A Modular 512-Channel Neural Signal Acquisition ASIC for High Density 4096 Channel Electrophysiology

Aikaterini Papadopoulou¹, John Hermiz², Carl Grace², and Peter Denes²

 $^{1}\mathrm{Lawrence}$ Berkeley National Laboratory $^{2}\mathrm{Affiliation}$ not available

February 27, 2024

Abstract

The complexity of information processing in the brain requires development of technologies that can provide spatial and temporal resolution by means of dense electrode arrays paired with high-channel-count signal acquisition electronics. In this work, we present an ultra-low noise modular 512-channel neural recording circuit that is scalable to up to 4096 simultaneously recording channels. The neural readout application-specific integrated circuit (ASIC) uses a dense 8.2 mm x 6.8 mm 2D-layout to enable high-channel-count, creating an ultralight 350 mg flexible module. The module can be deployed on headstages for small animals like rodents and songbirds, and can be integrated with a variety of electrode arrays. The chip was fabricated in a 0.18 µm 1.8-V CMOS technology and dissipates a total of 125 mW. Each DC-coupled channel features a gain and bandwidth programmable analog front-end, along with 14b analog-to-digital conversion at speeds up to 30 kS/s. Additionally, each front-end includes programmable electrode plating and electrode impedance measurement capability. We present both standalone and in vivo measurements results, demonstrating readout of spikes and field potentials that are modulated by a sensory input.

A Modular 512-Channel Neural Signal Acquisition ASIC for High-Density, 4096-Channel Electrophysiology

Aikaterini Papadopoulou, Member, IEEE, John Hermiz, Carl R. Grace, Senior Member, IEEE, and Peter Denes, Member, IEEE

Abstract—The complexity of information processing in the brain requires development of technologies that can provide spatial and temporal resolution by means of dense electrode arrays paired with high-channel-count signal acquisition electronics. In this work, we present an ultra-low noise modular 512-channel neural recording circuit that is scalable to up to 4096 simultaneously recording channels. The neural readout application-specific integrated circuit (ASIC) uses a dense 8.2 mm x 6.8 mm 2D-layout to enable high-channel-count, creating an ultralight 350 mg flexible module. The module can be deployed on headstages for small animals like rodents and songbirds, and can be integrated with a variety of electrode arrays.

The chip was fabricated in a 0.18 μ m 1.8-V CMOS technology and dissipates a total of 125 mW. Each DC-coupled channel features a gain and bandwidth programmable analog front-end, along with 14b analog-to-digital conversion at speeds up to 30 kS/s. Additionally, each front-end includes programmable electrode plating and electrode impedance measurement capability. We present both standalone and in vivo measurements results, demonstrating readout of spikes and field potentials that are modulated by a sensory input.

Index Terms—Brain-machine interface, biomedical electronics, in vivo, high-channel-count, neural readout.

I. INTRODUCTION

T HE brain is perhaps the most complex system we know of; multiple brain regions contribute to any given function through complex, anatomically distributed sub-circuits. We know that neurons generate electrical activity by means of action potentials which encode information, and that the timescales of brain activity range from milliseconds to years. However, the exact way that spiking patterns encode information is still a mystery.

As the neuroscience community attempts to translate these signals, the need for large-scale, high-density neural recording increases [1]. A large number of recording sites featuring high anatomical spatial coverage and millisecond temporal resolution is necessary for any new technology developed to tackle this problem. As a result, significant progress has been made in increasing the number of electrodes in silicon

Manuscript received October 15, 2022; revised Month Day, Year.

and polymer probes [2] [3], which in turn increases the requirement for high-channel-count neural readout electronics.

One of the most widely adopted commercial ASICs for neural readout features up to 64 recording channels [4] [5]. Each channel has an AC-coupled front-end and offers ultralow noise recording. The large capacitors required for a 1 Hz high-pass cutoff limit the scalability of the system, however, making it impractical for recording thousands of channels.

In [6] and [7] up to 384 readout channels are demonstrated on a single chip. The system is monolithically fabricated with electrodes and circuits on a silicon substrate and achieves small area and very low-power recording, at a moderate analog-to-digital converter (ADC) resolution. In addition, the probes can be used in multi-module assemblies of thousand of channels. Although monolithic fabrication allows increased channel count, the readout cannot be integrated with other probes or electrode technologies.

A massive 65,536-channel count recording system is demonstrated in [8]. The system consists of microwire electrode arrays bonded to readout electronics, and is the largest recording array to date. The readout ASIC does not include digitization, and power consumption can become a serious bottleneck since even a small temperature increase at the recording site can affect the measured potentials. Furthermore, the device weight and size are too large to be used in awake and free behaving experiments with small animals like rats.

This work achieves readout and digitization of 512 channels onto a single chip. The chip can be used in multi-module assemblies of up to 8 modules, therefore increasing the channel count to 4096. The prototype borrows a 2D layout approach that has previously led to major developments in particle physics and X-ray microscopy, allowing much higher density of electronics than standard 1D layouts. It features a DC-coupled programmable analog front-end and in-pixel 14bit digitization, as well as programmable clock distribution, and data encoding and serialization, making it a complete high-density neural readout solution, compatible with various high density electrode arrays, in standalone or multi-module configuration.

