30.1 Neurons to Silicon: Implantable Prosthesis Processor

Stephen O'Driscoll, Teresa Meng, Krishna Shenoy, Caleb Kemere

Stanford University, Stanford, CA

Each year hundreds of thousands of people suffer from neurological injuries and disorders, resulting in the permanent loss of motor function or even the ability to communicate. A few research groups have now demonstrated that monkeys and humans can learn to move computer cursors and robotic arms to various target locations simply by activating neural populations in the brain without actually performing the arm movements, e.g. [1]. More specifically, neurons in the parietal, pre-motor and motor cortex modulate the rate of emission of electrical impulses (action potentials) according to the particular arm movement that is desired, even in paralyzed individuals, thereby affording the opportunity to translate the emission rate into desired movement commands. As promising as these early demonstrations have been, achieving true clinical viability will require much faster and more accurate movement commands which, in turn, requires extracting as much information as possible from every neuron available. To this end, we propose an integrated silicon implant technology that combines research on cortical electrophysiology, algorithms and circuit design to achieve high levels of prosthetic performance while minimizing power consumption.

It is challenging to record, process and translate neural activity from the brain into control signals for guiding prostheses, as shown in Fig. 30.1.1, even in a lab setting without severe power constraints. Numerous low-power approaches have been pursued. One such approach is to use an analog comparator to simplify the signal processing task [2]. Another approach limited both the number of channels being recorded and the ADC resolution to save power [3]. Although these approaches offer an attractive compromise between power and information extraction, we sought an even lower-power architecture without any compromise on prosthesis performance.

Our proposed implantable prosthesis processor (IPP) consists of four major building blocks: an amplification stage and a variable-resolution ADC array, a digital spike sorter [4], a maximum-likelihood neural decoder [5], and a wireless data and power transceiver. The overall compression factor attained by employing the IPP is on the order of 10^6, translating raw neural data at a rate of 80Mb/s to less than 20b/s [1] indicating the intended movement. The total power budget of the IPP is limited to 1mW, which can be delivered through an external inductive coupling device.

Neural prostheses employ a permanently-implanted electrode array, typically containing 100 channels in an area of $4\times4mm^2$. An ADC array was developed that contains 100 ADC cells, each consuming 1μW on average at 100kSamples/s. The required resolution per channel varies with the neural activity displayed by that channel.

A successive approximation architecture was chosen for each ADC cell, as its resolution can be varied by selecting the branches in the capacitor array. As illustrated in Fig. 30.1.2, input sel [7:3] selects the ADC resolution. The circuit has been designed and laid out using 0.13μm CMOS. A 20fF, $4\times4\mu m^2$, MiM cap was chosen as the unit capacitance C, giving sufficient margin to process variations. Switches are necessary on both sides of the array capacitors to reduce the power consumed by parasitic capacitances at both the top and bottom plates. The comparator is a weak inversion latch.

Fig. 30.1.3 shows the ADC cell power versus ADC resolution, compared against the ADC figure-of-merit energy efficiency of 1pJ/stepsize. These results are based on layout extracted simulations with V_T = 3mV on all transistor pairs. INL and DNL remain within ±0.5LSB over the full scale. Fig. 30.1.3 confirms that a variable resolution (from 3 to 8 bits), 100kSamples/s ADC cell can be realized with an energy efficiency at 0.18pJ/stepsize. Even so, the power consumption of 100 such ADC cells operating at 8 bits per sample would be more than 400μW for the entire ADC array. An ADC array with variable resolutions is necessary to reduce its power further without degrading the neural decoding performance.

The resolution of each ADC cell is determined by taking advantage of the ensuing spike sorter as shown in Fig. 30.1.4. The spike sorter assigns each identified group of electrical impulses, called a spike, to a specific neuron. Reduced ADC resolution would generally increase the misclassification rate of the spike sorter. Let n_i be the resolution for the i^{th} ADC cell. The required ADC resolutions are estimated as shown in Fig. 30.1.5. Misclassification rates, $R(n_i)$ = #misclassifications(n_i) / #spikes, are calculated for each possible resolution. The lowest n_i for which $R(n_i) < R_{max}$ is chosen to be the ADC's resolution until its next training cycle. Resolution estimation uses 5 extra spike sorts for each electrode in a time window equal to the training period of the spike sorter, which is 120 seconds every 12 hours. For a spike sorter power consumption of 1.4μW per channel [4], this gives a power overhead for resolution estimation of 20nW per channel, or 2% of the ADC power with optimally assigned resolutions.

Fig. 30.1.6 shows an example resolution assignment. The spike sorter itself has a spike misclassification rate of approximately 5% based on measured neural data, so choosing R_{max} = 1% does not significantly compromise performance. The total power consumption is simulated to be 116μW, which represents a power saving on the order of 3.6 by using a variable-resolution ADC array.

This low power level makes it possible for it to be provided wirelessly through inductive coupling. However, to maximize this coupling, an alignment of the source magnetic field with the sensing coil that is implemented on the IPP chip is required. This alignment can best be accomplished through use of multiple source coils which are dynamically excited to optimize the magnetic field direction for maximum power transfer.

The IPP power budget of 10μW per channel, or 1mW total, is allocated as follows: less than 2μW for a multi-stage amplifier which provides 60dB gain and anti-aliasing filtering, 1μW for the variable-resolution ADC, and the rest for digital signal processing and wireless communication. The 100-channel ADC array occupies $2.6\times1.8mm^2$ in 0.13μm CMOS, a layout plot of which is shown in Fig. 30.1.7.

Acknowledgements:
This work was funded in part by the MARCO Focus Center for Circuit & System Solutions under contract 2003-CT-888, Stanford Graduate Fellowship, Burroughs Wellcome Fund, NSF, ONR and Whitaker Foundation.

References:
[1] Santhanam et al., "A High Performance Neurally-Controlled Cursor Positioning System," *Proc. IEEE Inter. Conf. on Neural Engineering*, pp. 494-500, Mar., 2005.
[2] Watkins et al., "Validation of Adaptive Threshold Spike Detector for Neural Recording," *Proc IEEE EMBS*, pp. 4079-4082, Sept., 2004.
[3] Olsson and Wise, "A Three-Dimensional Neural Recording Microsystem with Implantable Data Compression Circuitry," *ISSCC Dig. Tech. Papers,* pp. 558-559, Feb., 2005.
[4] Zumsteg et al., "Power Feasibility of Implantable Digital Spike Sorting Circuits for Neural Prosthetic Systems," *IEEE Trans. NSRE*, pp. 272-279, Sept., 2005.
[5] Yu et al., "Improving Neural Prosthetic System Performance by Combining Plan and Peri-Movement Activity," *Proc IEEE EMBS*, pp. 4516-4519, Sept., 2004.

