

## A Multi-Chip System Optimized for Insect-Scale Flapping-Wing Robots

Xuan Zhang\*, Mario Lok\*, Tao Tong, Simon Chaput, Sae Kyu Lee, Brandon Reagen, Hyunkwang Lee, David Brooks, Gu-Yeon Wei  
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, U.S.A

### Abstract

We demonstrate a battery-powered multi-chip system optimized for insect-scale flapping wing robots that meets the tight weight limit and real-time performance demands of autonomous flight. Measured results show open-loop wing flapping driven by a power electronics unit and energy efficiency improvements via hardware acceleration.

### Introduction

Recent demonstrations of an insect-scale flapping-wing robot (Harvard's RoboBee) achieved controlled flight—hovering and maneuvering along three axes—relying on an external motion capture system, a bench-top high-voltage amplifier to energize piezoelectric (PZT) actuators that flap its wings, and a desktop computer for computation [1]. The ultimate goal of autonomous flight requires converting the external test equipment into customized components that the robot can carry within its tight payload budget. Towards this end, the paper presents an energy-efficient BrainSoC to process sensor data and send wing flapping control signals to a power electronics unit that generates 200-300V sinusoids for driving a pair of piezoelectric actuators to individually flap each wing (Fig. 1).

### BrainSoC Architecture

Fig. 1 provides a detailed overview of the 3mg BrainSoC, which is the central controller required for flight. To minimize weight, the SoC integrates peripheral support circuits to obviate external components other than a battery. A fully-integrated, two-stage 4:1 switched-capacitor regulator (IVR) steps the 3.7V battery voltage down to 1.8V and 0.9V [2], resulting in three distinct voltage domains. On-chip oscillators generate multiple clock signals for a general-purpose core, accelerators, peripherals, and the IVR. A high-precision 10MHz relaxation oscillator sets the wing-flapping frequency. A 32-bit ARM Cortex-M0 handles general computing needs and is master of the AMBA high-speed bus that connects to various memories. An embedded ADC subsystem that comprises of four 8b SAR ADCs multiplexed between 16 analog inputs, combined with I2C, SPI, and GPIO controllers on the peripheral bus enable connection to a multitude of sensor chips (inertial measurement unit, vision, etc.). To meet real-time performance demands of autonomous flight in an energy efficient way, the M0 coordinates hardware accelerators (convolution filter, image interpolation optical flow (IIOF), DSP engines, and actuator controller) via memory-mapped registers.

### Power Electronics Unit

The 60mg power electronics unit (PEU) takes the 3.7V battery voltage and generates 200-300V sinusoidal signals to drive PZT actuators and flap its wings at 80-120Hz, the mechanical resonant frequency of the wing structure. The PEU comprises an off-chip tapped-inductor boost converter to generate a 200-300V VDDH. A 3.75mg driver IC integrates two high-voltage half-bridge drive circuits to produce two sinusoidal drive signals (one per wing) off of VDDH. An actuator control accelerator in the BrainSoC produces control signals to pulse the drive circuits, similar to [3,4]. Fig. 2 shows the drive signals needed for left and right wings to enable flight with three degrees of freedom: roll, pitch, and yaw. Asymmetric left and right wing stroke amplitudes and offsets lead to roll and pitch. Asymmetric up and down strokes yields yaw rotation. To save area and maximize power density, the drive circuits use n-type DMOS (nDMOS) transistors as high-side switches, which are 10x more area efficient than p-type DMOS (pDMOS) switches of equal resistance. Fig. 3 shows a fully integrated floating supply (VDDF) that moves with the driver's output ( $V_{SW}$ ), resolving the major problem of using n-type, high-side switches. An integrated regulator keeps VDDF above 4V and powers the buffers that drive the n-type, high-side switch. Measured transient waveforms verify the floating supply stays above 4V as driver output voltage rises. Fig. 4 shows video capture snapshots of wing flapping when driven by the PEU with open loop commands from the BrainSoC's M0 and actuator control accelerator. Although the M0 alone is sufficient for open-loop wing control, autonomous flight requires compute-intensive feedback control, motivating the need for hardware acceleration.