The system design is presented in II, with details on the circuit design for each channel discussed in Sections II-A and II-B. Bench and in-vivo measurement results are presented in Section III.

A. Papadopoulou, J. Hermiz, C. R. Grace, and P. Denes are with Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA.

Research reported in this publication was supported by the National Institute Of Neurological Disorders And Stroke of the National Institutes of Health under Award Number UF1NS107667. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.



Fig. 1. High-level system architecture of the proposed neural recording system (top) and readout ASIC (bottom). Top: One module of the system consists of a flexible ribbon cable which connects four electrode shanks with the substrate. The 512-channel ASIC is bump-bonded onto the substrate, which is bump-bonded onto the stiffened ribbon cable. Bottom: Block diagram of the proposed readout ASIC. The electrode signals are amplified, digitized, and encoded, before being sent to the next processing stage.

II. SYSTEM ARCHITECTURE AND DESIGN

The fabricated chip forms the basis of a modular, highchannel-count recording system. The proposed architecture of the neural data acquisition system is shown in the top part of Figure 1. Each module includes a 512-channel polymer probe consisting of 4 shanks of 128 electrodes, which is connected to the chip through a flexible ribbon cable. The chip is bumpbonded on a 10 mm \times 17 mm substrate, which is then bumpbonded to the ribbon cable. The chip I/O is routed to an ultralight 0.35mm-pitch connector, to be sent to an FPGA for processing. The whole system is designed to not exceed a weight budget of 350 mg for each module, therefore allowing for up to 8 modules to be stacked together inside the headstage. As a result, a total of 4096 electrodes can be simultaneously processed by a single FPGA.

The main focus of this work is the neural signal acquisition ASIC, shown in the bottom part of Figure 1. Each pixel consists of a front-end neural amplifier with programmable gain, a programmable anti-aliasing filter, a buffer, and a 14b $\Sigma\Delta$ analog-to-digital converter, including the digital decimation filter. Analog biasing is provided through programmable digital-to-analog converters (DACs). Digital control, programmable clock generation and distribution, as well as a serial communication protocol are implemented on chip. Low-voltage differential signaling (LVDS) is implemented for all high-speed inputs/outputs (I/O).

A. Analog front-end

A detailed block diagram of the analog front-end (AFE) is shown in Figure 2. In order to achieve both a programmable gain and the ultra-high input impedance required for neural recording, the first stage of the AFE is implemented as a



Fig. 2. Detailed block diagram of the analog front-end, including the electrode plating current scheme.

4-input operational transconductance amplifier, consisting of one main and one auxiliary input pair. The main amplifier has high-input impedance and is DC-coupled to the electrode pad. DC-coupling has the advantage of much lower area than ACcoupling, but it is more sensitive to electrochemical offsets. In order to compensate for such offsets while maintaining power, area, and noise requirements, a simple background offset calibration is implemented as feedback to the main amplifier; if the amplifier output exceeds a programmable threshold, a current charges a capacitor connected to the inverting input so that the amplifier output voltage is zero. This large capacitor is implemented as a MOS capacitor (MOSCAP) to minimize area while maintaining good charge retention. The offset correction scheme is implemented offline and does not otherwise interfere with the signal path. The auxiliary inputs of the amplifier are used to set the gain through an externally programmable resistive ladder. The available values for the amplifier gain

are 5, 10, 20, 50, 100, and 200. The final stage of the AFE is an anti-aliasing filter (AAF) with a programmable cutoff frequency.

In addition to gain, filtering, and offset calibration, the AFE features an electrode plating capability. Electrode plating currents are provided through gated current mirrors. A DAC common to all channels sets the value of the plating current. Each pixel can be individually programmed to source or sink the plating current. The electroplating feature can also be used for electrode impedance measurements. Both the electrode plating process and the impedance measurement process are discussed in more detail in Section III-A3.

B. Sigma-Delta ADC specification and design

To optimally determine the ADC specifications, a spike sorting algorithm was applied to publicly available 16-channel rat neural recordings, which were reconstructed to create a *golden* data set. The golden data was then digitized in software with various levels of non-idealities, including quantization, non-linearity, and noise. The golden data as well as the nonideal data was processed through a spike-sorting algorithm that produces neural clusters [9]. The results were visualized using a *confusion* matrix (Figure 3). Events that appear on the diagonal on the matrix are correctly matched between the two data sets, whereas off-diagonal events represent either missed or misidentified events. Additional analyses that explore spike sorting accuracy as a function of ADC specifications can be found in [10]. This process enables setting system specifications that are informed by spike-clustering algorithms.



Fig. 3. Simplified confusion matrices of two data sets of spike data. On the left, the confusion matrix shows no information loss between the two data sets. On the right, the confusion matrix shows both missing and misidentified events.

