The Neurochip: A complete system for long-term investigation of cultured neural network connectivity

Thesis by

Jonathan Christopher Erickson

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

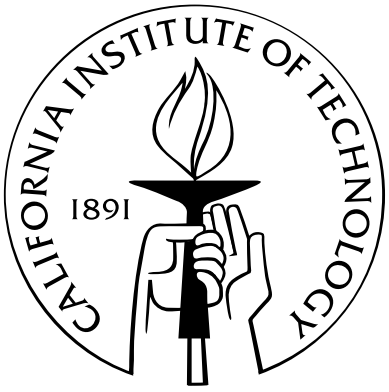

California Institute of Technology

Pasadena, California

2008

(Defended February 29, 2008)

Acknowledgements

There are a great many people to whom I am grateful for helping this work come to fruition: Prof. Jerry Pine for allowing me to make (and correct) mistakes and develop into a better scientist. From him I learned that a little bit of "freshman physics" can go a long way, and with this "basic" knowledge it is possible to build almost anything, be it an op-amp circuit or a cabin in the mountains. A man of incredible vigor, he would be down on the floor with me picking up the screws and BNC connectors at the end of the day. Also, Profs. Scott Fraser, Joel W. Burdick, and Y-C Tai who formed the rest of my thesis committee and offered helpful suggestions throughout.

The unsung heroes, the keystones of Caltech, contributed a huge amount to the success of this project, among them: Sheri McKinney for graciously providing neurons and smiles every week for 5+ years. Pat Koen and Jean Edens for producing masterful, beautiful SEMs. Mike Roy and Steve Olson, machine shop wizards, for manufacturing several custom-machined parts. Pat Perrone and Carole Worra for helping with a seemingly infinite amount of special purchase orders.

I am very grateful and appreciate having worked with excellent laboratory colleagues: Angela Tooker, for many, many rounds of microfabrication of neurochips; Gary Chow for many helpful suggestions and afternoon ice cream and humor breaks; Daniel Wagenaar for sharing a vast amount of knowledge and patience in the early years, and encouraging me to build and tinker.

Caltech pals: To all of you who provided laughs, diversions in the mountains and ocean and everywhere in between, thank you all.

Last, but certainly not least, I would like to express a special appreciation to my entire family, who let me chase my dreams from Colorado all the way out to California. And also to my beautiful wife, Sarah, for being a tireless listener and an amazing calming presence, especially during the endgame. I'll always remember the potato-leek gruyere gratin.

Financial support for this work was provided by the NIH Grant NS-044134.

Abstract

Traditional techniques for investigating cultured neural networks, such as the patch clamp and multi-electrode array, are limited by 1) the number of identified cells which can be in simultaneous electrical contact, 2) the length of time for which cells can be studied, and 3) the lack of 1:1 neuron-to-electrode specificity.

Here, I present a novel device—dubbed the "neurochip"— which overcomes these limitations. This micromachined device consists of 4x4 array of "neurocages" which mechanically trap a neuron near an extracellular electrode. While the cell body is trapped, the axon and dendrites can freely grow into the surrounding area to form a network. The electrode is bi-directional, capable of both stimulating and recording action potentials. This system is noninvasive, so that an entire network—all constituent neurons—can be studied over its lifetime with fixed 1:1 neuron-to-electrode correspondence. Proof-of-concept experiments have been completed to illustrate that functional networks do indeed form in the neurochip system, and that suprathreshold connectivity can be fully mapped over several weeks. The neurochip opens a new domain in neurobiology for studying small cultured neural networks.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Studying Cultured Neural Networks

The brain is an extremely versatile organ that performs an astonishing array of functions. Not only can it can store and recall the several thousand notes, sounds, and rhythms of Gershwin's *Rhapsody in Blue*, it can control ten fingers and two feet to play it. The brain instantly recognizes the faces of Mom, Dad, Brother, Sister, Fido, and (unfortunately) Paris Hilton. It contemplates the grand scale of the cosmos and the nanoscales of DNA. And it is the seat of invention that has devised everything from the diesel engine to Google's search engine.

The amazing functionality of the brain derives in part from sheer numbers. The human brain has been estimated to contain upwards of one hundred billion neurons. Each neuron makes, on average, about one thousand synaptic connections to other neurons. One hundred trillion connections, that's one heck of a wiring job! (And many a budding engineer thought building op-amp circuits in sophomore lab was bad.) Perhaps the brain's most astonishing feat is its ability, through a combination of nature and nurture, to correctly wire itself. What's more, new connections are constantly being formed, changed, broken, and reformed in response to external input and experience. This neural development and plasticity is precisely what I—and many other neurobiologists—are interested in studying. How are synapses made? How do neurons decide to keep or prune a connection? How do networks evolve over time?

It is fair to say that the brain is extremely complex. To begin to tackle such an enormous beast, many investigators have chosen to work with simpler, reduced systems. The leech, fruit fly, and grasshopper have all been extensively utilized as model neurobiological systems,

for example. Another popular reduced system is that of dissociated vertebrate neuronal cell-culture. Typically tens of thousands of cells are grown in a single culture—a factor of a million fewer neurons than the human brain. Conventional electrophysiology techniques for studying these—such as the patch clamp and multi-electrode array—have served as very important tools in neurobiology and much knowledge has been gained from experiments utilizing them. Each of these techniques, however, suffers from a serious shortcoming. Enter the neurochip. The neurochip project was initiated to provide a new tool for investigators, to enable them to study neural networks in new and more powerful ways. This thesis details the development of a *caged neuron multi-electrode array*, a novel micromachined device that allows studies of the individual neurons of *in vitro* (cultured) neural networks over their entire lifetime, stimulating and recording from all constituent neurons in the network.

1.2 Conventional Electrophysiology Techniques—Shortcomings

To put the neurochip in proper context, and to better appreciate its utility, a brief survey of current techniques is presented.

