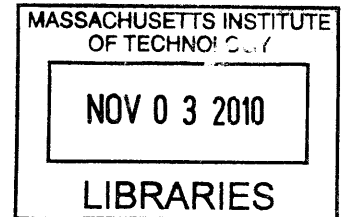**A Modular Neural Interface for Massively Parallel Recording and Control:**

Subsystem Design Considerations for Research and Clinical Applications.

by

Christian T. Wentz

S.B. Electrical Science & Engineering, M.I.T., 2009

Submitted to the Department of Electrical Engineering and Computer Science

In Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science

At the Massachusetts Institute of Technology

May, 2010
[June 2010]

Author _
Department of Electrical Engineering and Computer Science
May 21, 2010

Certified by
Dr. Edward S. Boyden, Assistant Professor
MIT Media Lab
Departments of Biological Engineering and Brain and Cognitive Science
Thesis Supervisor

Accepted by
Dr. Christopher J. Terman
Chairman, Department Committee on Graduate Theses

**A Modular Neural Interface for Massively Parallel Recording and Control:
Subsystem Design Considerations for Research and Clinical Applications.**

by
Christian T. Wentz

Submitted to the
Department of Electrical Engineering and Computer Science

May 21, 2010

In partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

The closed-loop Brain-Machine Interface (BMI) has long been a dream for clinicians and neuroscience researchers alike – that is, the ability to extract meaningful information from the brain, perform computation on this information, and selectively perturb neural dynamics in the brain for therapeutic benefit to the patient. Such systems have immediate application to treatment of paralysis, epilepsy and the amputated, and the potential for treatment of higher order cognitive dysfunction. Despite the promise of the BMI concept, the technology for bidirectional communication with the brain at sufficiently large scale to be truly therapeutically useful is lacking. Current state-of-the-art neuromodulation systems deliver open loop, 16-channel patterned electrical stimulation incapable of precisely targeting small numbers of neurons. Large-scale neural recording systems are limited to 16-128 electrodes, at the cost of several thousand dollars per channel. The ability to record from the awake behaving animal – let alone precisely modulate neural network dynamics in closed-loop fashion– presents a substantial challenge today.

In this thesis, I present decoupled design solutions for three critical subcomponents of the closed-loop BMI – (i) a highly miniature, wirelessly powered and wirelessly controlled implantable optogenetic neuromodulation system capable of selective neural network control with single neural subtype- and millisecond-timescale precision, (ii) a prototype, highly parallel and scalable bio-potential recording system for simultaneous monitoring of many thousands of electrodes, and (iii) a space- and energy-efficient battery charger for biomedical applications. In aggregate, these systems overcome many of the fundamental architectural problems seen in the research and clinical environment today, potentially enabling a new class of neuromodulation system capable of treatment of higher-order cognitive dysfunction. In the research setting, these systems may be scaled to enable whole-brain recording, potentially yielding insights into large-scale neural network dynamics underlying disease and cognition.

Thesis Supervisor: Edward S. Boyden
Title: Assistant Professor, MIT Media Lab

TABLE OF CONTENTS

TABLE OF FIGURES

LIST OF TABLES

1 Forward

Engineering is a discipline obsessed with abstraction. Compartmentalization of the design problem has proven wildly successful, in that it provides the engineer with a manageable set of variables and assumptions to optimize and test, working towards a larger, more meaningful goal. Such a design paradigm has reduced design the end-to-end design cycles of Intel to 18 months at the time of this writing, and this cycle time continues to drop. Yet, it seems that far too often engineers become obsessed with optimizing *for the sake of optimization*, pushing the limits of one minute problem without necessarily considering its importance relative to the larger goal.

Particularly in early, emergent applications of the engineering toolset to fields like neuroengineering though, the theme of this work, it is critical that the engineer work fluidly up and down the ladder of abstraction to identify areas in which most ground gain be gained, focusing on the problem to solve over the toolset itself. It is from this perspective that I tackle the design problems in this work – the first two projects, Wireless Optical Neural Control and Massively-Parallel Wired, focus on systems-level optimization to synthesize new tools for neuroscience research from predominantly existing technology solutions, with an eye towards long-term clinical impact. Schematics and code are therefore reserved almost exclusively for the appendix, though I make an effort to point out the hard-earned gritty details of the design solutions where appropriate. The third project, An Ultra-Efficient Biomedical Battery Charger, addresses one of the large problems in

existing technology solutions for biomedical applications, namely how to maximize the

energy available to implanted circuitry.

2   Acknowledgements

Many people deserve credit for helping to develop and implement the ideas presented in this work. I'd like to thank in particular Prof. Boyden – many of his thoughts went into the seeding of Wireless and Wired, both in the big picture sense and in the execution of the small details – he was bold enough to grant me the autonomy to take on such projects and provide resources wherever needed to ensure their success.

For Wireless, the LED arrays themselves are a product of Jake Bernstein, with some modifications from myself to allow compatibility with the system electronics. Alex Guerra also helped in the actual fabrication of these arrays. The animal work on Parkinson's Disease using Wireless was performed by Patrick Monahan with my oversight. The fabrication of the receiver and transmit coils for wireless power coupling, and the design and fabrication of the transmit coil driver was performed by Ferro Solutions, Inc. Jesse Simon was instrumental in this effort, along with Yiming Liu, JK Huang, Bob O'Handley and Matthew Farrell.

For Wired, Al Strelzoff provided great insight into high-speed analog signal acquisition and has been a great mentor overall. While not explicitly described in this thesis, the full-scale implementation of Wired, including GPU-accelerated closed loop analysis, is a join effort with Brian Allen.

The VLSI work described in section 5 was borne out of a class project for Prof. Rahul Sarpeshkar's Low Power Analog VLSI course, and the work was executed as a joint effort with Bruno Do Valle, now a PhD candidate in Prof. Sarpeshkar's lab. Work has continued on this project to further improve its efficiency.

I must finally thank my parents for nurturing my curiosity for such ideas and the ambition to undertake them, and for providing encouragement along the way.

3 Wireless Optical Neural Control of Freely-Moving Animals

3.1 Project Summary

The ability to control specific neural cell types or pathways with optical methods, utilizing 'optogenetic' sensitizers that make the electrical signaling in specific neural circuits controllable by light, is enabling the parsing of the causal substrates of normal and pathological brain functions in mouse and non-human primate models of neurological disease. In order to enable such studies in complex behaviors, longitudinal studies, and multi-animal studies, wireless control of and power delivery to implanted light sources in the awake behaving animal would be of great use.

Moreover, the ability to both optically modulate neural activity and record the resulting electrical activity of the neural network using such a wireless paradigm would enable systematic, high-throughput input-output screening of neural dynamics, providing critical insight into the mechanisms of action of neural disease and higher-order cognition.

To this end, I present a wirelessly powered and wirelessly controlled headborne system capable of simultaneously driving multiple LEDs and recording neural activity in the awake behaving animal. By utilizing robust commercially available components and taking advantage of chip-scale packaging, a prototype of the device presented here has demonstrated reliable optical stimulation in mice for greater than 3 months while weighing in total approximately 3 grams. This device negates the need for tethered stimulation systems and associated limitations on animal behavior.

Under continuous optical stimulation, the implant is capable of delivering >2W of power, sufficient to drive two 700um x 700um LEDs at 100% duty cycle at programmable frequency indefinitely (~500mW/mm^2 to cortical targets), while recording from as many as 4 electrodes simultaneously. Note that this is a very high level of power; for surface LEDs atop cortex, dozens can be run indefinitely. Intermittent LED input power can be increased to 5W for one second, enough to drive 5 LEDs at 100% duty cycle or 10 LEDs at 50% duty cycle. Intermittent high power delivery above the wirelessly supplied 2W is achieved using an onboard super capacitor energy storage cell with >4 Joules of reserve energy, thus a scaling law for burst mode power of 2W + 4W*seconds/duty ratio. Current implementations support a maximum of 2 driven LEDs at any time, though additional driver channels are easily added with a nominal increase in weight.

Wireless power is delivered at low magnetic field (400 A/m) and low frequency (120kHz) to reduce the possible side effects of magnetic field exposure. Power transfer is maintained at high efficiency (10%) by utilizing a synchronized bridge driving circuit and precision frequency tuning. The receiver is also optimized for power efficiency over a wide range of the animal's body orientation by utilizing an axial resonant LC receiver coil with high permeability core and under-cage transmit coils. The power receiver circuit is optimized to produce a maximum of 2W of received power, achieved at a height of approximately 2 cm above the cage floor.

The modular design of the implant is broken into 4 distinct, removable subsystems: (i) a skull-fixed LED array with thermal sink, optional fiber light guides and optional 4-channel electrode pre-amplifier with 16-pin universal connector, (ii) an implant motherboard with LED drivers, power management stages and ultra-low power 16 MHz integrated microprocessor (Texas Instruments MSP430), (iii) 2.4 GHz data telemetry link (TI/Chipcon CC2400), and (iv) wireless power rectifier. When not in use, components (ii-iv) can be removed from the head and subsystems can be upgraded, repaired or interchanged with other implanted LED arrays without significant disturbance to the animal.
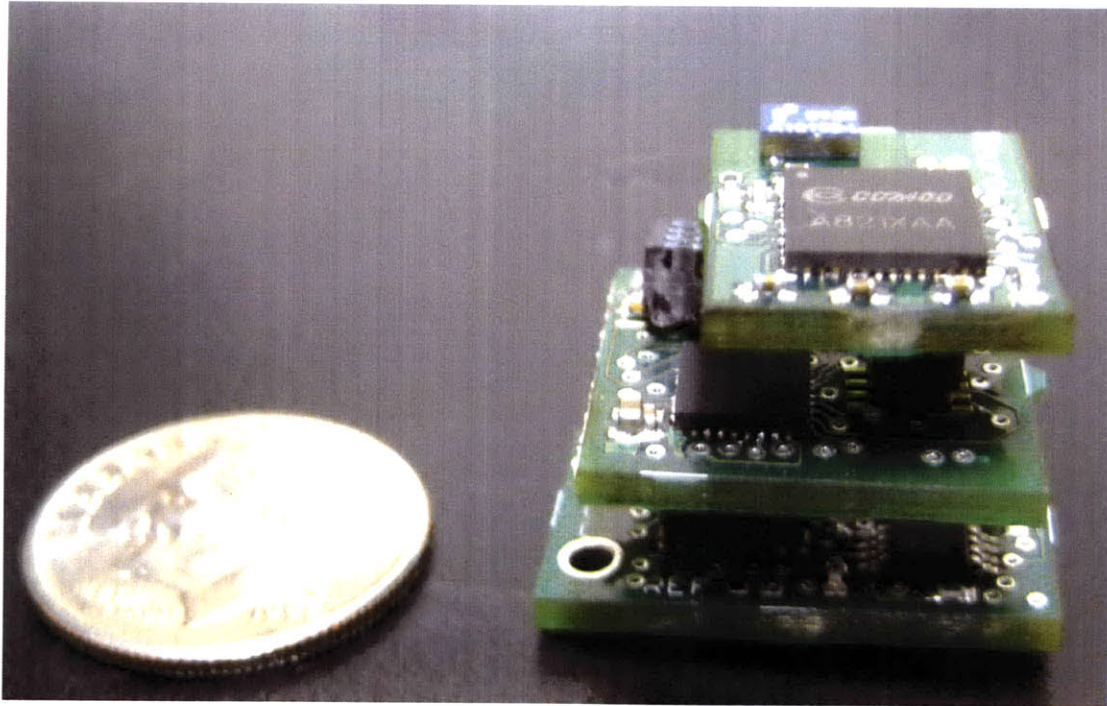
Figure 1: Side View of Complete Wireless Neural Control Hardware Stack
U.S. dime shown for size. Bottom: 4-channel 2-stage amplifier head stage (copper LED element not shown). Mounting hole for skull screws on far left. Middle: implant motherboard with integrated microprocessor, power converters and LED drivers. Top: 2.4 GHz transceiver with chip antenna. Power antenna and super capacitor not shown.

Implant board (i) consists of an array of bare LED dies embedded in a thermally dissipative ~1gram copper block, which is affixed to the skull over a cranial window using 3 standard skulls crews and anchoring dental cement. A miniature printed circuit board (PCB) features wire-bonded traces from the LED control terminals to a 10-16 pin connector accepting the implant motherboard stack.

In recording implementations, signal conditioning electronics are embedded in the implant PCB, composed of four 10X gain common-referenced differential preamplifiers, band-pass filter and 10-100X gain (hardware adjustable) second stage amplifier. Thus, the motherboard and telemetry stack is capable of simultaneously record from 4 dc-coupled

electrode channels at up to 25k samples per second (ksps) and 10 bit resolution, low-pass filtered at 10kHz.

Critically, as many as 83 separate implant-to-computer links can be simultaneously opened, by taking advantage of a Gaussian frequency-shift keying (GFSK) encoding scheme in the onboard radio and automatic frequency hopping algorithms in the Industrial, Scientific, and Medical (ISM) 2.4-2.4835 GHz band. A flexible software layer on the radio allows for addressing of 2^{16} distinct implants, making this paradigm suitable for institution-scale high-throughput testing. Wireless communication with an implanted device can be added to any PC using a small USB dongle, each dongle supporting one full-speed implant link. Rapid modification to stimulation, recording and communication protocols is possible due to a flexible and open source C/C++ based software architecture running on the implant and PC-side systems. Real-time remote triggering of stimulation and recording systems is supported, allowing for complex closed-loop behavioral tasks and high-throughput screening of large animal cohorts.

Figure 2: Wireless Neural Control System *in vivo* Test
Successful Deconstruction of Parkinson's Disease symptoms in a mouse model of PD using 130 Hz optogenetic stimulation of motor cortex.

3.2   Power Systems

Power delivery to the headborne device is achieved via precisely tuned resonant inductive coupling between matched LC networks. On the receiver side, AC signal is rectified using a full-wave passive rectifier and super capacitor energy storage element, serving to both filter AC ripple and provide reserve energy storage. This super capacitor element improves device reliability under varying coil-to-coil mismatch angle as the animal moves around the cage, and also allows for reserve energy to be used in high-powered short duration pulsing of the implanted LEDs. Passive Zener diode shunt and active pulse width modulated (PWM) rectifier load regulation are used to maintain safe voltage range on the super capacitor.

16

Following this first stage of AC/DC conversion, a switched mode buck/boost converter (Texas Instruments TPS61202) provides low ripple conversion from the 5.5V nominal super capacitor voltage to a regulated 3.3V digital supply (hardware adjustable, 1.8-3.6V), serving the onboard microcontroller (MCU) and both analog and digital supplies of the radio telemetry chipsets (CC2400/2500). An additional low-dropout 1.8V linear regulator, running off the 3.3V supply, provides analog power to the radio using a chip-scale packaged device (Analog Devices ADP121). This secondary voltage supply provides extreme flexibility in future radio chipset modifications or additional peripheral circuitry.

Finally, a second buck/boost converter is utilized as a parallel LED driver circuit to power the implanted LEDs. This circuit runs directly off of the super capacitor, with capability to operate at voltages as low as 0.3V, making it ideal for potentially variable input powers and output loading.

3.2.1 Resonant Power Coupling and Super capacitor Energy Storage

A demonstration system consisting of a 6" diameter under-cage transmit coil with series capacitive elements and 3mm diameter, 15mm long axially wound coil with ferrite core and parallel capacitive element on the head of the animal delivers 2W of output power to the headborne device. These power levels are achieved in a low strength magnetic field of 400 A/m oscillating at 120 kHz.