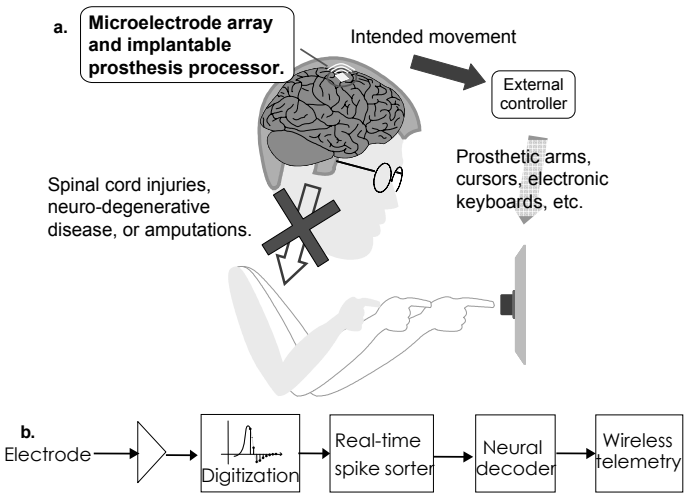

Figure 30.1.1: a. Implanted neural prosthetic system concept. b. Block diagram of implantable prosthesis processor.

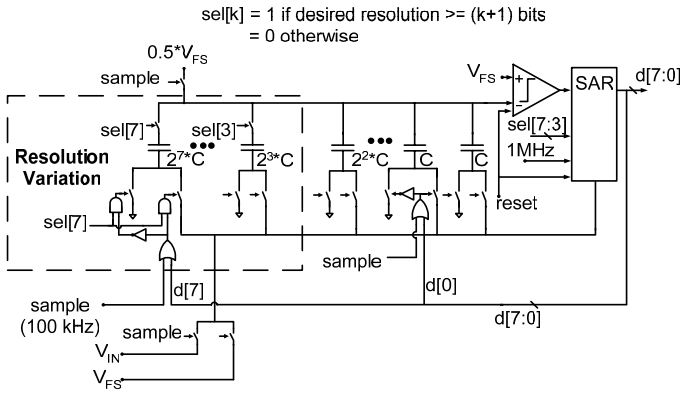

Figure 30.1.2: Successive approximation ADC cell with variable resolution.

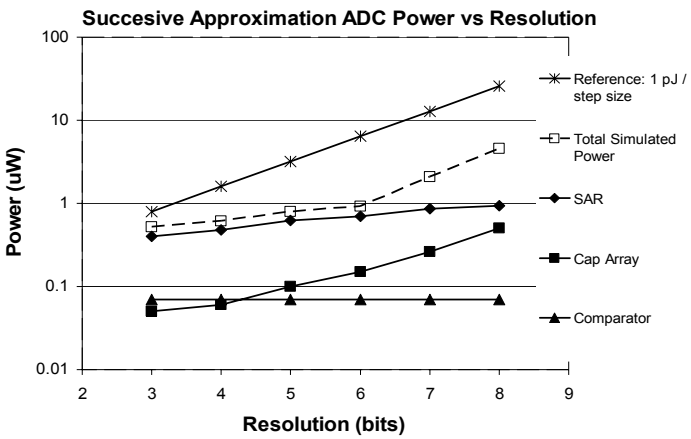

Figure 30.1.3: Power consumption of the ADC cell at different resolutions.

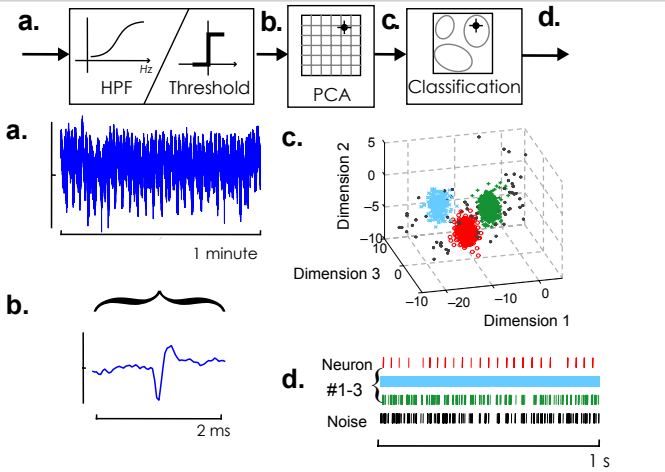

Figure 30.1.4: Real-time digital spike sorter. a. Digitized neural signals are first high-pass filtered and spike samples are selected. b. A spike window. c. Each spike is projected into a 3-dimensional discrimination space based on principle components analysis, PCA. d. Each identified spike is assigned to a specific neuron using a classification algorithm.

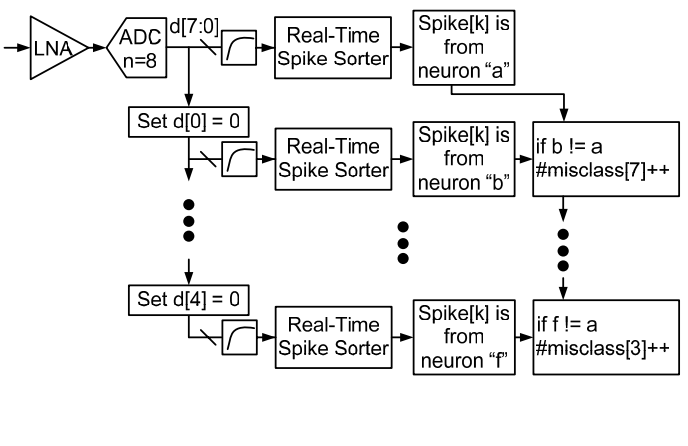

Figure 30.1.5: Method to determine the required ADC resolution for each channel.