1.2.1 Patch Clamp

The patch clamp electrode technique uses a micron-precision pulled-glass micropipette to interface with the neuron. A dizzying array of patch-clamp configurations is possible: perforated, outside-out, inside-out, and whole-cell. All of them rely on making a seal between the neuron cell membrane and a glass micropipette. This intimate contact invariably damages cell membrane, so the neuron is killed after a maximum of a few hours. The duration of an experiment, therefore, is very limited. These neurons can only be sampled at one time point—repeated recordings or chronic experiments are not possible. Additionally, the physical configuration of the patch-clamp system limits the number of cells that can be simultaneously patched. Manipulating two neurons simultaneously is common; a superstar patch-clamp performer might manage a third or forth cell, no more. In summary, the patch-clamp technique is limited both in time and in number of neurons manipulated.

1.2.2 Multi-Electrode Arrays

To overcome limitations inherent with the patch clamp method and the earlier sharp-electrode penetration, the multi-electrode array (MEA) was developed. The MEA essentially consists of an etched pattern of micron-scale metal electrodes deposited on a glass slide which doubles as the bottom of a culture dish. Neurons are cultured atop the electrode-containing glass.

In 1972, Thomas was the first to build and use MEAs. The Thomas lab fabricated MEAs using standard lithography methods to pattern nickel/gold electrodes. A second layer of photoresist served as an insulation layer. He and his colleagues extracellularly recorded electrical activity from cultured chick cardiac cells [62].

In 1977 Gross demonstrated that MEAs were capable of measuring neural activity, recording action potentials from snail (*Helix pomatia*) ganglia [28, 27]. To build his MEAs, Gross spin-coated an insulation layer over thin film metal electrodes. The insulation layer was subsequently "deinsulated" using a 600-nm, 0.5-mW laser. This photoetching method vaporized the insulation material over a small region to expose the metal electrode to saline.

This line of work culminated in 1980 with Pine demonstrating that MEAs could be used to both record and stimulate action potentials in dissociated neurons [52]. The fabrication of Pine's MEA was similar to that of Thomas, except that a silicon dioxide layer was used for insulation. Since this first demonstration that MEAs can *non-destructively* stimulate and record from dissociated neurons, MEAs have been used worldwide.

The main advantage of the MEA technique is that it is non-destructive: electrical activity can be recorded and stimulated without harming the nearby cell(s), so cultures can be tracked over time scales as long as months. The limitation with this tool, however, is two-fold: 1) There is no one-to-one correspondence between neurons and electrodes. Typically only a small fraction (on the order of 1/100) of cells are manipulated. Stimulation is not selective and can drive multiple neurons at unknown sites. 2) Neurons are migratory, so that the cells being manipulated may change over time. Both of these are undesirable for an ideal experiment to study cultured neural network development over time scales of several weeks.

1.3 Constrained Neural Networks

1.3.1 Patterned Substrates

To solve the problems with MEAs mentioned above, chemical patterning techniques were developed [72, 36, 6]. The basic idea is to chemically stamp adhesive/non-adhesive regions on top of the MEA so that neurons are allowed to anchor near the electrodes and grow along constrained paths. Typically poly-D-lysine is used as the adhesive substrate, with polyethylene glycol (PEG) as the non-adhesive substrate. Neural processes are constrained to grow along the geometry of the adhesive linear paths connecting the electrodes. This aspect of patterning is undesirable for studying development in cultured neural networks. Rather than being allowed to grow freely in two dimensions, axons and dendrites must follow a predefined geometry. In addition, chemical patterning is strong enough to spatially constrain neurons only for the first couple of weeks in culture.

1.3.2 Mechanical Traps

By 1988, mechanical traps were developed to constrain neurons in culture in an attempt to maintain a chronic one-to-one neuron-to-electrode correspondence. Regehr built micromachined cantilever "diving board" electrodes to interface to a neuron [55]. The essence of this technique was similar to the patch-clamping technique except that the diving board electrode was noninvasive. The diving boards had to be positioned by hand over a neuron, then glued to the bottom of the dish. This work demonstrated that snail and vertebrate neurons could be held in good electrical communication for up to four days. Nonetheless, this technique was not widely adopted due to the inherent awkwardness and limitation in placing the cantilever electrodes to interface with each cell.

Zeck and Fromherz [73] tried to trap neurons near an electrode by building a micromachined "picket-fences". Essentially six uniformly spaced post structures were built around a central electrode. They did manage to trap large snail neurons($\approx$ 100-μm radius), but never demonstrated their design or method could be adapted to much smaller vertebrate neurons ($\approx$ 10-μm radius). They don't explicitly say, but it can be inferred that the design and fabrication procedure are not adaptable to trap smaller cells.

Several years later, Maher and Wright [43, 44] developed bulk-micromachined "neurowells". Bulk silicon was etched to form a hole into which a single neuron was placed. A

mesh-like grill work was built atop the well to keep the neuron trapped. Maher demonstrated that this design was effective at trapping neurons, and that extracellular stimulation and recording were both possible. Fabrication of this device proved to be extremely challenging, and the yield was not good. Also, visibility of the neuron in the well was poor; the presence and health of a neuron were difficult to assess and could only be done by backtracking neurites to a particular location. This original neurowell-style microdevice, therefore, never garnered wide-spread popularity. However, the neurowell can be considered a direct predecessor to the neurocage-based neurochip described in this thesis. Note that the term *neurochip* appears to have been coined by Maher et al. and should not be confused with the current neurocage-based project.

