A Low-Power 32-Channel Digitally Programmable Neural Recording Integrated Circuit

Woradorn Wattanapanitch, *Student Member, IEEE*, and Rahul Sarpeshkar, *Senior Member, IEEE*

***Abstract*—We report the design of an ultra-low-power 32-channel neural-recording integrated circuit (chip) in a 0.18 μm CMOS technology. The chip consists of eight neural recording modules where each module contains four neural amplifiers, an analog multiplexer, an A/D converter, and a serial programming interface. Each amplifier can be programmed to record either spikes or LFPs with a programmable gain from 49–66 dB. To minimize the total power consumption, an adaptive-biasing scheme is utilized to adjust each amplifier's input-referred noise to suit the background noise at the recording site. The amplifier's input-referred noise can be adjusted from 11.2 $\mu\mathrm{V_{rms}}$ (total power of 5.4 μW) down to 5.4 $\mu\mathrm{V_{rms}}$ (total power of 20 μW) in the spike-recording setting. The ADC in each recording module digitizes the a.c. signal input to each amplifier at 8-bit precision with a sampling rate of 31.25 kS/s per channel, with an average power consumption of 483 nW per channel, and, because of a.c. coupling, allows d.c. operation over a wide dynamic range. It achieves an ENOB of 7.65, resulting in a net efficiency of 77 fJ/State, making it one of the most energy-efficient designs for neural recording applications. The presented chip was successfully tested in an *in vivo* wireless recording experiment from a behaving primate with an average power dissipation per channel of 10.1 μW. The neural amplifier and the ADC occupy areas of 0.03 mm^2 and 0.02 mm^2 respectively, making our design simultaneously area efficient and power efficient, thus enabling scaling to high channel-count systems.**

***Index Terms*—Analog-to-digital converters, brain-machine interfaces, digitally programmable, energy efficient, low power, neural amplifiers, neural-recording systems.**

I. Introduction

In the past few decades, direct recordings from the cortical area of the brain have enabled scientists to gradually understand and unlock the secrets of neural coding. With the aid of high-density microelectrode arrays, neural activities from a large population of neurons can be observed simultaneously with a spatial resolution down to that of a single cell [1], [2]. Many experiments in non-human primates [3]–[5] and a pilot clinical trial in a human subject [6] illustrated that control signals directly derived from spiking activities from a population of neurons in the cortical area of the brain can be used to successfully control and manipulate computer devices or robotic limbs. The study in [7] shows that cortical activities from a population of neurons can be used to control even a sophisticated device such as a robotic limb with multiple degrees of freedom. These studies have shown great promise for successful development of practical brain-machine-interface (BMI) systems to restore lost body functions to patients with disorders in the central nervous system such as those suffered because of spinal cord injuries. Practical BMI systems of the future will be portable and may enable the users to control dexterous robotic limbs or their limbs naturally.

Practical BMI systems require the use of neural-recording systems to amplify and digitize the neural signals. To avoid the risk of infection, the recording system should be entirely implanted under the skin while the recorded neural data and the power to operate the implant should be transferred through wireless means [8]. This implantability requirement poses major constraints on the size and total power consumption of the recording system [8]. To record from a large number of cortical neurons, high-channel-count recording systems are needed. Therefore, the area per recording channel must be small such that a high-channel-count recording system can be designed with a small form factor. Furthermore, the power consumption per channel must be minimal such that the total power dissipation of the recording system can be kept within feasible limits. To avoid excessive heat dissipation that may cause cell death in the surrounding tissues, the total power dissipation from the recording system should also be minimal [8]. For battery-operated recording systems, low power consumption could prolong the time between recharges, thus expanding the battery's life to avoid frequent surgeries for battery replacements.

Advances in integrated-circuit (IC) fabrication technologies have enabled engineers to increase the number of recording channels that can be put on a single chip by decreasing the size and power consumption per recording channel, while still significantly improving the recording system's performance. The first recording system reported in [9] contained 32 recording channels and data-reduction circuitry while consuming a total power of 5.4 mW (equivalent to 169 μW per channel). Another system reported in [10] contained 100 channels and also included wireless data transmission and power-transfer features. By counting only the power consumption from the recording channels and the analog-to-digital converter, the average power consumption of this system is approximately 140 μW/channel.

Manuscript received March 06, 2011; revised May 21, 2011; accepted July 11, 2011. Date of publication September 08, 2011; date of current version December 29, 2011. This work was supported by the National Institute of Health (NS056140 to R. Sarpeshkar). This paper was recommended by Associate Editor Wentai Liu.

W. Wattanapanitch was with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. He is now with the Department of Electrical Engineering, Kasetsart University, Bangkok 10900, Thailand (e-mail: fengwdw@ku.ac.th).

R. Sarpeshkar is with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: rahuls@mit.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TBCAS.2011.2163404

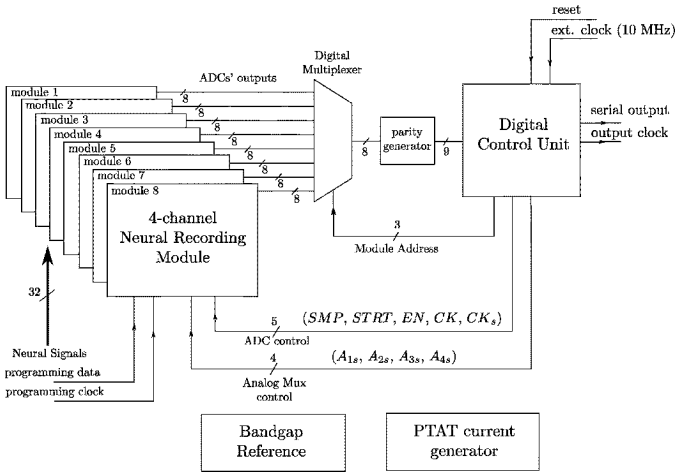

Fig. 1. Overall architecture of the 32-channel neural recording chip.

The systems reported in [11]–[14] represent examples of neural-recording systems with small area and power consumption per channel and thus are feasible for the development of high-performance BMI systems. The system reported in [15] achieved a very low average power consumption of 3.77 μW/channel. However, due to its very large area per channel, such a system may not be scalable to a high-channel-count system, and thus is not yet suitable for high-performance BMI systems.

Our goal is to develop an ultra-low-power implantable wireless neural recording system for practical use in neural prosthetic applications. For such applications, low power consumption and small area per recording channel are *both* very important. In this paper, we present the design and experimental results of an ultra-low-power 32-channel digitally-programmable neural-recording chip for use in such applications. Even though our recording chip is not yet optimized for total chip area, it achieves very small area and very low average power consumption per recording channel. The chip is thus suitable for scaling to include a large number of recording channels in subsequent generations. This paper is organized as follows: In Section II, we discuss the overall architecture of the 32-channel neural recording chip. In Section III, we present the design of the neural amplifier including an adaptive-biasing technique to optimize the total power consumption of the recording chip. In Section IV, we present the design and power-saving techniques in our energy-efficient ADCs, analog multiplexers, and control logic. In Section VI we present both benchtop characterizations of the components of the neural-recording chip and *in vivo* experimental results obtained wirelessly from the brain of an awake non-human primate. In Section VII, we conclude the paper by summarizing our contributions.

II. Overall Chip Architecture

Fig. 1 shows the overall architecture of our neural-recording chip. The chip contains a total of 32 recording channels which are grouped into eight 4-channel neural-recording modules. The schematic of one of the 4-channel neural-recording modules is shown in Fig. 2. Each neural-recording module contains

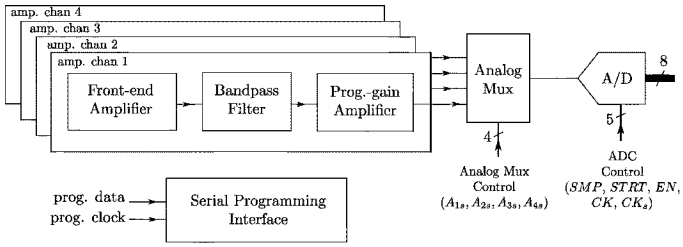

Fig. 2. Architecture of a 4-channel neural recording module.

four neural amplifiers, an analog multiplexer, an 8-bit ADC, and a serial-programming interface unit. The outputs from the four neural amplifiers in the neural-recording module are multiplexed to the ADC which digitizes its input signal at a rate of 125 kS/s. Effectively, each neural amplifier's output is sampled and digitized at a rate of 31.25 kS/s. The clock and control signals for the analog multiplexer and the ADC are generated from a centralized control-logic block which we call the Digital Control Unit. The output data from the ADCs are multiplexed to a parity bit generator before being sent to the Digital Control Unit. In the latter unit, the data are packetized and streamed off-chip for further processing. The configuration setting of each recording channel is achieved through the serial-programming interface unit via the programming data and clock pins.