It was a determined that 14-bit quantization causes a sufficiently small error in neural clustering produced by spike sorting. Linearity requirements are relaxed, and noise requirements allow the effective number of bits of the ADC to be as low as 12 bits. In addition, bandwidth requirements are limited to less than ≤ 10 kHz due to the nature of intracortical signals.

1) Sigma-Delta modulator: In order to achieve the desired specifications while maintaining low power and area, a 14bit $\Sigma\Delta$ topology was chosen. $\Sigma\Delta$ topologies have been increasingly attractive in the field of neural recording [11] [12], because they are uniquely appropriate for high-resolution, ultra-low power digitization, and offer advantages such as quantization noise shaping and relaxed filtering requirements.

A 2^{nd} -order $\Sigma\Delta$ loop was selected in order to provide good stability at a reasonable oversampling ratio. The oversampling ratio was set to 256 to achieve a good tradeoff between required clock frequency and capacitor area, such that the modulator is small enough to be integrated into the pixel. The modulator was implemented as a discrete-time, switchedcapacitor topology, in a single-ended-to-differential configuration, shown in the top of Figure 4. The input sampling capacitors employed a double-sampling scheme, which relaxes the noise requirements and therefore allows for a smaller input capacitor size. This allowed us to maintain reasonable linearity while maximizing dynamic range. The amplifiers were implemented using a fully-differential folded-cascode topology with capacitive common-mode feedback. The modulator draws a total of 60 μ A when the ADC is operated at 30 kS/s, including the input buffer current.



Fig. 4. Simplified schematic of the modulator (top) and block diagram of the decimator (bottom).

2) Digital decimation filter: In addition to a low-noise analog front end and a $\Sigma\Delta$ modulator, each pixel contains a digital decimation filter to convert the single-bit data stream from the $\Sigma\Delta$ modulator into a slower, multi-bit ADC output with improved resolution. The key objectives in the design of the decimation filter are minimum area (to fit the filter into a small pixel) and low power (to enable the inclusion of many pixels per chip). Typically, the decimation filter for a $\Sigma\Delta$ ADC is implemented using multiple stages and the final stage is usually a high-order finite impulse response (FIR) filter to reject close-in aliasing due to the decimation process. This final stage often drives the area and power of the decimation filter [14]. For this ADC, we separated the decimation into two components: a pre-filter to be implemented in the pixel and a post-filter implemented off-chip (Figure 4). This partitioning allowed a tradeoff between on-chip filter area and data communication bandwidth requirements because as the modulated signal is decimated in each stage bandwidth is traded for resolution. Partitioning the filter between on-chip hardware and off-chip hardware or software allows the decimation filter to be optimized for the specific application; both reducing area and improving performance. Integrating the decimation pre-filter on the chip reduces the data volume that must be transmitted from the prototype by a factor of 9 compared to transmitting the raw modulator data stream.



Fig. 5. Block diagram of Cascaded Integrator Comb filter. The output bits of the $\Sigma\Delta$ modulator are integrated k times, downsampled by a factor of N, and then differentiated an additional k times. This structure is highly computationally efficient as it allows the implementation of a high-order lowpass filter without multipliers, reducing power dissipation and area.

To minimize on-chip area, the decimation filter is implemented as a Cascaded Integrator Comb (CIC) filter. This filter structure is a computationally efficient implementation of a narrow-band FIR lowpass filter that does not require multipliers, which greatly relaxes the area and power dissipation required to implement the CIC filter [13]. A block diagram of the CIC filter is shown in Figure 5.

One consequence of using a CIC filter is that it has a strong sinc(x) response, which requires a droop compensation filter to recover the frequency response near the Nyquist band. Here, however, the spike-sorting routines that post-process the data acquired by the prototype include sharp lowpass filtering as part of their operations, so good high-frequency fidelity is not required in the decimation filter.

To balance performance and complexity, typically a CIC filter used in a $\Sigma\Delta$ ADC is implemented using an order one larger than the order of the modulator [14]. Since the ADC is using a second-order modulator, the implemented CIC is a third-order FIR filter. The CIC filter implemented as part of the $\Sigma\Delta$ ADC uses a two's complement data representation to simplify data flow. To ensure the final filter output does not overflow, the word width must be greater than:

$$W = Nlog_2(D) + 1$$

where W is the required word width to avoid overflow, N is the order of the CIC filter, and D is the decimation factor. In this case, the decimation factor for the pre-filter is 128. This leads to a required word width of 22 bits [13] in the prototype CIC filter. It is possible to shrink the word width as data progresses down the pipeline but this was not done here because we determined that the reduction in area possible was small relative to the additional effort required.

The key goal of the decimation filter is to minimize area to enable integration of a complete ADC inside each pixel. Because the full adder is the key circuit in the CIC filter, a number of full adder topologies were investigated to optimize area and power dissipation. We compared a conventional 28-transistor full adder and a more aggressive 18-transistor adder based on pass-transistor logic [15]. Each full adder was implemented using transistors with various thresholds. We also examined adders with even fewer devices but found the performance variation across corners was problematic. The results of this simulation study are summarized in Figure 6. We determined that the conventional 28-transistor full adder cell implemented using standard-threshold devices had the best balance of low power dissipation, small die area, and high reliability across corners given the expected workload.