1.4 Motivation for the Neurochip

The stated goal of this neurochip project was *to design and fabricate a micromachined device which remedies all of the problems with all of the techniques discussed above.*

In order to be successful, the neurochip must meet the following requirements:

1. Effectively trap neurons while allowing for network formation through unconstrained outgrowth of axons and dendrites.
2. Provide a bi-directional, noninvasive electrical interface with single-neuron specificity for every cell in a small (16-neuron) neural network.
3. Be simple and cost-effective to fabricate.
4. Potentially accommodate an arbitrary number of neurons.

This thesis describes the development and fabrication of such a device, and demonstrates that it meets all of the above requirements. Additionally, proof-of-concept experiments have been completed to demonstrate the full power of this neurochip for investigating network development with complete 1:1 correspondence over its lifetime.

1.5 Concept of the Neurocage

The neurochip presented here is a surface micromachined device consisting of a 4x4 array of *neurocages*. It is essentially a conventional MEA with micromachined structures—the

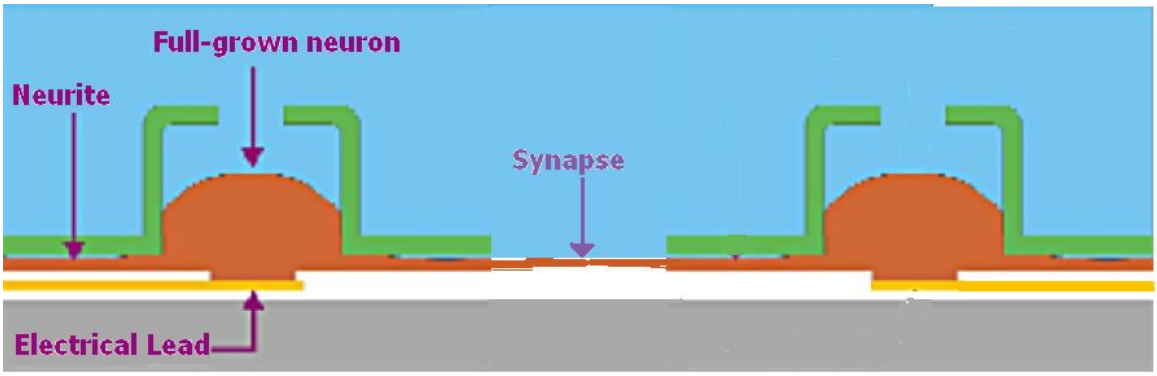

Figure 1.1: Cartoon drawing illustrating the concept of the neurocage. A neuron is trapped over its lifetime in the central cage region while axons and dendrites grow out through tunnels to synapse with other neurons. The nearby extracellular electrode is in bi-directional communication with the neuron.

neurocages—built over each electrode. Figure 1.1 illustrates the concept of the neurocage. This structure (green) mechanically traps a neuron (brown) in the central region near an extracellular electrode while still allowing axons and dendrites to grow outside to synapse and form networks with other neurons. Having trapped a neuron, a one-to-one correspondence is established over its lifetime. Furthermore, bi-directional communication—stimulation and recording—is maintained with *all* neurons in the network over the lifetime of the small neurochip culture. In this way, the investigator can gather a complete picture of the state of the neural network developing over time.

1.6 Organization of This Thesis

Chapter 2 provides a description of the final design and the fabrication process of the neurochip. Chapter 3 discusses electrical characteristics of the neurochip and how they relate to neural action potential measurement and stimulation. Chapter 4 demonstrates that survival on the neurochip is sustained at acceptable levels, and that neurons remain trapped in the neurocages. Chapter 5 covers stimulation issues, showing evidence that stimulation is safe and reliable. This chapter concludes with a brief theoretical note to compare with the experimental values. Chapter 6 describes extracellular measurement of action potentials. Theoretical considerations are presented, and I show that the actual recorded waveforms are explained well by the simple theory outlined. Chapter 7 is the climax: it demonstrates the ability of the neurochip to probe network connectivity over

several weeks. Chapter 8 presents results from initial experiments that were conducted to examine the natural development of synaptic connectivity in neurochip cultures—and to showcase the capabilites of this new tool. Chapter 9 concludes this thesis, highlighting the future of the neurochip and its potential uses.

Chapter 2

Neurochip Design and Fabrication

2.1 Introduction

This chapter provides a description of the final working design and fabrication methods of the neurochip and neurocage structure. Keep in mind when reading this chapter that the design and fabrication process were continually modified in smaller steps over a four-year period. For complete details of challenges faced during development please see Reference [63]. The following references also provide information on intermediate designs: [21, 19, 20, 18, 21, 65, 64, 47, 29]. The work was done in close collaboration with Angela Tooker, a graduate student in the Tai microfabrication lab at Caltech.

2.2 Neurochip and Neurocage Design

The final working design of the neurochip is presented here.

The heart of the neurochip consists of a 4x4 array of neurocages micromachined upon a standard, 0.5-mm-thick silicon wafer. The neurocage (or "cage", for short) is the basic repeated element, the structure that mechanically traps the neuron near the extracellular electrode.

There are three main parts to the neurocage structure. The final neurocage design includes a central cylindrical region, termed the *chimney*, for trapping the soma. Six surrounding *tunnels* serve as causeways between the interior of the cage and the exterior world (the rest of the culture substrate). Five *anchors* fasten the neurocage to the silicon substrate wafer.

