Electrode pooling: How to boost the yield of switchable silicon probes for neuronal recordings

Kyu Hyun Lee[1,†], Yu-Li Ni[1,†], and Markus Meister[1,2,*]

[1]Division of Biology and Biological Engineering, Caltech
[2]Tianqiao and Chrissy Chen Institute for Neuroscience
[*]Corresponding author: meister@caltech.edu

[†]These authors contributed equally to this work.

Abstract

State-of-the-art silicon probes for electrical recording from neurons have thousands of recording sites, but only a fraction of them can be used simultaneously due to the forbiddingly large volume of the associated wires. To overcome this fundamental constraint, we propose a novel method called *electrode pooling* that uses a single wire to serve multiple recording sites. Multiple electrodes are connected to a single wire through a set of controllable switches. Here we present the framework behind this method and an experimental strategy to support it. We show that under suitable conditions electrode pooling can save wires without compromising the content of the recordings. We make recommendations for the design of future devices to take advantage of this strategy.

1 Introduction

Understanding brain function requires monitoring the complex pattern of activity distributed across many neuronal circuits. To this end, the BRAIN Initiative has called for the development of technologies for recording "dynamic neuronal activity from complete neural networks, over long periods, in all areas of the brain", ideally "monitoring all neurons in a circuit" (BRAIN Working Group, 2014). Recent advances in the design and manufacturing of silicon-based neural probes have answered this challenge with new devices that have thousands of recording sites (Jun et al., 2017b; Dimitriadis et al., 2018; Rios et al., 2016; Torfs et al., 2010). But in many such devices only a small fraction of the recording sites can be used at once. The reason is that neural signals detected by the recording sites must be brought out of the brain via wires, which take up much more volume than the recording sites themselves. For example, in one state-of-the-art silicon shank, each wire displaces thirty times more volume than a recording site once the shank is fully inserted in the brain (Jun et al., 2017b). The current silicon arrays invariably displace or destroy more neurons than they record, and thus the goal of "monitoring all neurons" seems unattainable by simply scaling the present approach [1]. Clearly we need a way to increase the number of neurons recorded while avoiding a concomitant increase in the number of wires that enter the brain.

[1]But see Kleinfeld et al. (2019) for a wildly optimistic proposal.

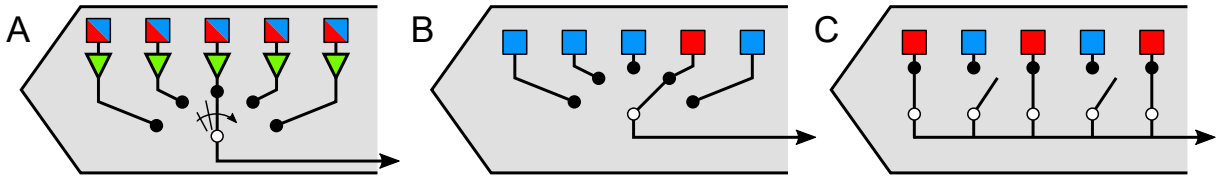

Figure 1: Strategies for using a single wire to serve many recording sites in switchable silicon probes. (A) Time-division multiplexing. Rapidly cycling the selector switch allows a single wire to carry signals from many recording sites interleaved in time. Green triangles represent anti-aliasing filters. (B) Static switching. A single wire connects to one of many possible recording sites through a selector switch. (C) Electrode pooling. Many recording sites are connected to a single wire through multiple controllable switches.

1.1 Time-division multiplexing

A common approach by which a single wire can convey multiple analog signals is "time-division multiplexing" (Obien et al., 2015). A rapid switch cycles through the N input signals and connects each input to the output line for a brief interval (Figure 1A). At the other end of the line, a synchronized switch or sampling system can demultiplex the N signals again. In this way a single wire carries signals from all its associated electrodes interleaved in time. The cycling rate of the switch is constrained by the sampling theorem (Shannon, 1949): It should be at least twice the highest frequency component present in the signal. The raw voltage signals from extracellular electrodes include thermal noise that extends far into the Megahertz regime. Therefore an essential element of any such multiplexing scheme is an analog low-pass filter associated with each electrode. This anti-alias filter removes the high-frequency noise above a certain cut-off frequency. In practice the cut-off is chosen to match the bandwidth of neuronal action potentials, typically 10 kHz. Then the multiplexer switch can safely cycle at a few times that cut-off frequency.

Given the ubiquity of time-division multiplexing in communication electronics, what prevents its use for neural recording devices? One obstacle is the physical size of the anti-alias filter associated with each electrode. When implemented in CMOS technology, such a low-pass filter occupies an area much larger than the recording site itself (Shahrokhi et al., 2010), which would force the electrodes apart, and thus prevent any high-density recording [2]. What if one simply omitted the low-pass filter? In that case aliasing of high-frequency thermal fluctuations will increase the noise power in the recording by a factor equal to the number of electrodes N being multiplexed. One such device with a multiplexing factor of $N = 128$ has indeed proven unsuitable for recording action potentials, as the noise drowns out any signal (Eversmann et al., 2003). A recent design with a more modest $N = 8$ still produces noise power 4-15 times higher than in comparable systems without multiplexing (Raducanu et al., 2016).

