to power grid resistance, as well as dynamic Ldi/dt -drop, which is the transient voltage fluctuation caused by the inductance and capacitance in response to load current changes. Also, for systems with integrated switching regulators, the intrinsic voltage ripple of the regulator contributes additional noise to the supply. The existence of supply noise can modulate the

Manuscript received December 24, 2013; revised February 23, 2014; accepted March 02, 2014. Date of publication March 31, 2014; date of current version July 24, 2014. This work was supported in part by the National Science Foundation (NSF) Expeditions in Computing Award #: CCF-0926148. This paper was recommended by Associate Editor J. M. de la Rosa.

The authors are with Harvard School of Engineering and Applied Sciences (SEAS), Cambridge, MA, 02138 USA (e-mail: xuanzhang@eecs.harvard.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/TCSI.2014.2312490

datapath delay, which may lead to setup time margin violation and eventually computation errors. In order to ensure sufficient delay margins under all operating conditions, the most straightforward approach is to lower the clock frequency and provide a “guardband” to tolerate even the worst supply noise scenario.

Such conservative design strategy may incur hefty performance loss and thus is highly undesirable. Instead, a number of alternative techniques have been proposed to mitigate supply noise with less performance penalty. The active management of timing guardband [3] in a prototype IBM POWER7 processor is an example of adaptive clocking: a digital phase-locked loop (DPLL) adjusts the processor core’s clock frequency based on the timing guardband sensed by a critical path timing monitor [6]. As resonant noise caused by the LC tank between the package inductance and the die capacitance has been identified as the dominant component of supply noise in high performance microprocessors [7], many studies have focused on this particular type of supply noise by proposing adaptive phase-shifting PLL [4] and clock drivers [8]. Following the duality between the clock frequency and the supply voltage in synchronous digital systems, the other approach to optimize performance in the presence of supply noise is adaptively adjusting the voltage level delivered by the power supply [1], [2], [9] at different desirable operating frequency. Despite their different implementations, both adaptive clocking and adaptive supply are along a similar vein of technical route that applies closed feedback loop to adjust frequency and/or supply based on monitored timing margin of the system, and therefore are subject to the bandwidth limitation of the feedback loop.

In addition to the above-mentioned systems and techniques, there exist other classes of logic implementations such as asynchronous logics and self-timed logics [10] that do not rely on a global clock for their operations. Unlike synchronous logics, these systems are intrinsically delay-insensitive and thus immune from the negative impact of supply noise. However, these logic implementations lack the full support of standard libraries, IPs, and EDA tools and thus are difficult to incorporate into the digital design flow of a sophisticated SoC.

B. RoboBee System

The RoboBee system is a platform developed by us and our collaborators to demonstrate autonomous flight in aerial micro-robots. It is a biologically inspired micro-robot [5], [11] that has similar size as a real honey bee as shown in Fig. 1(a). Its body weighs 250 mg and the lift force generated by its flapping wings is approximately 500 mg. Excluding its intrinsic weight, the payload RoboBee can carry is at most 250 mg, and this sets the weight budget for the entire electronic system employed by the RoboBee. Illustrated in Fig. 1(b), the RoboBee system encompasses many different components:

The power source of the entire system is provided by a lithium-ion battery of 3.7 V. The wings of RoboBee are propelled by multiple piezoelectric actuators to control the magnitude, frequency, and angle of the flapping wings. These actuators require high-voltage sinusoidal signals for efficient operation, and therefore a “Power IC” has been designed as controller and drivers for the actuators that converts the battery voltage up to 200 V [12]. Light-weight and low-power sensors are used to capture information about the environment, such as light intensity, optical images, magnetic field, and accelerations




Fig. 1. RoboBee is a bee-sized aerial microrobot developed to demonstrate autonomous flight. (a) Body of a RoboBee; (b) different components in the RoboBee system.

of the robot. Residing at the center of the system is the main “BrainSoC” for the microrobot that coordinates all the sensing, computation, and control functions of the system.

Since the Power IC requires special process technology to tolerate the high voltage, it cannot be integrated easily with the “BrainSoC”. Additionally, the Power IC requires extra external passive and magnetic components for its operation even after extensive optimization of weight and efficiency [12]. Therefore, the weight budget left for the BrainSoC after taking into account of the sensors, Power IC and its peripherals, and the battery, is extremely limited and has no room for any external components.

Although RoboBee is only one example of a microrobotic system, we believe it is representative of the kind of design constraints faced by the electronic systems for all autonomously operated microrobots, and thus can serve as a platform for more general studies on the design of SoC for microrobotic applications.

C. Microrobotic SoC

In many ways, the microrobotic SoC suffers similar setup time violations due to supply noise as other types of computing systems discussed in Section II-A, but its weight and form factor constraints present unique design challenges that require different supply-noise mitigation techniques from those employed in a microprocessor [3].

In our example, RoboBee uses a 3.7 V lithium-ion battery as its power source, while the central control SoC that acts as its ‘brain’ requires a digital supply voltage below 0.9 V. Due to the limited weight budget allocated for the SoC, neither external regulator module nor discrete components such as capacitors and inductors can be afforded, which leads to the integration of an 4:1 on-chip dc-dc converter as part of the SoC [13]. The switched-capacitor topology is selected for the integrated voltage regulator (IVR) because it does not need external inductors unlike the buck topology and it delivers much better conversion efficiency compared to a linear regulator, which suffers intrinsic low efficiency when the output voltage is only a fraction of the input voltage. The supply voltage generated by the fully integrated switched-capacitor voltage regulator (SC-IVR)

that is directly powered off of a battery can have different noise characteristics from the resonant- noise-dominated supply experienced by the microprocessor. In fact, for the IVR-enabled microrobotic SoC, the worst supply noise is often caused by sudden load current steps in a very small time scale instead of LC resonance. Therefore, the mitigation techniques developed earlier for slow-changing or periodic supply noise in microprocessors [3], [4], [8] are not applicable to the fast-changing supply noise induced by the load current in a microrobotic SoC.