The cage is made of parylene-C, a biocompatible polymer (for more details see next

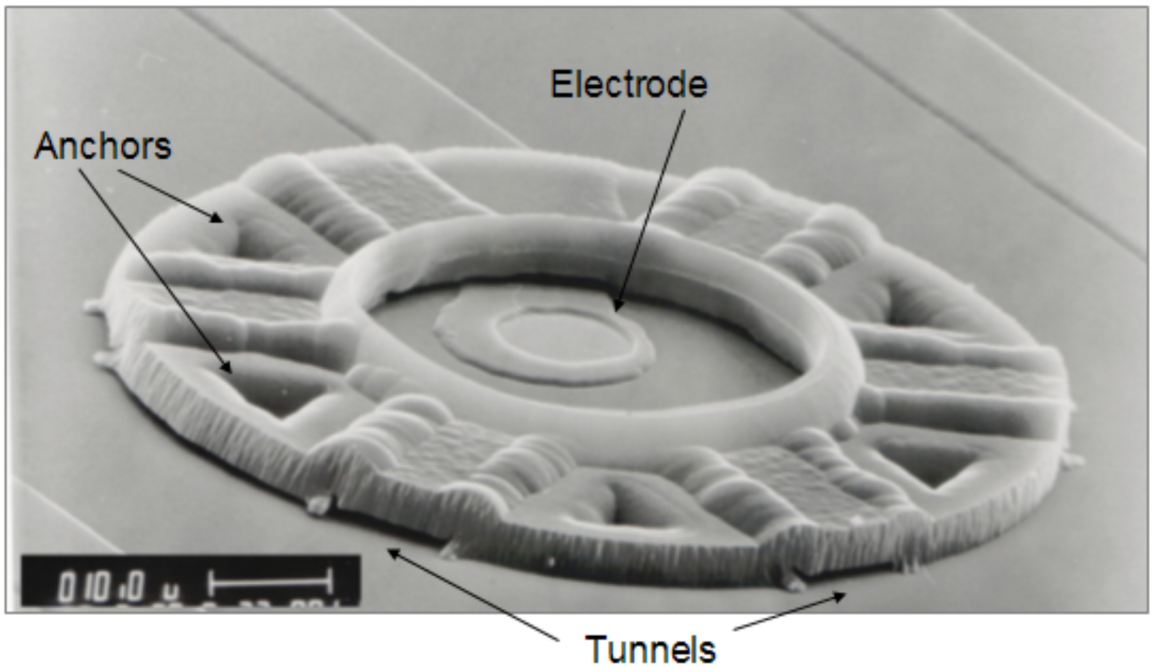

Figure 2.1: SEM of the final neurocage design. The major parts of the neurocage are labeled. A neuron is placed in the central chimney region, near the electrode. Axons and dendrites are free to grow though the tunnels to synapse with other neurons. The cage is made out of 4-μm-thick layer of parylene-C, a biocompatible plastic-like polymer. The electrode and lead are made of gold. Low-stress silicon nitride insulates the gold leads. Scale bar: 10 μm

section). The deposited parylene layer is 4 μm thick.

Two variants of cages have been fabricated, one with tunnels extending for 25 μm from the outer edge of the chimney, and the other with "no tunnels". In the latter case the effective length of the tunnels is 4 μm, equal to the parylene layer thickness of the chimney. In both cases the tunnel cross section is 10 μm wide by 1.0 μm high.

Figures 2.1 and 2.2 are scanning electron micrographs (SEMs) showing the anatomy of the 25 μm-long-tunnels variant from a side and overhead view, respectively. Figure 2.3 shows the anatomy of the no-tunnels variant.

The central cylindrical region (a "chimney", which became short over time) has a 40 μm inner diameter. It is 5 μm tall as measured from the silicon wafer to the lower lip of the access hole; or 9 μm tall measured to the top lip. A 30-μm-diameter access hole is cut in the top for loading a neuron.

The anchors stake the neurocage onto the silicon substrate. Anchors are formed by