Rectification of the coupled AC signal on the headborne device is performed with a full-wave passive rectifier using low forward voltage Schottky diodes and a super capacitor filtering element of 22-150mF.

The choice of energy storage method here is dominated by the need for rapid charging and discharging capability. Instantaneous LED currents in this system are on the order of 400mA or more. A typical stimulation waveform for treatment of a Parkinson's Disease model may require 5ms pulse width, 130 Hz stimulation, equating to a duty cycle of approximately 70%[citation]. Ultimately, these requirement suggests a preference for high power density devices rather than a strictly high energy density solution, mitigating temporary bursts of stimulation over the averaged power capability of the system, or in the case of intermittently poor power coupling.

Lithium Ion rechargeable cells are an attractive option at certain scales given their substantial energy density of ~620J/gram [1], typical peak discharge rates of approximately 5C and pulsed rates as high as 25C. However, several limitations to Li-Ion cells negate their use in this design. Li-Ion batteries must be carefully regulated to prevent deep discharge below 2.0V or overcharge above 4.2V. Moreover, higher power density batteries and smaller batteries in general have a higher volume fraction dedicated to non-storage contributing elements like current substrates [2], thus limiting their practical energy storage. 400mA discharge currents required by this design imply a 80mA-hour cell to maintain safe discharge rates of 5C, while only a few Joules of energy storage is required

to sustain surge currents experienced in photostimulation of neural networks in vivo. A one second pulsed sequence on an array of 5 LEDs operating at nominal 400mA, 3.6V load requires 7.2J of energy, less than 5% of the energy supplied by an 80mA-hour batter maintained between 4.2 and 2.0V. Thus, the Li-Ion cell is highly stressed in power density yet underutilized in energy density.

By contrast, super capacitors achieve on the order of 10J/gram energy density, but can sustain discharge currents of 1000C [2], implying that a one gram ideal super capacitor satisfies the demands of this system, while serving the additional purpose of filtering rectifier output. In practice, packaging dominates the weight and volume of these small super capacitors.

For these reasons, super capacitors were chosen as the storage elements of choice. A number of designs were explored, including thin-film, tantalum, electrolytic wet cell and others. Ultimately a prototype thin-film solution was chosen due to its low ESR and 5.5V limit (CapXX, 300mF, 5.5V, ESR 70 mOhm). The 5.5V limit was likely achieved by a series stacking of matched elements, as these processes typically do not see voltage limits greater than 3.0V.

Ironically, the ESR limit tends to be the limiting factor in terms of the maximum usable output current of most super capacitors at the small scale (a few grams), particularly at the 120kHz frequency operation of this circuit. Electrolytic cells tested (e.g.

Cooper Bussman PowerSTOR) yielded ESRs of nearly .5 ohms, resulting in capacitor voltage drops of several volts, effectively rendering the tank useless.

In the long term, a hybrid Li-Ion/Super capacitor storage system is extremely attractive, particularly for commercial neuromodulation systems subject to stringent FDA limitations on minimum time between device recharges, provided the system is operated within the limits of the Li-Ion chemistry. For an in-depth analysis of biomedical applications of Li-Ion cells and a novel battery charging circuit optimized for such operation, see Chapter 3.

3.2.2   Power Converters and LED Driver

Two DC/DC converter subsystems are present in this design: one system (Texas Instruments' TPS61202) is dedicated to providing efficiently regulated 1.8-3.6V digital supply for the onboard MCU and sensitive RF chipset, and another (either the Texas Instruments TPS61150 or another TPS61202) for high power up/down-conversion to drive the implanted LED array.

Both of these systems are straightforward, commercial off-the-shelf designs utilizing off-chip inductors of 2-10 uH and input/output filtering capacitors of approximately 10uF. Shielded inductors were chosen in both instances to minimize EMI to the RF and recording systems.

The use of a dedicated dual-channel LED driver (TPS61150) comes as a compromise between the need for a minimal design, such as a simple current source, and the need to maintain precise control over the I-V characteristics and rise/fall times of the LED waveform, given that it is the LED output signal which ultimately drives the neural network, as described above. In total, three iterations of the complete system were made, and in the final design (V1.3), the TPS61150 was scrapped for a second TPS61202 buck/boost converter. This decision was motivated twofold – first, the TPS61150 is rated to only 75mA per channel, making it unsuitable for driving sufficient LED power to overcome losses resulting from fiber-coupled optics systems (described in detail below, see section 2.6), though effective for cortical targeting. In practice, the device was robust when overdriven to several times this current limit. Second and most critically, a design switch from a magnetostrictive/electroactive (ME) power coupling system to resonant coil system resulted in ~20-fold increase in available power in V1.2/1.3, leaving a significant portion of the power delivered to the device unavailable. A parallel, high-side switched multiplexing topology was ultimately adopted using MCU-addressed BJTs to select which LEDs were to be driven at any time.

The move to a voltage-controlled LED driver allows for substantially great output currents to be achieved - with a 5.0V nominal super capacitor supply, efficiencies of 90% and maximal output currents of 1.8A are achievable at a fixed 3.5-3.9V LED forward voltage. This rating equates to 4.5 LEDs driven at 400mA, or 7W of output power. Voltage-mode control does, however, result in exponential output power variation with LED forward voltage (Figure 4). Moreover, LED output power varies significantly with

junction temperature (Figure 3), thus feedback is required if one desires to tightly regulate the output power of the implanted array; support for temperature sensing has therefore been built into the device. As can be seen from figure 3, an LED under continuous operation may shift its forward voltage down by as much as 400mV from ambient operating temperature. This 400mV shift correlates to a change of perhaps 200mA or >70% reduction in forward current. Thus, a nominal operating temperature must be considered when designing the LED driver circuit – in this system, a nominal 4V is chosen for drive voltage.
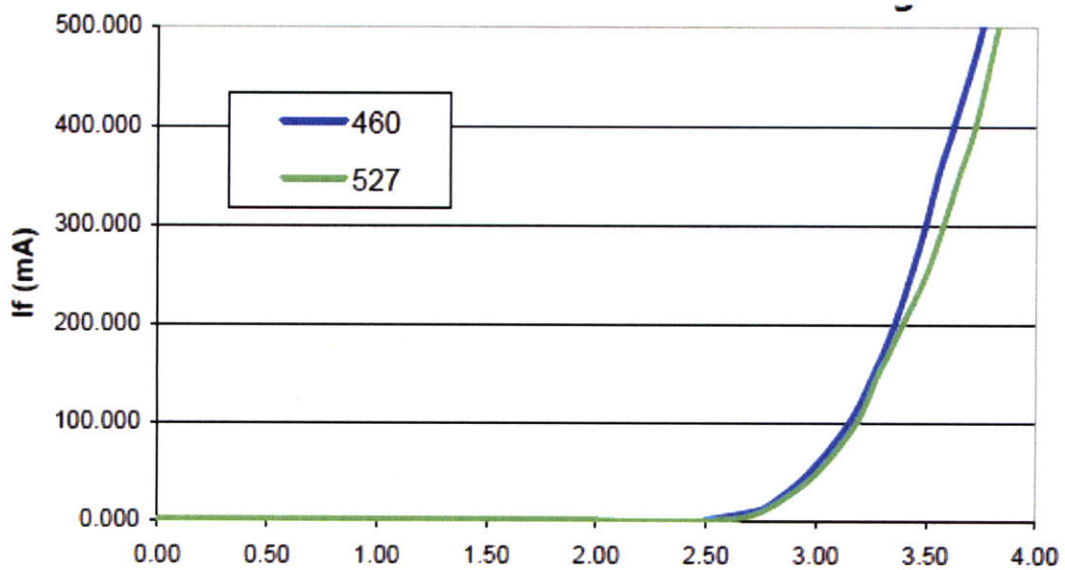
Figure 3: Forward Voltage vs Forward Current for 460-527 nm LEDs
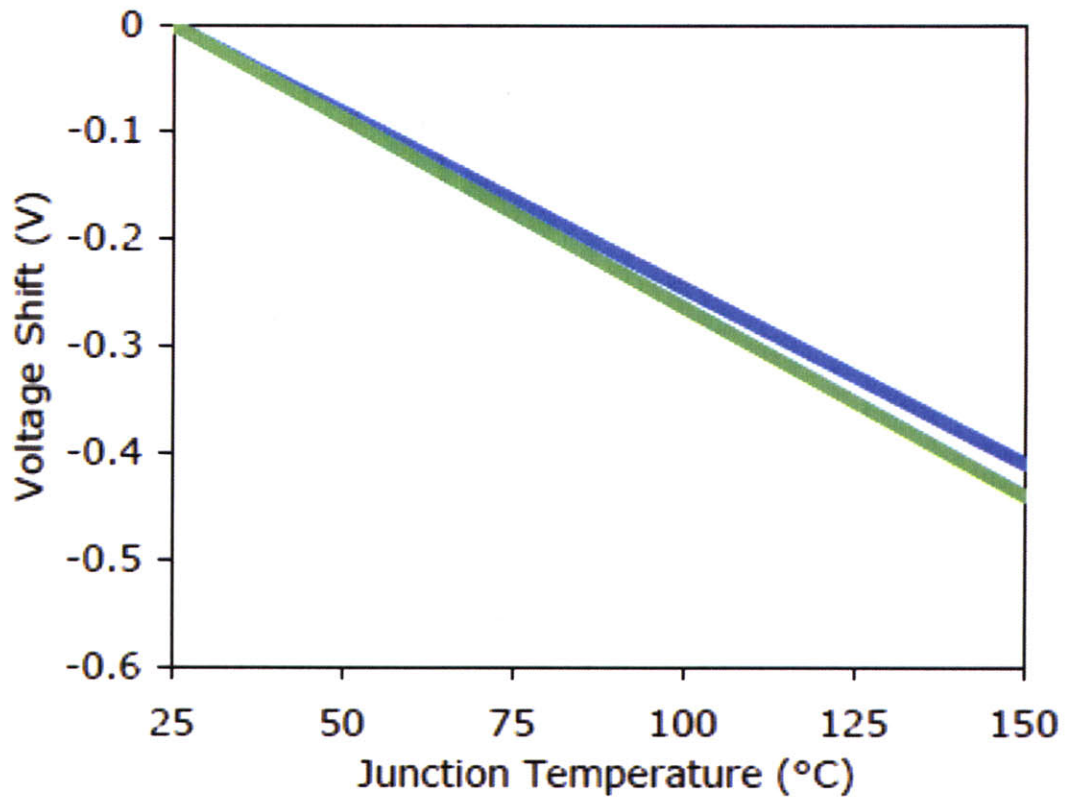
Figure 4: Forward Voltage Shift vs. LED Junction Temperature
Adapted from [3] © 2004 Cree LED corp.

In designing the overall power management system, an optimal value for energy storage exists such that the device is capable of performing a cold start – that is, the device may be placed in the presence of a 400 A/m drive field with all converter stages powered on, and sufficient storage capacitance exists to supply the initial startup energy for these buck/boost converters – yet does not have unnecessary amounts of reserve power (recall, the design specification was intended to provide 2J of bursting energy). An analytical solution based on the load currents of a typical LED, inductor and input and output capacitor sizing would suggest a nominal super capacitor of just a few tens of millifarads in addition to the amount of reserve power required (in the 2J case, roughly 66mF), but bench testing suggests 2-3X this amount of energy is necessary, largely due to the buck

converter's input surge current. Thus, a 300mF super capacitor was chosen. If a smaller capacitor is necessary, the system may be powered initially off of the 1.8V linear regulator. Once super capacitor voltage reaches full charge, the buck/boost stages can then be enabled without substantial capacitor voltage droop.

Waveforms of the capacitor voltage are shown below, with device operating in a 400 A/m field, 5cm above the drive coil, operating a 50% duty cycle, 100 Hz pulse waveform on a 400mA nominal LED. It can be seen that peak-to-peak ripple grater than 100mV is limited to high frequency (1MHz) switching caused by the boost converter. Overall, capacitor voltage is well-behaved.

Figure 5: 300mF Supercapacitor Voltage Waveform Under Load
Oscilloscope trace of super capacitor voltage during 130 Hz, 5 ms pulse width LED drive. 1 MHz peaking indicative of DC/DC converter switch operation.

3.3   Far-Field Wireless Link

The headborne device features a 1 Mbps 2.4-2.4835 GHz Industrial, Scientific, Medical (ISM) band wireless transceiver based upon the TI/Chipcon CC2400 integrated chipset with off-chip oscillator and differential to single-ended balun with chip antenna. An 8-pin header with 3-wire SPI interface to the MCU, 2 auxiliary I/O pins, 1.8,3.3V and ground terminals allows for independent re-design of the wireless transceiver. An alternative design has also been implemented using the highly popular 500ksps CC2500 chipset, also in the 2.4 GHz ISM band. Numerous open-source code examples exist for this chipset, such

that the end user may rapidly migrate to alternative polling strategies, ad-hoc mesh network topologies, etc.

To achieve low power in a design such as this where modification of the transceiver design itself is not possible, shutting down the oscillator during periods of inactivity is an effective dynamic power management technique, particularly in bursty BMI channels. Thus, rapid oscillator start-up time is of paramount importance in device selection.

Table 1: CC2400 Narrowband Transceiver Performance

RX Sensitivity	-87 dBm
TX Power	11 mA @ -25dBm
	15 mA @ -5 dBm
	19 mA @ 0 dBm
Channel Bandwidth	1 MHz, GFSK modulation, 1 Mbps
Operating Frequency Range	2.4-2.4835 GHz
Oscillator Start-up Time	15 uS
Stand-by Power Consumption	1.2 mA
	1.5 uA, Oscillator Core off

In addition to relatively lower power operation by commercial standards and fast oscillator startup, the CC2400 has a high level of autonomy and configurability. The radio handles preamble insertion, CRC generation, computation and check, and sync word insertion/detecton onboard, behind the scenes. Default settings for the CC2400 are transferred by the MCU once upon startup to initialize the radio in 44 data frames of 24 bits each, but can be modified online via a power-down of the radio's analog core.