To minimize power consumption of the recording chip, we utilize two power-supply domains. The neural amplifiers and analog multiplexer, requiring larger voltage headroom, operate from a 1.8 V supply voltage. The ADCs and the Digital Control Unit operate from a lower supply voltage of 1 V to save power. An energy-efficient DC-DC converter for light-load applications such as [16] achieves greater than 80% efficiency at our power level and can be included on chip to generate the lower supply voltage. With such conversion efficiencies, the power wasted for generating an extra supply voltage for the Digital Control Unit of our chip is only 10 μW. However, in this implementation of our chip, we merely provide two external power-supply voltages. Digital level translators are included to interface between the Digital Control Unit and the control switches in the analog multiplexers. The analog multiplexers also act as DC-level shifters between the neural amplifiers and the ADCs. A bandgap voltage reference circuit is included on chip to generate a temperature-independent 1 V reference (V_{ref}) for the ADCs and a 0.9 V reference (V_{mid}) for the mid-rail voltage of the neural amplifiers. The proportional-to-absolute-temperature (PTAT) current generator provides constant-gm biasing to all the subthreshold neural amplifiers and the analog multiplexers in the recording chip [8].

III. Neural-Amplifier Design

Fig. 3(a) shows the schematic of the neural amplifier, which consists of three stages: i) a front-end amplifier; ii) a bandpass filter; and iii) a programmable-gain amplifier. The neural amplifier is optimized for recording action potentials (spikes); however, it can be configured to record local field potentials (LFPs) if needed. The front-end amplifier uses a capacitively-coupled configuration to reject DC offset introduced at the electrode-

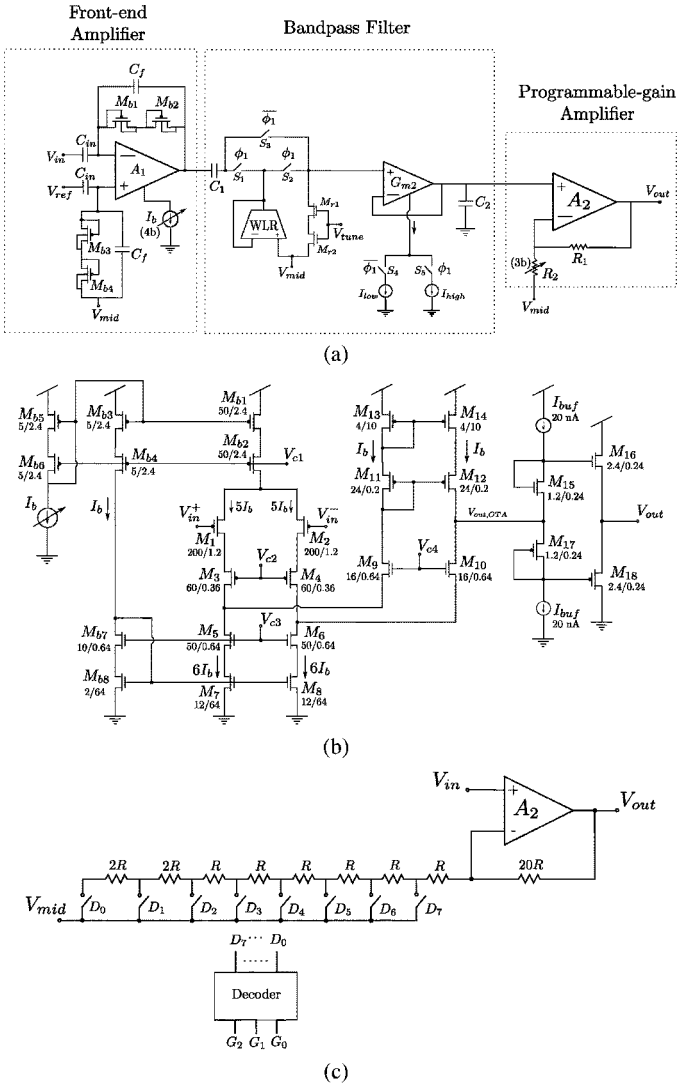

Fig. 3. (a) Schematic of the neural amplifier amplifier consisting of three stages: i) front-end amplifier ii) bandpass filter iii) programmable-gain amplifier. (b) Schematic of the amplifier A_1. (c) Schematic of the programmable-gain amplifier.

tissue interface along with high-resistance pseudoresistor elements to stabilize its DC operating point [17]. The midband gain of the front-end amplifier is set to 40 dB by capacitive feedback formed by C_{in} and C_f around the high-gain amplifier A_1. The front-end amplifier provides wide-band amplification from below 1 Hz to greater than 100 kHz, while the passband of the overall neural amplifier is determined by the bandpass filter, which can be chosen for one of the following two settings: i) a spike-recording setting (350 Hz–12 kHz) and ii) an LFP-recording setting (<1 Hz–300 Hz): Due to the difference in their amplitudes, spikes and LFPs are separated in the frequency domain before each signal type is amplified with programmably different gains [18]. The programmable-gain amplifier in Fig. 3(c) provides gain that ranges from 9 dB to 26 dB, adjustable in eight unequal steps. As a result, the overall gain of each channel can be adjusted from 49 dB to 66 dB based on a user-provided digital input.

To get clean neural recordings, while keeping the total power consumption of the entire amplifier array small, we utilize an adaptive-biasing scheme as proposed in [19]. In such a scheme, each neural amplifier's input-referred noise, and thus its power consumption, can be individually adjusted to suit the background noise level at its corresponding recording site. As a result, every neural amplifier in the array consumes just sufficient power to obtain clean recordings while the total power consumption of the recording chip is near optimal. Due to the high gain of the front-end amplifier, noise introduced by A_1 determines the noise performance of the overall amplification. By adjusting the input-referred noise per unit bandwidth of A_1 through changes in its bias current while keeping the bandwidth of the overall neural amplifier constant, we can adjust the input-referred noise of the neural amplifier. The bias current of A_1 is controlled by a 4-bit binary current DAC which is represented as the variable current source I_b in Fig. 3(a). The schematic of the amplifier A_1 including its transistor sizing is shown in Fig. 3(b). The amplifier A_1 consists of a folded-cascode operational transconductance amplifier (OTA), formed by M_{b1}, M_{b2}, M_1-M_{14}, followed by a class-AB output buffer, formed by M_{15}-M_{18}. The transistors M_{b3}-M_{b8} form the bias circuit that helps distribute current in the OTA. The class-AB output buffer helps minimize the output impedance of A_1 to ensure that, at every bias current level of the OTA, the bandwidth of the front-end amplifier is much wider than the bandwidth of the bandpass filter such that the bandwidth of the overall neural amplifier is constant and determined by that of the bandpass filter. The folded-cascode OTA is modified from our prior low-power, low-noise OTA presented in [20] which achieves nearly optimal noise efficiency at the fundamental limits. However, to minimize the layout area per recording channel, we replace the source-degeneration resistors in the OTA in [20] with the transistors M_{b8}, M_7, and M_8 at the expense of reduced noise efficiency. The input-referred noise power per unit bandwidth of A_1, $\overline{v_{n,A1}^2}$, which consists of both thermal and 1/f noise components can be expressed as

$$\overline{v_{\text{n,A1}}^2} = \frac{1}{g_{m1}}\left[\frac{4kT}{\kappa} + \frac{16}{3}kT\frac{g_{m7}}{g_{m1}} + \frac{1}{\alpha^2}\frac{16}{3}kT\frac{g_{m13}}{g_{m1}}\right] + \left[\frac{2K_p}{W_1L_1C_{\text{ox}}^2} + \frac{2K_n}{W_7L_7C_{\text{ox}}^2}\left(\frac{g_{m7}}{g_{m1}}\right)^2 + \frac{1}{\alpha^2}\frac{2K_p}{W_{13}L_{13}C_{\text{ox}}^2}\left(\frac{g_{m13}}{g_{m1}}\right)^2\right]\frac{1}{f} \quad (1)$$

where g_{mi} is the transconductance of the ith transistor in the OTA, κ is the inverse of the subthreshold slope [8], k is the Boltzmann constant, T is the absolute temperature, C_{ox} is the oxide capacitance per unit area, and K_p and K_n are the technology-dependent 1/f noise coefficients of the pFET and nFET respectively. The factor α represents the ratio of the effective transconductance of the OTA (G_m) and the transconductance of M_1 and M_2 (g_{m1}), and can be approximated as

$$\alpha \approx \frac{g_{s9}}{g_{s9} + g_{ds7}/g_{s5}r_{o5} + g_{ds1}/g_{s3}r_{o3}} \quad (2)$$

where g_{si} and r_{oi} represent the source admittance and the Early Effect drain-to-source resistance of the ith transistor respectively ($g_{dsi} = 1/r_{oi}$). Due to the use of cascode transistors M_3, M_4, M_5 and M_6, the factor α in (2) has a value close to 1. The noise optimization strategy for this OTA is similar to that described in [20].