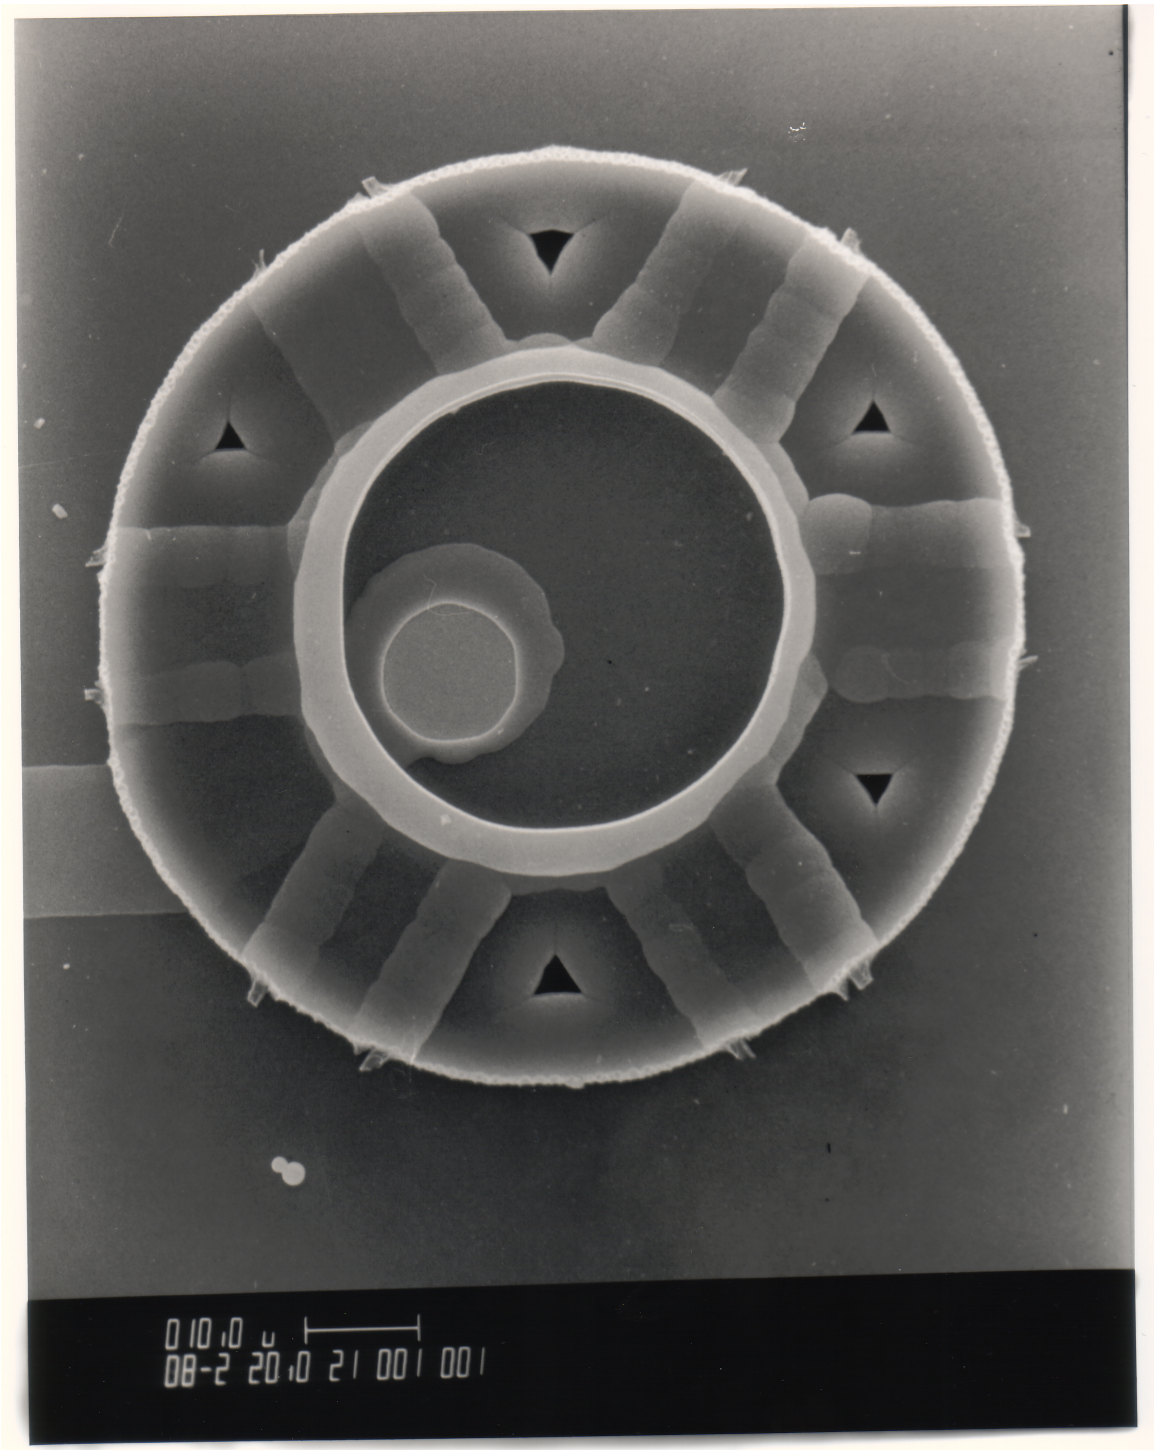

Figure 2.2: SEM of a neurocage in an overhead view. The electrode is offset from the center by 10 μm. Scale bar: 10 μm

filling cavities etched into the bulk-silicon with parylene. The cavities are 50 μm deep and shaped like an upside-down mushroom. At the silicon substrate surface they are roughly trapezoidal in shape and cover a total surface area of about 60 μm^2, as shown in Figure 2.1.

A gold electrode (later platinized), 10 μm in diameter, is positioned 10 μm off-center at the base of the neurocage. The electrical lead runs underneath the neurocage. It is 10 μm wide and extends for about 5 mm (on average, depending on the position of the neurocage), and terminates in a square-shaped bonding pad 0.5 mm on a side.

The insulation layer covering the leads is either a 1-μm-thick layer of low-stress silicon nitride, or 4-μm-thick layer of parylene-C. Both serve as good electrical insulators (see Section 3.5), but silicon nitride is highly preferable because it is compatible with fluorescence imaging. Parylene is not compatible with fluorescence imaging because it is autofluorescent at many wavelengths.