In addition to voltage conversion and regulation, clock generation is yet another function that must be entirely integrated on chip for microrobotic applications. As further described in Section II-B, RoboBee is a self-sustained autonomous system with no need to synchronize with other systems for communication or I/O transactions, therefore the timing jitter/phase noise requirement of its clock signal can be relaxed from the specifications in typical high-speed I/O interfaces that demand PLL-generated clean clock. Moreover, the lack of external frequency source as the reference signal renders the implementation of a PLL in the microrobotic SoC impractical. We therefore conclude that a free-running oscillator is a better candidate as the clock generator for the SoC.

The above discussion explains that the microrobotic SoC differs significantly from conventional digital systems in its integration of a battery-connected IVR and its ability to operate with a free-running clock. The former suggests distinctive supply noise characteristics dominated by fast load current changes and slow battery discharge, while the latter provides opportunity for an adaptive clocking scheme.

The main objective of this work is to use the RoboBee system as a platform to investigate effective supply-noise mitigation techniques that can be applied to more general microrobotic systems with minimal performance penalty. In the context of the RoboBee's BrainSoC, the goal is thus to optimize the flight time with respect to the total energy available in its battery and the associated battery discharge profile. Along this vein, we explore the relative merits of different operational modes offered by the supply regulation mechanisms and the clock generation schemes. Fig. 2 illustrates two possible modes of supply regulation with respect to a typical lithium-ion battery discharge profile. The first one is closed-loop operation at a fixed voltage. In this case, with the help of feedback control loop, the SC-IVR can provide a constant supply voltage that is resilient to input battery (V_{BAT}) and output load (I_{LOAD}) conditions. One advantage associated with the fixed voltage operation is that it provides a relatively constant operating frequency. However, for a target output voltage level (V_{REF}), the SC-IVR's operating range is limited to $V_{BAT} > 4V_{REF}$. In contrast, open-loop operation with variable unregulated voltage exhibits an entirely different set of attributes; with no feedback control, the SC-IVR's output voltage is roughly 1/4th the input battery voltage, but varies with both the discharge profile and load fluctuations. While open-loop SC-IVR mode allows the system to operate over a wider range, down to the minimum voltage limit of the digital load, performance and energy efficiency depend on the clocking strategy used. The choices are between two clocking schemes: *fixed-frequency* and *adaptive-frequency* clocking. Out of the four total combinations, this paper compares the following three: (1) fixed regulated voltage, fixed frequency; (2) fixed regulated voltage,




Fig. 2. Illustrations of two SC-IVR modes of operation versus typical battery discharge profile. (a) Closed-loop regulation; (b) open-loop operation.

adaptive frequency; and (3) variable unregulated voltage, adaptive frequency. Intuitively speaking, fixed-frequency clocking (F_{FIX}) requires extra timing margins to account for non-negligible worst-case voltage ripple, which is an intrinsic artifact of the SC-IVR's feedback loop; while, alternatively, an adaptive-frequency clocking (F_{ADP}) scheme that allows the clock period to track the changes in the supply voltage could offer higher average frequency. Adaptive-frequency clocking also works well for open-loop SC-IVR mode, because it maximizes performance with respect to battery and load conditions.

D. Timing Slack Analysis

So far, we have only given an intuitive explanation on the impact of supply noise on timing margin and the potential beneficial compensation effect when the clock period can match the datapath delay by tracking supply voltages. A more rigorous definition of timing slack and the "clock data compensation effect" can be found in previous work [4], [8]. However, the analytical treatment used is more applicable to resonant supply noise that has a narrow band frequency composition, which is not suitable for the supply noise in microrobotic SoC as discussed in Section II-C. In this section, we would like to extend the timing slack analysis to systems that employ adaptive free-running clock and exhibit broad-band supply noise with fast transients.

Fig. 3 shows one stage of a pipeline circuit that is clocked by a free-running digitally controlled oscillator (DCO). The clock signal (CLK) is generated by a clock edge propagating through the delay cells of the DCO, and is then buffered to trigger the flip-flops at the input and output of the datapath. The buffered

clock signals are labeled as CP1 and CP2. Using similar definition proposed by previous work [4], [8], the timing slack can be calculated as:

$$slack = t_{clk} + t_{cp2} - t_{cp1} - t_d \quad (1)$$

Here, $t = 0$ is the time the first clock edge is launched, and it takes t_{cp1} to travel through the clock buffers and reach the first flip-flop as CP1. In the meantime, the first clock edge propagates through the delay cells and, at $t = t_{clk}$, it completes the round trip in the DCO and the second clock edge is launched at CLK and takes t_{cp2} to reach the second flip-flop as CP2. Instead of resorting to small signals, we simply represent the supply voltage as a function of time with $v(t)$. Without loss of generality, let us assume each circuit block (X) has a unique function $f_x(v(t))$ that measures the rate of propagation delay accumulation as a function of the supply voltage, such that the propagation delay of the circuit t_x can be expressed as:

$$t_x = \int_0^{t_x} f_x(v(t)) dx \quad (2)$$

In this way, we can re-write the delay parameters in (1) as the following:

$$\begin{aligned} t_{clk} &= \int_0^{t_{clk}} f_{clk}(v(t)) dx \\ t_{cp2} &= \int_{t_{clk}}^{t_{clk}+t_{cp2}} f_{cp2}(v(t)) dx \\ t_{cp1} &= \int_0^{t_{cp1}} f_{cp1}(v(t)) dx \\ t_d &= \int_{t_{cp1}}^{t_{cp1}+t_d} f_d(v(t)) dx \end{aligned} \quad (3)$$