The bandpass filter is designed using a Gm-C bandpass filter topology with switches to configure the recording setting. The recording setting is controlled by the signal ϕ_1; when $\phi_1 = 1$, the filter is in a spike-recording setting, and when $\phi_1 = 0$, the filter is in an LFP-recording setting. The high-frequency cutoff f_h is set by G_{m2}-OTA connected in a unity-gain configuration and C_2. In the spike-recording setting ($\phi_1 = 1$), the switch S_5 is closed and the G_{m2}-OTA is biased with a higher bias current, $I_{\rm high}$, such that f_h is approximately 12 kHz. Similarly, in the LFP-recording setting ($\phi_1 = 0$), the switch S_4 is closed and G_{m2}-OTA is biased with a lower bias current, I_{low}, such that f_h is approximately 350 Hz. To set the low-frequency cutoff f_l in the spike-recording setting, we use a combination of C_1 and the unity-gain connected WLR-OTA [21]. In the LFP-recording setting, however, due to the difficulties of biasing a WLR-OTA at very low bias current to achieve $f_l < 1$ Hz, we simply use a series combination of M_{r1} and M_{r2} to provide a large effective resistance (denoted as R_p) to set f_l below 1 Hz. The gate voltage $V_{\rm tune}$ of M_{r1} and M_{r2} is set such that R_p is much larger than $1/G_{m,\rm WLR}$, where $G_{m,\rm WLR}$ is the transconductance of the WLR-OTA. In the spike-recording setting ($\phi_1 = 1$), the switches S_1 and S_2 are closed and the WLR-OTA appears in parallel with R_p. Since $R_p \gg 1/G_{m,\rm WLR}$, the effective resistance $1/G_{m,\rm WLR}$ dominates, resulting in $f_l = G_{m,\rm WLR}/2\pi C_1$. The WLR-OTA's bias current is set such that $f_l = 350$ Hz. In the LFP-recording setting ($\phi_1 = 0$), the switches S_1 and S_2 are open while the switch S_3 is closed. The WLR-OTA is disconnected from the signal path and C_1 appears in series with R_p through the switch S_3. In this case, the combination of C_1 and R_p determines the cutoff frequency f_l. By setting $V_{\rm tune} = V_{\rm mid}$, the effective resistance R_p is very large such that $f_l = 1/2\pi R_p C_1 < 1$ Hz. To operate OTAs at low frequencies, we utilize bump-linearization and source-degeneration techniques described in [21] or in [8] to reduce the G_m of WLR-OTA to achieve a low-frequency cutoff of 300 Hz in the spike-recording setting. However, only the bump-linearization technique is used in the G_{m2}-OTA to achieve a high-frequency cutoff of 350 Hz in the LFP-recording setting.

Fig. 3(c) shows the schematic of the programmable-gain amplifier. The gain of the programmable-gain amplifier can be programmed to any of the eight values and is given by $A_{v2} = 1 + (20R/R_v(D_i))$, $i \in \{0, \ldots, 7\}$, where $R_v(D_i)$ is the total resistance seen between the negative input terminal of A_2 and the node $V_{\rm mid}$ when the switch D_i is closed. The digital decoder ensures that only one of the switches D_i's can be closed at a given time depending on the decoder's inputs G_2, G_1, and G_0. The values of the gain A_{v2} for every combination of G_2, G_1, and G_0 are tabulated in Table I. The amplifier A_2 is designed using a standard two-stage amplifier topology with Miller compensation. The class-AB output buffer, similar to the one used in Fig. 3(b), is included to drive resistive loads (the feedback resistors) at the output of A_2.

TABLE I
GAINS OF THE PROGRAMMABLE-GAIN AMPLIFIER

G_2	G_1	G_0	switch closed	Gain (dB)
0	0	0	D_0	9.5
0	0	1	D_1	10.8
0	1	0	D_2	12.7
0	1	1	D_3	14
1	0	0	D_4	15.6
1	0	1	D_5	17.7
1	1	0	D_6	20.8
1	1	1	D_7	26.4

IV. NEURAL SIGNAL DIGITIZATION

In neural-recording applications, the sampling speed requirement per recording channel is quite modest. Since the amplified neural signal is bandlimited to 12 kHz, we choose a sampling rate per channel in the spike-recording setting of 31.25 kHz, which is slightly higher than the Nyquist frequency. The same sampling rate per channel is also used in the LFP-recording setting to ease the design of our control logic for the ADC. The ADC in each recording module operates at 125 kS/s to digitize the data from all the four amplifiers in the module. We choose to implement an 8-bit successive approximation register (SAR) ADC with a 1-V full-scale voltage due to its good energy efficiency and small area when implemented at 8-bit resolution. With 60 dB of gain from the neural amplifier, the ADC has sufficient resolution: The ADC's quantization noise referred to the input of the neural amplifier is only $V_{\rm LSB}/(1000 \times \sqrt{12}) = (1/2^8)/(1000 \times \sqrt{12}) = 1.12\,\mu{\rm V_{rms}}$, where $V_{\rm LSB}$ is the voltage corresponding to that between the adjacent digital codes of the ADC. This quantization noise is much less than the input-referred noise of the neural amplifier or noise from neural background activity and from high-impedance electrodes.

A. Analog-to-Digital Converter Design

The schematic of the SAR ADC is shown in Fig. 4(a). The high-level topology of the ADC is similar to the one presented in [22]. The ADC consists of a comparator, SAR logic, switching network, a capacitor DAC array, and a bootstrapped reference switch. To minimize dynamic power consumption of the ADC, the SAR logic is designed using dynamic logic techniques [23] to minimize internal capacitances. The capacitor DAC array is designed using a split-capacitor approach [24] to reduce the power consumed by the capacitor DAC array. The unit capacitance of the capacitor DAC array in Fig. 4(a) is $C = 8$ fF. The clock and control signals of the ADC are derived from an external 10 MHz clock by the Digital Control Unit which are common among all eight ADCs on chip. The timing diagram of the clock and control signals of the ADC is shown in Fig. 4(b). The clock signal CK is used for controlling the timing operation of the SAR logic, while $\rm CK_s$ is used for registering the outputs of the comparator to ensure synchronous operation. The control signals SMP, STRT, EN are used by the SAR logic for sampling the input voltage $V_{\rm in}$, initiating the conversion process, and duty cycling the comparator to reduce its static power consumption respectively.

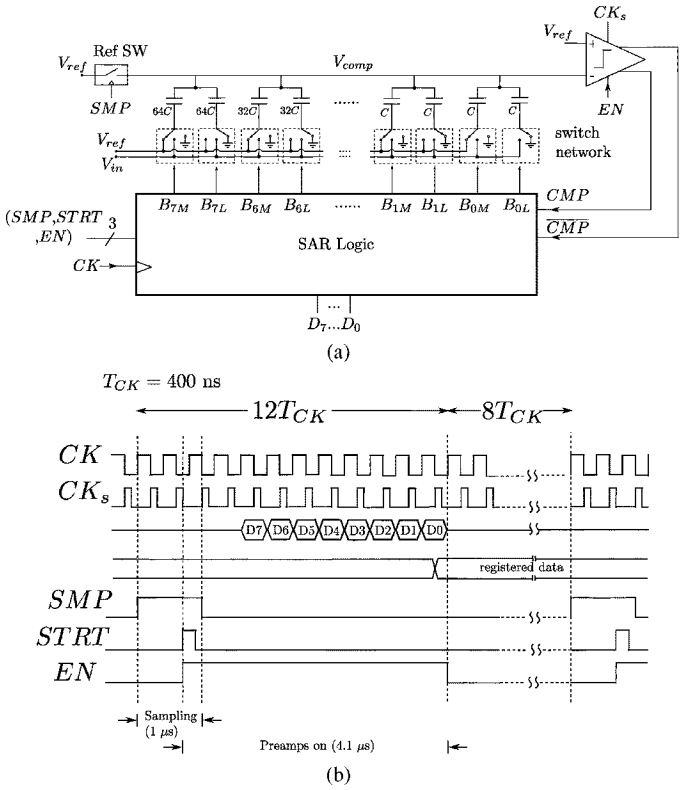

Fig. 4. (a) Schematic of the SAR ADC used in this neural recording chip. (b) Timing diagram of the ADC.