Other issues further limit the use of time-division multiplexing: The requirement for amplification, filtering, and rapid switching right next to the recording site means that electric power gets dissipated on location, which leads to local heating of exactly the neurons one wants to monitor. Furthermore, the active electronics in the local amplifier are sensitive to light. This can produce artifacts when combined with bright light flashes for optogenetic stimulation (Jun et al., 2017b; Kozai and Vazquez, 2015).

[2] One recent report claims to implement a 4 kHz low-pass filter with an electrode spacing of just 28 µm, but the underlying circuit has never been revealed (Angotzi et al., 2019)

1.2 Static selection

An alternative approach involves static electrode selection (Figure 1B). Again, there is an electronic switch that connects the wire to one of many electrodes. However the switch setting remains unchanged during the electrical recording. In this way the low-pass filtering and amplification can occur at the other end of the wire, outside the brain, where the base of the shank expands to offer virtually unlimited silicon space. The switch itself requires only minimal circuitry that fits comfortably under each recording site, even at a pitch of 20 µm or less. Because there is no local amplification or dynamic switching, the issues of heat dissipation or photosensitivity do not arise. This method has been incorporated recently into flat electrode arrays (Müller et al., 2015) as well as silicon prongs (Jun et al., 2017b; Lopez et al., 2017). It allows the user to choose one of many electrodes intelligently, for example because it carries a strong signal from a neuron of interest. However, it does not serve to increase the number of neurons per wire.

2 Electrode pooling

On this background we introduce a third method of mapping electrodes to wires: Select multiple electrodes with suitable signals and connect them to the same wire (Figure 1C). Instead of rapidly cycling the intervening switches, as in multiplexing, simply leave all those switches closed. This creates a "pool" of electrodes whose signals are summed and transmitted on the same wire. At first that approach seems counterproductive, as it mixes together recordings that one would like to analyze separately. How can one ever reconstruct which neural signal came from which electrode? Existing multi-electrode systems avoid this signal mixing at all cost, often quoting the low cross-talk between channels as a figure of merit. Instead, we will show how the pooled signal can be unmixed if one chooses the switch settings carefully during the recording session. Under suitable conditions this simple method can record many neurons per wire without appreciable loss of information. In this short paper we present the theory behind the approach and the predictions that can be derived from it. A more extensive report including experimental verification and extended simulations will appear elsewhere.

2.1 Spike trains are sparse in time

A typical neuron may fire $\sim$10 spikes/s on average (Attwell and Laughlin, 2001). Each action potential lasts for $\sim$1 ms. Therefore this neuron's signal occupies less than 1% of the time axis in an extracellular recording (Figure 2A). Sometimes a second neuron lies close enough to the same electrode to produce a large spike. That still leaves 98% of the time axis to transmit the action potentials of other neurons. Electrode pooling gives the experimenter the freedom to add more neurons to that signal by choosing other electrodes that carry large spikes (Figure 2). Eventually a limit will be reached when the spikes of different neurons collide and overlap in time so they can no longer be distinguished.

Fundamentally electrode pooling works because extracellular recordings are naturally sparse in time. The method can be seen as a variant of compressed sensing, which similarly relies on sparsity and randomly mixes multiple dimensions of sparse data into a common signal (Ganguli and Sompolinsky, 2012). However, one benefits greatly from setting the switches intelligently, rather than leaving them up to chance.