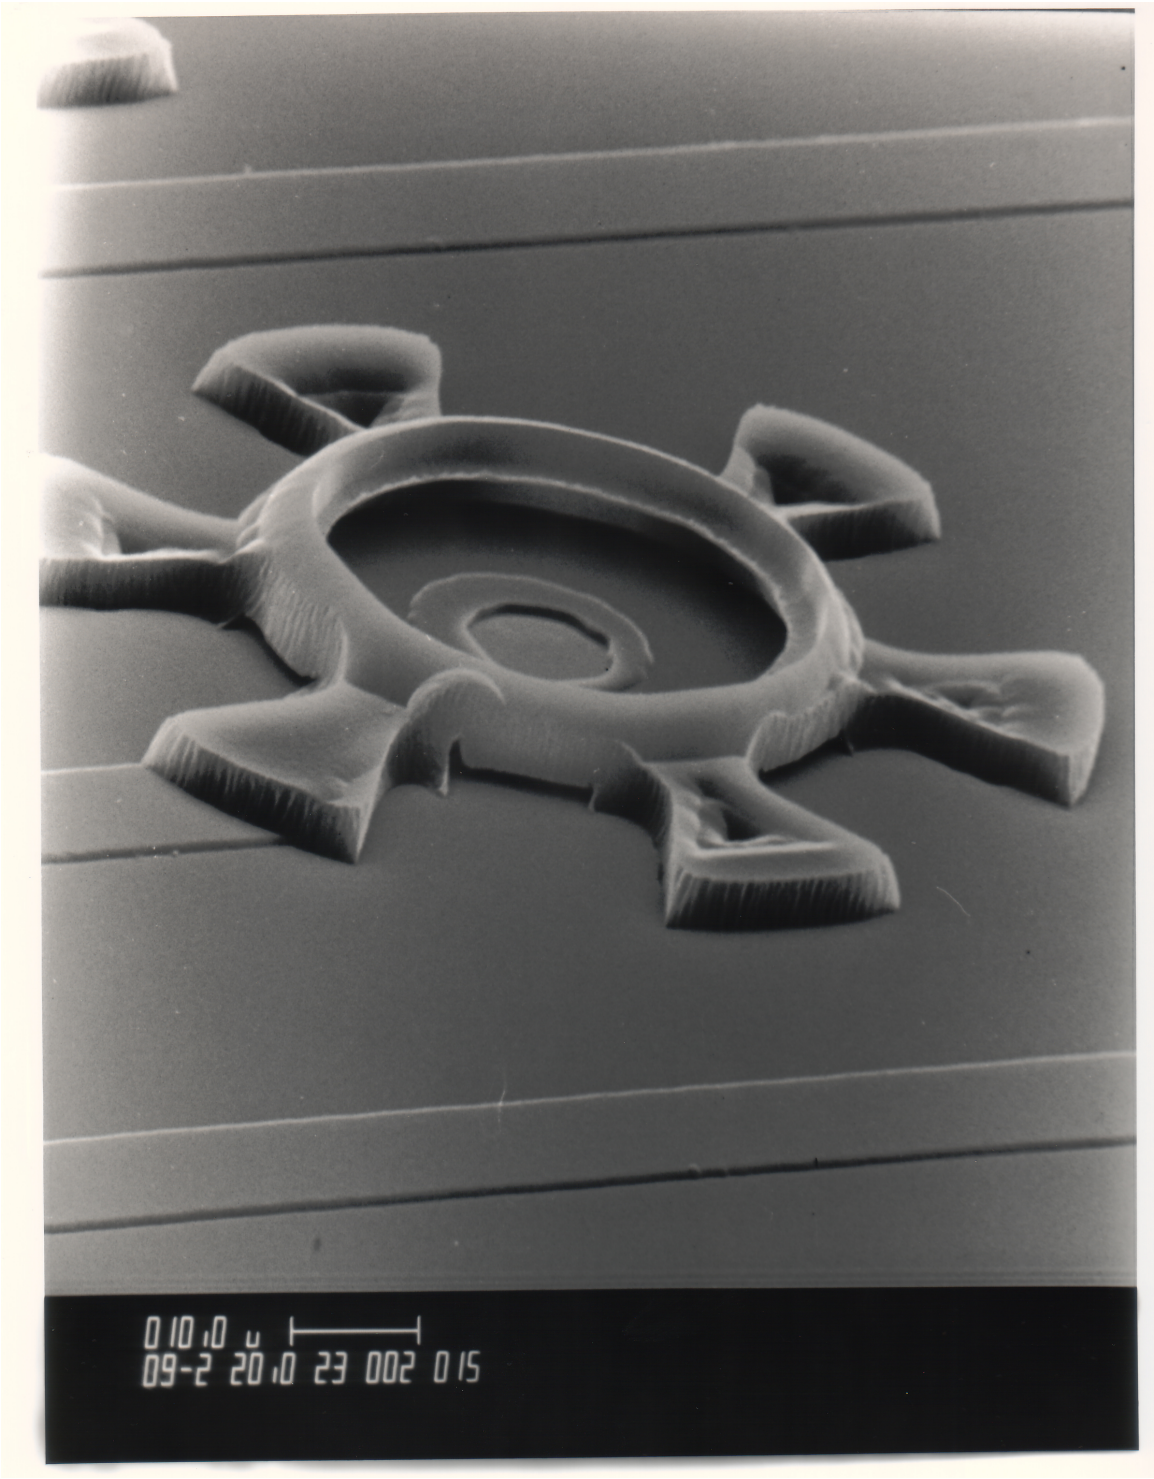

Figure 2.3: SEM of the "no tunnels" variant of the neurocage. Scale bar: 10 μm

2.3 The 4x4 Array

Neurocages are positioned on a square lattice, spaced 110 μm center to center. Figures 2.4 and 2.5 show the 4x4 array of neurocages. The array is positioned at the center of a 1-cm-square area of a silicon chip.