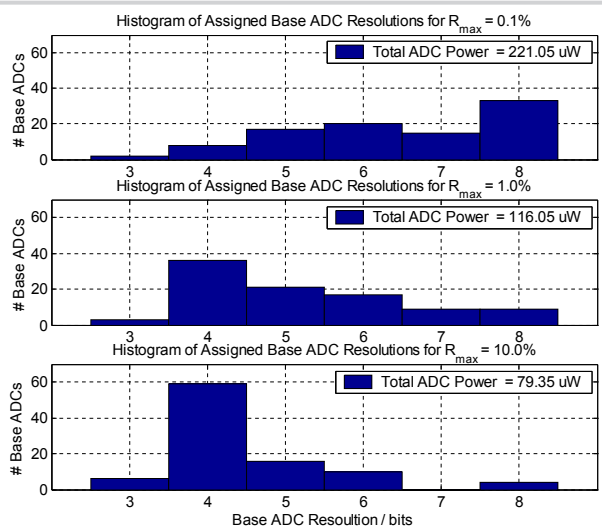

Figure 30.1.6: Histograms of ADC resolutions using sample neural data for R_{max} = 0.1%, 1%, and 10%.

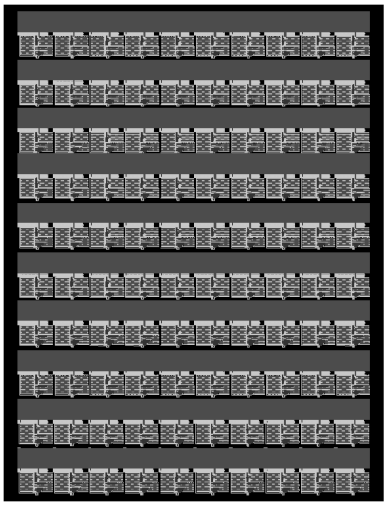

Figure 30.1.7: Layout plot of ADC array. Each row of the layout contains 10 ADC cells separated by input, signal and control lines. Dummy grounded lines are added between each signal line to avoid crosstalk.

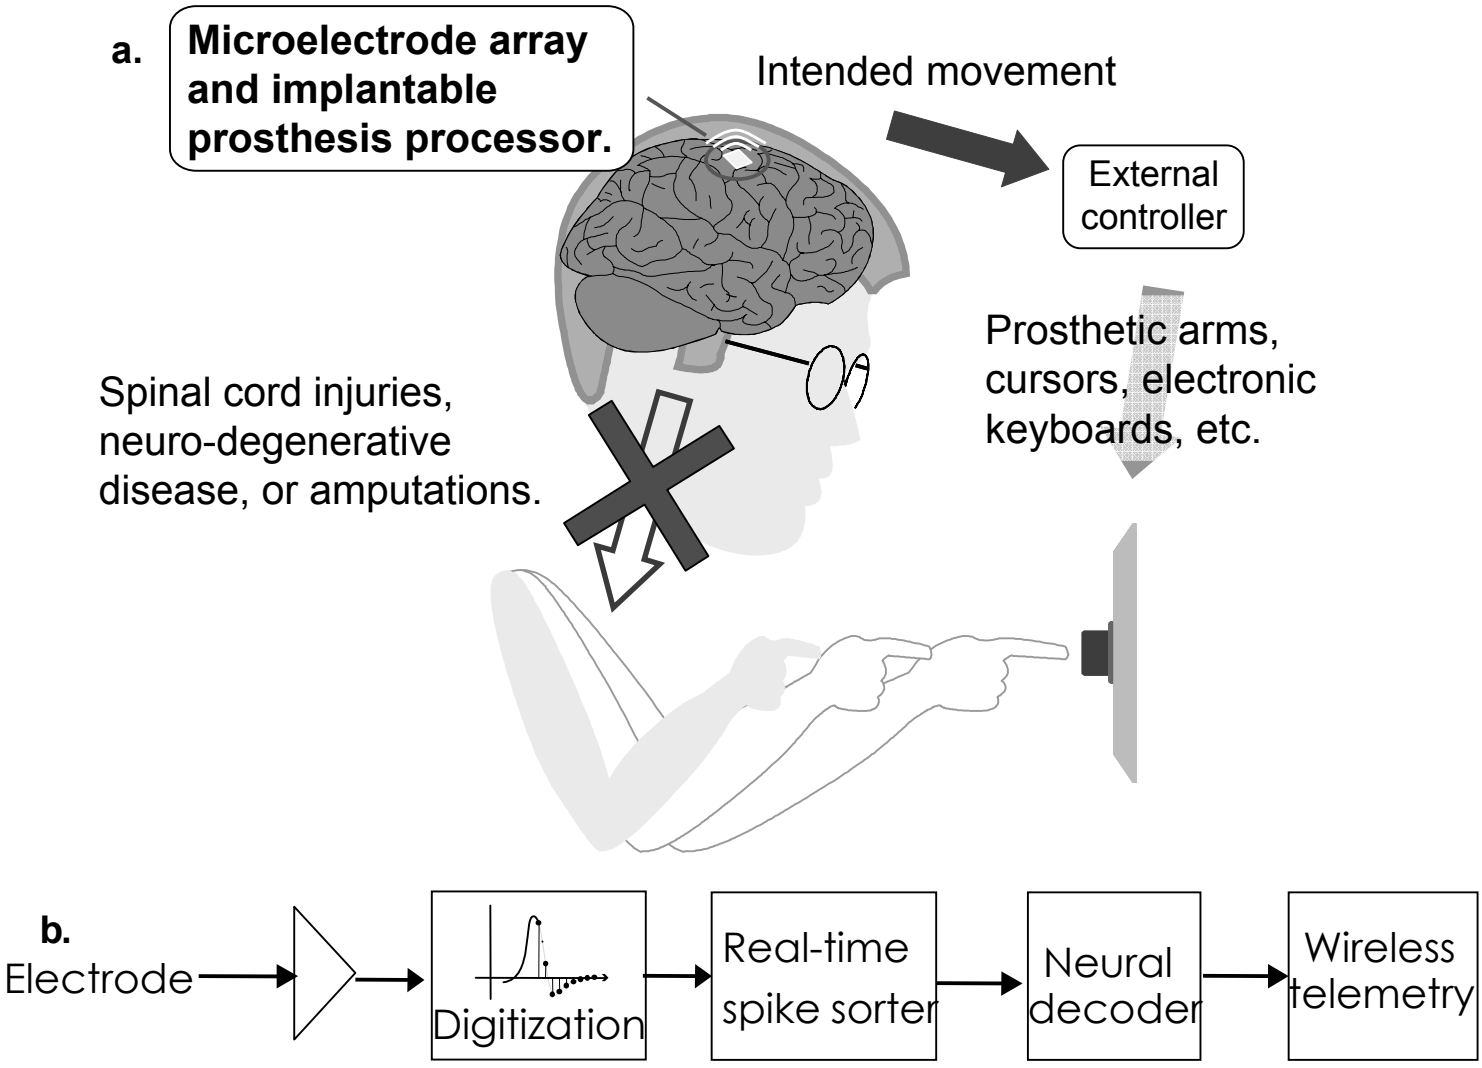

Figure 30.1.1: a. Implanted neural prosthetic system concept. b. Block diagram of implantable prosthesis processor.