To minimize the area of the digital filter, the layout was custom designed. The dimensions of the custom full adder, shown in Figure 7, are 12.1 μ m by 4.8 μ m (59.1 μ m²) and the full adder area was reduced by over 30% compared to a commercial standard cell full adder. The area of the other cells in the filter was reduced by a similar factor. In addition, because of the use of minimum-size devices throughout the layout (to minimize area), the power dissipation of the custom full adder was reduced by a factor of approximately 3 compared to a commercial standard cell full adder.

In addition to full adders, digital latches are required to pipeline the data as it flows through the filter. To minimize area, the latches are implemented using 2-phase logic. This is possible to do in a simple way by reusing the 2-phase clocks required for the switched-capacitor circuits in the $\Sigma\Delta$ modulator. The integrators operate at the same speed as the modulator, while the differentiators (after the downsampler) operate at 1/128th the rate of the modulator. This implicit clock division is implemented by masking every 128th phase of the higher-frequency clock.

The entire filter is implemented using a full-custom layout style, without including any standard cells in the design. The filter has an area of approximately 200 μ m by 100 μ m or 0.02 mm².

The physical implementation of the CIC filter in each pixel required 5850 transistors. The breakdown of the per-pixel CIC filter device usage is shown in Table I. The coder block converts the single-bit output of the $\Sigma\Delta$ modulator into a 22-bit two's complement representation.

C. Digital services

Two 14-bit words, representing the sample from two channels, are encoded in a 32-bit DC-balanced output word. The chip serializes the data output on two LVDS channels, each



Fig. 6. Comparison of power dissipation of several full adder topologies simulated at 4 MHz and implemented using various transistor flavors. The conventional 28-transistor full adder cell implemented using standardthreshold devices had the best balance between power dissipation, die area, and reliability across power, voltage, and temperature variation. "svt" refers to standard-threshold devices, "mvt" references to medium-threshold devices, and "native" refers to zero-threshold devices.



Fig. 7. Custom full-adder layout. This layout of a standard 28-transistor full adder consumes 30% less die area compared to the full adder included in a commercial standard cell library.

TABLE I TRANSISTOR USAGE IN IMPLEMENTED CIC FILTER

Circuit	Instances	Transistors	Total Transistors
Coder	1	42	42
Integrator	3	880	2640
Downsampler	1	132	132
Differentiator	3	1012	3036
Complete Filter	1	5850	5850

carrying the data from 16 rows \times 16 columns of pixels, at a rate of $32 \times 128 \times f_s$, where f_s is the sampling frequency. Chip programming is performed via a 2-wire interface, with token passing to control multiple chips in multi-module assemblies.

III. SILICON MEASUREMENT RESULTS

The design was fabricated in a 0.18 μ m 1.8-V CMOS process. Figure 8 shows a microphotograph of the prototype, along with the floorplan and power distribution scheme. Each chip contains 512 channels, organized as 32 rows by 16 columns. Each channel includes the electrode pad, and one column of additional power pads is placed between every two electrode pad columns. These additional power pads achieve ultra-low resistance power routing, easy decoupling, and reduced on-chip power regulation requirements, at the cost of chip area.

Each channel occupies 0.099mm², including the area of the power pads. The channel layout is shown in Figure 9. The total chip area is 55.8 mm². Figure 10 shows the chip bump-bonded on the substrate board.

Both standalone bench testing and *in vivo* testing were performed, and the results are presented in Sections III-A and III-B, respectively.

A. Standalone Measurements

1) Chip performance: For standalone testing, the peripheral power pads and 32 electrodes are wirebonded onto a custom testboard. Initial testing was done using the programmability and testability features of the chip. Figure 11 shows measurements of the read-back analog biasing current compared to simulated values. Measurement reveals currents slightly higher than expected, but still within the desired range. Figure 12 shows excellent bandwidth programmability of the AAF filter.



Fig. 8. Die photo of the chip. The die size is 8.2 mm x 6.8 mm. Power pad columns (blue) are placed between every two electrode pad columns. Digital I/O and test pads are the right-most column.



Fig. 9. Layout of the complete pixel. The dimensions of the of pixel are approximately 205 μ m \times 477 μ m, with more than half of the area dedicated to the analog-front-end. The modulator and decimation filter occupy approximately 200 μ m \times 100 μ m or 0.02 mm² each.



Fig. 10. Chip prototype bump-bonded on the substrate board.



Fig. 11. Measured bias current vs. biasing DAC value.



Fig. 12. Measured system bandwidth vs. AAF DAC value.



Fig. 13. Simulated amplitude for a 0.2 mV 10 kHz input sinusoid, compared to simulated AAF response. The AFE gain setting is 50. ADU size is 61μ V.



Fig. 14. Simulated and measured plating current programmability.