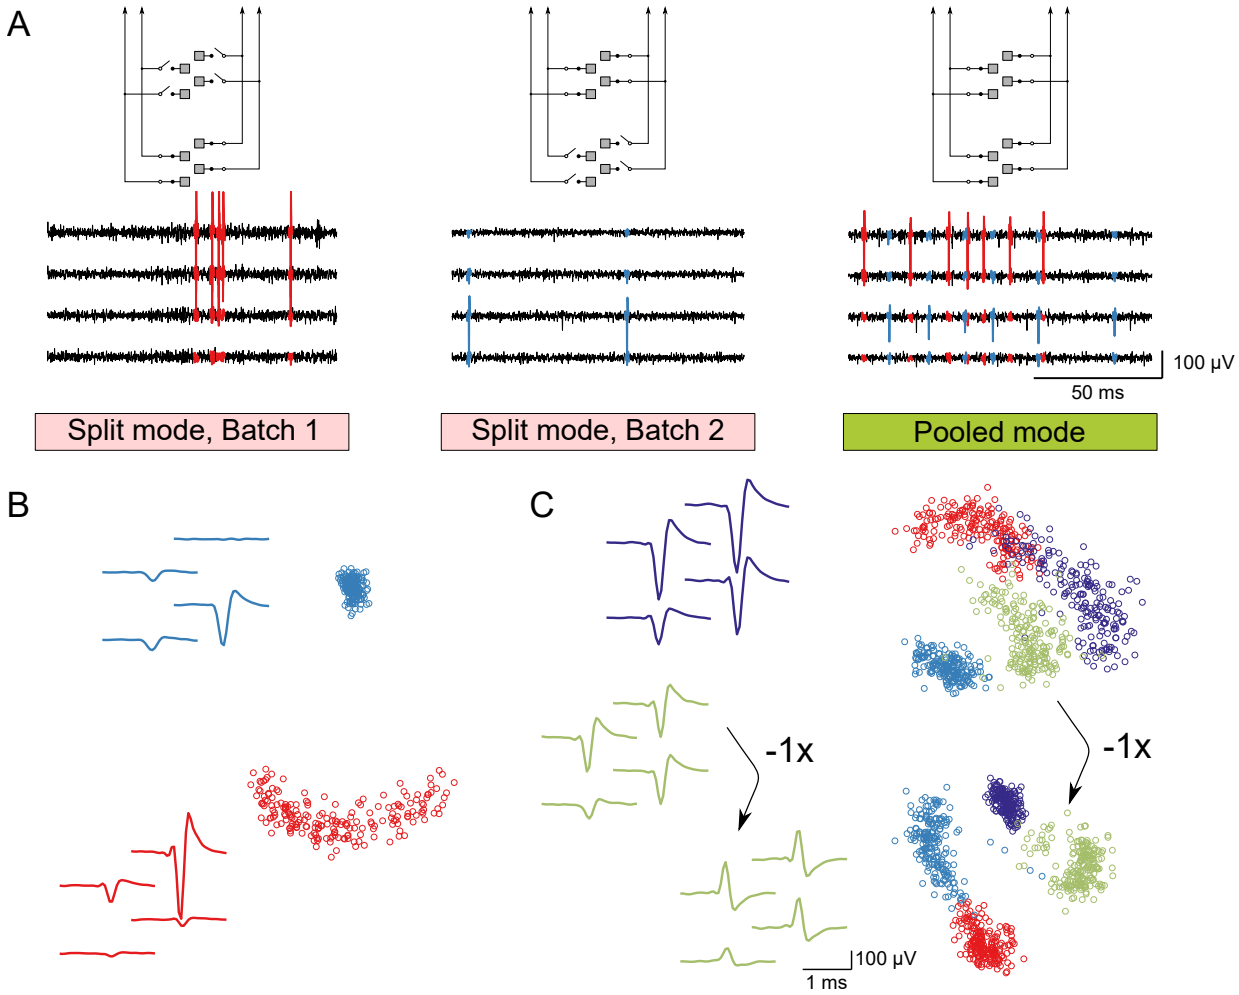

Figure 2: Pooling signals from multiple electrodes. (A) Successive split-mode recordings, illustrated with signals from one group of 4 electrodes (tetrode) (left) and later another (middle). In the subsequent pooled-mode recording (right) the two tetrode signals are merged. Units recorded from mouse superior colliculus with a passive non-switched electrode array. Pooling (right) simulated by simply averaging two electrode signals. (B) Spike waveforms of the units in (A) and their projection onto the first two principal components. The waveforms were recovered from the simulated pooled data by spike sorting with KiloSort2 (Pachitariu, 2019) without any manual curation and identified by comparing to sorted units in split-mode recordings. Note that the spikes from the two units in the pooled signal form two clusters that are easily separated. (C) Adding two additional tetrodes to the pool, each with a prominent spiking unit (purple and olive). Left: waveforms of the additional units. Right top: projections of the spikes on the first two principal components. Right bottom: same with inversion of the signal from the fourth tetrode before pooling. Note that the clusters are more separable.

2.2 Principles of electrode pooling

A key requirement for this method is that the experimenter can control all the switches that map electrodes to wires (Figure 1C). The commercial Neuropixels device (Jun et al., 2017b) has this capability (Figure 5A), but as of today it does not offer the user an option to activate multiple electrodes per wire. If given that kind of control, the user can adjust the electrode-to-wire map to the specific contingencies of any particular neural recording. In fact the experimenter will benefit from using different switch settings during the same session.

In Figure 2 we propose an overall workflow for experiments using electrode pooling. We envision that a recording session begins with a short period of acquisition in "split mode" with only one electrode per wire. The purpose is to acquire samples of the spike waveforms from all neurons that might be recorded by the entire array. If the device has E electrodes and W wires, this sampling stage will require at least E/W segments of recording to cover all electrodes. For each segment the switches are reset to select a different batch of electrodes. Each batch should cover a spatially compact group of electrodes on the array, perhaps with some overlap between batches. Because a neuron often produces signals on several adjacent electrodes this ensures that the entire "footprint" of each neuron can be captured (Figure 2A).

For the main phase of the experiment the user creates electrode pools, by combining the signals from several electrodes onto a common wire. We elaborate the criteria for these pooling decisions below. After allocating all the available wires to effective electrode pools one begins the main recording session in pooled mode (Figure 2A). Ideally this phase will capture all the neurons whose spike signals are within reach of the electrode array.

Finally the pooled data are analyzed by making use of the earlier recordings in split mode. By comparing the spike shapes encountered in the pooled data with those predicted from the split mode recordings and the electrode-to-wire map one can identify each recorded unit with the electrodes of origin. This process effectively demixes the pooled signal (Figure 2B-C).

2.3 The effects of pooling on spikes and noise

To predict the efficacy of this method, one needs to understand what signal actually results when one connects two electrodes to the same wire. Figure 3 shows the relevant circuit for an electrode array that allows electrode pooling. Here the common wire is connected via programmable switches to two recording electrodes. The CMOS switches themselves have low impedance, typically $R_{\mathrm{swi}} \approx 100\,\Omega$ (Seidl et al., 2012). Each electrode and the external electrolyte can be modeled as an RC element with a total impedance on the order of $R_i \approx 100\,\mathrm{k}\Omega$ (Seidl et al., 2012; Robinson, 1968). Thus the voltage on the shared wire (V_{P}) is the average of the signals at the recording sites, weighted inversely by the electrode impedances,