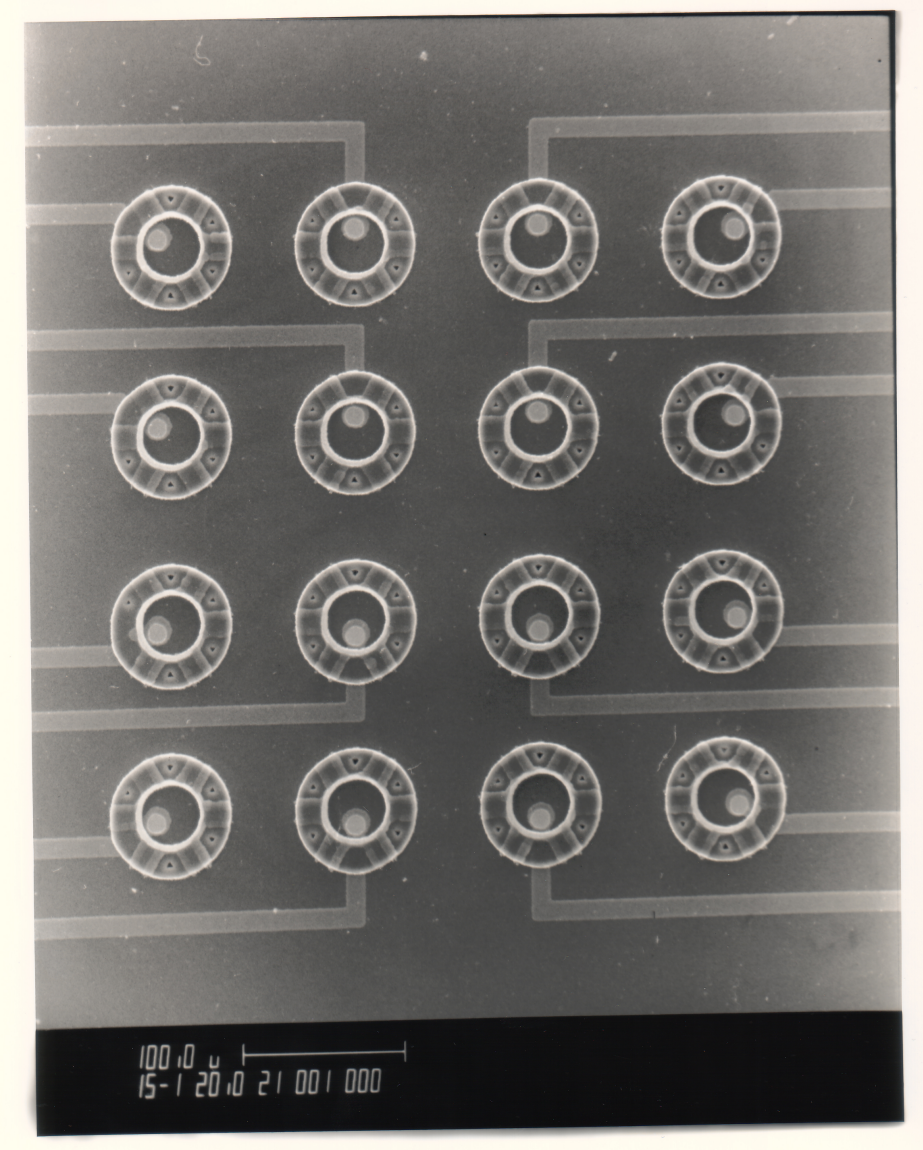

Figure 2.4: SEM of the 4x4 array of neurocages from an overhead view. Cages are spaced 110 μm apart. Scale bar: 100 μm