Figure 13 shows the measured signal amplitude compared against the simulated AAF response, for various AAF DAC settings. For a AAF DAC value of 20, the measured gain is 53.4 V/V, very close to the expected value. The plating current DAC was also measured to be comparable with simulation results (Figure 14). Figure 15 shows signal amplitude over frequency for a 0.2 mV input sinusoid. The response is compared to the theoretical response of the AAF and digital CIC filter with gain scaling, and follows the expected roll-off. It also reveals a 15% lower gain comparing to the expected value, in this particular gain setting.



Fig. 15. Measured amplitude for a 0.2 mV input sinusoid. The AFE gain setting is 100. ADU size is 61 $\mu V.$

Figure 16 shows the input-referred noise (IRN) histogram. Measured noise is 5.4 μ V in the 0.3-10 kHz action potential (AP) band, and 3.1 μ V in the 0.5 Hz-1 kHz local-field potential (LFP) band. For our application, the AP band of interest is the 0.3-6 kHZ band, which yields an IRN of 4.8 μ V.



Fig. 16. Measured noise spectrum in AP band of interest (0.3-6 kHZ).

At a 1.8 V supply voltage, the total power consumption is 244 μ W/channel when sampling at the maximum sampling frequency of 30 kS/s. This includes all on-chip components; the AFE, buffer and complete ADC, as wells as all programmability features, digital communication protocol implementation, and LVDS I/O circuitry. The power breakdown of the chip is shown in Figure 17.

2) Data post-processing: Larger than anticipated leakage currents cause the offset compensation capacitor to charge up,



Fig. 17. Measured power breakdown (top) and measured power at various sampling frequencies (bottom).

imposing a sawtooth-shaped artifact on the signal. The primary source of this leakage is gate leakage of the capacitor. To remove this sawtooth background, we take advantage of the consistent shape of the leakage current-induced patterns. This involves: 1) detecting the up and down-phases of the sawtooth signal 2) fitting a line to the up-phases and 3) subtracting away the linear fit while zeroing out changes in the negligibly short down-phases. Let the acquired signal be denoted x, nbe the index of the discretely sampled time series, and dbe a sampling-rate dependent delay parameter. In step 1, the down-phase of the sawtooth is detected when both a delayed amplitude threshold (1) and a first order difference threshold (2) are met.

$$x[n-d] > x_{min} \tag{1}$$

$$\Delta x[n] < \Delta x_{max} \tag{2}$$

Since, the leakage current is always in one direction, this phase detection criteria assumes that sawtooth signal always rises positively; however, this can easily be generalized to sawtooth signals of the opposite sign. Samples not detected as down-phases are classified as up-phases. A least squares linear model is fit to the up-phases of x, ignoring the downphases by concatenating only the up-phases of the signal xtogether. Let c be the slope of the linear fit. The final step involves subtracting away the linear sawtooth component of the signal which can be done with the first order difference signal and accumulating $\Delta x[n]$:

$$\Delta x[n] \leftarrow \begin{cases} \Delta x[n] - c & \text{if } up - phase \\ 0 & \text{if } down - phase \end{cases}$$
(3)



Fig. 18. The top left panel shows raw data collected with a 1 Hz sine wave injected into the test chip. The top right shows the signal and fitted sinusoid after sawtooth removal. The bottom left panel shows raw data collected with a 0.1 Hz signal injected into the test chip. The bottom right panel shows the signal and fitted sinusoid after sawtooth removal. In both the 1 Hz and 0.1 Hz case, there is excellent agreement ($R^2 = 0.96$ and $R^2 = 0.94$ for the 1 and 0.1 Hz respectively) between the injected sinusoids and the signals after post-hoc sawtooth removal.

As shown in Figure 18, post-processing eliminates the sawtooth background in signals of frequencies both above and below the sawtooth frequency.

3) Electroplating and electrode impedance measurements: As outlined in Section II-A, the AFE is equipped with a digitally-controlled electroplating DAC. There are two main steps in electroplating process: 1) electrochemically cleaning the metal surface of the electrodes, and 2) coating with poly(3,4-ethylenedioxythiophene) polystyrene sulfonate (PE-DOT:PSS). Both steps are performed using a 2-electrode configuration where the reference/ground electrode is a pure silver wire, and the microelectrodes are 20 µm-diameter platinum contacts. Cleaning involves submerging the electrodes in a sulfuric acid bath and passing a current that is swept between -30 to 150 nA at an average rate of 2 nA/s for 10 minutes. After cleaning, the electrode array is rinsed with 70% alcohol and deionized water. Subsequently, the electrode array is submerged in a solution of PEDOT:PSS. Finally, a constant current of 10 nA is supplied for 45 seconds through each electrode, simultaneously plating all electrodes. Figure 19 shows the bare metal electrodes and the plated PEDOT:PSS electrodes. As a result of plating, the impedance decreased from 5.4 (90% confidence interval [CI] = 2.4 to 8.5) $M\Omega$ to 0.2 (90% CI = 0.04 to 0.88) $M\Omega$, as shown in Figure 20.