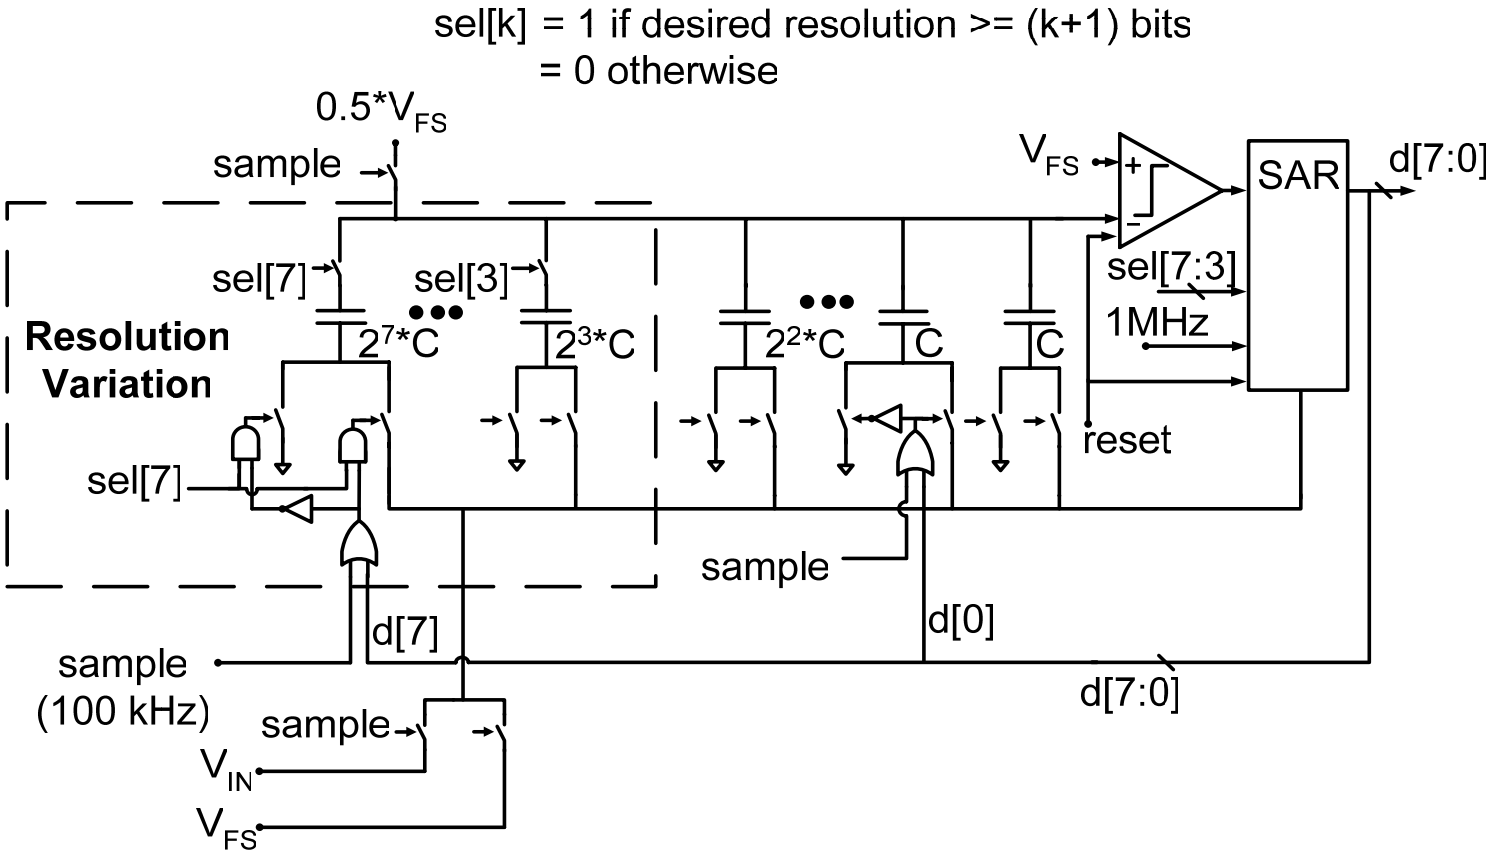

Figure 30.1.2: Successive approximation ADC cell with variable resolution.