The important finding derived from (3) is that to ensure constant positive timing slack under all supply noise conditions is to match all the delay accumulation functions ($f_{clk}, f_{cp1}, f_{cp2}, f_d$), rather than simply the DCO and the datapath. Intuitively, this is because the delay is accumulated first at the DCO and then at the clock driver for CP2, whereas it is first at the clock driver and then at the datapath for CP1. If there is any mismatch between the clock driver and the DCO or the datapath, the impact of the supply noise cannot be fully compensated. Therefore, our design uses the fanout-of-4 delay tracking DCO as a reasonably good approximation to track the typical delay in both the datapaths and the clock drivers, rather than a precise implementation to perfectly match either the datapaths or the clock drivers.

The remainder of the paper presents the design of a prototype ‘brain’ SoC implemented in TSMC’s 40 nm process and the experimental results in order to explore the relative trade-offs and merits associated with the different operational modes described above. We are able to evaluate the reliability enhancement and the performance advantage of adaptive clocking with both regulated and unregulated voltages generated by the SC-IVR.

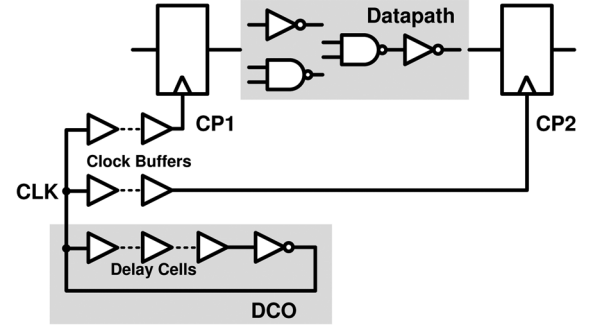


Fig. 3. Simplified diagram of a pipeline circuits.

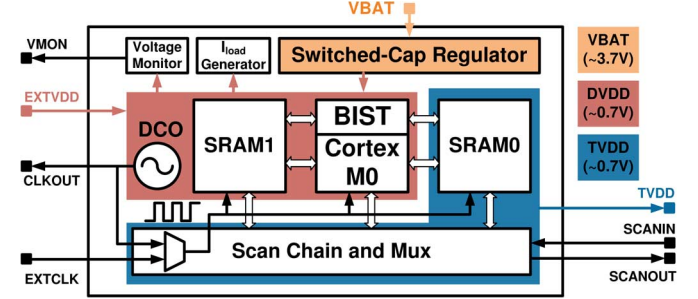


Fig. 4. Block diagram of the fully integrated SoC.

III. SYSTEM ARCHITECTURE

The prototype microbotic SoC designed for the RoboBee is not a full-fledged implementation of all the functionalities required of a ‘brain’ SoC, but it captures the most essential components in such SoC for our investigation of the interaction between supply noise and clocking scheme. Shown in Fig. 4, the prototype SoC contains a fully integrated two-stage 4:1 switched-capacitor voltage regulator (SC-IVR), a 32-bit ARM Cortex-M0 general-purpose processor, two identical 64 KB memories, and a programmable digitally controlled oscillator that generates the voltage-tracking adaptive-frequency clock. To gather measurement data for performance evaluation, the chip also includes numerous blocks for testing and debug purposes: a built-in self test (BIST) block allows thorough testing of the two memory blocks; a scan chain configures the digital blocks; a voltage monitor block probes internal voltages to record fast transients on the supply line; and a current-load generator enables different testing scenarios for load current-induced supply noise. Lastly, the prototype chip provides direct interfaces for external power and clock sources to set up different operating modes.

IV. CIRCUIT IMPLEMENTATION

A. Switched-Capacitor Integrated Voltage Regulator

The SC-IVR converts the battery voltage ($V_{BAT} \approx 3.7$ V) down to the digital supply ($DVDD \approx 0.7$ V). It consists of a cascade of two 2:1 switched-capacitor converters that are respectively optimized for high input voltage tolerance and fast load response and individually tuned to maximize conversion efficiency, and each converter stage employs a 16-phase topology to reduce voltage ripple. A low-boundary feedback control loop can regulate DVDD to a desired voltage level. A thorough discussion of the SC-IVR and its implementation details can be found in [13].

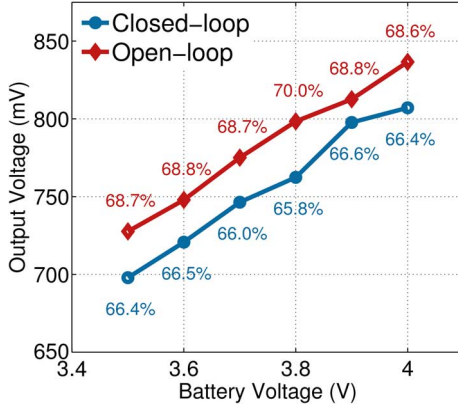


Fig. 5. Output voltages at peak conversion efficiency versus battery voltage for open- and closed-loop operation.