### Hardware Acceleration

Fig. 5 lays out the computation required to close the feedback loop with sensors and actuators for autonomous flight. There are four types: 1) *image processing* acquires images from a vision sensor one pixel at a time, uses a convolution filter for edge sharpening, and optical flow for position estimation; 2) *rotation estimation* uses low-pass filtered gyroscope readings from an inertial measurement unit (IMU); 3) *body control's* feedback loop takes outputs from image processing and rotation estimation to feed an adaptive PID controller that determines the rotation action needed to stabilize and move the robot; 4) *actuator control* takes the roll, pitch, and yaw rotation commands from body control and generates pulse signals for the PEU. Initial tests identified minimum speed requirements for these computations, labeled in Fig. 5, which are not easily met by the M0. Fig. 6 presents the aggregate number of cycles needed for the above computations as a percentage of M0 CPU cycles running at 220MHz. Image processing and body control dominate and the M0 maxes out for image size 32x32. Finer image resolution, higher frame rates, or higher control loop update rates, all of which can improve flight performance, are not possible with the M0 alone. To alleviate this problem, we designed multiple hardware accelerators. Given their fixed and well-defined algorithms, convolution filter (Conv.) and IIOF are dedicated monolithic accelerators. Parameters in the body control algorithm require constant tweaking and are best if flexible. Therefore, the M0 implements body control, but offloads atomic arithmetic operations, such as dot product, matrix multiply, and filtering, to the DSP engines.

Measured results confirm the computation efficiency advantages of these hardware accelerators. Fig.7 compares dynamic power across frequencies for synthetic workloads run on the M0 versus the accelerators. The M0's power does not change with workload due to the simplicity of its in-order 3-stage pipeline. Accelerators not only consume less dynamic power across the board, but they also reduce execution time, significantly. This advantage is best seen via dynamic energy comparisons. Simple arithmetic functions in the DSP engine yield ~10x energy improvements (Fig. 8). Monolithic convolution filter and IIOF accelerators offer over two orders of magnitude improvement (Fig. 9). While difficult to implement monolithically, we can accelerate body control using the DSP engines to gain more than 2x energy improvement.

To evaluate flight worthiness of the BrainSoC's computing capabilities, we constructed a synthetic autonomous flight workload that represents a flight experiment by combining the computations outlined in Fig. 5. Executing this synthetic workload with and without accelerators further highlights the advantages of acceleration. In these measurements, the frame rate is 100 frame/s, the image size is 32x32 pixels, and we sweep the body control's feedback loop frequency. Fig. 10 plots the system clock frequency versus control loop frequency. With accelerators, the system requires a 60MHz system clock frequency to achieve the 1500Hz loop frequency limit compared to 190MHz needed when computing with only the M0. Fig.11 plots the resulting power consumption of the BrainSoC for different voltage scaling scenarios. Reduction in system clock frequency down to 60MHz with accelerators allows for much more aggressive voltage scaling, reducing the BrainSoC's digital power from 24.8mW at 0.84V down to 4.2mW at 0.63V. In addition to the quadratic reduction in dynamic power, leakage current reduces from 17.2mA at 0.84V down to 5.2mA at 0.63V. Using the IVR to power the SoC increases the overall power due to IVR losses and further degradation in conversion efficiency when regulating to lower voltage levels [2]. However, this higher power cost is offset by the reduction in weight. In summary, this multi-chip system fulfills the weight budget and real-time computation demands for future autonomous RoboBee flight experiments.

### Acknowledgements

Authors thank TSMC's university shuttle program for chip fabrication. This work was supported in part by NSF (IIS-0926148 and CCF-1218298) and DARPA (HR0011-13-C-0022).

\* These authors have equal contributions.

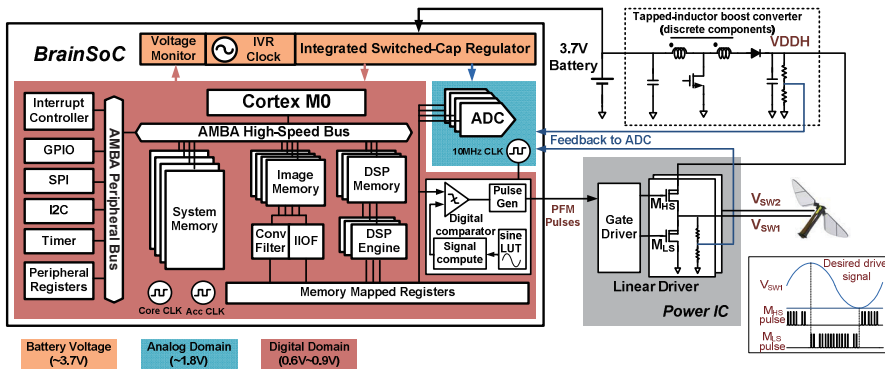


Fig. 1: The robotic multiple chip module consists of the BrainSoC and the power electronic unit (Power IC + tapered-inductor boost converter).