B. Analog Multiplexer

The outputs from four amplifiers are multiplexed through an analog multiplexer and input to the ADC in the recording module. Since the input voltage range of the ADC is from 0 to $V_{\text{ref}} = 1$ V while the DC level of the amplifier's output is at $V_{\text{mid}} = 0.9$ V, the analog multiplexer performs DC level shifting such that the input to the ADC is centered near the midpoint of the ADC's input range (0.5 V). Fig. 5(a) shows the schematic of the analog multiplexer. The core of the analog multiplexer consists of four source-follower drivers, which are formed by the transistors $M_1, M_{b1}, \ldots, M_4, M_{b4}$. The source-follower drivers buffer the amplifiers' outputs $V_{\text{out}1}$-$V_{\text{out}4}$ to provide low output impedance necessary for driving the input capacitance of the ADC and to perform DC level shifting. Multiplexing of the amplifiers' outputs is achieved through the switches S_1-S_4 which are controlled by the control signals A_{1s}-A_{4s} respectively. When A_{is} $(i \in \{1, \ldots, 4\})$ is high, the switch S_i is closed and $V_{\text{out}i}$ is buffered and multiplexed to the input of the ADC. The timing diagram of the control signals $A_{1s}, \ldots, A_{4s}$ is shown in Fig. 5(b). Note that only one of the switches S_1-S_4 is closed at a given time. The duration for which each switch is closed is 1.1 μs and suffices to span the ADC's sampling duration (when SMP is high for 1 μs). As a result, the ADC's input is driven by only one low-impedance source at any sampling instant of the ADC (at the negative edge of SMP). According to Fig. 5(b), each source-follower driver only needs to be active when it is driving the input capacitance of the ADC, which is only 1.1 μs for the whole 32 μs duration of the sampling period. Therefore, to save power, we duty cycle the bias current

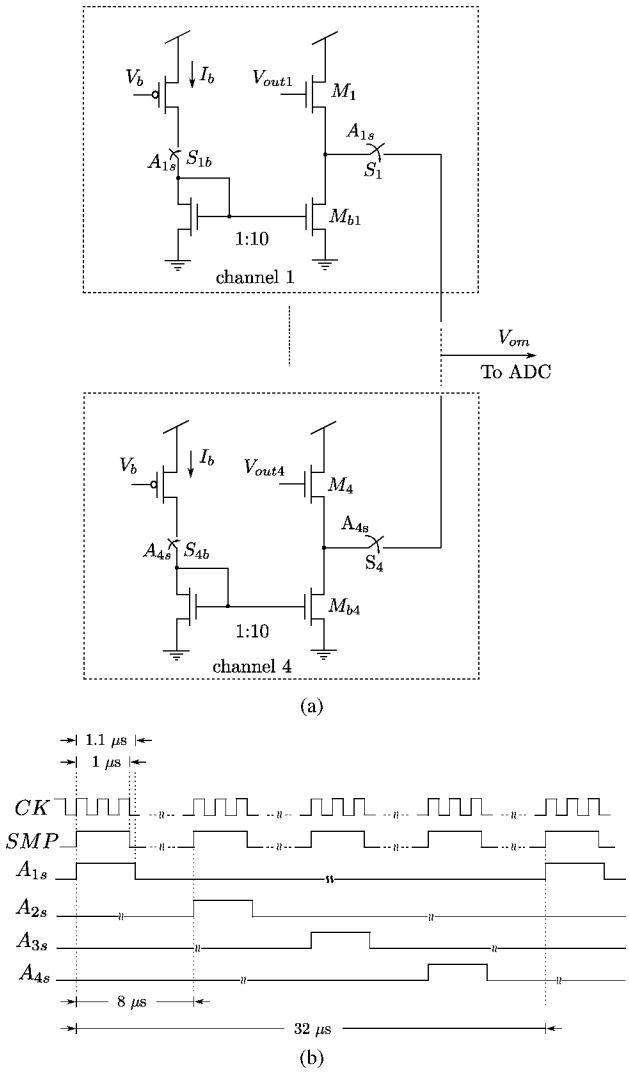

Fig. 5. (a) Schematic of the analog multiplexer. (b) Timing diagram of the analog multiplexer's control signals.

in the ith source-follower driver by its corresponding control signal A_{is}. A power savings of a factor of 32x is thus achieved compared to the case when all drivers are constantly powered.

V. Configuration Commands and Outgoing Data Packet

A. Receiving Configuration Settings

Configuring the 32-channel neural-recording chip is achieved on a module-by-module basis. As shown in Fig. 2, each recording module contains a dedicated serial-programming interface unit. In each recording channel, the front-end amplifier's bias current for adjusting its input-referred noise, the programmable-gain amplifier's gain, and the recording setting (spike or LFP) are configured with 4-bit, 3-bit, and 1-bit control inputs respectively. The bias currents of comparator's preamplifiers in the ADC can also be controlled by a 3-bit control input for calibrating against process variations. To configure each recording module, the user provides a 56-bit programming packet in the format shown in Fig. 6(a). The data packet is loaded into a 56-bit shift register on the positive

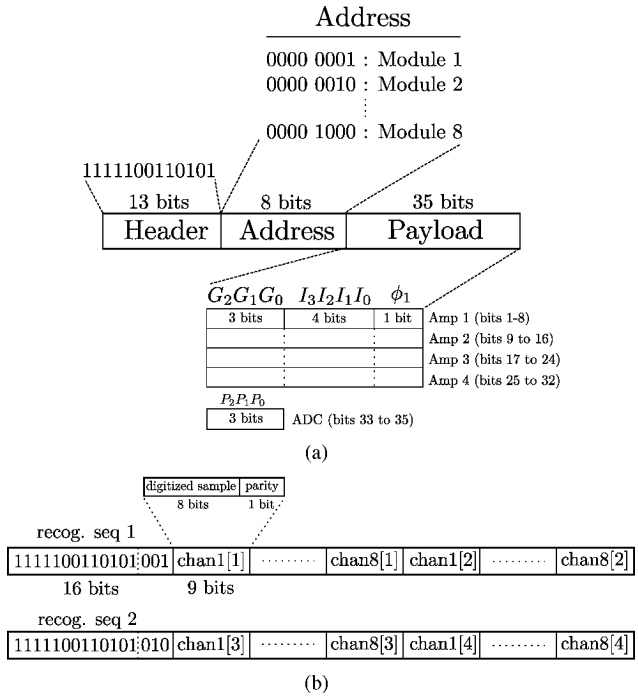

Fig. 6. (a) Format of the programming packet for configuring each neural-recording module. (b) Format of the outgoing data packet of the 32-channel neural-recording chip.

edges of the programming clock. The 56-bit programming data packet consists of a 13-bit header field, an 8-bit module address, and a 35-bit payload field. The header field is a fixed sequence "1111100110101". The module address is an 8-bit binary value that specifies the address (from 1 to possibly 256 modules although we only use 8) of the recording module to be programmed. Once the header field and the module address are recognized by the internal logic, 35 payload bits are latched into storage registers that provide parameter inputs for the recording channels and the comparator's preamplifiers.

B. Outgoing Data Packet

With a 32 μs sampling period per channel and no storage mechanism on chip, the data from each recording channel must be transmitted off-chip within the next 32 μs. To reduce the number of output wires, the digitized raw data from all 32 recording channels is streamed out in a serial-data format. To provide easy synchronization between the 32-channel neural-recording chip and an external FPGA, the data are packetized into a 320-bit data frame as shown in Fig. 6(b). A data frame contains two 16-bit recognition sequences and is streamed out serially at 10 Mbps. For error-correction purposes, an even parity bit is appended to each digitized sample to create 9-bit sample data. In Fig. 6(b), the data frame is divided into two halves with each half containing one 16-bit recognition sequence and sixteen 9-bit samples from 16 recording channels. In Fig. 6(b), the 9-bit sample chan i[j] represents the sample from the recording module i ($i \in \{1, \ldots 8\}$) and channel j ($j \in \{1, \ldots, 4\}$) respectively.

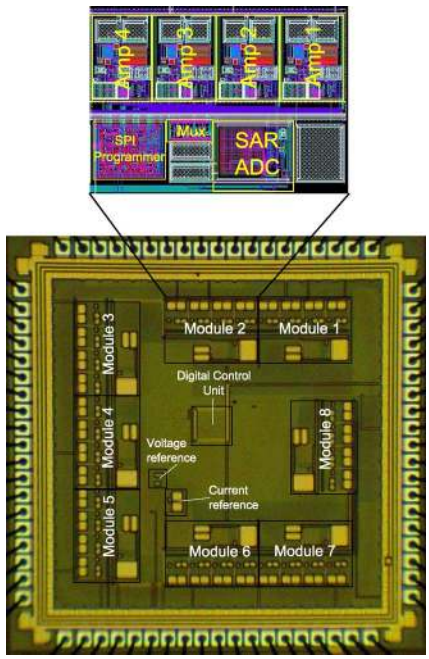

Fig. 7. The micrograph of the 32-channel neural recording chip.