While this SC-IVR can achieve high conversion efficiency, 70% at its optimal operating point, it is subject to the inherent limitations of any switched-capacitor based dc-dc converter; efficiency varies with respect to input and output voltages and the load current level. As an example, Fig. 5 plots the SC-IVR's output voltages at different battery voltage levels for both open- and closed-loop operation, and the peak efficiency is labeled at each point. These efficiency numbers and output voltage levels are derived by sweeping the output voltage to find the peak conversion efficiency and its corresponding output voltage level at each fixed input battery voltage, and hence each point presents the best possible efficiency for each input and output voltage combinations. The data in Fig. 5 clearly shows that open-loop operation consistently offers 2 ~ 3% higher conversion efficiency and 16 ~ 30 mV higher output voltage levels than closed-loop operation, resulting in both higher efficiency and higher performance for the SoC.

Fig. 6 plots the transient behavior of the IVR with respect to load current steps between 3 mA and 50 mA for both open- and closed-loop operation. According to the transient waveform of the output voltage, the closed-loop operation quickly responds to avoid the steep voltage droop otherwise seen in the open-loop case; however, it exhibits larger steady-state voltage ripple, especially for higher load currents, due to the control loop implementation and feedback delay.

The voltage ripple in closed-loop operation is caused by the delay of the feedback control that prevents the regulator from instantaneous reaction when its output voltage drops below the reference voltage and thus can result in both undershoot and overshoot in the output voltage. Even with sophisticated feedback strategy and fast transistors in advanced technology node, it is difficult to reduce the feedback delay beyond 1 ns, and this fundamentally determines the magnitude of the steady-state closed-loop ripple. On the other hand, the voltage ripple in open-loop operation is due to charge sharing between the flying capacitor and the decoupling capacitor at the output node. It can be reduced by phase interleaving, and therefore does not experience the same fundamental limit as the closed-loop ripple. This is why in our 16-phase interleaved implementation of the IVR, the open-loop ripple magnitude is significantly smaller than the closed-loop ripple.

B. Digitally Controlled Oscillator (DCO)

Our proposed adaptive-frequency clocking scheme needs a clock generator whose frequency tracks closely with changes




Fig. 6. Transient response of SC-IVR output voltage to load current steps.

in supply voltage. This allows the operating frequency of the digital load circuitry to appropriately scale with voltage fluctuations, providing intrinsic resilience to supply noise. There are numerous examples of critical-path-tracking circuits for local timing generation [6], [10], [14]–[17]. Instead, we use a programmable digitally controlled oscillator to generate the system clock. The DCO contains a ring of programmable delay cells comprising of transmission and NAND gates that approximate a typical fanout-of-4 inverter delay. Such designs have been found to deliver decent tracking accuracy, acting as a proxy for critical path delay in complex digital logic [1], [9], [18].

As shown in Fig. 7(a), a 7-bit digital code, $D = D_6 \dots D_1 D_0$, sets the DCO frequency by selecting the number of delay cells in the oscillator loop. While our implementation uses 7 bits of control code, measurement results show that the lower 4 bits are sufficient for the normal operating range of the digital system. Fig. 7(b) plots DCO frequency versus supply voltage (DVDD) across a range of the digital control codes. Over the measured voltage range (0.6 to 1 V), frequency scales roughly linearly with supply voltage, but slightly flattens out for voltages below 750 mV. The uneven frequency spacing with respect to the control code D results from the delay cell's asymmetric design and should be improved for future implementation.

It is worth noting that the delay accumulation function introduced in Section II-D is not only a function of the supply voltage, but can also vary with other variation-inducing factors, such as process and temperature. The fanout-of-4 delay is a reasonably good approximation of typical critical path, which means it can track the delay sensitivity to process and temperature to some extent. However, since the main objective of our proposed adaptive clocking scheme is to improve supply resilience, we focus on verifying the matching quality of our DCO with respect to the supply voltage.

C. Cortex-M0 and Memory

Both the Cortex-M0 microprocessor and the memories used in the prototype SoC are IP blocks provided by our collaborators. To facilitate our testing strategy for different operation modes, the Cortex-M0 and one of the 64 KB memories (SRAM1) are intentionally designed to share the voltage

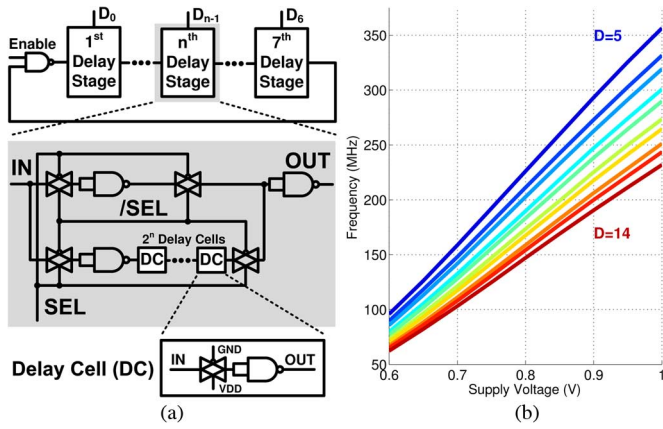


Fig. 7. Digitally controlled oscillator schematic and measured characteristics. (a) DCO schematic; (b) DCO frequency versus DVDD at digital control code D.

domain (DVDD) with the DCO, while the other memory (SRAM0) sits in another separate voltage domain (TVDD) with the rest of the test peripheral circuits. Except for being in a different voltage domain, SRAM0 shares the same physical design as SRAM1 as both are using the same hard memory IP.

A built-in self-test (BIST) module performs at-speed test for both SRAMs. To make sure full test coverage of all the memory cells, the BIST performs a modified MARCH-C routine that writes to and reads different data patterns from each memory address in the SRAMs successively. At the conclusion of the MARCH-C routine, the BIST module would raise a pass/fail flag if any read or write error has been detected. Due to the limited storage for testing vectors, only the address and data of the last failure are recorded for post-test analysis.