The MSP430F2132 MCU runs a simple three state interrupt-based Finite-State Machine to minimize idle transceiver on-time. Transition delay between the three states, Transmit, Receive and Idle is adjustable in software in this implementation. In future

implementations, a channel-adaptive delay may be programmed onto the MCU to further improve performance.
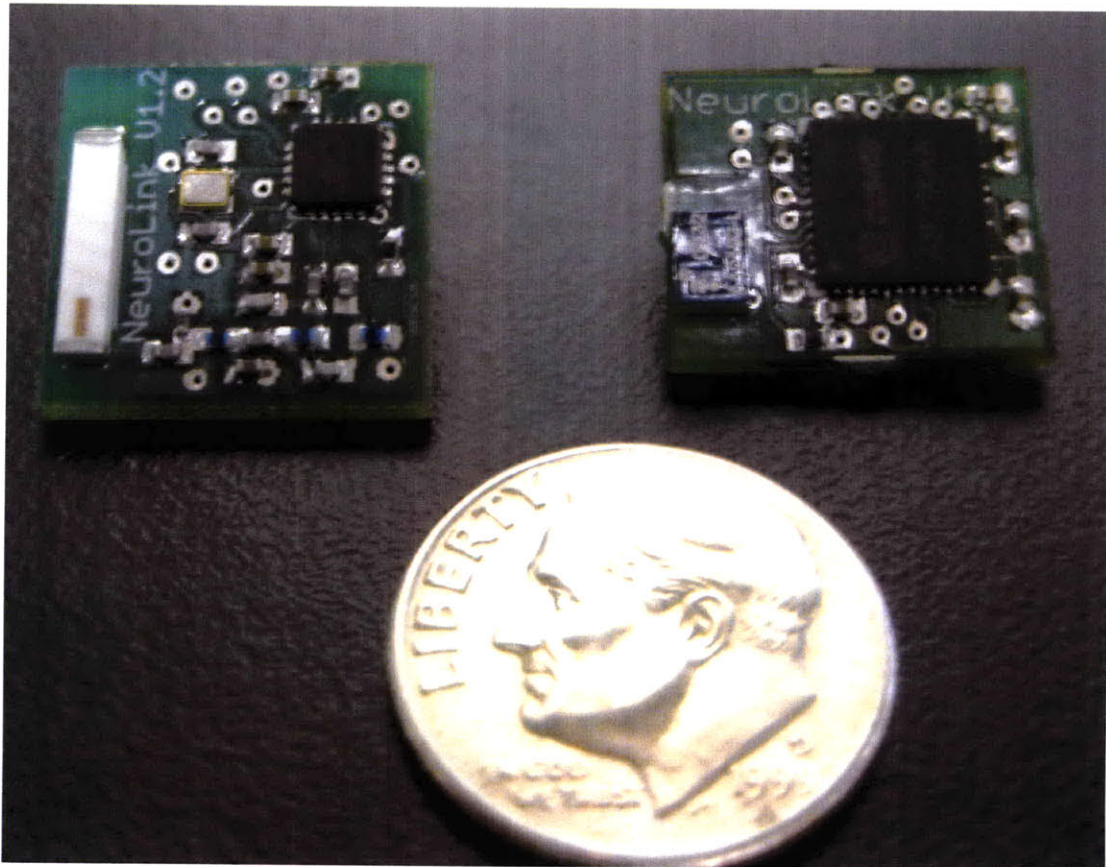
Figure 6: Hot-Swappable Wireless Telemetry Units
At left: Rev 1.2, based on the 500 ksps, highly supported CC2500 2.4GHz chipset. At right: Rev 1.1, based on the higher bandwidth 1 Msps CC2400, also operating in the 2.4GHz ISM band. Both devices are supported in hardware and software.

3.4   Link Performance

To perform system level testing for power consumption during wireless transmission, two of the prototype wireless boards were set up in the laboratory with three meters of open air separation, each lying flat on bench surface. One wireless board was pre-programmed

using a JTAG programmer and analog and digital V_{DD} were powered separately using bench-top power supplies. A second wireless board was interfaced to a Windows machine via the USB JTAG programmer, which also supplied power.

With both radios operating at 2.400 GHz and +/- 500kHz offset GFSK modulation, the CC2400's onboard random number generator on the USB-connected board was used to generate a sample data stream at 1 Mbps with CRC calculation enabled. The second wireless board was set up to receive this signal, decode the CRC, and transmit a flag if bit errors were detected.

Finally, the USB-connected wireless board calculated overall bit error rate (BER) and provided a flag via an open Hyperterminal in Windows if BER reached a threshold (10e-3 in this test). At no point in the test did a flag arise. After 60 seconds at a given power setting, the USB-connected wireless board's CC2400 was then increased in transmit power. A multimeter was used to measure CC2400 current consumption, and the results are plotted below:
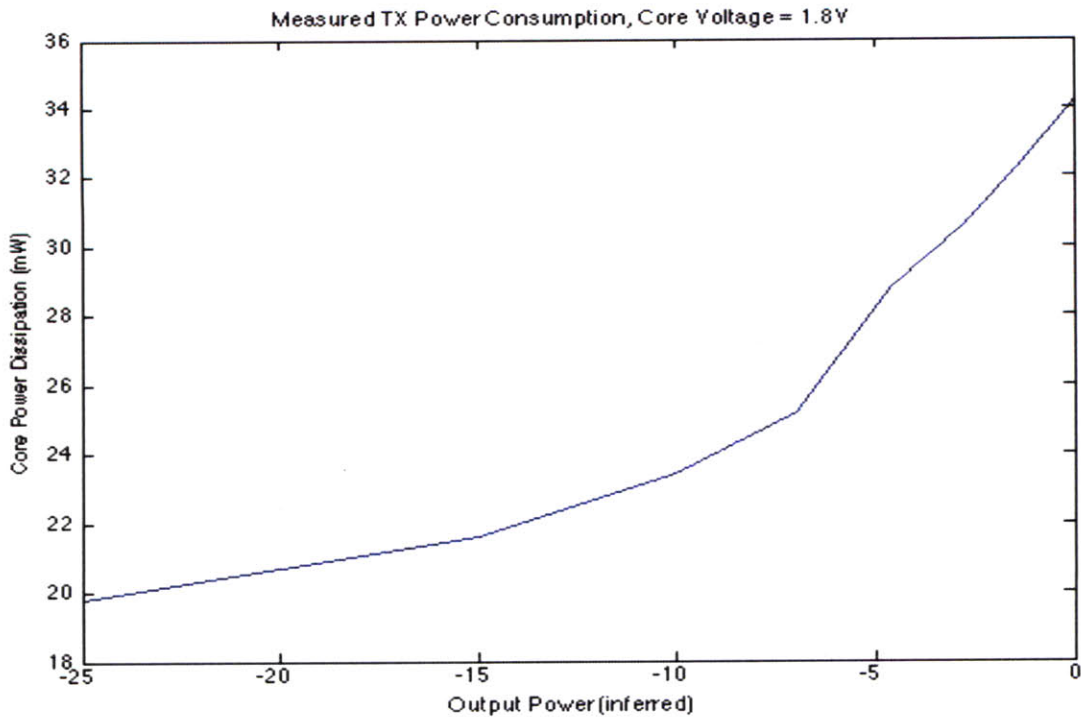
Figure 7: Measured CC2400 Transceiver Power Consumption
Data taken at 1 Mbps, GFSK modulation, 1 MHz channel width.

LNA power was inferred from the CC2400 data sheet and register settings. This rudimentary test indicates that the PCB's RF connections, balun and antenna performed within design tolerances at an operating range similar to that used in BMI applications.

3.4.1 Alternative Commercial, Off-the-Shelf Solutions

The commercially available far-field wireless landscape is summarized briefly below. Areas considered in the design process were Narrowband, Broadband, Ultra-Wide Band, and active RFID. Differential voltage signaling via twisted pair cabling is also considered on an energy per bit basis for comparison in the quantitative analysis.

3.4.1.1 Narrowband

By far the cheapest and most numerous radio transceivers fall in this category. The 900 MHz, 2.4 and 5.8 GHz ISM bands are the most common operating frequencies encountered. Because of the ubiquitous nature of these devices, bare-bones transceivers are easy to come by, thus reducing the size and power demands of the system as a whole. However, as a fundamental limitation of a narrow bandwidth link, these types of transceivers are in practice over-the-air data rate limited to less than 1 Mbps.

Narowband TX Wattage Comparison
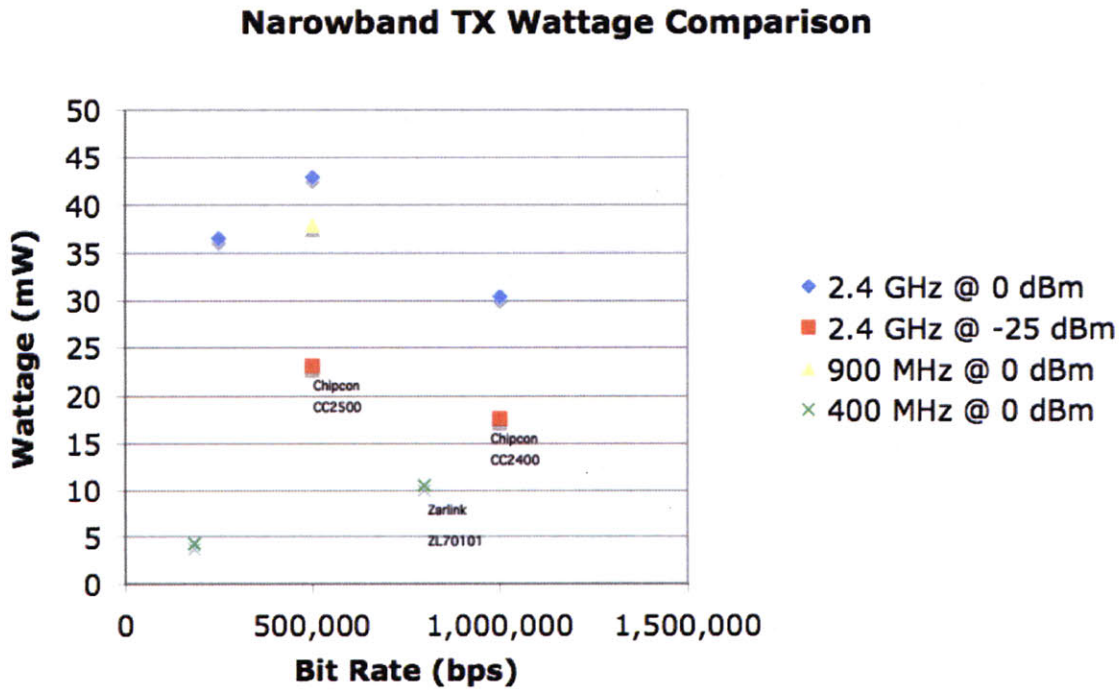
Figure 8: Performance of Slected Narrowband Radio Chipsets

3.4.1.2 Broadband and Ultra-Wide Band (UWB)

While the broadband wireless device market has resulted in a vast array of RF systems in the 10's of Mbps range, these systems typically incorporate unneeded and power-hungry features, most often resulting in a bulky general-purpose chip with extensive hardware abstraction (e.g. integrated media access controller (MAC), algorithmic frequency hopping for noisy environments, etc). As such the 802.11x standard devices are poorly suited to this application.

One emerging technology of great interest for lower power sensor networks such as BMI is the Ultra-Wide Band radio, which achieves extraordinarily low power levels in one implementation through impulse-based transmission. Unfortunately, commercially available options are limited for UWB, particularly those with power-conscious designs. Antenna options in the COTS domain are unsuitable for implant at the moment due to the very wide spectrum of the transmitted signal, but compact designs have been demonstrated in several research publications. This will be an attractive option when very high bandwidth wireless links are needed in future biomedical systems.

3.4.1.3 Active RFID

In the ultra-low power realm, active RFID technology is appealing for medium bandwidth transmission, particularly in the 5.8 GHz band. Again, however, most designs encountered are focused on industrial applications where size is less of a constraint, and as such these systems are often many centimeters in size. A long-term design strategy

focused on ultra-low-power design (i.e., one in which radio budget is the dominant power budget) has many attractive options in the sub-milimeter scale RFID tag space.

3.5   System Control Architecture and Onboard Processor

An onboard integrated 16-bit, 16 MHz microcontroller (TI MSP430F2132) with 8-channel, 200ksps 10-bit analog-to-digital converter (ADC), 512 byte RAM and 264 kB flash memory supervises on-the-head computation. This chipset is used to sample analog signal from the 4-channel neural recording amplifiers and optional temperature and super capacitor voltage sensors, enable data transmission to and from the wireless telemetry, and handle pulse width modulation (PWM) of the LED drivers to in turn modulate neural network activity *in vivo.*

The overall control architecture is as follows:

- Upon startup, the MCU transmits 44 data frames of 24 bits each to initialize the radio chipset. By default this initialization is at 2.4000 GHz transmit frequency, Gaussian Pulse Shift Key (GPSK) modulation, 1 MHz channel width, and 1 Mbps data rate via unbuffered transmit mode. A CRC is, however, calculated.

- The remote device transmits its unique channel ID periodically at default frequency of 2.4000 GHz, awaiting instruction from the PC-side controller, implemented as a simple Hyperterminal interface to a USB-connected 2.4GHz transceiver.

- Control signals from the PC-side controller trigger state change in the remote device – 4 Neuromodulation parameters, namely LED address, pulse frequency, pulse width and on/off toggle – and 3 Recording parameters, namely electrode channel select, sample rate and on/off toggle are all transmitted to the remote device with unique device ID. In addition, optional LED temperature and super capacitor voltage monitoring can be toggled, to have the remote device transmit real-time state variables to the PC. These 9 control variables may be updated at any time by text entry into the Hyperterminal window. Finally, a frequency hopping routine can be initiated to find the nearest clear channel in the ISM band.

- Neuromodulation Programming - On the remote device, these pulse frequency and pulse width parameters are stored as 16-bit comparator trigger values. An up/down waveform generator operating at 1.2 MHz is toggled on and off by the on/off toggle control signal, such that the comparator output triggers the particularly addressed LED on/off with a square wave as demonstrated in the figure below.
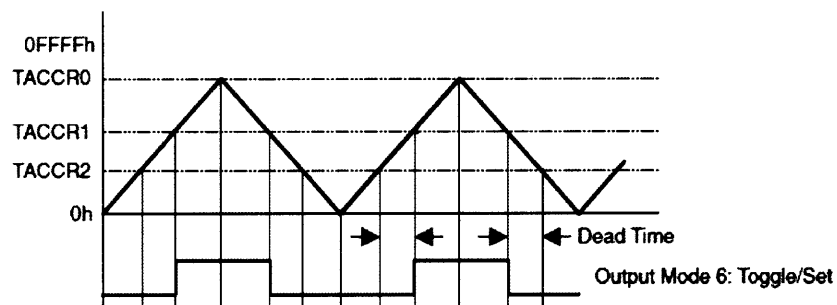
Figure 9: PWM Waveform Generation
Adapted from [4], © 2008, Texas Instruments.

This allows for independent control of two LED waveforms simultaneously, and 8-LED addressing using a direct port-to-port mapping of pins. A more sophisticated addressing scheme is easily implemented using a binary addressing scheme.

Pulse Frequency Programming: Frequency is transmitted to the device as a decimal number (e.g. 'FR=130' = 130 Hz). On the device, decimal to hex conversion is performed, and the upper trigger register (TACCR0, see figure above) is set as TACCR0 = Clock frequency / 2*Pulse frequency.

Pulse Width Programming: Pulse width parameter is transmitted as a decimal number in milliseconds (e.g. 'PW=5' = 5ms). On device this is converted to hex, and stored in the Toggle/Set register (TACCR1 for LED channel 1, TACCR2 for LED channel 2), as TACCR1/2 = PW*Clock frequency*500.

- Recording Programming – On the remote device, the electrode channel select and sample rate are programmed into a direct memory access (DMA) controller which automates the sampling of data from the 10-bit ADC to memory or flash. By default, data is written temporarily to memory and then output unbuffered to the radio for streaming recording.

- Super capacitor Overvoltage Control - In implementations utilizing active PWM control of super capacitor voltage, the MCU's comparator's interrupt flag is enabled to alert the system in the event that voltage on the super capacitor exceeds safe 5.5V

threshold. In this instance, the MCU open circuits the rectified in a back-off-and-wait manner for a programmed period of time.

- Channel Frequency Hopping – In modes of operation where either multiple animals are under simultaneous control, an optional channel hopping protocol can be initiated between PC-side and remote device to find the nearest clear channel.

3.6 Optics

The headborne device described above is intended to be easily disconnected from an implanted optics array. This array consists of a set of bare die LEDs affixed to a small copper block, which serves as thermal sink to draw lost energy away from the brain. LEDs are wire bonded to a small PCB also connected to the copper block, which serves as an anchor point for a 10-16 pin Samtec connector allowing easy disconnect of the headborne electronics for repair or replacement, to ease animal housing needs, etc.