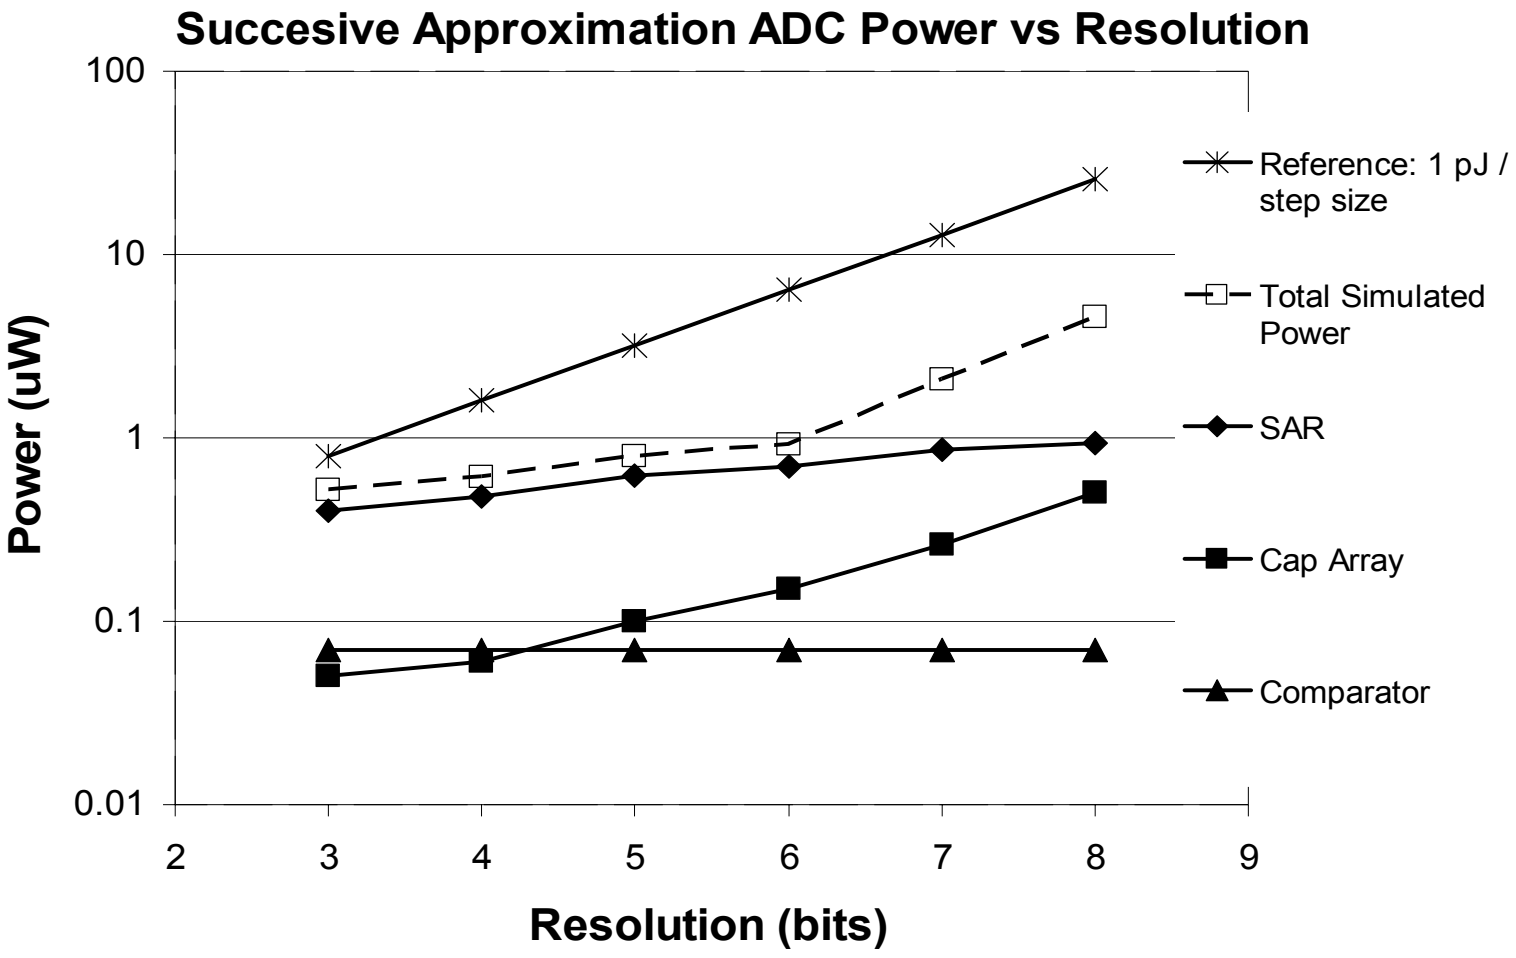

Figure 30.1.3: Power consumption of the ADC cell at different resolutions.