$$V_{\mathrm{P}} = \sum_{i=1}^{M} c_i V_i \tag{1}$$

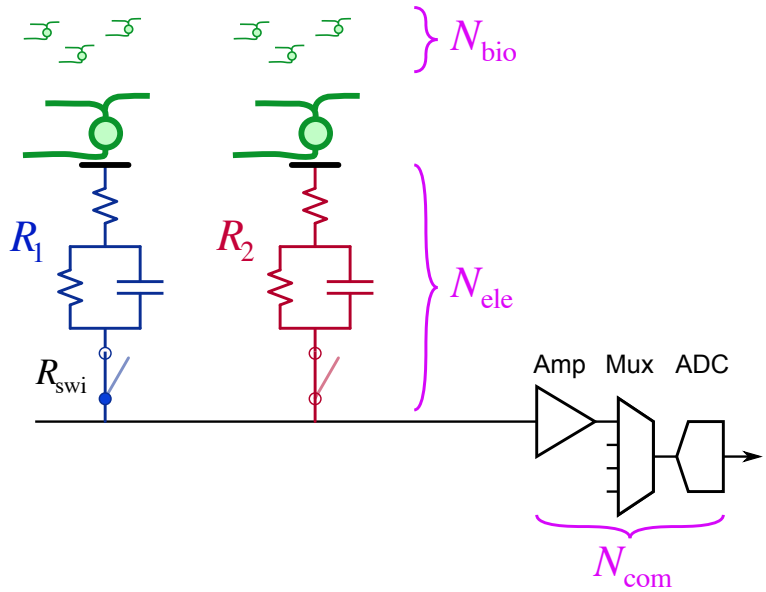

Figure 3: Pooling of signal and noise. An equivalent circuit model for two electrodes connected to a common wire along with downstream components of the signal chain, such as the amplifier, multiplexer, and digitizer. R_1, R_2: total impedance for electrodes 1 and 2; R_{swi}: switch impedance. N_{bio}: biological noise from distant neurons; N_{ele}: thermal noise from the electrode impedance; N_{com}: electronic noise common to all electrodes sharing the same wire.

where

$$c_i = \frac{1/R_i}{\sum_{j=1}^{M} 1/R_j} \tag{2}$$

is defined as the pooling coefficient for electrode i. If all electrodes have the same size and surface coating, they will have similar impedance, and in that limit one expects the simple relationship

$$V_{\text{P}} = \frac{1}{M} \sum_{i=1}^{M} V_i. \tag{3}$$

Thus an action potential that appears on only one of the M electrodes will be attenuated in the pooled signal by a factor $\frac{1}{M}$.

In order to understand the trade-offs of this method, we must similarly account for the pooling of noise (Figure 3). There are three relevant sources of noise: (1) thermal ("Johnson") noise from the impedance of the electrode; (2) biological noise ("hash") from many distant neurons whose signals are too small to be resolved; (3) electronic noise resulting from the downstream acquisition system, including amplifier, multiplexer, and analog-to-digital converter. The thermal noise is private to each electrode, in the sense that it is statistically independent of the noise at another electrode. The biological noise may be similar on neighboring electrodes that observe the same distant populations, but becomes independent for well-separated electrodes. Finally the noise introduced by data acquisition is common to all the electrodes that share the same wire.

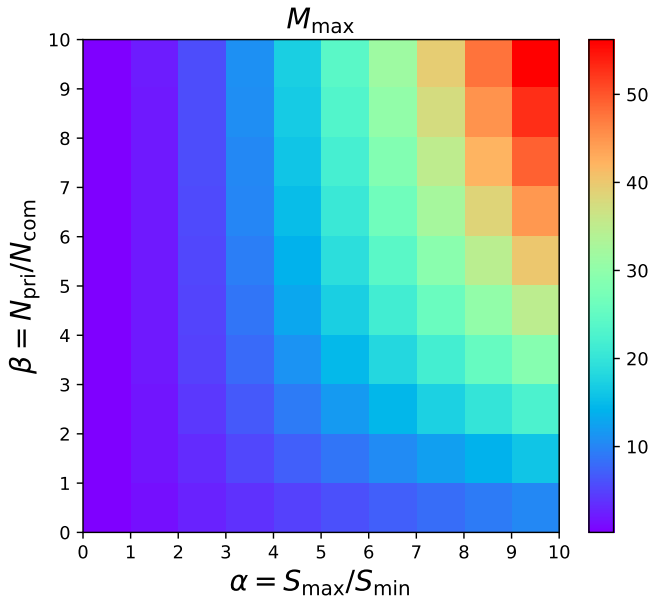

Figure 4: Maximal pool size $M_{\max}$ as a function of the parameters α and β that characterize spike signals and noise.

In the following we assume that the noise sources along the signal path are additive. For each electrode the Johnson noise and the hash are statistically independent, and therefore sum in quadrature to a total private noise with RMS amplitude

$$N_{\text{pri},i} = \sqrt{N_{\text{ele},i}^2 + N_{\text{bio},i}^2}. \tag{4}$$

In the course of pooling this signal gets attenuated by the pooling coefficient c_i (Eq 2). Then it gets added to the common noise from data acquisition, which again is statistically independent of the other noise sources. Thus the total noise at the output has RMS amplitude

$$N_{\text{tot}} = \sqrt{N_{\text{com}}^2 + \sum_{i=1}^{M} c_i^2 N_{\text{pri},i}^2}. \tag{5}$$