VI. Experimental Results

The 32-channel neural recording chip was fabricated in an IBM 0.18 μm CMOS (7RF) technology through MOSIS. The micrograph of the chip is shown in Fig. 7. Excluding the area for I/O pads, the chip has dimensions of 2.1 mm $\times$ 2.1 mm. Due to the number of I/O pads included for testing purposes, the layout of the whole chip was not optimized for the total chip area. However, the neural amplifier, analog multiplexer, and the ADC were laid out as compactly as possible for further scaling to higher channel counts in subsequent generations. The active areas of the neural amplifier, the analog multiplexer, and the ADC are 0.03 mm^2, 0.006 mm^2, and 0.02 mm^2 respectively. The remaining area of the recording module is occupied by the serial programming interface unit and by power-supply decoupling capacitors.

A. Benchtop Testing of the Neural Amplifier

In this section, we describe experimental measurements of our neural amplifier. For frequency-response measurements, swept-sine techniques were used with an input amplitude of 100 μV. Fig. 8(a) shows the measured magnitude responses of the neural amplifier in the spike-recording setting for four different gain settings. The lower and upper -3 dB cutoff frequencies were measured to be $f_l = 350$ Hz and $f_h = 11.7$ kHz respectively, and were constant across different gain settings. Fig. 8(b) shows the magnitude response of the neural amplifier when configured for an LFP-recording setting. At this particular setting, the amplifier exhibits a midband gain of 53.3 dB with -3 dB cutoff frequencies of $f_l = 126$ mHz and $f_h = 293$ Hz. The common-mode-rejection-ratio (CMRR) and the power-supply-rejection-ratio (PSRR) were measured

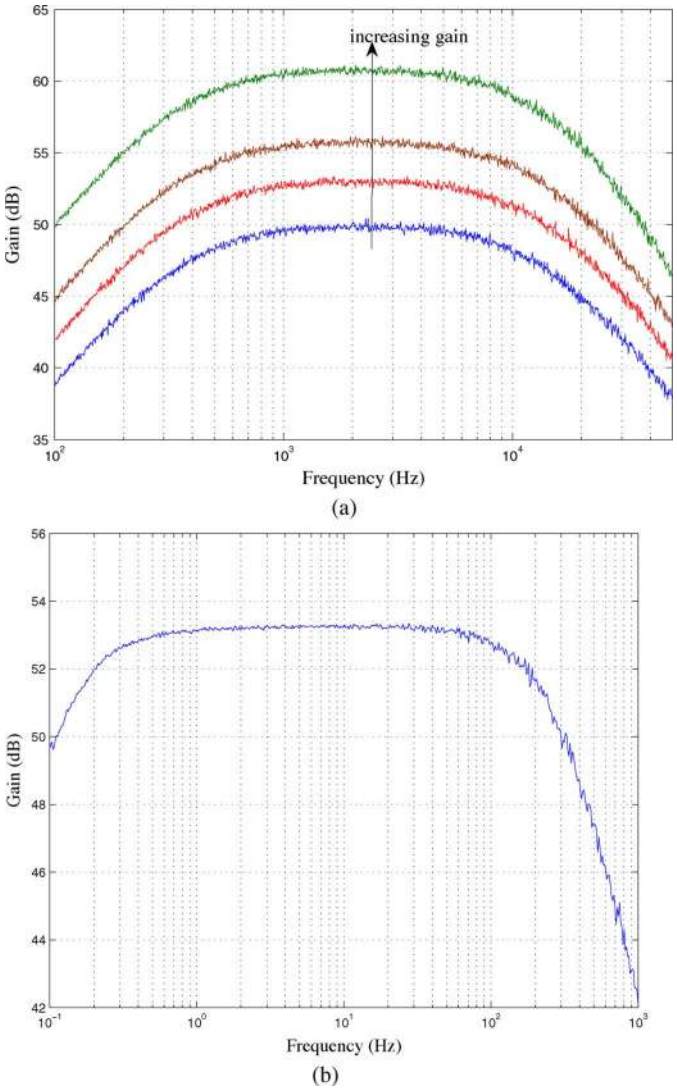

Fig. 8. (a) Magnitude responses of the amplifier at different gain settings in the spike-recording setting. (b) Magnitude Response of the amplifier in the LFP-recording setting.

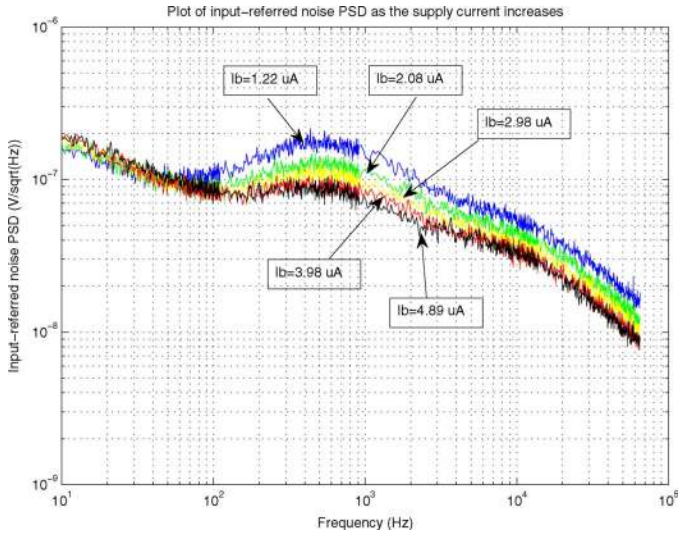

Fig. 9. Input-referred noise density of the neural amplifier for various amplifier's supply currents.

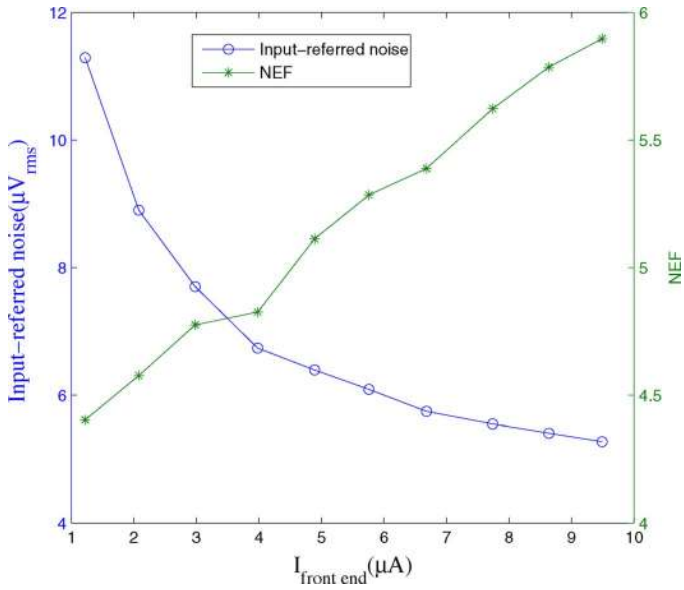

Fig. 10. Integrated input-referred rms noise and the NEF of the neural amplifier versus front-end amplifier's bias current.

to be 62 dB and 72 dB respectively. Fig. 9 shows the input-referred noise spectral density of the neural amplifier in the spike-recording setting as we increase the front-end amplifier's bias current [I_b in Fig. 3(a)]. Notice that 1/f noise can be observed in the passband (350 Hz–12 kHz). This 1/f noise results from making our neural amplifier small for further scaling to higher channel-count systems and is a consequence of our choosing not to implement resistive degeneration as in our lower-noise design [20]. As we increase the front-end amplifier's bias current, the input-referred noise spectral density in the frequency range above 100 Hz decreases as expected. However, at frequencies below 100 Hz, the input-referred noise spectral density is invariant to the front-end amplifier's bias current. At such low frequencies, the noise of the front-end amplifier is no longer the dominant factor as the gain from the front-end amplifier to the output drops significantly due to the filtering effect of the bandpass filter ($f_l = 350$ Hz). As a result, the noise from the programmable gain amplifier (especially its 1/f noise) becomes the dominant noise source such that the overall noise is independent of the front-end amplifier's bias current.