Finally, the chip has the ability to perform electrode impedance measurements at arbitrary frequencies. Each channel can individually be programmed to enable a positive or negative plating current. By continuously re-programming the DAC that sets the plating current for the chip, current patterns can be implemented including sinusoids. Electrode impedance can be measured by passing a sinusoidal current through an electrode and measuring the voltage across the electrode. In a

Fig. 19. Bare metal electrodes and the PEDOT:PSS plated electrodes.

sample of 6 channels, the resistance measured across a 47 k Ω resistor at 250, 500, and 1000 Hz was $\leq \pm 10\%$, as shown in Figure 21.

Fig. 20. Impedance of a baseline and a PEDOT:PSS plated set of electrodes. The impedance decreases from 5.4 (90% CI = 2.4 to 8.5) $M\Omega$ to 0.2 (90% CI = 0.04 to 0.88) $M\Omega$ after electrode plating. The impedance was measured at 1kHz.

Fig. 21. Resistance measurements using a 47 k Ω resistor. The measured resistance is within 10% of the reference.

B. In vivo electrophysiology

We performed an acute craniotomy experiment on a Sprague Dawley rat, which was anesthetized by intraperitoneal injections of Ketamine and Xylazine. All rat procedures were performed in accordance with established animal care protocols approved by the LBNL Institutional Animal Care and Use Committees (IACUC). A commercial silicon laminar probe was inserted into the primary auditory cortex, and a platinum

Fig. 22. Photo of the in vivo test setup. A custom testboard was made in order to interface the chip prototype with the commercial silicon laminar probe. Auditory stimulus was provided through a speaker, and results were post-processed using an FPGA.

reference wire was inserted into a contralateral frontal region. A testboard was fabricated to connect the silicon laminar probe with the chip. Figure 22 shows a photo on the test setup, including the custom testboard, prototype, and silicon laminar probe.

Auditory stimulus included a white noise burst lasting 100 ms played every 1 second for 60 repetitions [17]. The digital output was sent to an FPGA and main controller unit for digital processing and finally sent to a computer for visualization and data saving¹.

The recordings were post-processed with a similar background subtraction technique described in Section III-A2, and then passed through spectral and spike sorting analysis pipelines. The spectral analysis involves computing the constant-Q wavelet transform for each trial for center frequencies ranging from 8.3 Hz to 1200 Hz [18] [19]. The magnitude of the transform is taken and then normalized by z-scoring relative to baseline. The baseline period lasts 200 samples or ~6.67 ms, and starts 100 ms prior to the upcoming stimulus presentation.

For a separate spike analysis, high pass filtering at 300 Hz, whitening, and automated spike clustering were performed using the publicly available spike-sorting algorithms spikeinterface [20] and MountainSort [9]. Finally, the produced units or clusters were manually curated to identify putative single units.

The prototype was able to readout *in vivo* electrophysiological signals including action potentials measured from laminar polytrodes inserted into cortex. The filtered measured signals from 4 channels are shown in Figure 23(a). Evoked potentials were strongly driven by auditory stimuli across the neural frequency spectrum as expected (Figure 23(b)). Spike sorting revealed isolated putative single units (Figure 23(c)). These results demonstrate that the proposed design can readout spikes and field potentials that are modulated by a sensory input.

¹FPGA developement and post-processing was performed by our collaborators at SpikeGadgets.

Fig. 23. Chip readout of *in vivo* electrophysiological signals. a) High pass filtered ($f_c = 300$ Hz) signals from 4 channels during the presentation of an auditory stimulus indicated by the red vertical line at time 0. b) Median spectrogram across 60 trials of auditory presentation showing a broadband increase in amplitude relative to a baseline window. The auditory presentation is indicated by the red vertical line at time 0. The auditory presentation is indicated by the red vertical line at time 0. The amplitude is normalized by z-scoring (Z) relative to baseline. c) Four putative single unit waveforms generated by using an automated spike-sorting algorithm. The average waveform is plotted in black and the 95% standard error is plotted as a gray shaded region about the average.

The chip performance summary is shown in Table II, compared to the two most widely adopted state-of-art commercial neural signal acquisition systems.

TABLE II Performance summary

	[5]	[6]	[12]	[21]	This work
Channels	64	384	128	16	512
Tot. area [mm ²]	28.7	45.2	0.005	5.8	55.8
Area/ch. [mm ²]	0.448 ^a	0.12	0.0045	0.16	0.099
ADC bits	16	10	14	8	14
ADC f_s [kHz]	30	30	30	31.25	30
IRN (LFP) $[\mu V]$	2.4 ^b	10.32	11.9	-	3.1
IRN (AP) [µV]	2.4 ^b	6.36	7.71	5.4 ^c	5.4
Power/ch. [µW]	351	49.06	8.34	0.96	244
Supply [V]	3.0	1.2/1.8	0.8	0.5	1.8
In vivo results	Yes	Yes	No	Yes	Yes
Technology [um]	0.35	0.13 SOI	0.022	0.18	0.18

^aIncludes I/O and digital interface. ^bUnspecified frequency range.