If all electrodes have similar noise properties and impedances this simplifies to

$$N_{\text{tot}} = \sqrt{N_{\text{com}}^2 + N_{\text{pri}}^2/M}. \tag{6}$$

2.4 Benefits and limits of pooling

Now we are in a position to estimate the maximal possible benefits from electrode pooling. Suppose that using the conventional split-mode recording mode (1 electrode per wire, $M = 1$) we find a range of spikes on the electrode array: from the largest, with spike amplitude $S_{\max}$, to the smallest that can still be sorted

reliably from the noise, with amplitude $S_{\min}$. To create an effective pool, we find the electrodes with the largest spikes and add them to the pool. Eventually the spikes in the pooled signal will be so attenuated that they are no longer sortable from the noise. Pooling is beneficial as long as the signal-to-noise ratio of spikes in the pooled signal is larger than that of the smallest sortable spikes in a split-mode recording, namely

$$\frac{S_{\max}/M}{\sqrt{N_{\text{com}}^2 + N_{\text{pri}}^2/M}} > \frac{S_{\min}}{\sqrt{N_{\text{com}}^2 + N_{\text{pri}}^2}}. \tag{7}$$

This leads to a limit on the pool size M,

$$M < M_{\max} = \sqrt{\left(\frac{\beta^2}{2}\right)^2 + (1+\beta^2)\alpha^2} - \frac{\beta^2}{2} \tag{8}$$

where

$$\alpha = S_{\max}/S_{\min},\ \beta = N_{\text{pri}}/N_{\text{com}} \tag{9}$$

So the maximal beneficial pool size depends on two parameters: the ratio of private to common noise, and the ratio of largest to smallest useful spike amplitudes. These parameters vary across applications, because they depend on the target brain area, the recording hardware, and the spike sorting software. Users can estimate those parameters from experience with conventional recordings, and find $M_{\max}$ from the lookup table in Figure 4. For example, assuming a modest 5-fold range between the largest and smallest sortable spikes, and supposing that the "hash" is 4 times larger than the electronic noise, one finds that pooling can be beneficial up to 14 electrodes.

Note these calculations focus on the discrimination of spikes from noise, not on distinguishing different spikes. As more neurons join the pool, two different cells may produce spikes with similar waveforms that are difficult to separate. Furthermore, as the combined firing rate of all neurons increases, action potentials start to overlap in time, which further hampers their separation. We address these challenges below, but for now the above expression for $M_{\max}$ (Eq 8) must be considered an upper bound on the beneficial pool size.

2.5 Workflow for experiments and analysis of pooled recordings

With these insights one can derive more specific instructions for the electrode pooling method:

1. Split mode: The purpose of this sampling phase (Figure 2A) is to acquire a catalog of single units that exist on each batch of electrodes. This requires a quick pass at spike sorting, yielding the spike waveforms and firing rates for neurons on each electrode. Finally, for each electrode one also samples the private noise N_{pri}. The common noise N_{com} can be assessed ahead of time, because it is a property of the recording system.

2. Choosing electrode pools: The experimenter now has all the information needed to select useful electrode pools. Some principles one should consider in this process:

- Pool electrodes that carry large signals.
- Pool electrodes with different spike waveforms.
- Don't pool neighboring electrodes that share the same hash noise.
- Don't pool electrodes that carry dense signals with high firing rates.

3. Pooled mode: This is the bulk of the experiment that samples all the conditions required by the study, maximizing the number of units recorded (Figure 2A). It probably pays to monitor the development of spike shapes during the pooled-mode recording. If they drift substantially, for example because the electrode array moves in the brain (Jun et al., 2017a), then a recalibration by another split-mode sampling at the end of the session may be productive.

4. Analysis: The goal is to detect spikes in the pooled signals and assign each spike correctly to its electrode of origin. This step uses the split-mode recordings from the early sampling stage of the experiment, from which one can predict how the corresponding spike will appear in the pooled signals (Figure 2B-C). Here it helps to know all the electrode impedances R_i so the weighted mix can be computed accurately (Eq 1). This prediction serves as a search template for spike sorting the pooled recording. Because these spike templates are obtained in split mode, they are less affected by noise than if one had to identify them *de novo* from the pooled recordings.

By its very nature electrode pooling produces a dense neural signal with more instances of temporal overlap between spikes than the typical split-mode recording. This places special demands on the methods for spike detection and sorting. The conventional cluster-based algorithm (peak detection - temporal alignment - PCA - clustering) does not handle overlapping spikes well (Lewicki, 1998). It assumes that the voltage signal is sparsely populated with rare events drawn from a small number of discrete waveforms. Two spikes that overlap in time produce an unusual waveform that cannot be categorized. Recently some methods have been developed that do not force these assumptions (Yger et al., 2018; Pachitariu et al., 2016). They explicitly model the recorded signal as an additive superposition of spikes and noise. The algorithm finds an efficient model that explains the signal by estimating both the spike waveform of each neuron and its associated set of spike times. These methods are well suited to the analysis of pooled recordings.

3 Recommendations for future neural recording systems

3.1 Hardware

The ability to service multiple recording sites with a single wire opens the door for much larger electrode arrays that nevertheless maintain a slim form factor and don't require any onboard signal processing. In the current version of the Neuropixels device (Jun et al., 2017b) the ratio of electrodes to wires is only 2.5, and thus there is little practical benefit to be gained from electrode pooling. In most circumstances the user can probably use static selection to pick 40% of the electrodes and still monitor every possible neuron. However, devices are already under development with an electrode:wire ratio of 10 or greater. With pooled recording in mind one should contemplate future devices with much higher ratios still (see Figure 4 for guidance).