The total input-referred rms noise of the neural amplifier for different front-end amplifier's bias current levels (including the current from the local bias circuits) is calculated by integrating the corresponding input-referred noise density curve from 10 Hz to 65 kHz. The Noise Efficiency Factor (NEF) [25] is calculated for each front-end amplifier's bias current level. The plot showing the input-referred rms noise and the NEF versus the front-end amplifier's bias current level is shown in Fig. 10. It is interesting to note that as the front-end amplifier's bias current increases, its input-referred noise decreases as expected, however, at a slower rate than expected from purely thermal-noise considerations alone. As we significantly increase the front-end amplifier's bias current, its thermal noise component decreases and its 1/f noise becomes the limiting component. Our neural

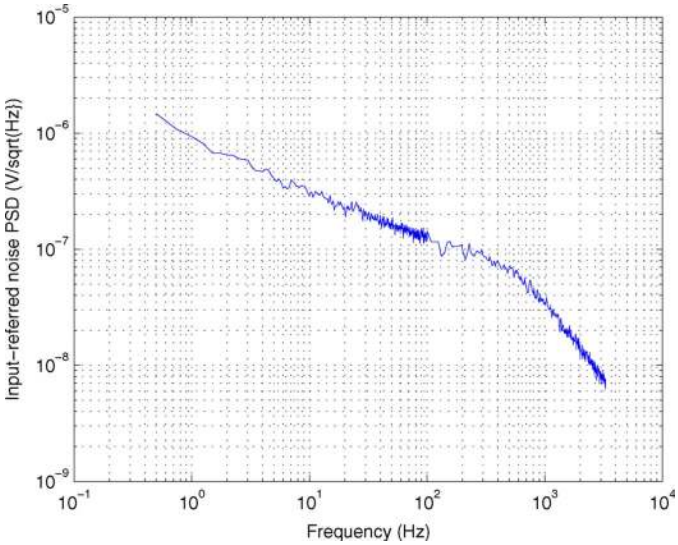

Fig. 11. Input-referred noise density in the LFP-recording setting.

TABLE II
PERFORMANCE SUMMARY OF THE NEURAL AMPLIFIER

Performance Metric	Value
Supply Voltage	1.8 V
Programmable Gain	49-66 dB
Bandwidth:	
spike-recording setting	350 Hz-11.7 kHz
LFP-recording setting	126 mHz-293 Hz
Input-referred noise (spike-recording)	5.4 μV_{rms}-11.2 μV_{rms}
CMRR	62 dB
PSRR	72 dB
Power Consumption per Channel:	
Bandpass filter & Current DAC	700 nW
Programmable-gain amplifier	2.6 μW
Front-end amplifier	2.1 - 16.6 μW
Total	5.4 - 20 μW
Active Area	0.03 mm^2
NEF of 1^{st} stage	4.4-5.9

amplifier achieves an NEF in the range of 4.4–5.9 which is only slightly higher than our NEF of 2.67 reported in [20]. However, this design is significantly more area efficient than the design in [20] such that the higher NEF represents a good system tradeoff. The input-referred noise spectral density of the neural amplifier in the LFP-recording setting is shown in Fig. 11. Integrating the input-referred noise spectral density curve from 500 mHz to 3.3 kHz yields a total input-referred noise of 3.14 μV_{rms} with a front-end amplifier's bias current of 3.98 μA. At this low frequency, 1/f noise of the front-end amplifier dominates the overall input-referred noise of the neural amplifier making our adaptive biasing strategy not useful. Table II summarizes the measured performance of the neural amplifier.

B. Benchtop Testing of the Analog-to-Digital Converter

In this section, we describe experimental characterization of our ADC. Static measurements were performed with a histogram test to obtain integral nonlinearity (INL) and differential nonlinearity (DNL) plots as shown in Fig. 12(a). A least-squared approximation was used to calculate the INL.

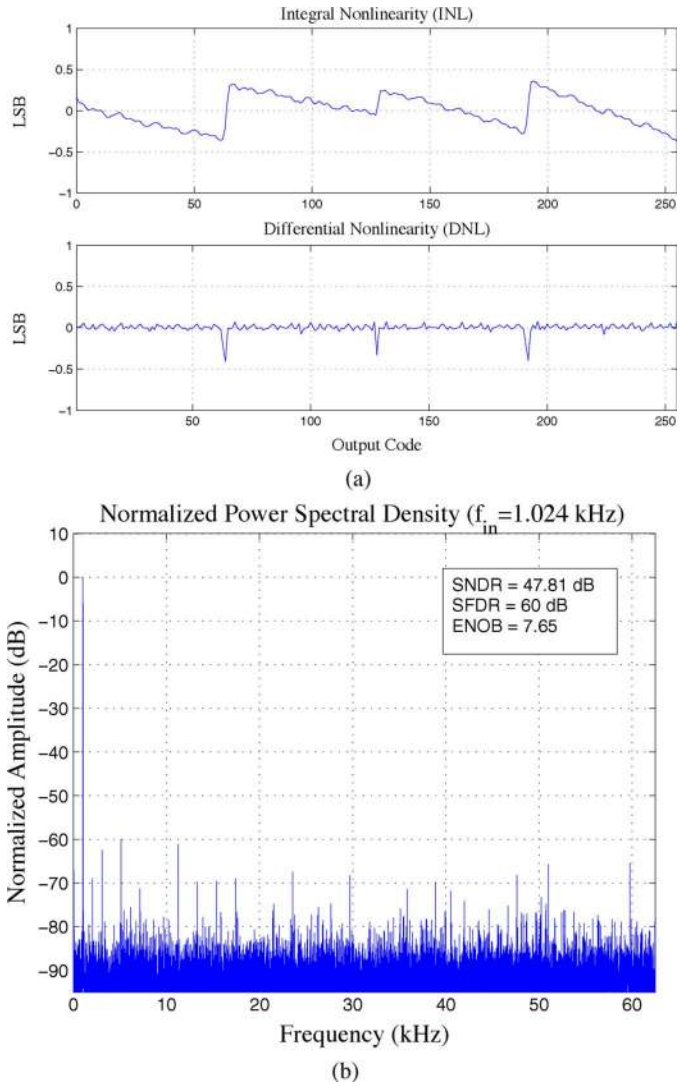

Fig. 12. (a) Low-frequency INL and DNL plots of our ADC. The INL is obtained using a least-squared approximation. (b) Measured output spectrum of the ADC with a rail-to-rail input sine wave of 1.024 kHz.

Both the INL and DNL of our ADC are within ± 0.4 LSB. For dynamic measurements, we input a full-scale (1 V_{pp}) 1.024-kHz input sine wave and sampled the signal at a rate of 125 kHz. Fast Fourier-Transform (FFT) analysis applied to 31,982 sample points yields a plot of the power-spectral density shown in Fig. 12(b). The signal-to-noise-and-distortion-ratio (SNDR) is calculated from this plot to be 47.81 dB. The effective number of bits (ENOB) is calculated from the SNDR to be 7.65 bits. The spurious-free dynamic range is 60 dB and is limited by the fifth harmonic. The total power consumption of the ADC from a 1-V supply voltage for this particular measurement is 1.93 μW which can be apportioned as follows: 592 nW from the comparator, 1.13 μ W from the custom SAR logic, and 210 nW from the capacitor DAC array. Note that this power consumption is for converting the output signals from four amplifiers in a recording module. The average power consumption per recording channel of the ADC is only 483 nW. One important figure of merit (FOM) of the ADC that is widely used to compare the ADCs across a wide

TABLE III
PERFORMANCE SUMMARY OF THE ANALOG TO DIGITAL CONVERTER

Performance Metric	Value
Supply Voltage	1 V
Full-scale voltage	1 V
Precision	8 bits
ADC's input range	0-1 V
Sampling rate	125 kHz
INL	$< \pm$ 0.4 LSB (8 bits)
DNL	$< \pm$ 0.4 LSB (8 bits)
SNDR	47.8 dB
SFDR	60 dB
ENOB	7.65 bits
Power Dissipation:	
Analog(comparator & Capacitor DAC)	802 nW
Digital (SAR logic)	1.13 μW
Energy per quantization level	77 fJ/State
Active Area	0.02 mm^2

range of bandwidths and precisions is the energy consumption per quantization level [8]. The FOM can be calculated from the formula $\mathrm{FOM} = P_{\mathrm{total}}/(2^{\mathrm{ENOB}} \times f_{\mathrm{samp}})$, where P_{total} and f_{samp} are the total power consumption and the sampling frequency of the ADC respectively. The FOM of our ADC is calculated to be 77 fJ per quantization level and is among the most energy-efficient ADCs reported to date especially for its compact 0.02 mm^2 area. The performance summary of our ADC is provided in Table III.