In the event that deep brain structures are targeted for optogenetic control, small fiber waveguides are affixed directly to the surface of the LED die using optics glue, and the remaining die surface coated with reflective epoxy. An additional reflector is placed behind the die for secondary redirection. Significant loss of energy results from this arrangement due to the uncollimated nature of LED light. Future systems, particularly for implantable medical devices in clinical application, must adopt a Laser based approach. Still, output

power at the tip of the fiber is approximately 5-10mW/mm^2, sufficient for local activation of ChR2 targeted neurons[5].
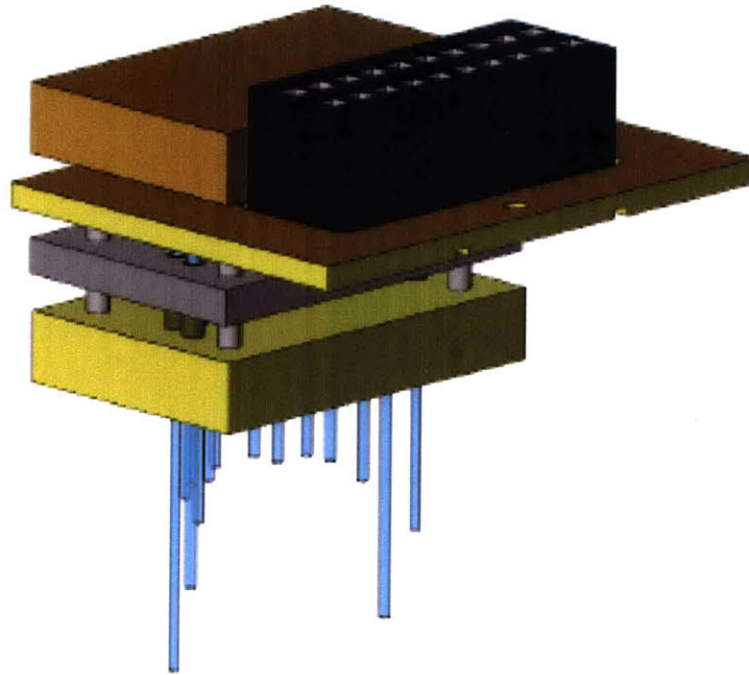
Figure 10: CAD Rendering of Multi-LED Optical Implant for Hippocampal CA1 Targeting Design and rendering courtesy of Jake Bernstein. From bottom, an array of variable-depth fibers is held in alignment with milled plastic aligner (yellow) and reflector (grey). A routed PCB (yellow and copper) serves as mounting platform for copper block housing LEDs (copper, above left) and plastic 2x16 electrical connector (black).

3.7   Sensing

The integrated 8-channel 10-bit ADC affords the ability to monitor critical device parameters remotely. This feature is useful for diagnostic purposes in the event of aberrant behavior on the part of the animal, and also adds an element of safety to the device by allowing remote shutdown in the event of super capacitor overvoltage, over-

current (indicating a short circuit), or over temperature events. Many of these features are also hard-wired into the existing circuitry, e.g. over-temperature shutdown in the TPS61202 and over-current limiters. Specific designs are listed below.

3.7.1 Rectifier Current Monitor

The ability to monitor rectifier current is useful for diagnostic purposes, but is also critical in the event that alternative power transfer techniques are employed. The use of Magnetostrictive/Electroactive (ME) sandwich materials for power conversion was explored using this design. In such designs, resonant frequency is tightly coupled to the temperature of the core material. As such, a constant monitor of rectifier current can be used to optimize drive coil tuning frequency over a wide range of implant temperatures.

A high-side current shunt monitor is used to measure small voltage drops across a precision sense resistor placed in series with the super capacitor load. The shunt monitor (Texas Instruments INA193) amplifies the voltage drop across a 0.3 Ohm sense resistor by a factor of 20-100V/V and outputs this voltage to one of the eight ADC channels with operating range of Vdd. This implies a maximal sensed current of 500mA for 3.3V supply, sufficient for a 2.5W rectifier with 5.0V nominal capacitor voltage. Power loss due to the sense resistor is $I^2*Rsense = 75$ mW. Reduction of sense resistor below 0.3 Ohms linearly reduces power loss, though values much below this require very careful layout to achieve accurate readings. In higher power settings or where dissipation is critical, one might consider a current sense FET.

3.7.2 Super Capacitor Voltage Monitor

As already described, super capacitor peak voltages are generally limited to only a few volts before breakdown occurs. Catastrophic breakdown is very easily achievable with the power figures developed in this system. Thus, the first order protection circuit ought to imply a current shunt. The 5.5V peak allowed by the super capacitors used in this design were convenient for use of a Zener diode at 5.1 or 5.6V, though care must be taken to ensure that the diode is rated for peak power dissipation of several watts. This results in unwanted bulk, as can be seen in the Parkinson's Disease implant which utilizes a 5W Zener diode. More creative approaches have also been taken in the design process, such as an SCR-based overvoltage shunt ("crowbar" circuit) to protect lower voltage capacitors, but the simplicity of the Zener cannot be beat.

Alternatively (or additionally, in the case of a clinical device), an active modulation scheme described in section 2.6 can be used, in which a comparator circuit is used to detect over-voltage limit and open circuit the rectifier input (leaving a discharge path to ground through the power converter load) with a simple series power FET (Fairchild FDMA430NZ). To minimize system resources, a lower threshold for FET turn-on is set at 4.5V. A simple version of this is implemented using a buffered voltage divider input to a spare ADC channel.

Power loss due to sensing current in the voltage divider is minimized with large value resistors, such that loss is negligible. Conduction losses due to the use of a power

FET are $I^2*R_{ds,on}$. Efficient designs at room temperature have $R_{ds,on}$ values of tens of milliohms, such that less than 1 mW of power is lost in conduction (e.g. Fairchild Semiconductor FDMA430NZ 30V, 5A 40mOhm N-FET). One must be careful to consider the voltage gain of the unloaded LC resonant circuit in selecting a FET with sufficient $V_{ds,\,max}$ headroom. Recall that voltage gain in a resonant LC-tank system is the square of the ratio of receiver to transmitter inductances – in this design a low voltage gain was used, roughly 3, such that 15V open circuit voltages are achievable. In the ME-based system, however, open circuit voltages may approach 30V or more.

3.7.3   LED Temperature Monitor

Thermistor-based temperature monitoring circuits are easily incorporated into the existing design. Due to the temperature sensitivity of the brain, FDA limits typically allow no more than one degree Celsius temperature rise in neural tissue due to device heating. In the research setting, temperature is an important factor to control for in experimental procedures. Finally, future implementations of LED drive systems ought to incorporate junction temperature compensation. To serve these needs, a simple resistor divider circuit using a remotely located thermistor epoxied to the LED heat sink, in series with a bias resistor, is placed across the input to one of the ADC channels. As in the case of all analog voltage-based sensing used in this system, capacitive filtering is added, using the input resistance of the ADC to create a high-pass filter pole.

3.8 Recording

A recording amplifier with 4 input channels, one serving as common mode ground, is built into the implant headstage motherboard, top view seen in the figure below. The first stage amplifier is a unity to 20X gain, capacitively coupled single rail amplifier based upon the Texas Instruments TLC2264 quad operational amplifier (on underside of PCB in the figure). All four channels including the reference channel are buffered in this manner. With the first channel serving as common ground, the outputs of the first stage are fed into a 10-1000X programmable gain instrumentation amplifier (INA333) with first order low pass filter set to approximately 10 kHz. The input impedance of the system is approximately 10^{12} Ohms. Adjustment of the second stage gain must be done in hardware, however, by replacing the gain setting resistor on each of the INA333's. Importantly, these amplifiers are easily swappable given the modular design, such that any preamplifier producing a controllable gain output voltage can be utilized.
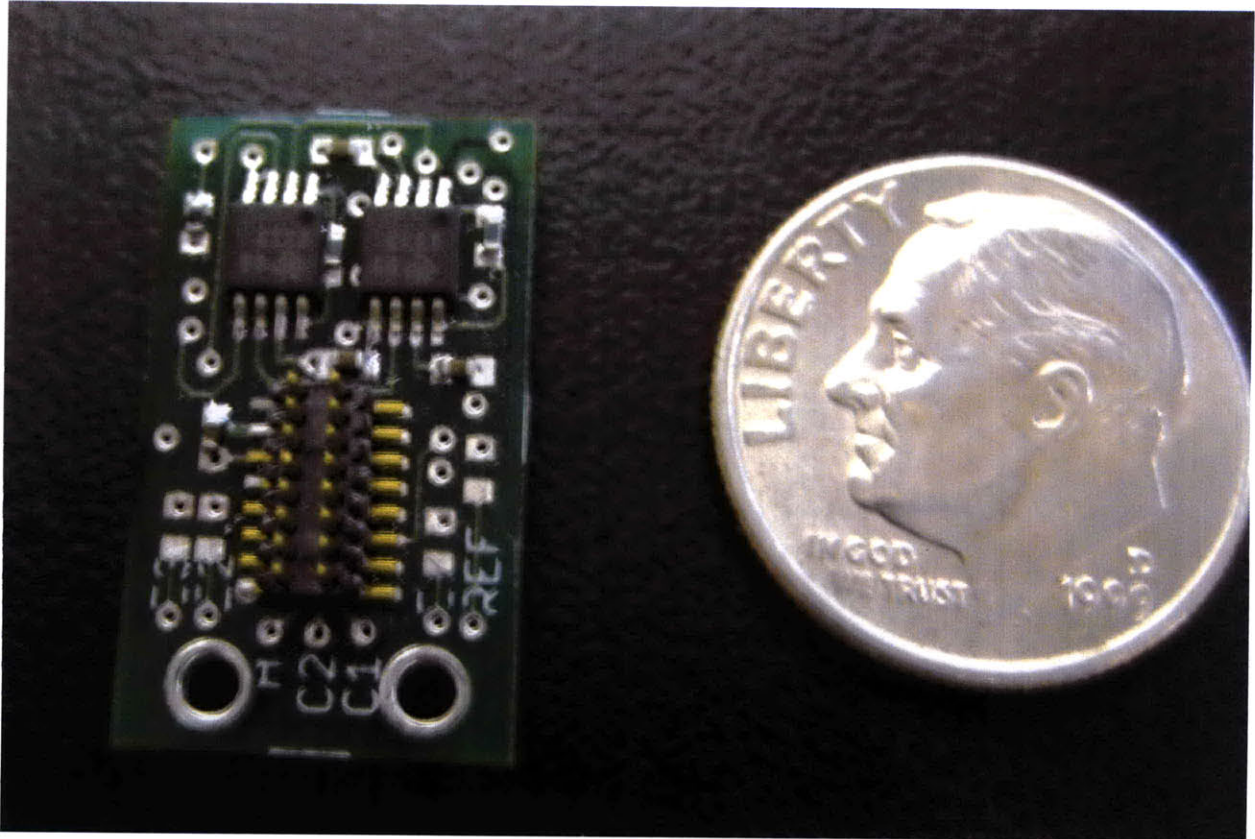
Figure 11: 4-Channel Prototype Headstage Amplifier

3.9 Thermal Considerations

As mentioned above, the thermal dissipation of electronic components on the head is a significant concern. Numerous design principles have been employed to limit the amount of unwanted heat generation. Where possible, dissipative elements (sense resistors, diodes, pass transistors) have been chosen for their smallest possible conduction loss. The microcontroller and radio elements have been chosen to maximize low power mode operation, with rapid startup times on the radio oscillator core and the MCU.

The choice of TI's MSP430 and other subthreshold-based circuit designs should also be considered carefully in biomedical applications, as their operation has an exponential dependence on temperature. Thus, the quoted power dissipation levels of the MSP430 are substantially below what is seen at body temperature, and power budgets should reflect this. Mounting the MCU on a secondary PCB elevated from the two dominant thermal sources (LEDs and rectifier coil) helps to minimize unwanted heating.

The power receiver coil is also a significant source of heat, with roughly equal amounts of energy being produced electrically as thermally. Thus, at nominal operation in a 400 A/m field, the receiver is generating almost 2W of thermal energy. Low conductivity mounts, as seen in the prototype in section 2.11, mitigate this issue.

Finally, protection circuit power dissipation must be considered carefully. Zener shunts are excellent protection at these modest powers, but the diode should be chosen at a thermal rating suitable for long-term dissipation of full load. In the case of the Parkinson's Disease therapy implant pictured in section 2.11, a 5W Zener diode is used. Careful design, both from a circuit design and layout perspective, can minimize the influence of device heat on the neural network itself.

3.10 PC-based Device Interface

A modified USB-to-serial JTAG programmer serves as the base for a wireless base station for communication with the remote device. The complete prototype system is depicted below.
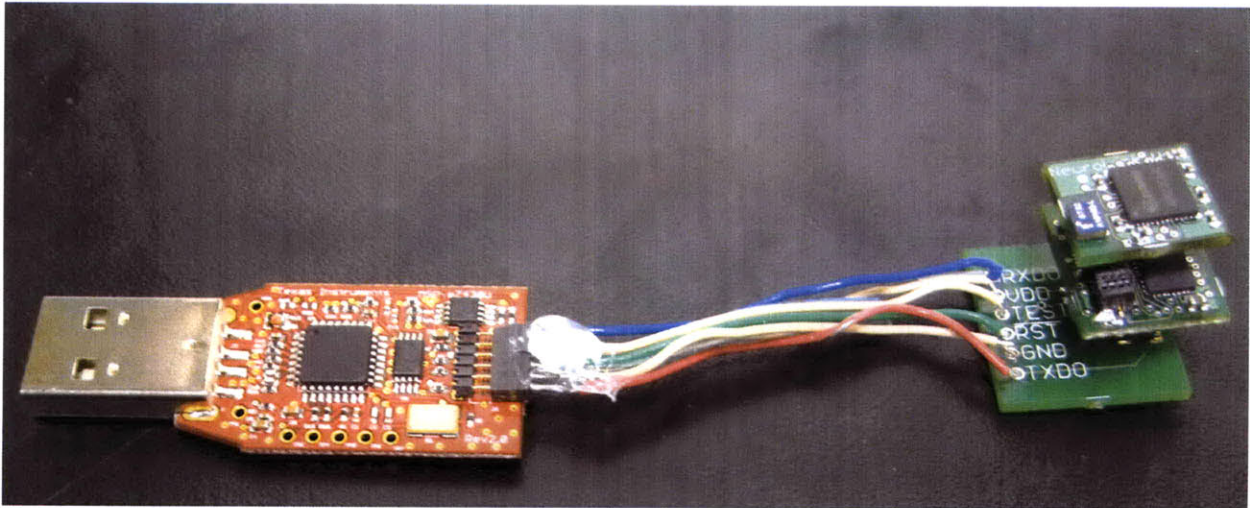
Figure 12: USB wireless base station tether
Tethered implant docked in the receptacle at right.

The USB programmer/wireless base station has a 16-pin Samtec header for mounting an implant motherboard and radio PCB. With a simple re-flashing of memory, the implant becomes a 1Mbps tether. The original TI MSP430-UIF programmer (red PCB) has been modified to allow for full speed communication with the USB port, as default units are limited to 9600 baud backchannel communication.