The design of effective electrode pools requires some flexibility in how recording sites are connected to wires. In the current Neuropixels each electrode has only one associated switch, and thus only one candidate wire. Electrode pools are limited to sites spaced 3.8 mm apart. The CMOS switch itself is small, but the local

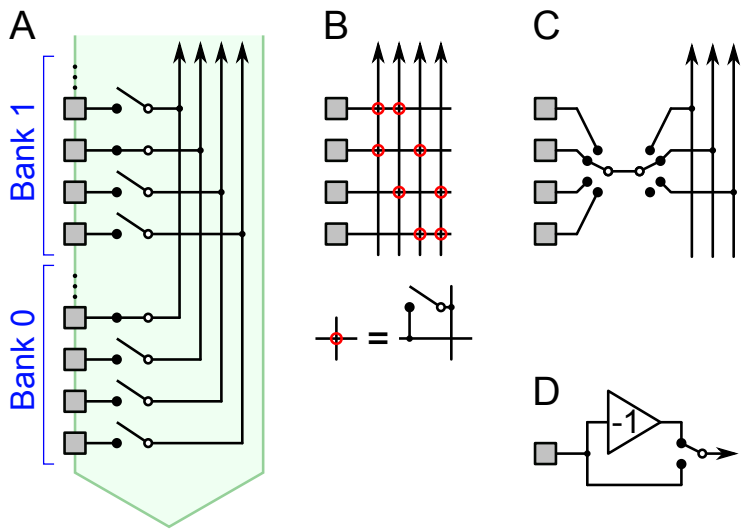

Figure 5: Hardware schemes for flexible connection between electrodes and wires. (A) In the current Neuropixels array each electrode can be connected to just one wire with a controllable switch. (B) Two switches per electrode would allow a choice of 2 wires, enabling many more pooling configurations. (C) Because neighboring electrodes often carry redundant signals, one may want to choose just one from every group of 4. This switch circuit matches that choice with one of 3 (or no) wires. (D) An optional inverter for each electrode, controlled by a local switch.

memory to store the switch state occupies some silicon space (Seidl et al., 2011). Nonetheless one can implement 3 switches per electrode even on a very tight pitch (Dragas et al., 2017). When arranged in a hierarchical network (Müller et al., 2015) these switches could effect a rich diversity of pooling schemes adapted to the specifics of any given experiment (Figure 5). For example, one could route any one electrode among a group of four to any one of three wires with two 1:4 switches (Figure 5C). This requires just 1 bit of storage per electrode, as in the current Neuropixels probe (Jun et al., 2017b).

Another hardware design feature could greatly increase the capacity for electrode pooling: An optional inverter at each electrode (Figure 5D). This is a simple CMOS circuit that changes the sign of the waveform (Bae, 2019) depending on a local switch setting. If half of the electrodes in a pool use the inverter, that helps to differentiate the spike shapes of different neurons. Because extracellular signals from cell bodies generally start with a negative voltage swing, this effectively doubles the space of waveforms that occur in the pooled signal. In turn this aids the spike-sorting analysis, ultimately allowing even more electrodes to share the same wire (Figure 2C).

3.2 Software

Electrode pooling will also benefit from further developments in spike sorting algorithms. For example, as we discuss in section 2.5, a promising strategy is to acquire all the spike shapes present on the electrode array using split-mode recordings, compute the expected pooled-mode waveforms, and use those as templates in sorting the pooled-mode signals. We expect that this capability could be readily added to existing sorting algorithms, such as KiloSort (Pachitariu et al., 2016). The challenges of pooled recordings also place great value on improving the resolution of temporally overlapping spike waveforms.

Another interesting software challenge lies in the optimal design of electrode pools. As one considers experiments with 10,000 or more recording sites, it becomes imperative to automate this process, so that the

user wastes no time before launching into pooled-mode recording. In section 2.5 we lay out some heuristic rules one might follow, and one can envision turning these into an effective algorithm that makes use of the full noise and impedance specifications of the device.

4 Discussion

4.1 Summary of results

This work presents the concept of flexible electrode pooling as a way to multiply the yield of large electrode arrays. We show how the signals from many recording sites can be combined onto a small number of wires, and then recovered by a combination of experimental strategy and spike-sorting software. We developed the theory behind electrode pooling, analyzed the trade-offs of the approach, derived a mathematical limit to pooling, and developed a recipe for experiment and analysis that implements the procedure (section 2). By these methods the signals from different neurons can be reliably disambiguated and assigned back to their electrodes of origin.

Future electrode array devices should take full advantage of this new method, and we propose a number of hardware and software developments to pursue that goal (section 3). Electrode pooling can dramatically reduce the number of wires needed per recorded neuron. This will allow the manufacture of slender array devices that cause less damage to the brain circuits they are meant to observe.

4.2 General strategies for using switchable probes

The switching circuitry in today's silicon probes opens up new ways of operating these devices. Static switching, time-division multiplexing, and electrode pooling are just some of the possible modes of operation (Figure 1). Although this paper focused on the benefits of electrode pooling, experimenters and hardware designers could in principle mix and match these strategies. Indeed they are fully complementary: static switching takes advantage of the redundancy in the neural signal across neighboring electrodes; time-division multiplexing uses the fact that a wire has much higher bandwidth than a neuron; and electrode pooling exploits the sparseness of spiking neural signals on the time axis. Exploring the best combination of these strategies for each use case may dramatically enhance the yield of extracellular recordings in the future.