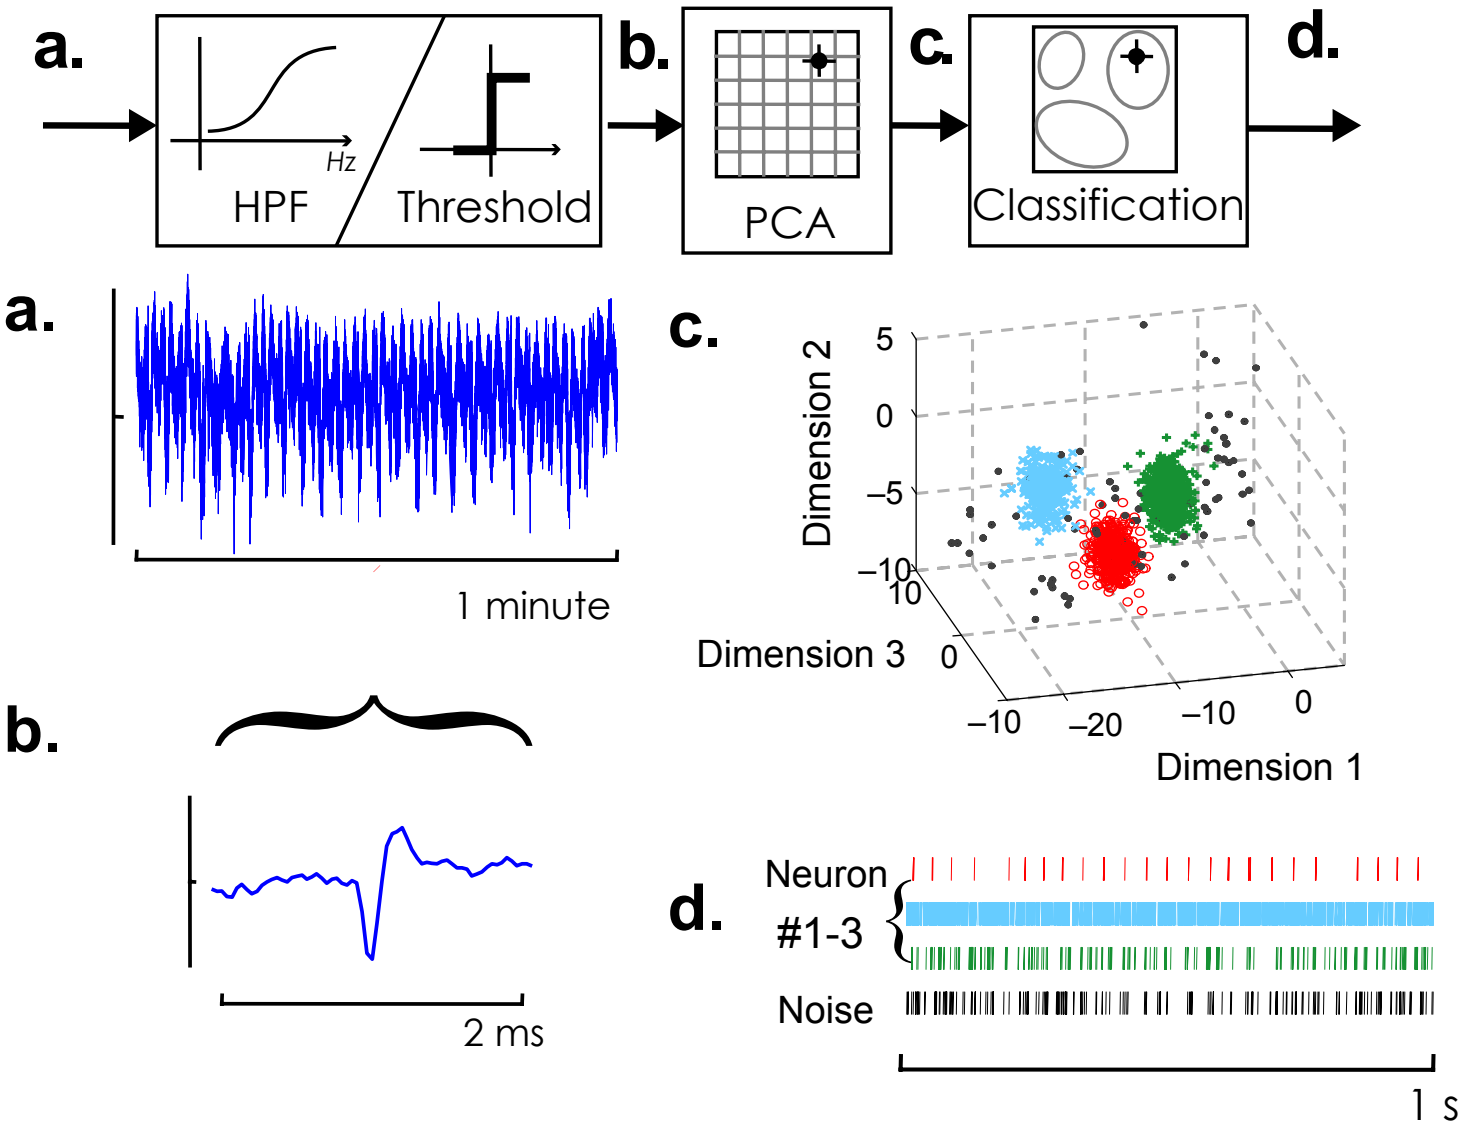

Figure 30.1.4: Real-time digital spike sorter. a. Digitized neural signals are first high-pass filtered and spike samples are selected. b. A spike window. c. Each spike is projected into a 3-dimensional discrimination space based on principle components analysis, PCA. d. Each identified spike is assigned to a specific neuron using a classification algorithm.

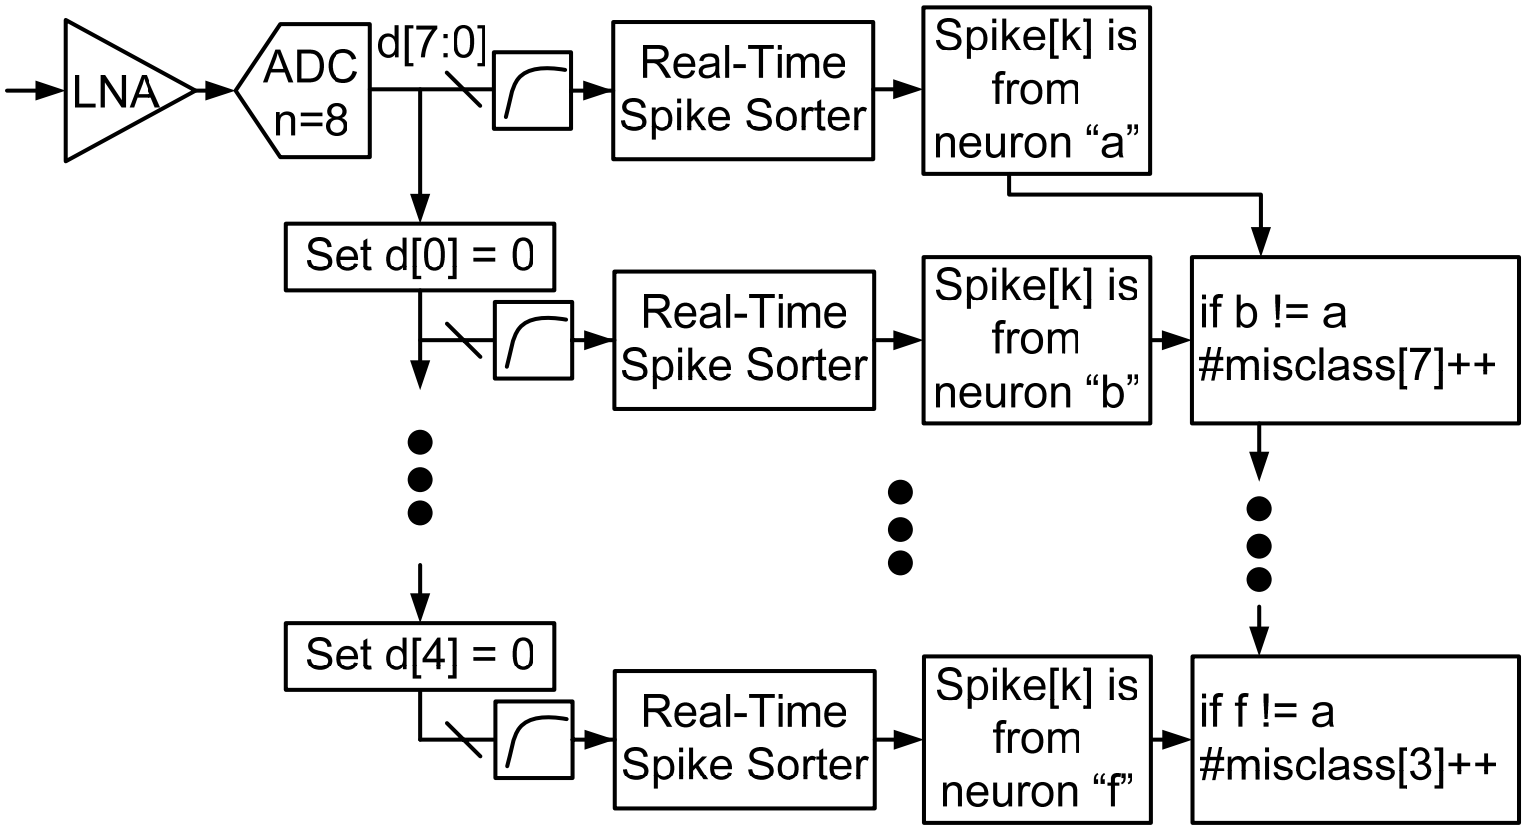

Figure 30.1.5: Method to determine the required ADC resolution for each channel.