All PC-side interface software is maintained in firmware in the tethered device, thus, no software other than a Windows Hyperterminal is required. The USB device reports itself as a serial interface with a dedicated COM port. The user connects the USB device,

opens a new Hyperterminal session using the COM port associated with the USB port on their PC, and the tethered device automatically initializes with a splash screen listing the command-line options for communication with tethered devices. All keystrokes are parroted back to the user, such that the link appears to be native software.

3.11 Behavioral Demonstration

To demonstrate the high power capacity of the headborne wireless optical stimulator, a Parkinson's Disease (PD) mouse model was wirelessly treated using a simple, autonomous version of the full featured device. The headborne system, depicted in the figure below, was pre-programmed to generate a 130 Hz, 5ms pulse width waveform previously reported to successfully halt Parkinson's like bahavior in the PD mouse model [6]. The device was targeted toward right hemisphere M2 motor cortex, some of whose afferent axons projection to the subthalamic nucleus, a target known to be effective in deconstruction of Parkinsonian essential tremor. The PD behavioral phenotype is modeled as a rightward tendency to rotate on the part of the animal. When stimulated with the wireless system, the animal shows a halting of rotational behavior. This behavior has recently been demonstrated in fiberoptically tethered animals [6], however, the power requirements of stimulation, roughly 1W input power to the LED, previously made wireless operation inaccessible.
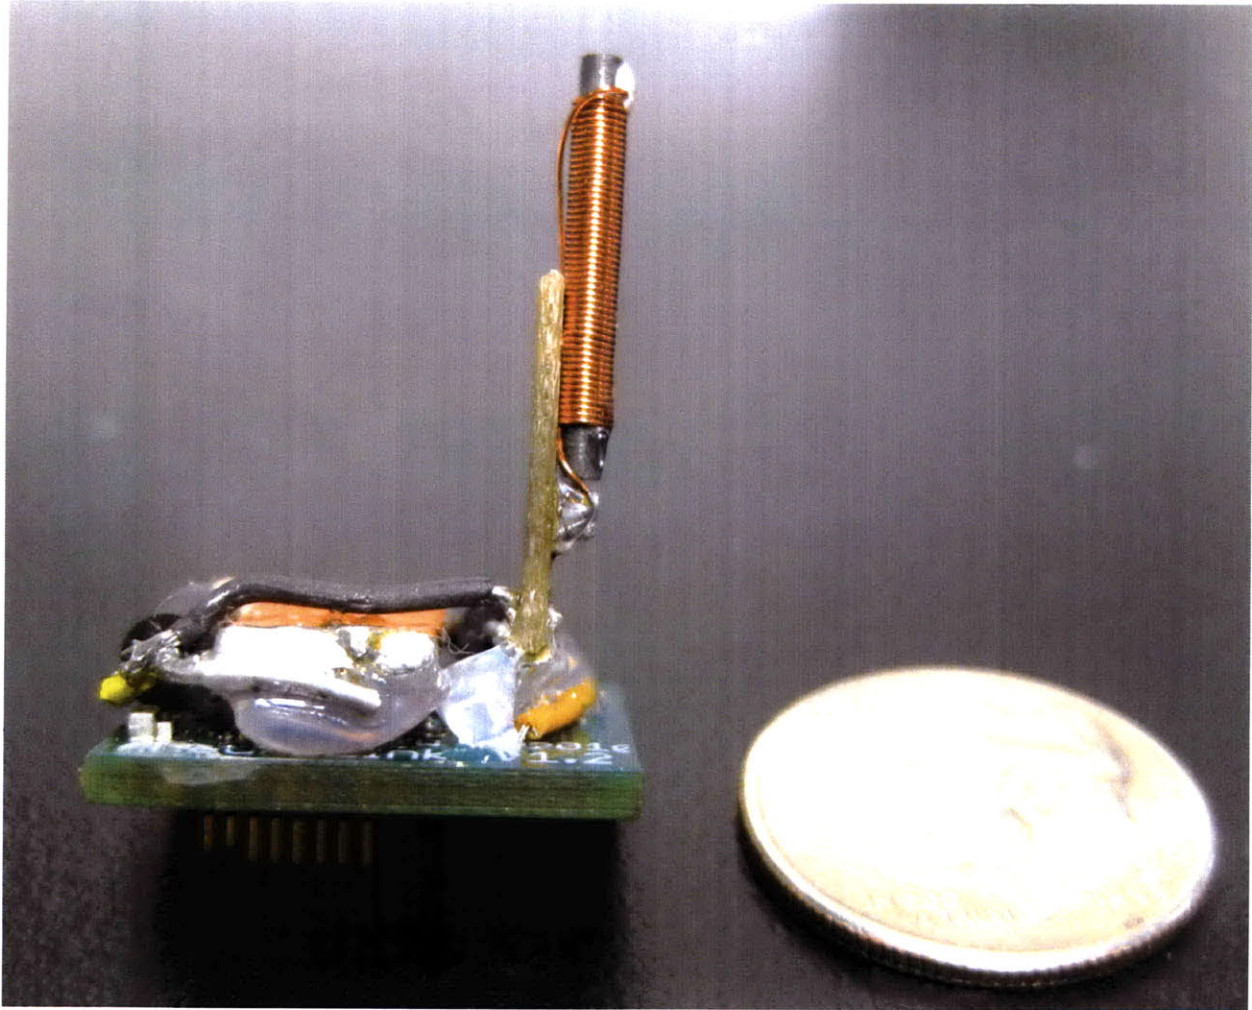
Figure 13: Wirelessly Powered Implant for Parkinson's Disease Therapy
Device functions autonomously using onboard MCU, for mouse model of PD.

3.12 Future Applications

Potential applications of the wireless optical neural control system are significant for *in vivo* neuroscience research. The freedom to explore complex environments while maintaining recording and optical neuromodulation capability has potential in fear behavior research, complex social behavior and other not yet considered. The ability to remotely address dozens of device simultaneously presents an opportunity to perform

high-throughput screening of complex behaviors. Whereas previously researchers were required to select the animal, tether the animal singly, perform the experiment and begin anew, such research could in principle become semi-automated, with only the observational task to be done manually. Additionally, closed-loop paradigms in complex environments, in which an optogenetic control event is triggered based upon either neural network activity or behavioral event, could be easily implemented without modifying the hardware or software interface already developed.

4 A Modular, Massively-Parallel Neural Recording and Control Architecture

4.1 Project Summary

Electrophysiology has been a mainstay of neurobiology for more than 50 years, yet relatively little technological innovation has unfolded to improve upon the methodology or scope. This project seeks to develop a scalable recording architecture for simultaneous capture of many thousands of channels of neural data, with the long-term goal to perform real-time computation on this large-scale data set and derive control signals to perturb the neural network. Such large-scale systems are necessary to elucidate the causal substrates of higher order cognition in the healthy and diseased brain.

A single extracellular recording electrode placed in brain tissue might integrate the spiking behavior of perhaps 100 nearby neurons within a 50 um radius, each of these with peak-to-peak voltages of 60-100 uV [7, 8]. Use of multisite recording techniques and currently available clustering methods can provide further spatial isolation of neural activity within this 50 um sphere, e.g. through use of tetrodes and multi-electrode arrays (MEAs), by triangulating in space, time and frequency domains the firing of clusters of neurons. It is predominantly at this level of network activity that systems neuroscience operates today – arrays of tens to perhaps 100 microelectrodes in a 100 um pitch, 2-dimensional grid. Critically, however, emerging evidence suggests that it is not just the local neuronal activity that plays a role in cognition but that large-scale ensemble activity

giving rise to synchrony (e.g. theta, gamma oscillations) is critical to an understanding of the neuronal network as a whole.

Understanding large-scale neuronal network dynamics will require several orders of magnitude expansion in electrode channel count, further necessitating improvements in large-scale electrode design, massively-parallel neural recording architectures to capture the data, and automated, parallelized clustering of information in the spatiotemporal and frequency domain.

A proof-of-concept system is provided here that rides upon the state-of-the-art in existing technology trends for critical subsystems, including communications protocols (USB 2.0), computational hardware (90 nm Field Programmable Gate Arrays [FPGAs], multicore graphics processing units [GPUs]), and very-large scale integrated (VLSI) circuit techniques. The overarching design principle is therefore one of modularity, such that as individual components are improved, subsystems are swapped out without further redesign.

4.2   Systems Level Description

The end-to-end design of a single recording module from electrode to PC is depicted in the graphic below. Each of the three core blocks of this system are modularly constructed, such that they can be thought of as largely independent elements. "Level 0",

not depicted in the graphic, is some means of connecting from a discrete electrode element to an amplifier channel. In the implemented design described in detail below, such connection is a simple array of 32-channel Omnetics connector interfaced microelectrodes, though higher density systems utilizing multiplexed electrode-to-preamplifier interfaces are likely to be required given the spatial granularity required for whole-brain recording. Level I (pre-amplification) consists of a moderate number of serially multiplexed, amplified analog data streams collected near to the electrode-tissue interface, where space and thermal dissipation constraints are most significant. Level II (Signal Conditioning) is placed as far as several centimeters from the neural tissue. This section band-pass filters the serialized data stream using very high slew rate buffering amplifiers in an effort to minimize the digitization rate limitations implied by input capacitance of the following stage. Level III(a) is a set of high-speed, moderate bit precision analog-to-digital converters (ADCs). These elements dominate the power budget of the entire signal conditioning chain, thus their spatial isolation from neural tissue is of concern.

The general system description up to this point is based upon a first-order optimization of noise, area and energy constraints. The mixed-signal engineer's first thought in optimizing for power in such a scheme as this is likely to be a more highly parallel system – i.e., multiplex in the analog domain to lower speed, moderate precision ADCs in parallel very near to the electrode. Frequency splitting of the neural signal into a so-called local field potiential (LFP) channel from ~DC-300Hz and a spiking channel from 300Hz-8KHz is a more energy efficient solution, and at moderate channel counts of a few dozen to a few hundred, this approach is preferred. This design asks the question, "how do

we achieve one million channels?" As designs scale from a few hundred to several thousand and beyond, the interconnects between electrode and amplifier, and from amplifier to digitizer become the bottleneck. As the vast majority of VLSI fabrication techniques are fundamentally two-dimensional, the area available for circuits *in silico* grows as the square of the peripheral pad count. An extremely efficient design presented in the literature utilizes approximately 0.2 x 0.5 mm of die area in 0.18 um process to amplify and digitize with 13 bit equivalent precision 2 electrode channels. A 10 x 10 mm die with two stacked rows of such modules therefore achieves ~500 electrode channels. This amounts to maximal ~450 um pitch on a land grid array type package, which is already less then 20% of the spatial precision required for whole-brain recording! Moving to a deep sub-micron fabrication process provides only incremental improvements to the argument. Therefore, three-dimensional integration of VLSI wafers, likely in addition to an analog-multiplexed scheme at levels "0" and "1" in the figure below are required.

Returning to the high-level description, Level III(b) of the system is realized as an FPGA performing pre-processing of the digitized data stream. This pre-processing may include lossless compression in very large scale implementations. In the current modular architecture, pre-processing consists of adding a unique channel ID tag and time stamp to each chunk of data. Hardware-level time stamping of data relaxes the timing constraints of the upstream network interconnect between Wired Neural Recording modules and the PC or PC network. Modules are synchronized with a global digital clock signal and channel ID tags are preconfigured to ensure no double-labeling of channels. Assuming the network protocol is a no-drop protocol (e.g. USB and some Ethernet implementations), we

guarantee collection of the neural signal. Closed-loop control of the system, however, remains a constraint to drive low-latency networking.

Level III(c) is a FPGA-based network configuration engine handling low-level interaction with an off-FPGA physical layer chip supporting a network protocol (USB 2.0 in the current implementation). Finally, Level III(d) is the physical layer itself, mediating data transmission between the FPGA and PC or PC network.
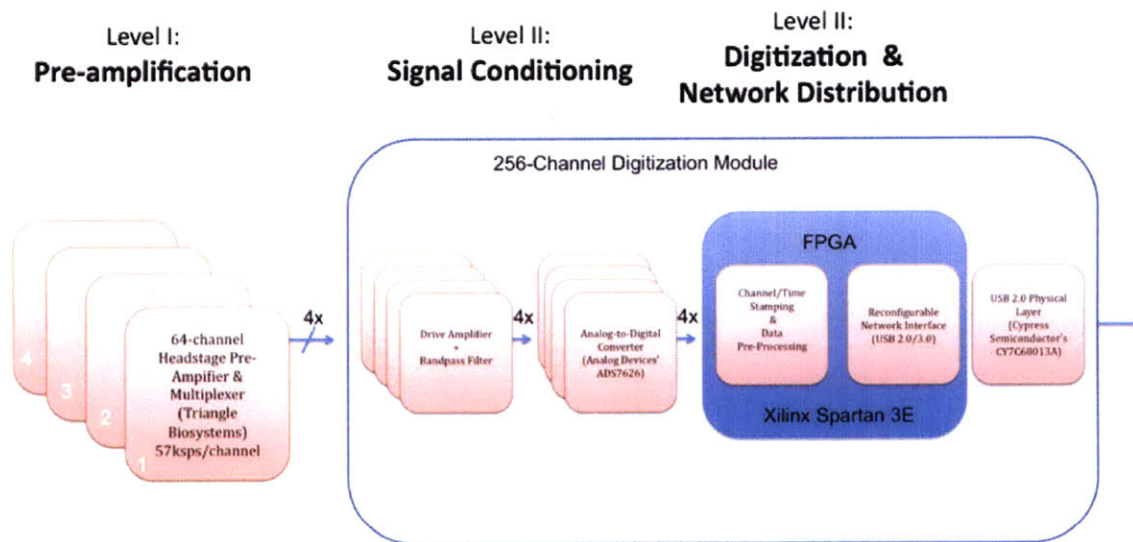
Figure 14: Single 256-channel Module of Wired Neural Recording System

A single FPGA handles four 64-channel headstages. Each headstage signal chain (see section 4.3.2) is handled by a separate high-speed ADC. Total data load over USB 2.0 is 32.816 MHz (see section 4.3.4.1).

4.3   Implementation of a Prototype Wired Module

A single 64-channel module based upon the high-level design principles outlined above has been constructed. The complete description of the prototype is elaborated below.

4.3.1   Electrode Interface and Pre-Amplification

Prototype 64-channel neural amplifier headstages from Triangle Biosystems (TBSI) serve as pre-amplifiers. Capacitively coupled inputs and low pass filtering onboard result in a 0.7Hz – 8 kHz channel bandwidth. Amplifiers are set with a fixed gain of 100X. An internal oscillator-driven multiplexer samples each channel at 57 kHz in serial, preceding each 64-channel block with a low-swinging sync pulse to signal the beginning of a new time frame. Measured power consumption is approximately 20mA from a 3.6V supply.

Connection from headstage to electrode array is made by two 36-channel Omnetics connectors, with 4 unused channels per connector. Output from multiplexer to Level II of the system is a single wire shielded cable. In total , only thee cables (V_{DD}, ground and multiplexer output) make up the headstage-to-digitizer connections.

These sample rates imply an analog output signal switching at 3.648 MHz. While channel-to-channel swing is only perhaps 10 millivolts, such signals place significant demands on signal conditioning and digitization systems downstream. As such, it appears

the 64-channel multiplexing is near the upper limit of what is practical for these sample rates. Of course, 57kHz is more than twice the sample rate required for spike analysis of recorded neural signals (on the order of 30 kHz for tetrode recording [9]) – reduction of sample rate will increase analog multiplexing bandwidth to perhaps 128 channels without increasing the burden on clock recovery circuitry. An ideal solution will have multiplexer clock signal triggered externally, in order to synchronize ADC conversion with the middle of the analog signal pulse. To accommodate next-generation serially multiplexed headstages with higher channel count, the ADC signal chain has been designed to handle roughly 3X the current sample rate specification.