^c 1kHz-12kHz frequency range.

IV. CONCLUSION

We report a massive 512-channel neural signal acquisition ASIC designed to target high-density electrophysiology. Modularity and scalability enable addressing mulliple brain regions and were key components of the design, as well as integration with commercial high-density probe systems. We briefly discuss our complete system headstage design which targets 4096-channel recording, and focus on the chip design, testing, and data processing. The ASIC features programmable gain, filtering, and 14b $\Sigma\Delta$ digitization, including digital decimation filtering in each channel. It occupies a 55.8 mm^2 area, measures an IRN of 4.8 µV in the AP band of interest (0.3-6 kHz), and dissipates 244 μ W/channel from a 1.8 V supply. The chip also provides electrode plating and electrode impedance measurement capability. Finally, we present in vivo measurements of action potentials using silicon laminar probes on anesthetized rats. This work demonstrates an ultra-low

noise flexible signal acquisition modular system with potential for ultra high-density neural recording.

ACKNOWLEDGMENT

The authors are members of the Kavli Institute for Fundamental Neuroscience. The authors would like to thank Loren Frank and his group at the University of California, San Francisco Center for Integrative Neuroscience and Kristofer Bouchard at LBNL for their scientific input, Magnus Karlsson and Mattias Karlsson at SpikeGadgets for the digital postprocessing backend development, and Alison Yorita, Travis Massey, and Razi Haque at Lawrence Livermore National Laboratory for probe development.

REFERENCES

- E. Litvina et al., "BRAIN Initiative: Cutting-Edge Tools and Resources for the Community," J. Neurosci., vol. 39, no. 42, pp. 8275–8284, Oct. 2019, doi: 10.1523/JNEUROSCI.1169-19.2019.
- [2] G. Buzsáki, E. Stark, A. Berenyi, D. Khodagholy, D. R. Kipke, E. Yoon, and K. D. Wise "Tools for probing local circuits: high-density silicon probes combined with optogenetics," Neuron, vol. 86, no. 1, pp. 92-105, Apr. 2015, doi: 10.1016/j.neuron.2015.01.028
- [3] J. E. Chung et al., "High-Density, Long-Lasting, and Multi-region Electrophysiological Recordings Using Polymer Electrode Arrays," Neuron, vol. 101, no. 1, pp. 21-31.e5, Jan. 2019, doi: 10.1016/j.neuron.2018.11.002
- [4] J. Du, T. J. Blanche, R. R. Harrison, H. A. Lester, and S. C. Masmanidis, "Multiplexed, High Density Electrophysiology with Nanofabricated Neural Probes," PLoS ONE, vol. 6, no. 10, p. e26204, Oct. 2011, doi: 10.1371/journal.pone.0026204.
- [5] Intan Technologies. "RHD Electrophysiology Amplifier Chips," intantech.com: https://intantech.com/products_RHD2000.html, Oct 2021.
- [6] C. M. Lopez et al., "A 966-electrode neural probe with 384 configurable channels in 0.13µm SOI CMOS," in 2016 IEEE International Solid-State Circuits Conference (ISSCC), Jan. 2016, pp. 392–393. doi: 10.1109/ISSCC.2016.7418072.
- [7] N. A. Steinmetz et al., "Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings," Science, vol. 372, no. 6539, p. eabf4588, Apr. 2021, doi: 10.1126/science.abf4588.
- [8] K. Sahasrabuddhe et al., "The Argo: A high channel count recording system for neural recording in vivo," J. Neural Eng., Dec. 2020, doi: 10.1088/1741-2552/abd0ce.
- [9] J. E. Chung et al., "A Fully Automated Approach to Spike Sorting," Neuron, vol. 101, no. 1, pp. 1381-1394, 2017, 10.1016/j.neuron.2017.08.030.