The test peripheral circuits operated off of TVDD include a voltage monitor block that captures fast nanosecond-scale transient changes of the supply line and a programmable load current generator that is made up of differently weighted switched current sources. Although external supply and clock interfaces such as TVDD, EXTVD, and EXTCLK are included, they are for testing purposes only, as the SoC is capable of operating directly off of the battery without any external supply or clock reference.

V. EXPERIMENTAL RESULTS

To demonstrate improved resilience and performance of the proposed adaptive clocking across a wide range of supply voltage, measurement results were obtained from a prototype SoC chip (Fig. 8) fabricated in TSMC's 40 nm CMOS technology. We use the maximum error-free operating frequency of the memory performing built-in self-test as a proxy metric, because it is often the on-chip SRAM sharing the same voltage domain with the digital logic that limits the system performance at lower supply levels. Also, the retention voltage of the SRAM cells typically determines the minimum operating voltage of the system [19].

This section presents the following set of experimental results: First, we characterize the voltage versus frequency relationship of the SRAMs using external sources in order to determine the efficacy of using the DCO for adaptive-frequency clocking. Then, we compare the fixed- and variable-frequency clocking schemes with a regulated voltage generated by the




Fig. 8. Die photo of the fully integrated system-on-chip.

SC-IVR in closed-loop operation. Lastly, we present the advantages of combining adaptive clocking with a variable voltage provided by operating the SC-IVR in open loop.

A. Frequency versus Voltage Characterization

The on-chip SRAMs were characterized at static supply voltage levels provided externally via EXTVD, in order to determine the SRAM's voltage to frequency relationship under quiet supply conditions. Using an external clock (EXTCLK) at different fixed frequencies, we obtained the Shmoo plot in Fig. 9(a). It shows (1) the minimum retention voltage of SRAM cell is between 0.6 V and 0.65 V; (2) the maximum SRAM frequency scales roughly linear with supply voltage and ranges from 68 MHz at 0.65 V to 256 MHz at 1.0 V; and (3) the maximum SRAM frequency closely correlates with DCO frequency plot for control code $D = 10$. This correlation suggests the FO4-delay-based DCO tracks the critical path delay in the SRAM across a wide supply range and should enable error-free memory BIST for control word D above 10. We experimentally verified this by turning on the internal DCO to provide the system clock instead of using EXCLK and sweeping the same voltage range via EXTVD. The resulting Shmoo plot in Fig. 9(b) further demonstrates the DCO's ability to track SRAM delay at different static supply voltage levels. Considering process and temperature variation, the same control word $D = 10$ may not apply to chips from all process corners and over all temperature ranges. However, since the process and temperature conditions are relatively static, and our results show that a fixed control word can cover the range of fast-changing supply variation, it is possible to determine this fixed control word during the calibration phase before the normal operation of the robotic system starts.

B. Fixed versus Adaptive Clocking With Regulated Voltage

Having verified the DCO, we now compare fixed- and adaptive-frequency clocking schemes for a system that operates off of a regulated voltage with the SC-IVR operating in closed loop. We also emulate noisy operating conditions using the on-chip I_{LOAD} generator that switches between 0 and 15 mA at 1 MHz. Measurements made via the on-die voltage monitoring circuit showed approximately ± 70 mV worst-case ripple about a mean voltage of 0.714 V.

For the conventional fixed-frequency clocking scheme, the maximum operating frequency ought to depend on the worst-case voltage droop, measured to be 0.647 V. Using the measured relationship in Fig. 9(a), the maximum frequency cannot exceed 68 MHz. To measure the actual maximum error-free frequency, we performed 100 independent BIST runs using the external clock, EXTCLK, set to a fixed frequency and recorded the failure rate. Fig. 10(a) summarizes the measured failure rates across different externally driven operating frequencies. These

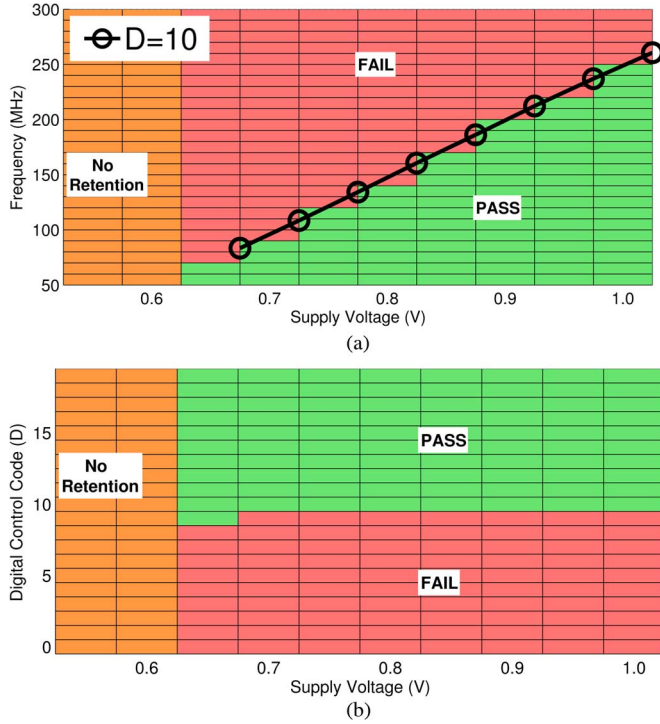


Fig. 9. Shmoo plots for two different clocking schemes. (a) External clock at fixed frequencies; (b) internal adaptive clock generated by the DCO at different control code (D).