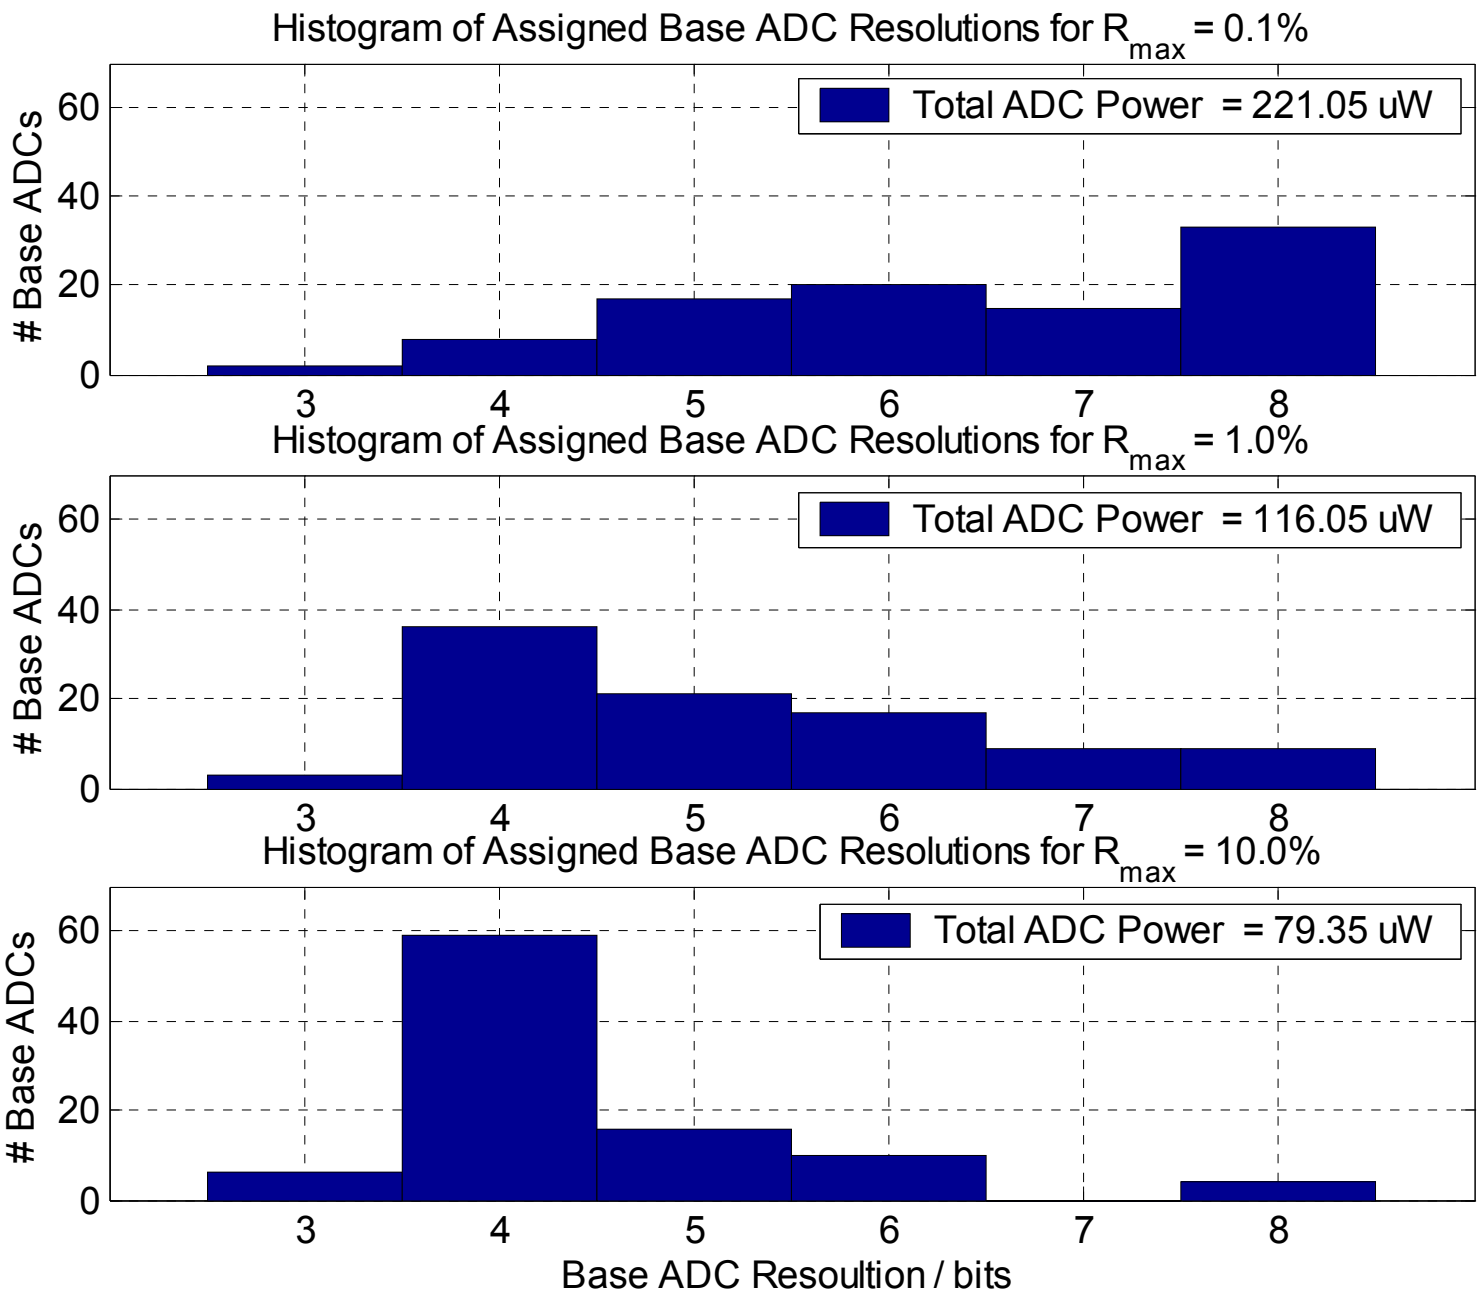

Figure 30.1.6: Histograms of ADC resolutions using sample neural data for R_{max} = 0.1%, 1%, and 10%.