4.3.2   Signal Conditioning and Digitization

Input from the TBSI headstage's multiplexed 3.648 MHz output channel is buffered via wide-band, ultra-low distortion operational amplifier (Analog Devices ADA4899-1) with RC feedback loop achieving approximately 20X gain. The -3 dB point of this stage is approximately 60 MHz, or an order of magnitude beyond the switching frequency of the analog signal, thus ensuring precision operation. Since the downstream ADC is fully-differential, a single-ended to differential conversion is made using a low-distortion bipolar process (SiGe) differential amplifier (Analog Devices ADA4932-1) with 50 ohm terminated inverting input. A standard 4-resistor feedback network sets this third-stage gain to 2X, resulting an end-to-end amplification of approximately 1000X. Large-signal -3 dB bandwidth at 2X gain is greater than 200 MHz.

Digitization is performed using a 16-bit, 10 Msps successive approximation register ADC (Analog Device, AD7626 PulSAR) achieving 91 dB signal-to-noise ratio and 68 dB common-mode rejection. Buffered external reference circuitry sets the input voltage range of the 1000X amplified signal to 4.096V differential peak-to-peak, allowing a large degree of input variation and overhead for clock pulse. This voltage range can be adjusted by modifying the voltage gain of the buffer reference circuitry, or replacement of the reference source. Output from the ADC and conversion enable input signals are low-voltage differential and 2.5V CMOS signaling, respectively.

4.3.3 Clock Recovery Circuitry

An unfortunate consequence of the internally clocked headstage multiplexer is the need to perform clock recovery on the multiplexed signal to ensure accurate and efficient timing of ADC conversion triggering. A brute-force approach, oversampling at more than twice the Nyquist rate (7.296 MHz) is a less than desirable alternative, as it requires digital pre-processing to identify the electrode channel data and thresholding to remove the low-swinging clock pulses.

Thus for maximal efficiency, the differential driver's high side output is additionally sampled by a high speed, minimum delay comparator circuit (Analog Devices AD8561). This comparator achieves <10 ns worst-case propagation delay and <5ns rise/fall times ensuring essentially zero latency compared with the 57 kHz clock pulse signal to be

detected. Additionally, we check that input capacitance is sufficiently low at 3pF, ensuring that the addition of the comparator circuit will not have significant effect on the ADC input.

Low-swinging 57 kHz clock signal is compared with a potentiometer-adjustable voltage reference (Analog Device ADR130) in the prototype system (nominally 0.5V). The comparator output swings high when the clock signal crosses this reference threshold, which latches a divide-by-64 circuit implemented on the FPGA (Xilinx Spartan 3E, 1.2M gates). A Digital Frequency Synthesizer (DFS) block on the FPGA is responsible for this 3.648 MHz clock generation. To ensure minimal phase delay between the extremely low latency comparator output and the synthesized clock output, a Digital Clock Manager (DCM) module compares these two signals and phase-aligns the their output. Thus, following these processing blocks, the input to the conversion enable signal on the ADC is aligned to the middle of the multiplexer's sample pulse. Clocking of the digital I/O between ADC and FPGA is conversion-triggered, that is, the ADC synthesizes the I/O clock signal and the FPGA's simple finite state machine (FSM) waits for a 0-1-0 transition to precede a 16-bit data word.

4.3.4 FPGA Control Architecture and Digital Pre-Processing

Each Wired Neural Recording module consists of two distinct functional blocks synthesized as independent, synchronous blocks – (i) an ADC interface implemented as an FSM, (ii) a USB interface FSM. This system is thus a simple example of a GALS architecture, or globally-asynchronous, locally synchronous digital system. Interface between different

clock domains is achieved using generous first-in first-out buffers. The prototype system implements minimal data pre-processing, though less than 10% of the FPGA's 1.2 million gates are utilized, allowing for substantial expansion within the existing hardware platform. All FPGA code has been written and tested in VHDL, using Xilinx ISE 11.0.

4.3.4.1 Analog-to-Digital Converter Finite State Machine (ADC-FSM)

As described above, the ADC-FSM is clocked by the output of the ADC itself at 250 MHz nominally, well within the 666 MHz clock limit of the Spartan 3E FPGA. The FPGA's LVDS module converts this serial bit stream into two parallel bytes, which are latched into a single 16-bit buffer. The ADC-FSM takes as input the 57kHz CMOS-level output from the comparator, and appends a unique 16-bit channel ID tag to the beginning of every 64 channel sample, thus resulting in a 64x16 bit data block representing one instance in time at one 64-channel electrode array.

In sequence with the ID tagging, the register value of a 16-bit up counter, synchronized across all modules, is latched into the data block, padded with 7 trailing words of value zero. Thus, the total processed data block is 72x16 bits. This block is fed into an asynchronous FIFO to be read by the USB-FSM. The 57 kHz sample rate multiplied by 72-bit deep data block implies a minimum steady state network data rate of 2*4.104 MHz = 8.204 megabytes per second to prevent overflow. This rate is well within the limits of the USB 2.0 data stream. To mitigate PC-side latency-induced overflow of data, the FIFO

is initialized by default to a depth of 1024k, using block RAM modules on the FPGA allowing 17 ms of lag time before data is overwritten for a 57 ksps sampled headstage.

4.3.4.2 USB Interface FSM (USB-FSM)

The prototype system implements a USB-FSM built upon the Digilent/Cypress Semiconductor USB 2.0 reference design. USB I/O is constructed as a parallel array of 56 8-bit registers on the FPGA. Handshaking signals, including port direction and idle/busy state of the PC are mirrored from the Cypress Semiconductor CY7C68013A-56 EZ-USB chipset to the FPGA's state machine as inputs. A communication event is initiated from the P.C. by requesting one of the 56 registers, the direction of the operation (read or write), and the data to be written in the case of a write operation. Other low-level operations are abstracted away from the FPGA and handled exclusively by the USB chipset and PC-side dynamic link library (DLL).

This USB chipset operates at a clock frequency of 48 MHz with two separate 8-byte channels, allowing for 96 megabytes per second theoretical link capacity. In practice, PC-side latency limits actual throughput to approximately 50% of this figure. The FPGA is configured to act as slave to the PC, given the tendency towards PC-side latency. In this abstraction layer, the Data I/O channel consists of a pair of one byte registers. A second pair of one byte registers is used to hold control variables transmitted from the PC. In the prototype system, the lower control byte is limited to an on/off toggle for initialization of

ADC sampling (bit 0, LSB) and ADC reset (bit 1), FIFO clear (bit 2) and a global synchronous reset (bit 2). The upper control byte is set prior to initializing sampling, and is used as a master byte to set the ID tag word (all zeros for upper byte) uniquely identifying the blocks of data coming from this FPGA's associated headstage.

4.4 PC-side Software

A simple Windows terminal program was written for debugging and single-module operation. Built upon the Digilent/Cypress Semiconductor USB 2.0 reference design, written in Visual C/C++ and compiled as a standalone executable file, the terminal program allows the user to manually initialize sampling to disk of neural signals from the headstage. To begin a session, the user simply plugs the USB device into a Windows PC which, provides power to the Wired Neural Recording module, and runs the executable at command prompt. The user may sample a single register of the data output port to screen or disk, or continuously stream data to disk for later analysis.

4.5 Applications and Future Work

The system described here is an implementation of a single 64-channel module of the Wired Neural Recording System. Design principles demonstrated in this prototype will allow low-cost scaling to many thousands of channels of recording. Calculations show that each FPGA+USB2.0 link can handle four 64-channel headstages, with the USB 2.0 protocol being the limiting factor. The FPGA-level design has been structured in such a way that

each data stream, from headstage amplifier to network I/O register, is completely independent and parallelizable. Overhead has been designed in wherever possible to accommodate incremental improvements to each of the subsystems to ease scaling of the system (higher channel count headstage amplifiers, faster network interface, additional parallel signal chains, etc.). Critically, the data path from module to PC is bidirectional, allowing implementation of future closed-loop neural network perturbation systems.

Software-side visualization of incoming data will be necessary to perform significant electrophysiology, and is currently under development. Additionally, the extraordinary amount of incoming data allowed by this front end will place incredible demands on PC-side data analysis. A back end data redistribution architecture will likely be required for such analysis, and the use of GPU-accelerated online computation is being explored.

Finally, the current software is implemented in Windows – USB chipset drivers are readily available for Linux/Unix platforms, and movement from Windows to Linux will likely reduce the PC-side latency for acquisition of data over USB.

5   An Ultra-Compact and Efficient Li-Ion Battery Charger for Biomedical Applications

5.1   Project Summary

This project describes a novel, all-analog Li-ion battery charging circuit intended for operation in a wirelessly rechargeable medical implant. Lithium-ion (Li-ion) batteries are a popular choice for implants due to their ability to provide relatively high performance in both energy and power densities, of 158 Wh/Kg and 1300 W/Kg, respectively [1].

Previous Li-ion charger designs, however, often suffer from two significant problems. First, unnecessarily complex control circuitry [10, 11] is often employed to manage battery charging at the expense of circuit area and power consumption. Additionally, many circuits require a sense resistor in order to detect end-of-charge [12, 13]. This latter point is especially problematic for battery longevity due to the challenges of precision on-chip resistor fabrication, as undercharging the battery can drastically reduce its capacity [14].

The circuit presented here addresses both of these issues. By utilizing the tanh output current profile of an OTA, the circuit naturally transitions between constant current (CC) and constant voltage (CV) charging regions without the need for complex control circuitry. As a result, this circuit is an order of magnitude smaller than previous designs, while achieving an efficiency of greater than 75% in this proof-of-concept design. This design does not require sense resistors to determine end-of-charge, as our control circuitry operates in the current domain. Upon startup the device is capable of monitoring battery

voltage levels and providing charging current during periods of power coupling, as in the case of a wireless power link. This design represents a simple, analog, power- and area-efficient version of previous, more complicated and power-hungry designs.

5.2   Background

Battery longevity is a primary concern in implanted medical devices due to the significant cost and risk of resurgery. Battery longevity, in turn, is highly sensitive to the accuracy of the final charging voltage on the battery. Previous reports indicate that undercharging a Li-ion battery by 1.2% of the 4.2V target value results in a 9% reduction in capacity [14]. Conversely if the Li-ion battery is overcharged, dangerous thermal runaway can occur. During discharge, deeply discharging the Li-ion battery below 3V can permanently reduce the cell's capacity [15].

The charging profile of a Li-ion battery can be divided into four distinct regions as illustrated by Fig. 15: trickle-charge, constant current, constant voltage, and end-of-charge. Trickle charging is required only if the battery is deeply discharged (voltage is less than 3 V). During trickle-charge, the battery is charged with a small amount of current, typically no more than 0.1 times the rated capacity of the battery, or (0.1C) [14]. C represents the battery capacity expressed in terms of amp-hours (Ah). Charging currents greater than 0.1C may be hazardous as the battery has a high internal impedance at these low voltages. Above 3.0 V, the battery may be charged at higher currents; this is the constant current region.
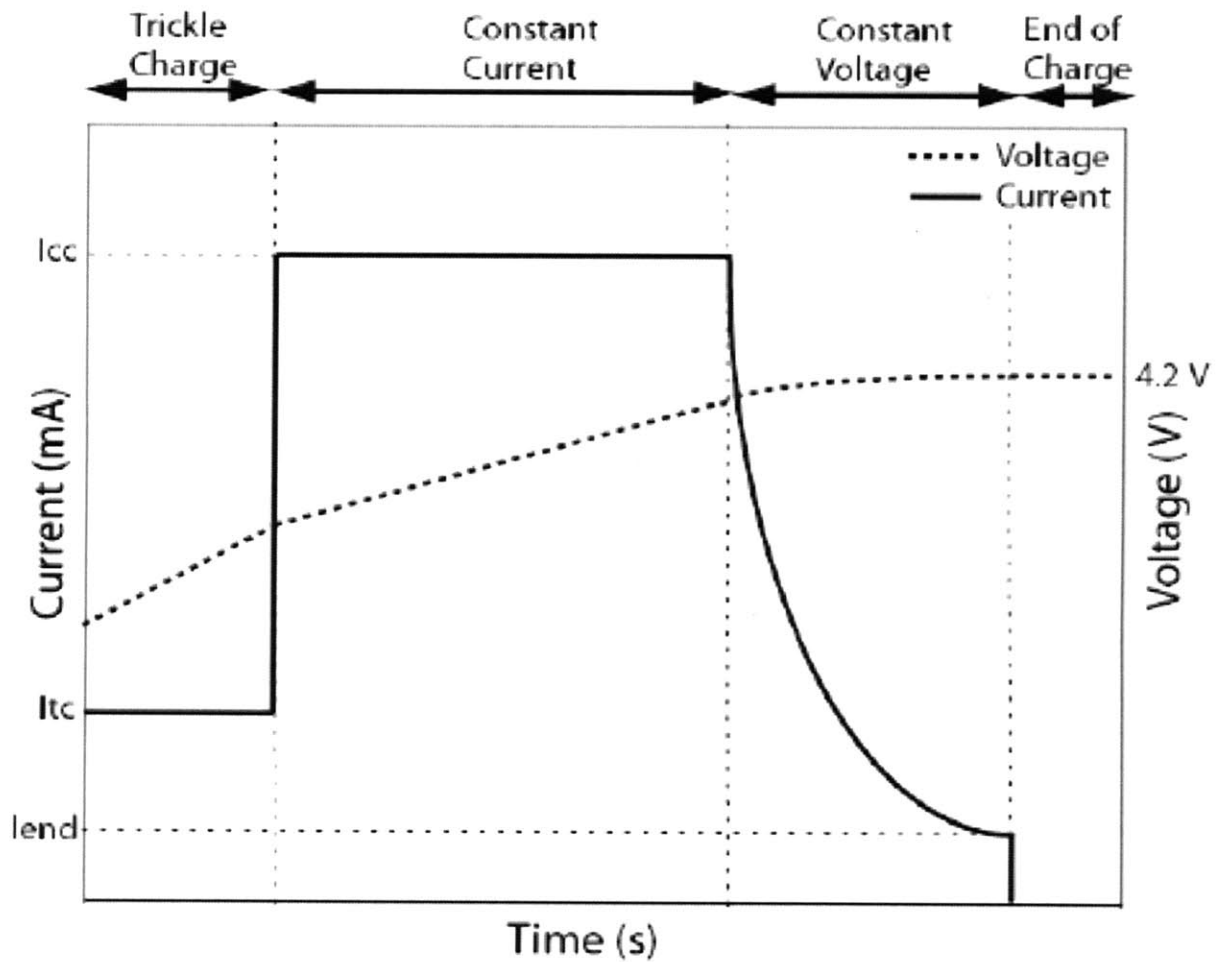
Figure 15: Theoretical Li-ion Charging Profile

As the battery voltage approaches 4.2V, the charging profile enters the constant

voltage region. In this region, the charging current should be progressively decreased as

the battery voltage approaches 4.2V. The constant voltage region is required in order to

compensate for internal battery voltage drop; as the charging current decreases, the

battery output voltage also decreases due to lower voltage drop across its internal

impedance. Charging current should be decreased until a certain threshold is met, which is

usually about 2% of the rated battery capacity [14]. Once this charging current is reached,

the charger enters the end-of-charge region.