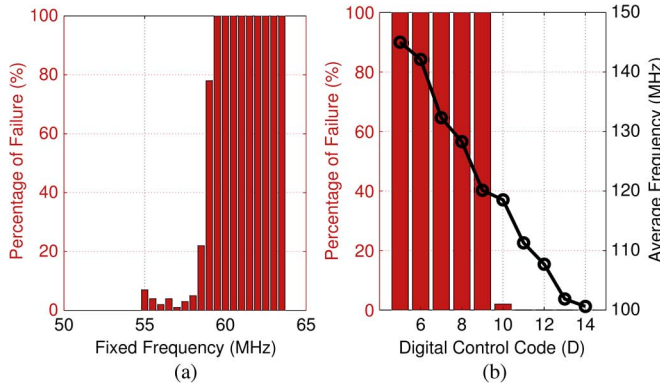


Fig. 10. Comparison of memory BIST failure rates for fixed- versus adaptive-frequency clocks, but under the same IVR closed-loop regulation. (a) Fixed-frequency externally; (b) adaptive-frequency from DCO.

results show that the maximum error-free frequency is below 55 MHz for the fixed-frequency clocking scheme, which is even lower than the anticipated 68 MHz, perhaps attributable to the additional noise injection.

Using the same IVR configuration and test conditions, Fig. 10(b) plots the failure rates versus different digital control codes of the DCO. At $D = 10$, there were intermittent failures, attributable to the additional noise not present in the prior experimental results of Fig. 9(b). The adaptive-frequency clocking scheme delivers consistent and reliable operation at $D = 11$. Based on the average DCO frequency measured during the tests and also plotted in Fig. 10(b), $D = 11$ corresponds to an average frequency of 111 MHz, which is $2\times$ the fixed-frequency clocking scenario. Here we use the average frequency instead of the worst frequency as a measure of system performance, because the operating frequency, which tracks the noisy supply, changes at much faster time scale than the task performance requirement of the robot, which means that

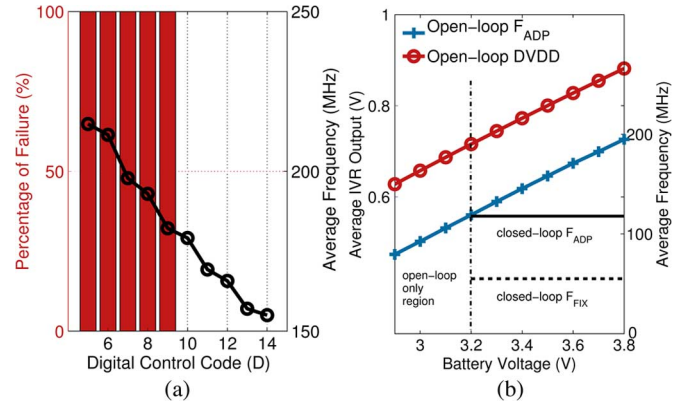


Fig. 11. Performance under unregulated voltage from open-loop SC-IVR operation. (a) Failure rate and average frequency versus DCO settings; (b) measured average clock frequency and output voltage.

the aggregated throughput measured by the average frequency is the more meaningful metric and occasional drop to lower frequency does not have destructive effect on the system.

We attribute this large frequency difference between the two clocking scenarios to a couple of factors. Fixed-frequency clocking requires sufficiently large guardbands to guarantee operation under the worst-possible voltage droop condition. In contrast, adaptive-frequency clocking allows both the clock period and load circuit delays to fluctuate together as long as both vary with voltage in a similar manner. Hence, the guard-band must only cover voltage-tracking deviations between the DCO and load circuit delay paths across the operating voltage range of interest, and can be built into the DCO. Another factor that penalizes the performance of the fixed-frequency clocking comes from the additional noise on the external clock signal for crossing the TVDD to DVDD boundary.

C. Adaptive-Frequency Clocking With Unregulated Voltage

We now turn our attention to how adaptive-frequency clocking performs with an unregulated voltage generated by the SC-IVR operating in open loop. We used the same test setup with $V_{BAT} = 3.7$ V and noise injection via the on-chip I_{LOAD} generator. The failure rates and the average frequencies are captured in Fig. 11. Compared to the measured results in Fig. 10(b), average frequencies are much higher, because DVDD settles to higher values (≈ 0.8 V) when the SC-IVR operates in open loop. Despite the high susceptibility to fluctuations on DVDD to load current steps as seen in Fig. 6, Fig. 11(a) shows zero errors occurred even for $D = 10$. The higher DVDD voltage provides more cushion to avoid intermittent retention failure.

In order to illustrate the extended operating range offered by running the SC-IVR in open loop, Fig. 11(b) plots the average DCO frequency and average DVDD voltage for error-free operation versus battery voltage. These measurements were again made with 0 to 15 mA current load steps. As expected, the open-loop SC-IVR's average output voltage scales proportional to the battery voltage. Moreover, the system can operate error-free even for battery voltages below 3 V, which approaches the 2.5–2.7 V lower discharge limit of Li-ion batteries. In comparison, assuming a target SC-IVR regulated voltage of 0.7 V, the system would only operate down to a battery voltage of 3.2 V and at a lower frequency across the battery discharge profile even with the adaptive-frequency clocking scheme. A fixed-frequency clocking scheme would lead to even lower performance.