C. Wireless In Vivo Testing of the Neural-Recording Chip in an Awake Behaving Primate

We performed an *in vivo* biological experiment in a rhesus macaque monkey using our chip. We used our chip as the core of a wireless neural-recording system. Our wireless neural-recording system consisted of a recording unit and a receiver unit. For wireless data transmission, we used the impedance-modulation data-telemetry system described in [26], which transmits data using an inductive-coupling scheme at a carrier frequency of 25 MHz. The 32-channel neural recording chip (in a QFP80 package) was integrated onto the recording unit on a custom PCB with dimensions of 8.5 cm $\times$ 9.5 cm. The recording unit also contained a low-power FPGA (IGLOO series, AGL060 from Microsemi) for data selection, power regulators to generate specific supply voltages for different subsystems, a 10-MHz crystal oscillator for system clock generation, the internal unit of the impedance-modulation data telemetry system, and a transmitting RF coil. The receiver unit was constructed on another PCB and housed the external unit of the impedance-modulation data telemetry system and another receiving RF coil. The recording unit was powered by a 3.7 V supply voltage and the receiver unit was powered by a 2.5 V supply voltage from Keithley source meters. The 3.7 V supply voltage of the recording unit was chosen to emulate the nominal output voltage of a rechargeable Li-ion battery, which is normally used for powering implantable electronics [8]. During the experiment, the monkey was seated in an experimental chair, with its head restrained. The recording unit was mounted on a platform placed near a shoulder of the monkey, and eight input channels of the recording unit were wired to Tungsten microwire electrodes (Alpha Omega Co., Alpharetta, GA) with an electrode impedance of approximately 1 MΩ at 1 kHz. The

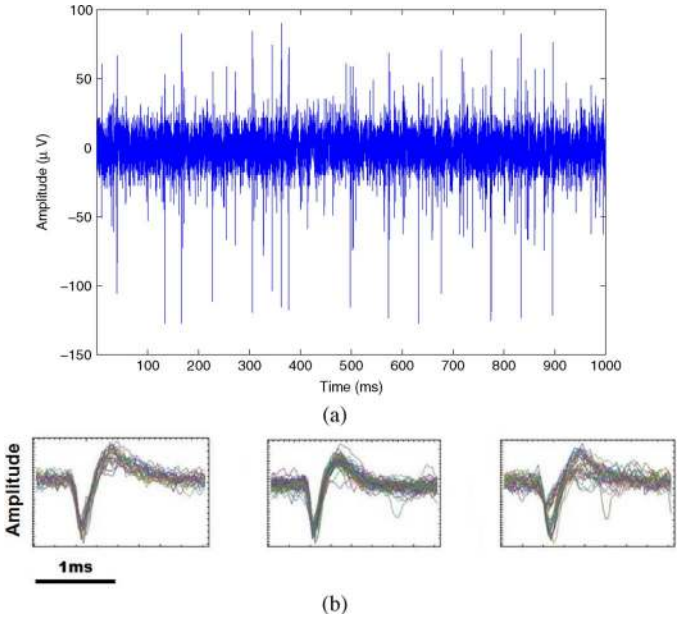

Fig. 13. Electrode-referred neural signals recorded from the brain of a rhesus macaque and transmitted wirelessly: (a) 1-second long raw neural data from 1 channel. (b) Spikes extracted by our system from 3 channels over a 1-minute period.

TABLE IV
SYSTEM LEVEL PERFORMANCE

Technology	0.18 μm CMOS
Voltage Supply:	
Amplifier array & reference circuits	1.8 V
ADC array & Digital Control Unit	1 V
Channel Count	32
Dimensions of the Neural Amplifier	215 μm $\times$ 155 μm
Dimensions of the Recording Module	680 μm $\times$ 450 μm
Die Dimensions	3.15 mm $\times$ 3.15mm
ADC's Input Range	0-1 V
ADC's Sampling Rate per Channel	31.25 kHz
INL of ADC	$< \pm$ 0.4 LSB
DNL of ADC	$< \pm$ 0.4 LSB
ENOB of ADC	7.65 bits
Power Dissipation:	
Neural Amplifier Array	207 μW (biased at NEF = 4.5)
ADC Array	15 μW
Digital Control Unit	42 μW
Voltage and Current References	61 μW
Total	325 μW

electrodes were lowered into the brain tissue of the monkey just before the experiment. The FPGA on the recording unit was programmed to select digitized data from eight input channels of the 32-channel neural recording chip. The receiver unit was interfaced with a PC via a USB data acquisition system (XEM3010 from Opal Kelly). Throughout the experiment, the transmitting and receiving coils of the data-telemetry system were placed concentric to each other and spaced approximately 1 cm apart.

Prior to each recording session, configuration commands were sent wirelessly from the PC to the 32-channel neural-recording chip on the recording unit. All 32 amplifiers were configured with a spike-recording setting, with a gain of 60 dB, and a nominal front-end amplifier's bias current of 2.08 μA. The power consumption of the 32-channel neural-recording

TABLE V
COMPARISON TO OTHER STATE-OF-THE-ART NEURAL RECORDING SYSTEMS

Reference	[14]	[27]	[28]	[29]	This work
Channel Count	128	128	16	8	32
Supply Voltage	3 V	± 1.65 V	1.8 V	1.5 V	1.8 V(analog) 1 V (digital)
Technology	0.35-μm CMOS	0.35-μm CMOS	0.18-μm CMOS	0.35-μm CMOS	0.18-μm CMOS
Mid-band gain	54-73 dB	57-60 dB	70 dB	51.9-65.6 dB	49-66 dB
Low freq. cutoff	0.5-50 Hz	0.1-200 Hz	100 Hz	1.1-525 Hz	0.126 Hz, 350 Hz
High freq. cutoff	500 Hz-10 kHz	2 kHz-20 kHz	9.2 kHz	5.1-12 kHz	293 Hz, 12 kHz
Input-referred noise	6.08 μV_{rms} (10 Hz-10 kHz)	4.9 μV_{rms} -	5.4 μV_{rms} -	3.12 μV_{rms} (0.5 Hz-50 kHz)	5.4-11.2 μV_{rms} (10 Hz-65 kHz)
NEF	5.5	-	4.9	2.68	4.4 - 5.9
ADC resolution	8 bits	6-9 bits (adjustable)	8 bits	10 bits	8 bits
ADC sampling rate/channel	14 kHz	40 kHz	30 kHz	35.7 kHz	31.25 kHz
ADC INL & DNL (LSB)	-	-	0.5/0.5	0.8	0.4/0.4
ADC ENOB	6.5 bits	-	7 bits	9.2 bits	7.65 bits
Total Power Dissipation	2.43 mW	3 mW	680 μW	211 μW	325 μW (@NEF=4.5)
Average Power /channel	19 μW/chan.	23.4 μW/chan.	42.5 μW/chan.	26.4 μW/chan.	10.1 μW/chan. (@NEF=4.5)

chip for this experiment was 325 μW. The power consumption of the internal unit of the impedance-modulation wireless data telemetry system was approximately 100 μW, while the receiver unit's power consumption was approximately 3.0 mW for a 2.5 Mbps data rate [26]. Fig. 13(a) shows a 1-second long raw neural data from one of the input channels recorded with our system. Fig. 13(b) shows superimposed neural spikes extracted by our chip from 3 channels over a 1-minute period. The performance summary of the neural-recording chip during the wireless recording experiment is shown in Table IV.

Table V compares the performance of our design with some designs in the literature that achieve low power consumption and small area per channel. The design in [14], [29] includes both recording and stimulation features while the design in [27] also included digital signal processing (DSP) and an ultra-wide-band (UWB) transmitter. In Table V, only the recording features that include signal amplification and digitization are compared.

VII. CONCLUSION

We described the operation and measured performance of an ultra-low-power 32-channel neural-recording chip. The chip can amplify and convert neural signals from 32 input channels into 8-bit digital representations and transmit this data off chip in a serial format. An adaptive-biasing technique was utilized in the design of the neural amplifier to help minimize the total power consumption of the overall recording chip. Our neural amplifiers are highly programmable: their gain, recording mode (spike or LFP), and input-referred noise can be digitally programmed to suit the recording environment. Our recording chip was successfully tested in an *in vivo* wireless recording experiment from a behaving primate while dissipating only 10.1 μW/channel. Due to its very small area and power consumption per recording channel, our recording chip is suitable for high-channel-count recording systems.