Acknowledgements

This research was supported by a grant from the Tianqiao and Chrissy Chen Institute for Neuroscience.

References

Angotzi, G. N., Boi, F., Lecomte, A., Miele, E., Malerba, M., Zucca, S., Casile, A., and Berdondini, L. (2019). SiNAPS: An implantable active pixel sensor CMOS-probe for simultaneous large-scale neural recordings. *Biosensors and Bioelectronics*, 129:355–364.

Attwell, D. and Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab*, 21:1133–45.

Bae, W. (2019). CMOS Inverter as Analog Circuit: An Overview. *Journal of Low Power Electronics and Applications*, 9:26, DOI: 10.3390/jlpea9030026.

BRAIN Working Group, . (2014). BRAIN 2025: A Scientific Vision. Technical report, https://braininitiative.nih.gov/strategic-planning/brain-2025-report.

Dimitriadis, G., Neto, J. P., Aarts, A., Alexandru, A., Ballini, M., Battaglia, F., Calcaterra, L., David, F., Fiáth, R., Frazão, J., Geerts, J. P., Gentet, L. J., Van Helleputte, N., Holzhammer, T., van Hoof, C., Horváth, D., Lopes, G., Lopez, C. M., Maris, E., Marques-Smith, A., Márton, G., McNaughton, B. L., Meszéna, D., Mitra, S., Musa, S., Neves, H., Nogueira, J., Orban, G. A., Pothof, F., Putzeys, J., Raducanu, B., Ruther, P., Schroeder, T., Singer, W., Tiesinga, P., Ulbert, I., Wang, S., Welkenhuysen, M., and Kampff, A. R. (2018). Why not record from every channel with a CMOS scanning probe? *bioRxiv*, DOI: 10.1101/275818, http://biorxiv.org/lookup/doi/10.1101/275818.

Dragas, J., Viswam, V., Shadmani, A., Chen, Y., Bounik, R., Stettler, A., Radivojevic, M., Geissler, S., Obien, M. E. J., Müller, J., and Hierlemann, A. (2017). In Vitro Multi-Functional Microelectrode Array Featuring 59 760 Electrodes, 2048 Electrophysiology Channels, Stimulation, Impedance Measurement, and Neurotransmitter Detection Channels. *IEEE Journal of Solid-State Circuits*, 52(6):1576–1590, ISSN: 0018-9200, DOI: 10.1109/JSSC.2017.2686580.

Eversmann, B., Jenkner, M., Hofmann, F., Paulus, C., Brederlow, R., Holzapfl, B., Fromherz, P., Merz, M., Brenner, M., Schreiter, M., Gabl, R., Plehnert, K., Steinhauser, M., Eckstein, G., Schmitt-Landsiedel, D., and Thewes, R. (2003). A 128 x 128 cmos biosensor array for extracellular recording of neural activity. *IEEE Journal of Solid-State Circuits*, 38(12):2306–2317, ISSN: 0018-9200, DOI: 10.1109/JSSC.2003.819174, http://ieeexplore.ieee.org/document/1253878/.

Ganguli, S. and Sompolinsky, H. (2012). Compressed Sensing, Sparsity, and Dimensionality in Neuronal Information Processing and Data Analysis. *Annual Review of Neuroscience*, 35(1):485–508, DOI: doi:10.1146/annurev-neuro-062111-150410, http://www.annualreviews.org/doi/abs/10.1146/annurev-neuro-062111-150410.

Jun, J., Mitelut, C., Lai, C., Gratiy, S. L., Anastassiou, C. A., and Harris, T. D. (2017a). Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. *bioRxiv*, pages 1–29, DOI: https://doi.org/10.1101/101030.

Jun, J. J., Steinmetz, N. A., Siegle, J. H., Denman, D. J., Bauza, M., Barbarits, B., Lee, A. K., Anastassiou, C. A., Andrei, A., Aydin, [U+FFFD], Barbic, M., Blanche, T. J., Bonin, V., Couto, J., Dutta, B., Gratiy, S. L., Gutnisky, D. A., Häusser, M., Karsh, B., Ledochowitsch, P., Lopez, C. M., Mitelut, C., Musa, S., Okun, M., Pachitariu, M., Putzeys, J., Rich, P. D., Rossant, C., Sun, W. L., Svoboda, K., Carandini, M., Harris, K. D., Koch, C., O'Keefe, J., and Harris, T. D. (2017b). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, ISSN: 14764687, DOI: 10.1038/nature24636, http://dx.doi.org/10.1038/nature24636.

Kleinfeld, D., Luan, L., Mitra, P. P., Robinson, J. T., Sarpeshkar, R., Shepard, K., Xie, C., and Harris, T. D. (2019). Can One Concurrently Record Electrical Spikes from Every Neuron in a Mammalian Brain? *Neuron*, 103(6):1005–1015, ISSN: 08966273, DOI: 10.1016/j.neuron.2019.08.011, https://linkinghub.elsevier.com/retrieve/pii/S0896627319306956.

Kozai, T. D. Y. and Vazquez, A. L. (2015). Photoelectric artefact from optogenetics and imaging on microelectrodes and bioelectronics: New challenges and opportunities. *Journal of Materials Chemistry B*, 3(25):4965–4978, ISSN: 2050-7518, DOI: 10.1039/C5TB00108K.

Lewicki, M. S. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network*, 9 4:R53–78.