- [10] J. Hermiz et al. "The impact of reducing signal acquisition specifications on neuronal spike sorting." 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). IEEE, 2021.
- [11] H. Kassiri et al., "Rail-to-Rail-Input Dual-Radio 64-Channel Closed-Loop Neurostimulator," IEEE Journal of Solid-State Circuits, vol. 52, no. 11, pp. 2793–2810, Nov. 2017, doi: 10.1109/JSSC.2017.2749426.
- [12] X. Yang et al., "A 128-Channel AC-Coupled 1st-order Δ - $\Sigma\Delta$ IC for Neural Signal Acquisition," in 2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits), Jun. 2022, pp. 60–61. doi: 10.1109/VLSITechnologyandCir46769.2022.9830236.
- [13] E. Hogenauer, "An economical class of digital filters for decimation and interpolation," IEEE Transactions on Acoustics, Speech, and Signal Processing, vol. 29, no. 2, pp. 155-162, April 1981, doi: 10.1109/TASSP.1981.1163535.
- [14] S. Pavan, R. Schreier, and G. Temes, Understanding Delta-Sigma Data Converters, 2nd ed., New York: Wiley-IEEE Press, 2017, doi: 10.1002/9781119258308
- [15] N. Weste, D. Harris, CMOS VLSI Design: A Circuits and Systems Perspective, 4th ed. Boston: Addison-Wesley, 2011
- [16] V. Lanzio, V. Gutierrez, J. Hermiz, K. Bouchard, and S. Cabrini, "Neural optoelectrodes merging semiconductor scalability with polymeric-like bendability for low damage acute in vivo neuron readout and stimulation," Journal of Vacuum Science & Technology B, vol. 39, no. 6, p. 063001, Dec. 2021, doi: 10.1116/6.0001269.
- [17] V. L. Baratham, M. E. Dougherty, J. Hermiz, P. Ledochowitsch, M. M. Maharbiz, and K. E. Bouchard, "Columnar Localization and Laminar Origin of Cortical Surface Electrical Potentials," J. Neurosci., vol. 42, no. 18, pp. 3733–3748, May 2022, doi: 10.1523/JNEUROSCI.1787-21.2022.
- [18] C. Schörkhuber and A. Klapuri "Constant-Q transform toolbox for music processing," 7th sound and music computing conference, Barcelona, Spain, July 2010.
- [19] J. Livezey. "process_nwb: Functions for preprocessing (ECoG) data stored in the NWB format," Github: https://github.com/BouchardLab/process_nwb, Oct. 2021
- [20] A. P. Buccino et al., "SpikeInterface, a unified framework for spike sorting," eLife, vol. 9, p. e61834, Nov. 2020, doi: 10.7554/eLife.61834.
- [21] S.-J. Kim et al., "A Sub-μW/Ch Analog Front-End for Δ-Neural Recording With Spike-Driven Data Compression," IEEE Transactions on Biomedical Circuits and Systems, vol. 13, no. 1, pp. 1–14, Feb. 2019, doi: 10.1109/TBCAS.2018.2880257.

John Hermiz earned a bachelors in Electrical Engineering and Computer Science from the University of Michigan, Ann Arbor. He earned his M.S. and PhD in Electrical and Computer Engineering from the University of California, San Diego. He is a Horatio Alger National Scholar and Dennis Washington Graduate Fellow.

Currently, he is a postdoctoral researcher at Lawrence Berkeley National Lab working in the Neural Science and Data Science Lab. His research focuses on developing advanced neurotechnologies,

especially high-channel-count electrophysiological systems. Generally, he is interested in developing and utilizing advanced tools to answer basic neuroscience questions as well as translating neurotechnologies and discoveries to clinical applications.

Carl R. Grace (S'98-M'04-SM'09) received the B.S., M.S., and Ph.D. degrees in Electrical Engineering from the University of California, Davis, in 1997, 2001, and 2004, respectively.

From 2004 to 2006, he was with ClariPhy Communications, Irvine, CA, where he developed analog integrated circuits for optical communications. From 2006 to 2010, he was with Analog Devices, Raleigh, N.C., where he designed integrated circuits for wireless communications. Since 2010, Carl has been with Lawrence Berkeley National Laboratory,

Berkeley, CA, where he is currently a Staff Scientist. His research is focused on the development of integrated circuits for the instrumentation of subatomic particle detectors in extreme environments.

Peter Denes received a B.S. in Physics in 1980 and a Ph.D. in Physics in 1984 from the University of New Mexico. From 1985 to 2000, Dr. Denes was a senior research physicist at Princeton University, NJ. Dr. Denes has been with Lawrence Berkeley National Laboratory in Berkeley, California, since 2000. His work focuses on high-speed electron and soft Xray imaging detectors for in-situ microscopies. The high-speed electron microscopy detectors he has developed have been pivotal in enabling ultra-high resolution structural biology. His current research

efforts aim to extend those techniques to massively parallel neural recording. In 2009, Dr. Denes was awarded the Secretary of the Department of

Energy's Excellence Award for his work on the Transmission Electron Aberration-Corrected Microscope project. He is the recipient of the 2015 Berkeley Laboratory Lifetime Achievement Award for pioneering development of direct detectors for electron and soft X-ray microscopy, and of the 2017 Joseph F. Keithley Award For Advances in Measurement Science. Dr. Denes is a member of the American Physical Society, American Association for the Advancement of Science and the Institute of Electrical and Electronics Engineers.

Aikaterini Papadopoulou (S'09-M'17) received the diploma in Electrical Engineering and Computer Science from the University of Patras, Greece, in 2009, and the M.S. and PhD degrees in Electrical Engineering from the University of California, Berkeley, in 2011 and 2017 respectively.

She has worked on variability characterization of high-speed wireline receivers at Altera Corporation, San Jose, CA, and on high-speed and low-power wireline links for backplanes and cable interconnects at Kandou Bus, Lausanne, Switzerland. In 2017 she

joined the Detectors, Integrated Circuits and Electronic Systems group at Lawrence Berkeley National Lab, where she focuses on sensor readout circuits and analog-to-digital conversion for imaging, neuroscience, and high-energy physics applications.