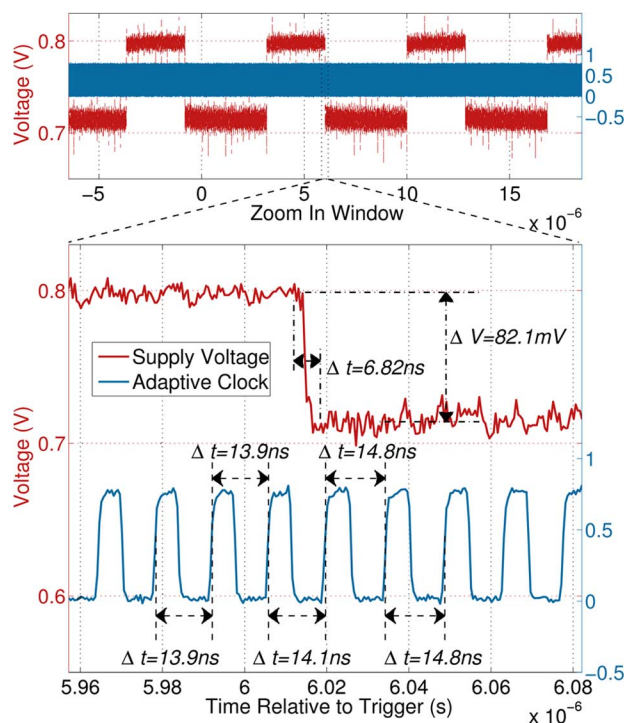


Fig. 12. Waveforms of the supply voltage (DVDD) and a divide-by-2 clock signal (CLKOUT) during IVR open-loop operation with load current (I_{load}) periodically switching between 5 mA and 30 mA.

Finally, we look at the transient waveform of the supply voltage captured by the internal voltage monitor when the SC-IVR operates in open-loop with the same periodic load step condition used in Fig. 11(a). Fig. 12 shows both CLKOUT, which is a divide-by-2 signal of the internal system clock, and the supply voltage (DVDD). The DCO's control code is set to 10 as suggested by Fig. 11(a). In the zoom-in window of the waveforms, it clearly illustrates that the DCO frequency can respond promptly to a 82.1 mV supply droop within 6.82 ns, as the load current steps from 5 mA to 30 mA, so that no memory access error is recorded by the BIST module, suggesting superior supply-noise resilience of the SoC system due to the adaptive clocking scheme. Moreover, the waveform of the supply voltage demonstrates that supply noise in a typical microrobotic SoC with IVR is characterized more by the voltage droop and ripple in response to the load current steps and the IVR switching, instead of resonant noise cause by the LC tank in the power delivery path.

In addition to validating the resilience and the performance advantages of adaptive-frequency clocking, our experimental results also reveal the synergistic properties between the clocking scheme and the IVR design in a battery-powered microrobotic SoC. The supply-noise resilience provided by an adaptive clock alleviates design constraints imposed by voltage ripple and voltage droop. Therefore, the IVR can trade off its transient response for better efficiency or smaller area when co-designed with adaptive-frequency clocking.

VI. CONCLUSION

An adaptive-frequency clocking scheme offers several advantages when combined with an IVR in a battery-powered microrobotic SoC. We are able to thoroughly explore the relative trade-offs and merits associated with different mode

configurations of the IVR and the clock using a prototype SoC chip fabricated in 40 nm CMOS and an aerial microrobot as the application platform. For regulated voltage operation via closed-loop SC-IVR, adaptive-frequency clocking enables $2\times$ performance improvement, compared to conventional fixed-frequency clocking scheme. Combining adaptive-frequency clocking with an unregulated voltage via open-loop IVR extends the operating range across a wider portion of the battery's discharge profile. The noise resilience demonstrated by the adaptive-frequency clocking scheme calls for co-design and co-optimization of clock generation and voltage regulation in weight-and-power constraint microrobotic systems that exhibit distinctive supply noise characteristics dominated by fast load current changes and slow battery discharge.


ACKNOWLEDGMENT

The authors thank the TSMC university shuttle program for chip fabrication.

REFERENCES

- [1] G.-Y. Wei and M. Horowitz, "A fully digital, energy-efficient, adaptive power-supply regulator," *IEEE J. Solid-State Circuits*, vol. 34, no. 4, pp. 520–528, 1999.
- [2] Y. K. Ramadass and A. P. Chandrakasan, "Minimum energy tracking loop with embedded dc/dc converter enabling ultra-low-voltage operation down to 250 mV in 65 nm cmos," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 256–265, 2008.
- [3] C. R. Lefurgy, A. J. Drake, M. S. Floyd, M. S. Allen-Ware, B. Brock, J. A. Tierno, and J. B. Carter, "Active management of timing guard-band to save energy in power7," in *Proc. 44th Annual IEEE/ACM Int. Symp. Microarchitecture*, New York, USA, 2011, ser. MICRO-44 '11, pp. 1–11 [Online]. Available: <http://doi.acm.org/10.1145/2155620.2155622>
- [4] D. Jiao, B. Kim, and C. H. Kim, "Design, modeling, test of a programmable adaptive phase-shifting pll for enhancing clock data compensation," *IEEE J. Solid-State Circuits*, vol. 47, no. 10, pp. 2505–2516, 2012.
- [5] M. Karpelson, J. P. Whitney, G.-Y. Wei, and R. J. Wood, "Energetics of flapping-wing robotic insects: Towards autonomous hovering flight," in *Proc. IROS*, 2010, pp. 1630–1637.
- [6] A. Drake, R. Senger, H. Deogun, G. Carpenter, S. Ghiasi, T. Nguyen, N. James, M. Floyd, and V. Pokala, "A distributed critical-path timing monitor for a 65 nm high-performance microprocessor," in *Proc. ISSCC*, 2007, pp. 398–399.
- [7] Y. Kim, L. K. John, S. Pant, S. Manne, M. Schulte, W. L. Bircher, and M. S. S. Govindan, "Audit: Stress testing the automatic way," in *Proc. MICRO*, Dec. 2012, pp. 212–223.
- [8] D. Jiao, J. Gu, and C. H. Kim, "Circuit design and modeling techniques for enhancing the clock-data compensation effect under resonant supply noise," *IEEE J. Solid-State Circuits*, vol. 45, no. 10, pp. 2130–2141, 2010.
- [9] 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, 2000.
- [10] M. E. Dean, "STRIP: A Self-Timed RISC Processor," Ph.D. dissertation, Stanford Univ., Stanford, 1992.
- [11] K. Y. Ma, P. Chirarattananon, S. B. Fuller, and R. J. Wood, "Controlled flight of a biologically inspired, insect-scale robot," *Science*, vol. 340, no. 6132, pp. 603–607, 2013.
- [12] M. Lok, D. Brooks, R. Wood, and G.-Y. Wei, "Design and analysis of an integrated driver for piezoelectric actuators," in *Proc. ECCE*, 2013, pp. 2684–2691.
- [13] T. Tong, X. Zhang, W. Kim, D. Brooks, and G.-Y. Wei, "A fully integrated battery-connected switched-capacitor 4:1 voltage regulator with 70% peak efficiency using bottom-plate charge recycling," in *Proc. CICC*, 2013.
- [14] M. Nomura, Y. Ikenaga, K. Takeda, Y. Nakazawa, Y. Aimoto, and Y. Hagiwara, "Delay and power monitoring schemes for minimizing power consumption by means of supply and threshold voltage control in active and standby modes," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 805–814, 2006.