Lopez, C. M., Putzeys, J., Raducanu, B. C., Ballini, M., Wang, S., Andrei, A., Rochus, V., Vandebriel, R., Severi, S., Van Hoof, C., Musa, S., Van Helleputte, N., Yazicioglu, R. F., and Mitra, S. (2017). A Neural Probe With Up to 966 Electrodes and Up to 384 Configurable Channels in 0.13 µm SOI CMOS. *IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS*, 11(3):510–522.

Müller, J., Ballini, M., Livi, P., Chen, Y., Radivojevic, M., Shadmani, A., Viswam, V., Jones, I. L., Fiscella, M., Diggelmann, R., Stettler, A., Frey, U., Bakkum, D. J., and Hierlemann, A. (2015). High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. *Lab on a Chip*, 15(13):2767–2780, ISSN: 1473-0197, 1473-0189, DOI: 10.1039/C5LC00133A, http://xlink.rsc.org/?DOI=C5LC00133A.

Obien, M. E. J., Deligkaris, K., Bullmann, T., Bakkum, D. J., and Frey, U. (2015). Revealing neuronal function through microelectrode array recordings. *Frontiers in Neuroscience*, 8, ISSN: 1662-453X, DOI: 10.3389/fnins.2014.00423, http://journal.frontiersin.org/article/10.3389/fnins.2014.00423/abstract.

Pachitariu, M. (2019). Mouseland/kilosort2. https://github.com/MouseLand/Kilosort2.

Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M., and Harris, K. D. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems*, 29(Nips):4448–4456, ISSN: 10495258, https://papers.nips.cc/paper/6326-fast-and-accurate-spike-sorting-of-high-channel-count-probes-with-kilosort.pdf.

Raducanu, B. C., Yazicioglu, R. F., Lopez, C. M., Ballini, M., Putzeys, J., Wang, S., Andrei, A., Welkenhuysen, M., Van Helleputte, N., Musa, S., Puers, R., Kloosterman, F., Van Hoof, C., and Mitra, S. (2016). Time multiplexed active neural probe with 678 parallel recording sites. *European Solid-State Device Research Conference*, 2016-Octob:385–388, ISSN: 19308876, DOI: 10.1109/ESSDERC.2016.7599667.

Rios, G., Lubenov, E. V., Chi, D., Roukes, M. L., and Siapas, A. G. (2016). Nanofabricated Neural Probes for Dense 3-D Recordings of Brain Activity. *Nano Letters*, 16(11):6857–6862, ISSN: 1530-6984, 1530-6992, DOI: 10.1021/acs.nanolett.6b02673, https://pubs.acs.org/doi/10.1021/acs.nanolett.6b02673.

Robinson, D. (1968). The electrical properties of metal microelectrodes. *Proceedings of the IEEE*, 56(6):1065–1071, ISSN: 0018-9219, DOI: 10.1109/PROC.1968.6458, http://ieeexplore.ieee.org/document/1448388/.

Seidl, K., Herwik, S., Torfs, T., Neves, H., Paul, O., and Ruther, P. (2011). CMOS-Based High-Density Silicon Microprobe Arrays for Electronic Depth Control in Intracortical Neural Recording. *Journal of Microelectromechanical Systems*, 20:1439–1448.

Seidl, K., Schwaerzle, M., Ulbert, I., Neves, H. P., Paul, O., and Ruther, P. (2012). CMOS-Based High-Density Silicon Microprobe Arrays for Electronic Depth Control in Intracortical Neural Recording–Characterization and Application. *Journal of Microelectromechanical Systems*, 21(6):1426–1435,

ISSN: 1057-7157, 1941-0158, DOI: 10.1109/JMEMS.2012.2206564, http://ieeexplore.ieee.org/document/6249712/.

Shahrokhi, F., Abdelhalim, K., Serletis, D., Carlen, P. L., and Genov, R. (2010). The 128-Channel Fully Differential Digital Integrated Neural Recording and Stimulation Interface. *IEEE Transactions on Biomedical Circuits and Systems*, 4(3):149–161, ISSN: 1932-4545, 1940-9990, DOI: 10.1109/TBCAS.2010.2041350, http://ieeexplore.ieee.org/document/5471738/.

Shannon, C. E. (1949). Communication in the Presence of Noise. *Proceedings of the IRE*, 37(1):10–21, ISSN: 0096-8390, DOI: 10.1109/JRPROC.1949.232969.

Torfs, T., Aarts, A., Erismis, M. A., Aslam, J., Yazicioglu, R. F., Puers, R., Van Hoof, C., Neves, H., Ulbert, I., Dombovari, B., Fiath, R., Kerekes, B. P., Seidl, K., Herwik, S., and Ruther, P. (2010). Two-dimensional multi-channel neural probes with electronic depth control. In *2010 Biomedical Circuits and Systems Conference (BioCAS)*, pages 198–201, Paphos, Cyprus. IEEE, ISBN: 978-1-4244-7269-7, DOI: 10.1109/BIOCAS.2010.5709605, http://ieeexplore.ieee.org/document/5709605/.

Yger, P., Spampinato, G. L., Esposito, E., Lefebvre, B., Deny, S., Gardella, C., Stimberg, M., Jetter, F., Zeck, G., Picaud, S., Duebel, J., and Marre, O. (2018). A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *eLife*, 7, ISSN: 2050-084X, DOI: 10.7554/eLife.34518.