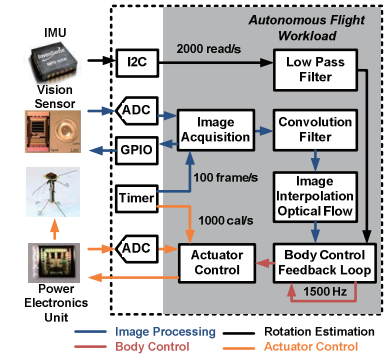


Fig. 5: The autonomous flight experiment requires running multiple computational workloads concurrently.




Fig. 2: (a) Roll, pitch, and yaw rotation for the "RoboBee". (b) Drive signal waveforms for roll, pitch, and yaw rotation.




Fig. 6: Aggregated cycle count of the computation workload as a percentage of M0 cycles at 220MHz.




Fig. 7: The dynamic power of running different workloads in Cortex M0 (CM0) and accelerators (Acc) as frequency scales.

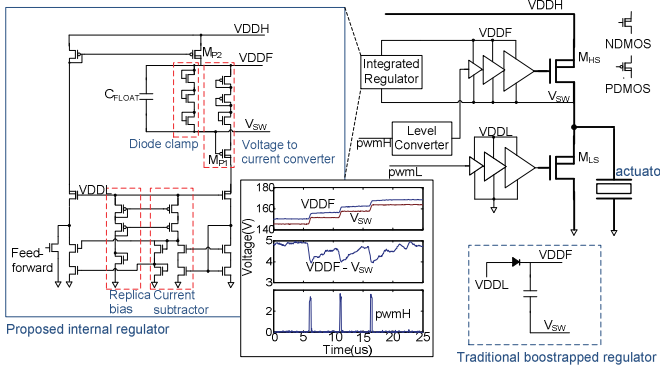


Fig. 3: The floating supply powers the high-side NMOS gate driver. An internal regulator ensures the floating supply voltage is above 4V, while a traditional bootstrapped regulator does not work for capacitive loads.




Fig. 4: Slow motion video capture snapshots of RoboBee's right wing (left wing is stationary) flapping driven by the PEU with open loop commands from the BrainSoC's M0 and actuator control accelerator.

### References

- [1] K. Ma, et al., "Controlled flight of a biologically inspired, insect-scale robot," *Science*, May 3, 2013.
- [2] T. Tong, et al., "A fully integrated battery-connected switched capacitor 4:1 voltage regulator with 70% peak efficiency using bottom plate charge recycling," *CICC*, 2014.
- [3] M. Lok, et al., "Design and analysis of an integrated driver for piezoelectric actuators," *ECCE*, 2013.
- [4] M. Karpelson, "Low power control IC for efficient high-voltage piezoelectric driving in a flying robotic insect," *VLSI*, 2011.




Fig. 8: The dynamic energy of different DSP workloads.




Fig. 9: The dynamic energy of Conv., IIOF, and body control




Fig. 10: The minimum system clock frequency required at different body control loop frequency.




Fig. 11: SoC power consumption.




Fig. 12: Annotated die photos

Table I CHIP SUMMARY

BrainSoC Characteristics	
Technology	IPI10M 40nm CMOS
VDD: Core, I/O	0.6-0.9V, 1.8V
Clock Frequency	30-220MHz
Die Size	1.85mm x 2.9mm
Logic Cell Count	200K
Memory	100KB
Power	M0+Accelerators 7.5mW @ 0.63V
	M0 Only 41.6mW @ 0.84V
Leakage Current	5.2mA @ 0.63V
	17.2mA @ 0.84V
Weight	3mg (die thinning)
PEU Characteristics	
Technology	0.8µm 300V BCD
Active Area	4mm <sup>2</sup>
DC-DC efficiency	63% for 300V output
Power	200mW
Driver Stage Weight	3.75mg (die thinning)
Projected Total Weight	60mg