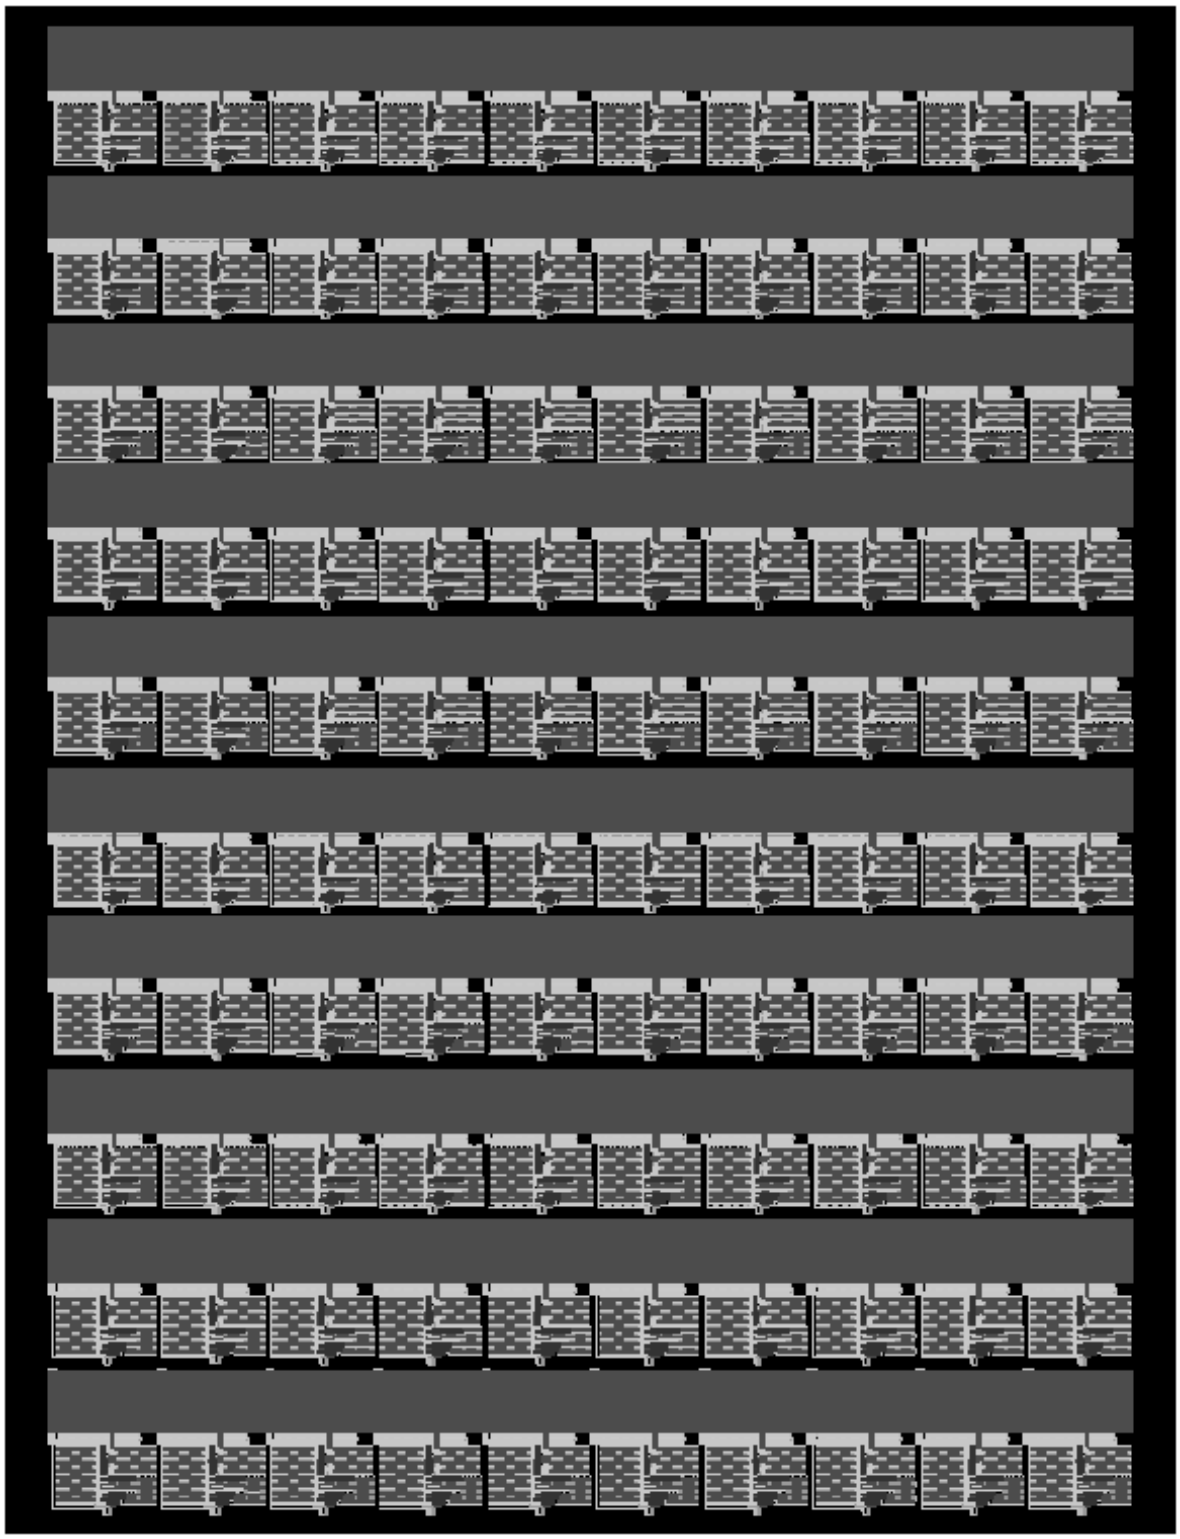

Figure 30.1.7: Layout plot of ADC array. Each row of the layout contains 10 ADC cells separated by input, signal and control lines. Dummy grounded lines are added between each signal line to avoid crosstalk.