- [15] X. Wang, M. Tehranipoor, and R. Datta, "Path-ro: A novel on-chip critical path delay measurement under process variations," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, Piscataway, NJ, 2008, ser. ICCAD '08, pp. 640–646 [Online]. Available: <http://dl.acm.org/citation.cfm?id=1509456.1509597>
- [16] I. J. Chang, S. P. Park, and K. Roy, "Exploring asynchronous design techniques for process-tolerant and energy-efficient subthreshold operation," *IEEE J. Solid-State Circuits*, vol. 45, no. 2, pp. 401–410, 2010.
- [17] J. Park and J. A. Abraham, "A fast, accurate and simple critical path monitor for improving energy-delay product in dvs systems," in *Proc. ISLPED*, 2011, pp. 391–396.
- [18] Y. Ikenaga, M. Nomura, S. Suenaga, H. Sonohara, Y. Horikoshi, T. Saito, Y. Ohdaira, Y. Nishio, T. Iwashita, M. Satou, K. Nishida, K. Nose, K. Noguchi, Y. Hayashi, and M. Mizuno, "A 27% active-power-reduced 40-nm CMOS multimedia SoC with adaptive voltage scaling using distributed universal delay lines," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 832–840, 2012.
- [19] S. Jain, S. Khare, S. Yada, V. Ambili, P. Salihundam, S. Ramani, S. Muthukumar, M. Srinivasan, A. Kumar, S. K. Gb, R. Ramanarayanan, V. Erraguntla, J. Howard, S. Vangal, S. Dighe, G. Ruhl, P. Aseron, H. Wilson, N. Borkar, V. De, and S. Borkar, "A 280 mV-to-1.2 V wide-operating-range 1A-32 processor in 32 nm CMOS," in *Proc. ISSCC*, 2012, pp. 66–68.



Xuan Zhang (S'08) received the B.Eng. degree from Tsinghua University, Beijing, China, and the Ph.D. degree from Cornell University, Ithaca, NY, USA. She is currently a Postdoctoral Fellow in computer science at the Harvard School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA.

Her research interest includes energy-efficient and highly reliable computing systems for embedded and high performance applications. She is the recipient of Intel PhD Fellowship in 2008. In the summers of 2008 and 2010, she interned at Broadcom Central Engineering Center and Schlumberger Research Center respectively, where she worked on reference buffer design and wireline communication system prototyping.



Tao Tong (S'10) received the B.E. degree from Tsinghua University, Beijing, China, and the M.S. degree from Oregon State University. He is currently a Ph.D. student in electrical engineering at the Harvard School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA.

He worked at MediaTek Wireless Inc. and Lion Semiconductor Inc., designing analog-to-digital converters and fully integrated dc-dc converters for mobile applications. His research interests include integrated voltage regulators and their applications in energy efficient computing systems.



David Brooks (M'02) received the B.S. degree in electrical engineering from the University of Southern California, Los Angeles, CA, USA, and the M.A. and Ph.D. degrees in electrical engineering from Princeton University, Princeton, NJ, USA.

He is a Gordon McKay Professor of Computer Science in the School of Engineering and Applied Sciences at Harvard University, Cambridge, MA. He joined Harvard in 2002 after spending one year as a research staff member at IBM T. J. Watson Research Center. His research interests include resilient and

power-efficient computer hardware and software design for high-performance and embedded systems.

Prof. Brooks is a member of ACM.



Gu-Yeon Wei (M'00) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1994, 1997, and 2001, respectively.

He is a Gordon McKay Professor of Electrical Engineering and Computer Science in the School of Engineering and Applied Sciences (SEAS) at Harvard University. His research interests span multiple layers of a computing system: mixed-signal integrated circuits, computer architecture, and runtime software for automatic code parallelization.

Particular efforts focus on research opportunities across these layers to develop energy-efficient solutions for a broad range of systems from flapping-wing microrobots to large-scale servers.