5.3   Circuit Description

The simplified block diagram of the circuit topology is illustrated in Fig. 16. The circuit consists of four major blocks: a 4.2V reference, OTA, current gain stage, and end-of-charge detector. The 4.2V reference was designed using a bandgap reference followed by a non-inverting op-amp to produce a stable output voltage over a range of temperatures. This design is intended to be used in an implantable device, so the expected temperature variation is limited. Nevertheless, the design presented here is robust enough for charging applications where temperature varies significantly.

Figure 16: Simplified Charger Block Diagram

The OTA compares the battery voltage to the 4.2V bandgap reference in order to determine the charging current. For battery voltages less than approximately 4.1V, the OTA output is saturated. As the battery voltage reaches 4.1V, the difference in input terminal voltages becomes small enough that the OTA enters the linear region and the output current begins to decrease. The OTA was designed to operate in subthreshold to save power and also to reduce its linear range. In order to account for the trickle-charge region the OTA topology was slightly modified. Fig. 17 shows the schematic of the OTA with the trickle-charge modification, which is the addition of transistors M1 and M2. If the battery voltage is less than 3V, the *Trickle Charge Flag* is low enabling M1.

In this case, transistor M2 conducts some current, which reduces the OTA output via current stealing of the bias current. The reduction in charging current during trickle-charge is proportional to the ratio of W/L of M2 to the W/L of M6. Once the battery voltage crosses the 3V threshold, the *Trickle Charge Flag* goes high disabling the current path through M1 and M2. As a result, the current output of the OTA is increased to its maximum value.
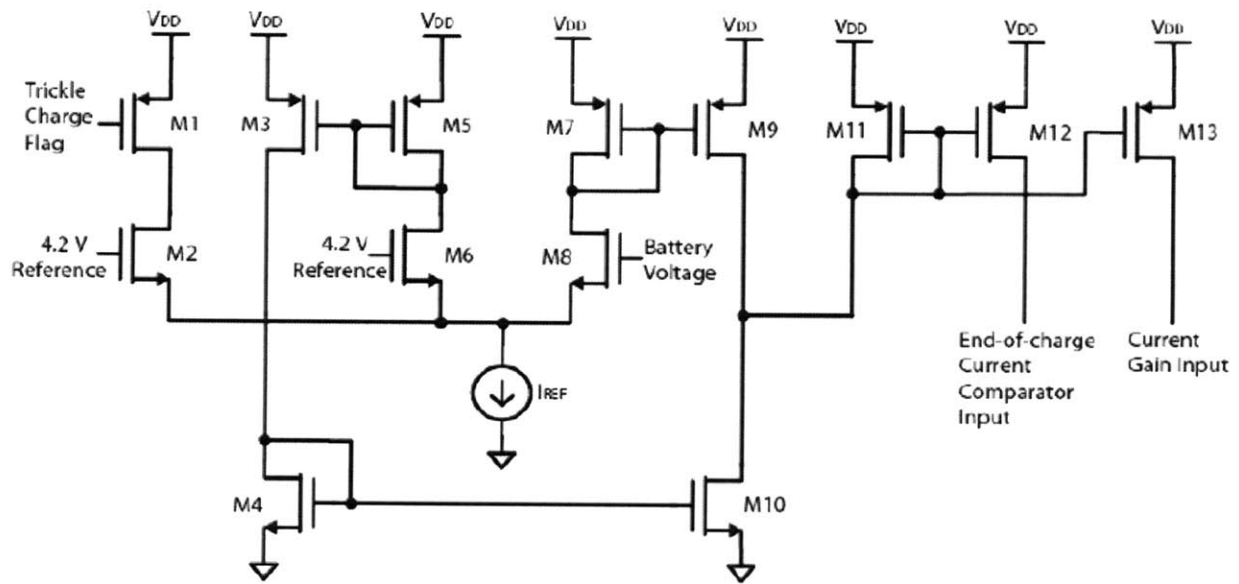
Figure 17: OTA Circuit Schematic

The current gain stage is simply composed of current mirrors to increase the current output of the OTA, from a few hundred nano-amps to whatever charging current is required in the design. The initial design requirements constrained the power consumption to 10 mW, so the charging current was limited to 2 mA. All current mirrors in this design including those in the OTA are of the Wilson Current Mirror type in order to reduce channel length modulation error.

The end-of-charge is detected by comparing the output of the OTA to a reference current; this reference current is proportional to the reference current used to bias the OTA in order to minimize error. Fig. 18(a) shows the schematic of the current comparator [16]. The *End-of-charge Output* signal goes low when the OTA output is higher than I_{REF}, otherwise it equals V_{DD}. When the *End-of-charge Output* signal is high, the last stage of current mirrors in the current gain block is disabled, reducing the charge current to zero.
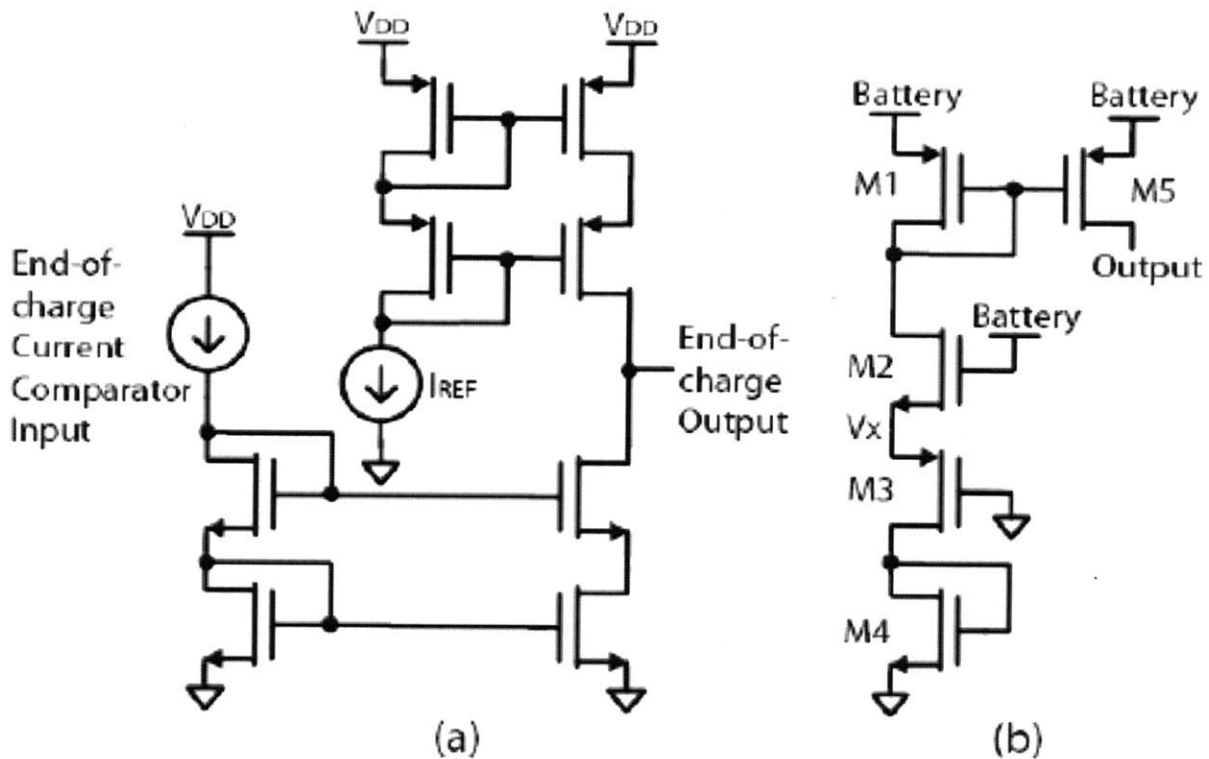
Figure 18: End-of-charge Detector
(a) End-of-charge Current Detector, (b) Trickle Charge Threshold Detector.

In order to detect the battery reaches the 3 V threshold for the trickle-charge region, a simple low-power detector circuit was designed, shown in Fig. 18(b). This circuit is used to detect critically low battery voltage, in order to prevent any damage to the battery due to deep discharge; when critical threshold is reached, the detector circuit cuts off power to the load. As the battery voltage decreases, the voltage at the node V_x between transistors M2 and M3 decreases. The relationship between the voltage at this node and the battery is linear, so the current flowing through transistor M5 reduces quadratically when M2 and M3 are in saturation and exponentially when they enter sub-threshold. The current output

of M5 goes through another current comparator similar to the one shown in Fig. 18(a), in order to detect when the battery voltage falls below 3 V. Transistors M1 through M4 were designed with large widths and lengths in order to minimize process variation. This strategy also minimizes power consumption such that the threshold detector may be run off the battery voltage directly for constant protection against deep discharge. The designed threshold detector consumes only 3 µW when the battery voltage is approximately equal to 3.7 V.

5.4   System Performance

The battery management chip was fabricated in an AMI 0.5 µm CMOS process, consuming 0.15 mm2 of chip area. Fig. 19 shows the die micrograph of the test chip. Fig. 20(a) shows the measured results of the battery management IC charging a 25 mAh battery during tricklecharge and a portion of the constant current region. The battery was charged with 1.5 mA and 2.2 mA during trickle charge and constant current, respectively. Although trickle-charge is not strictly needed in this case since the constant current charging rate is already less than 0.1C for the 25 mAh battery, it is included here to demonstrate circuit functionality. Further, while the proof of concept circuit was limited to about 2 mA maximum charging current, the design can easily be modified if a higher charging current is required by adjusting the current gain in the last stage of current mirrors.
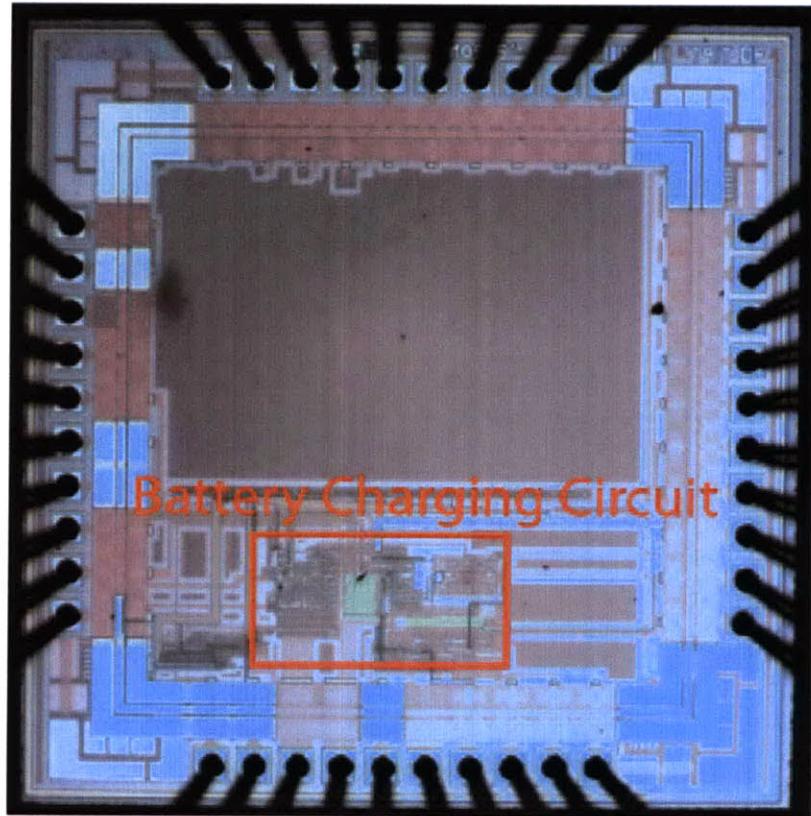
Figure 19: Charger Die Photograph

Fig. 20(b) shows the remaining regions of the charging profile: constant current, constant voltage, and end-of-charge. The constant voltage region begins when the battery reaches approximately 4.1 V. The transition between constant current and constant voltage is continuous since the control loop is based on a simple tanh function. As shown in Fig. 20(b) the charging current decreases as the battery voltage goes from 4.1 V to 4.2 V, reaching the end of charge when the current is approximately 0.26 mA. At the end of charge the battery voltage is 4.21 V, providing an accuracy of 99.8%. In this test with a 25 mAh battery the total charging time was about 800 minutes. This long charging time is

purely due to the maximum charge current of 0.1C, which was determined by the power consumption requirement of 10 mW. If the current mirrors are adjusted to provide 1C during constant current, a charging time of a few hours can be attained with this 25 mAh cell.
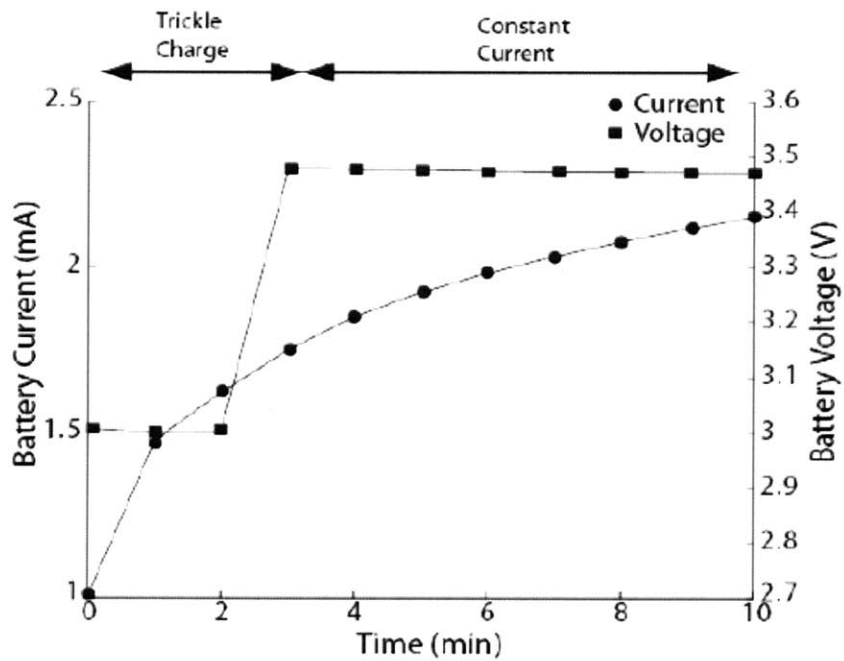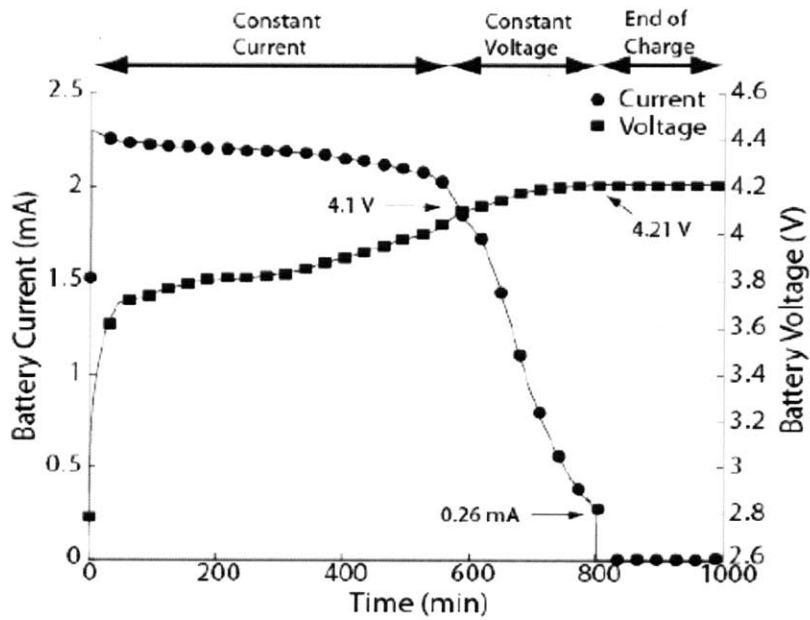
Figure 20: Experimentally Derived Charging I-V Profile
(a) Trickle Charge Profile of Prototype Charger, (b) Charge Profile During Constant
Current and Constant Voltage Modes.