ACKNOWLEDGMENT

The authors would like to thank Prof. R. A. Andersen for providing facilities for the animal experiments, B. I. Rappoport and E. J. Hwang for tremendous help with the rhesus macaque recording experiment, D. Kumar for help in designing the on-chip current reference and digital standard cells, L. Turicchia for developing a software interface between the data-telemetry system and a computer, B. Do Valle for help in designing the wireless setup, and S. K. Arfin for valuable discussions and comments during manuscript preparation.

REFERENCES

[1] D. R. Kipke, R. J. Vetter, J. C. Williams, and J. F. Hetke, "Silicon-substrate intracortical microelectrode arrays for long-term recording of neuronal spike activity in cerebral cortex," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 11, no. 2, pp. 151–155, Jun. 2003.

[2] E. Maynard, C. Nordhausen, and R. Normann, "The utah intracortical electrode array: A recording structure for potential brain-computer interfaces," *Electroencephalogr. Clin. Neurophysiol.*, vol. 102, no. 3, pp. 228–239, Mar. 1997.

[3] J. Wessberg, C. Stambaugh, J. Kralik, P. Beck, M. Laubach, J. Chapin, J. Kim, J. Biggs, M. Srinivasan, and M. Nicolelis, "Real-time prediction of hand trajectory by ensembles of cortical neurons in primates," *Nature*, vol. 408, no. 6810, pp. 361–365, Nov. 16, 2000.

[4] D. Taylor, S. Tillery, and A. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, no. 5574, pp. 1829–1832, Jun. 7, 2002.

[5] M. Serruya, N. Hatsopoulos, L. Paninski, M. Fellows, and J. Donoghue, "Instant neural control of a movement signal," *Nature*, vol. 416, no. 6877, pp. 141–142, Mar. 14, 2002.

[6] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue, "Neuronal ensemble control of prosthetic devices by a human with tetraplegia," *Nature*, vol. 442, no. 7099, pp. 164–171, Jul. 13, 2006.

[7] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz, "Cortical control of a prosthetic arm for self-feeding," *Nature*, vol. 453, no. 7198, pp. 1098–1101, Jun. 19, 2008.

[8] R. Sarpeshkar, *Ultra Low Power Bioelectronics: Fundamentals, Biomedical Applications, and Bio-Inspired Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2010.

[9] R. Olsson and K. Wise, "A three-dimensional neural recording microsystem with implantable data compression circuitry," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2796–2804, Dec. 2005.

[10] R. R. Harrison, P. T. Watkins, R. J. Kier, R. O. Lovejoy, D. J. Black, B. Greger, and F. Solzbacher, "A low-power integrated circuit for a wireless 100-electrode neural recording system," *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 123–133, Jan. 2007.

[11] A. A. Sodagar, K. D. Wise, and K. Najafi, "A fully integrated mixed-signal neural processor for implantable multichannel cortical recording," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 6, pt. 1, pp. 1075–1088, Jun. 2007.

[12] M. Chae, W. Liu, Z. Yang, T. Chen, J. Kim, M. Sivaprakasam, and M. Yuce, "A 128-channel 6 mw wireless neural recording ic with on-the-fly spike sorting and uwb tansmitter," in *Proc. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 146–603.
[13] M. Yin and M. Ghovanloo, "A flexible clockless 32-ch simultaneous wireless neural recording system with adjustable resolution," in *Proc. Solid-State Circuits Conf. Dig. Tech. Papers*, 2009, pp. 432–433, 433a.
[14] F. Shahrokhi, K. Abdelhalim, D. Serletis, P. Carlen, and R. Genov, "The 128-channel fully differential digital integrated neural recording and stimulation interface," *IEEE Trans. Biomed. Circuits Syst.*, vol. 4, no. 3, pp. 149–161, Jun. 2010.
[15] W.-S. Liew, X. Zou, L. Yao, and Y. Lian, "A 1-V 60-mu W 16-channel interface chip for implantable neural recording," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2009, pp. 507–510.
[16] Y. Ramadass and A. Chandrakasan, "Voltage scalable switched capacitor dc-dc converter for ultra-low-power on-chip applications," in *Proc. IEEE Power Electronics Specialists Conf.*, Jun. 2007, pp. 2353–2359.
[17] R. Harrison and C. Charles, "A low-power low-noise CMOS amplifier for neural recording applications," *IEEE J. Solid-State Circuits*, vol. 38, no. 6, pp. 958–965, Jun. 2003.
[18] Y. Perelman and R. Ginosar, "An integrated system for multichannel neuronal recording with spike/LFP separation, integrated A/D conversion and threshold detection," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 1, pp. 130–137, Jan. 2007.
[19] R. Sarpeshkar, W. Wattanapanitch, S. K. Arfin, B. I. Rapoport, S. Mandal, M. W. Baker, M. S. Fee, S. Musallam, and R. A. Andersen, "Low-power circuits for brain-machine interfaces," *IEEE Trans. Biomed. Circuits Syst.*, vol. 2, no. 3, pp. 173–183, 2008.
[20] W. Wattanapanitch, M. Fee, and R. Sarpeshkar, "An energy-efficient micropower neural recording amplifier," *IEEE Trans. Biomed. Circuits Syst.*, vol. 1, no. 2, pp. 136–147, Jun. 2007.
[21] R. Sarpeshkar, R. Lyon, and C. Mead, "A low-power wide-linear-range transconductance amplifier," *Analog Int. Circuits Sig. Proc.*, vol. 13, no. 1–2, pp. 123–151, May–Jun. 1997.
[22] M. Scott, B. Boser, and K. Pister, "An ultralow-energy ADC for smart dust," *IEEE J. Solid-State Circuits*, vol. 38, no. 7, pp. 1123–1129, Jul. 2003.
[23] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 2003.
[24] B. Ginsburg and A. Chandrakasan, "An energy-efficient charge recycling approach for a SAR converter with capacitive DAC," in *Proc. IEEE Int. Symp. Circuits and Systems*, 2005, pp. 184–187.
[25] M. Steyaert, W. Sansen, and Z. Chang, "A micropower low-noise monolithic instrumentation amplifier for medical purposes," *IEEE J. Solid-State Circuits*, vol. 22, no. 6, pp. 1163–1168, Dec. 1987.
[26] S. Mandal and R. Sarpeshkar, "Power-efficient impedance-modulation wireless data links for biomedical implants," *IEEE Trans. Biomed. Circuits Syst.*, vol. 2, no. 4, pp. 301–315, Dec. 2008.
[27] M. S. Chae, Z. Yang, M. R. Yuce, L. Hoang, and W. Liu, "A 128-channel 6 mW wireless neural recording IC with spike feature extraction and UWB transmitter," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 17, no. 4, pp. 312–321, Aug. 2009.
[28] B. Gosselin, A. E. Ayoub, J.-F. Roy, M. Sawan, F. Lepore, A. Chaudhuri, and D. Guitton, "A mixed-signal multichip neural recording interface with bandwidth reduction," *IEEE Trans. Biomed. Circuits Syst.*, vol. 3, no. 3, pp. 129–141, Jun. 2009.
[29] M. Azin, D. Guggenmos, S. Barbay, R. Nudo, and P. Mohseni, "A battery-powered activity-dependent intracortical microstimulation ic for brain-machine-brain interface," *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 731–745, Apr. 2011.

Woradorn Wattanapanitch (S'07) received the B.S. degree (Summa Cum Laude) in electrical and computer engineering from Cornell University, Ithaca, NY, and the M.Sc. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2005, 2007, and 2011, respectively,

At MIT, he worked on developing ultra-low-power electronics for biomedical applications. In 2011, he joined the Department of Electrical Engineering, Kasetsart University, Bangkok, Thailand, as a faculty member where he continues his research interests in the areas of low-power analog and mixed-signal VLSI for biomedical applications and control theory.

Rahul Sarpeshkar (SM'07) received the B.S. degrees in electrical engineering and physics from the Massachusetts Institute of Technology (MIT), Cambridge, and the Ph.D. degree in electrical engineering from the California Institute of Technology, Pasadena.

He then joined Bell Labs as a member of the technical staff in the Department of Biological Computation within its physics division. Since 1999, he has been on the faculty of MITs Electrical Engineering and Computer Science Department where he heads a research group on analog very-large-scale and biological systems. He holds more than 25 patents and has authored many publications, including one featured on the cover of *Nature*. He has authored *Ultra Low Power Bioelectronics: Fundamentals, Biomedical Applications and Bio-inspired Systems*, which provides a broad and deep treatment of the fields of low-power electronics and bioelectronics. He is an Associate Editor of the IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSYTEMS.

Dr. Sarpeshkar has received several awards, including the National Science Foundation Career Award, the Office of Naval Research Young Investigator Award, the Packard Fellows Award, and the Indus Technovator Award for his interdisciplinary bioengineering research. He serves on the program committees of several technical conferences.