A power efficiency of approximately 75% was obtained during constant current mode. The limiting factor in efficiency is the fact that the test circuit was designed for a 5 V supply. One can easily design for a lower supply voltage, increasing the overall power efficiency of the system. By simply reducing the supply voltage from 5 V to 4.5 V, the efficiency of this circuit can be increased to approximately 83%. In this chip it was not possible to reduce the supply voltage to 4.5 V because of the Wilson current mirrors in the OTA. Nevertheless, if these mirrors are replaced with current mirrors that require less voltage headroom, the supply voltage can be easily reduced to 4.5 V.

Table 2 compares this design with previous Li-ion charger circuits in the literature. While the design presented here has yet to be optimized for supply voltage, it nevertheless achieves competitive power efficiency while consuming at least an order of magnitude less area than other designs.

Most of the literature uses the maximum power efficiency during charging as a figure of merit for battery chargers. However, power efficiency is not constant during a charge, as the battery voltage varies from 3 V to 4.2 V. A suggestion for a better figure of merit is the total energy delivered to the battery divided by the total energy consumed. Using this figure of merit, the design presented achieves energy efficiency close to 70%.

Table 2: Comparison of Li-ion Charger Performance

Design	Power Efficiency	Layout Area
[10]	67.9%	1.96 mm^2
[11]	82%	2.6 mm^2
[13]	83%	PCB
[17]	72.3%	Not Specified
This Work	**75%**	**0.15 mm^2**

5.5 Applications and Future Work

A novel design for a Li-ion battery charger that simplifies the control circuit by using the tanh output current profile of an OTA has been presented and experimentally verified. This design does not require the use of sense resistors to determine the end-of-charge point, reducing layout area and charging errors due to resistor variability. The layout area required for this chip is more than an order of magnitude smaller than previous designs, as Table 2 illustrates.

Without optimization, the proof of concept design achieved a power efficiency of 75%, which is comparable to previous designs. This efficiency can be further improved if one designs the circuit to operate with a lower supply voltage or with an adaptive supply that varies with battery voltage. If the supply voltage is reduced to 4.5 V, or an adaptive power supply is introduced, a power efficiency greater than 83% can be obtained. The

circuit presented here achieves excellent energy efficiency with potential for further improvement, and consumes the smallest layout area of any design thus far presented in the literature.
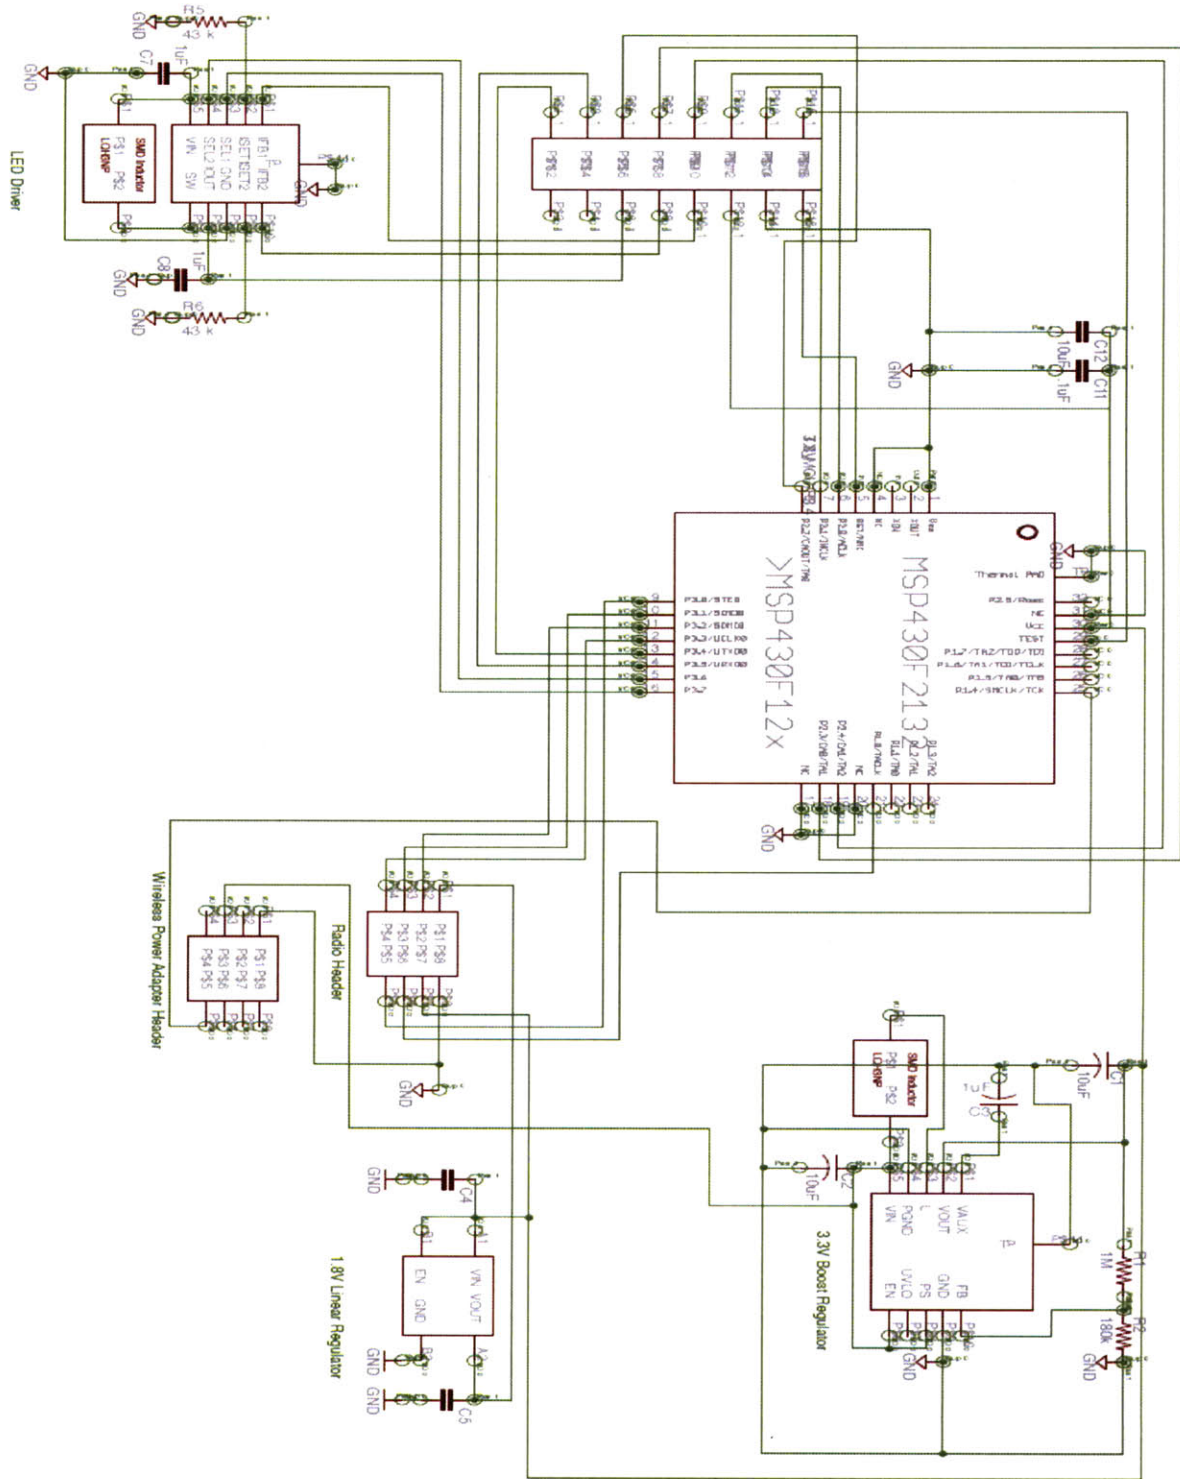
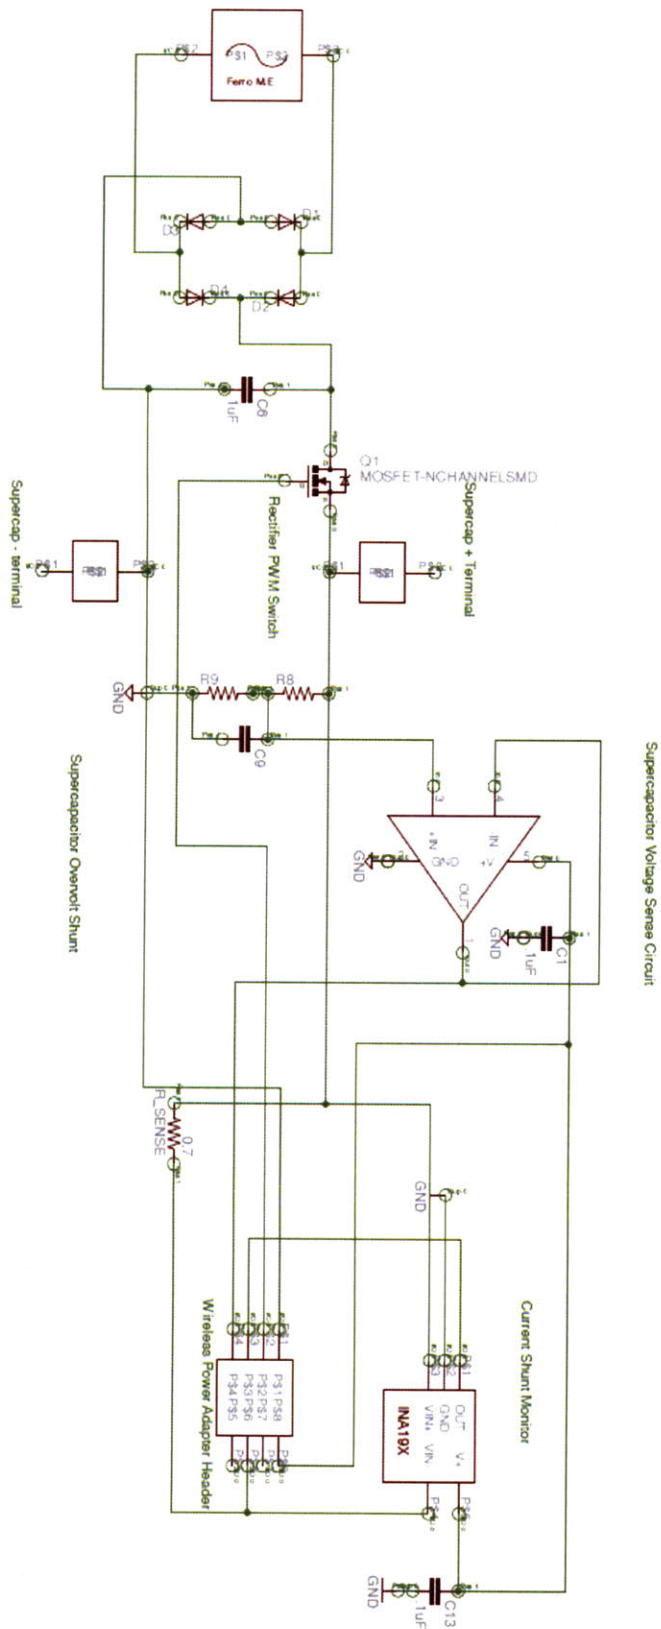# 6 Appendix

Figure 21: Wireless Motherboard Schematic

Figure 22: Wireless Rectifier and Voltage/Current Sense Schematic

Figure 23: CC2400-based Telemetry Schematic (1 Mbps)

Figure 24: CC2500-based Telemetry Schematic (500 kbps)

7 Bibliography

[1]   D. Linden, *Handbook of Batteries.* New York: McGraw-Hill, 2001.

[2]   R. Sarpeshkar, *Ultra Low Power Bioelectronics.* New York: Cambridge University Press, 2010.

[3]   Cree Corp., "EZ700 LED Data Sheet," ed, 2004.

[4]   Texas Instruments, "MSP430x2xx User Guide," E ed, 2008.

[5]   E. S. Boyden, *et al.,* "Millisecond-timescale, genetically targeted optical control of neural activity," *Nat Neurosci,* vol. 8, pp. 1263-8, Sep 2005.

[6]   V. Gradinaru, *et al.,* "Optical deconstruction of parkinsonian neural circuitry," *Science,* vol. 324, pp. 354-9, Apr 17 2009.

[7]   G. Buzsaki, "Large-scale recording of neuronal ensembles," *Nat Neurosci,* vol. 7, pp. 446-51, May 2004.

[8]   D. A. Henze, *et al.,* "Intracellular features predicted by extracellular recordings in the hippocampus in vivo," *J Neurophysiol,* vol. 84, pp. 390-400, Jul 2000.

[9]   T. Jochum, *et al.,* "Integrated circuit amplifiers for multi-electrode intracortical recording," *J Neural Eng,* vol. 6, p. 012001, Feb 2009.

[10]   Y.S. Hwang, S.C. Wang, F.C. Yang, J.J. Chen, ""New Compact CMOS Li-ion Battery Charger Using Charge-Pump Technique for Portable Applications," *IEEE Trans. on Circuits and Systems - Part I: Regular Papers,* vol. 54, pp. 705-712, 2007.

[11]   F.C. Yang, C.C. Chen, J.J. Chen, Y.S. Hwang and W.T. Lee, ""Hysteresis-current-controlled buck converter suitable for Li-ion batter charger"," *Proceedings of the IEEE International Conference on Communications, Circuits and Systems (SCCCAS),* pp. 2723-2726, 2006.

[12]   P. Li and R. Bashirullah, "A Wireless Power Interface for Rechargable Battery Operated Medical Implants," *IEEE Trans. Circuits Syst. II, Exp. Briefs,* vol. 53, pp. 912-916, 2007.

[13] M. Chen and G.A. Rincon-Mora, "Accurate, compact, and power-efficient Li-ion battery charger circuit," *IEEE Tans. Circuits Syst. II, Exp. Briefs,* vol. 53, pp. 1180-1184, 2006.

[14] S. Dearborn, "Charging Li-ion batteries for maximum run times," *Power Electron. Technol. Mag.,* pp. 40-49, 2005.

[15] F. Hoffart, "Proper Care Extends Li-ion Battery Life," *Power Electron. Technol. Mag.,* pp. 24-48, 2008.

[16] D. A. Freitas and K.W. Current, "CMOS current comparator circuit," *Electronics Letters,* vol. 19, pp. 695-697, 1983.

[17] C.C. Tsai, C.Y. Lin, Y.S. Hwang, W.T. Lee and T.Y. Lee, "A multi-mode LDO-based Li-ion battery charger in 0.35 um CMOS technology," *IEEE Asia-Pacific COnference on Circuits and Systems,* pp. 49-52, 2004.