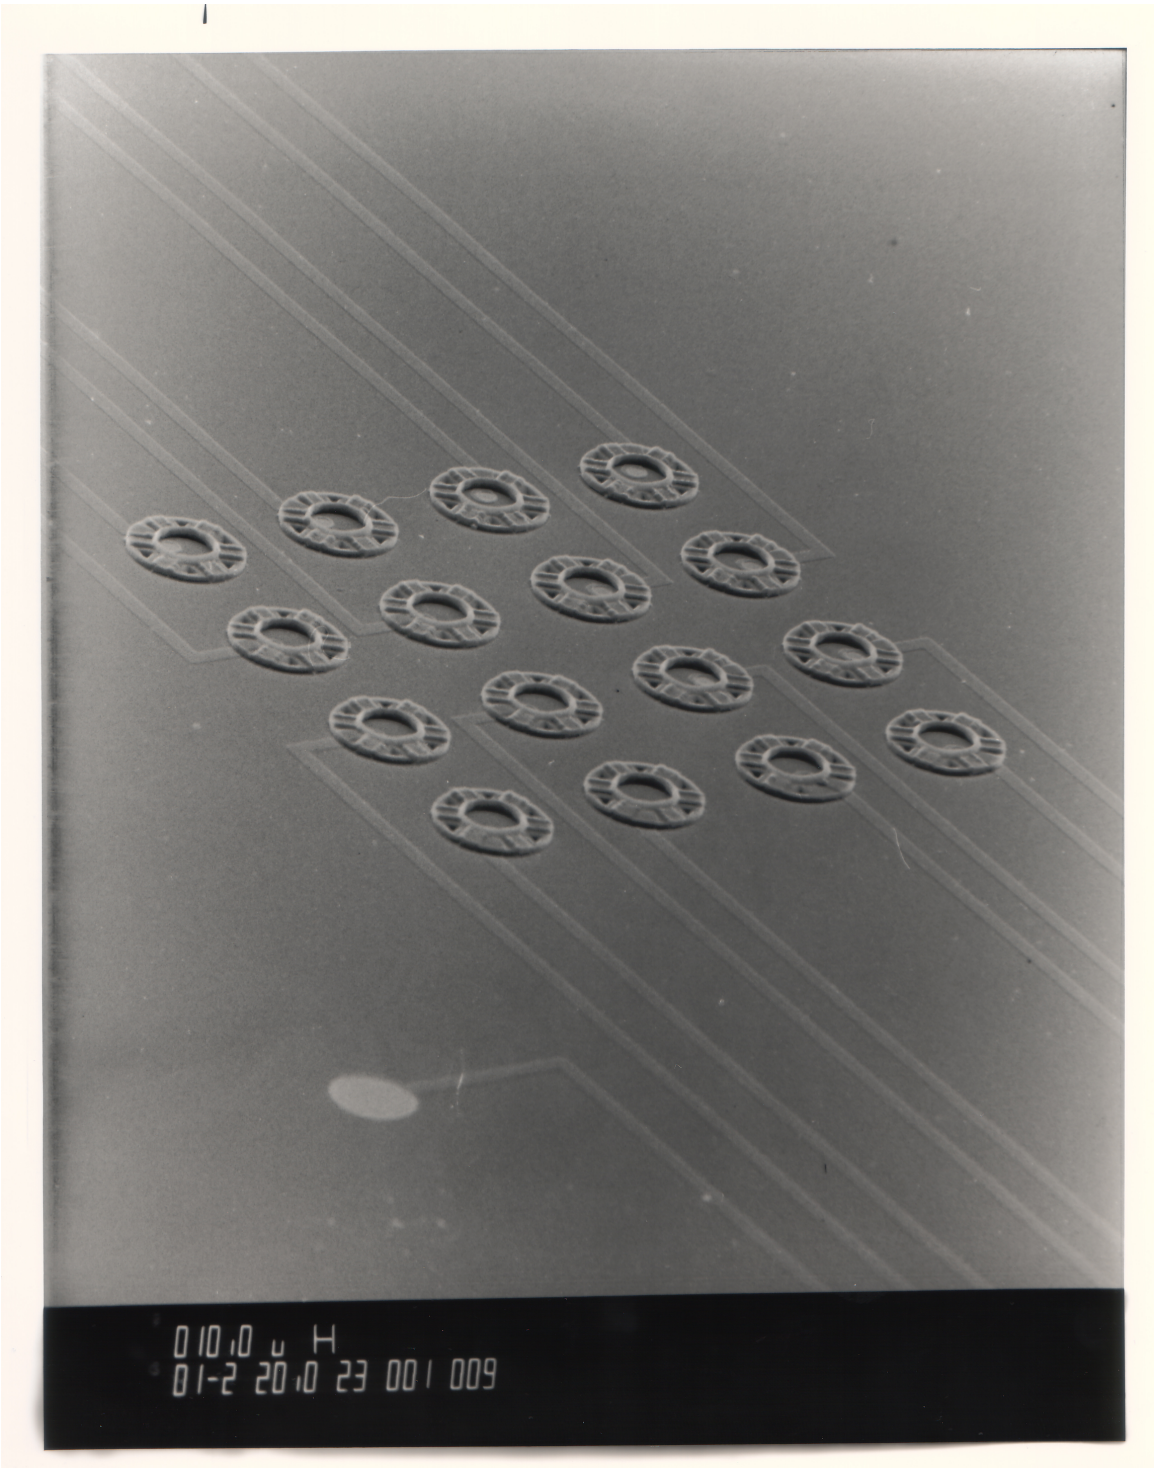

Figure 2.5: SEM of the 4x4 array of neurocages from a side view. The electrical leads run parallel out toward bonding pads. Cages are spaced 110 μm apart. The circular structure seen at lower left is a fabrication remnant of an unused reference electrode. Scale bar: 10 μm

2.4 Rationale for Dimensions of the Neurocage

The cage design was optimized over several years for trapping rat hippocampal neurons. I have not tried trapping other types of neurons, so it is not known if they are also successfully trapped in this design. Given the dimensions of the cages presented above, it is natural to ask the question "Why that design?" Ask and ye shall receive:

Cage Height: The chimney height dimension is not critical to keep a neuron trapped, but it appears to be critical for survival of neurons. Early experiments attempted with a much taller, 20 μm chimney structure never sustained neuronal growth (for unknown reasons). Sustained growth was noted in the neurocages only after decreasing the chimney height to the current, final height of 5μm. No neuron, however, has ever been seen to escape from a cage by traveling over the top. As the neuron grows, it flattens out along the base of the cage. Since axons and dendrites adhere and extend laterally along the neurochip substrate, no vertical tension develops to pull a migrating neuron up and out of the cage. The lateral tension sometimes pulls the soma snug against the cage sidewall.

Cage and Electrode Diameter: The cage diameter was chosen to be accommodate the size scale of cultured rat hippocampal neuron. The soma of a cultured hippocampal neuron is typically about 20 μm in diameter. With the electrode occupying 10 μm diameter in cage real estate, the neuron can spread out and grow in a wedge shape on roughly a 30 μm size scale. The electrode is positioned off-center to enhance visibility of growing neurons, or of their death.

Tunnel Height: Tunnel height is critical in keeping a neuron trapped. Early designs employed taller tunnels, 1.8-μm-high. About 20% of neurons loaded into these cages escaped by migrating through the tunnels. Subsequent reduction to 1.3-μm-high tunnels still allowed about 10% of neurons to escape in the same manner. The 1.0-μm-high tunnels reduced escape to levels essentially zero.

Tunnel Length: Survival has not been found to depend on tunnel length. Both variants have sustained cultures for several weeks at the same survival rates (see Section 4.5). Perhaps surprisingly, both variants trap neurons equally effectively. Regardless of tunnel length, cages with 1.0-μm-high tunnels trap neurons over time at a > 99 % level.

2.5 Microfabrication Summarized

A summarized description of the microfabrication of the neurochip is presented in this section. The entire process is depicted in cartoon form in Figures 2.6–2.8. Note that the cartoon drawings are not to scale. The basic idea can be broken down into the following steps:

1. Build a standard MEA on a silicon wafer.

2. Cast a form for the neurocages out of aluminum and photoresist sacrificial layers.

3. Deposit the parylene, etch it to form the cages, and dissolve the form.

A single wafer yields about twenty-two neurochips.

More specifics follow below. Complete details can be found in [63].

Electrodes: Thin film gold electrodes and leads are patterned using the "lift-off" method on a silicon-dioxide-coated substrate wafer. Sixteen electrodes are placed on a square lattice, spaced at 110 μm, with leads extending to the bonding pads at the edge of the chip. Next, an insulation layer (1 μm thick) of low-stress silicon nitride is deposited and selectively etched with SF_6 to expose the gold electrodes.

Anchors: Once the electrodes and leads have been patterned, anchor holes are chemically "drilled" into the silicon substrate in the shape of inverted mushrooms with a repeated Bosch process. They will later be filled with parylene. The parylene-filled head of the inverted mushroom mechanically stakes the neurocage to the substrate silicon wafer.

Cages: A 1-μm-thick sacrificial layer of aluminum is thermally evaporated and patterned with a photoresist mask to form the star-burst shape of the tunnels. A layer of photoresist is spun-on and patterned with standard photolithography to form the central chimney structure. A layer of parylene—biocompatible polymer—is *conformally* deposited over the aluminum and photoresist sacrificial layers. This property of parylene is the key to the fabrication method. During this deposition, the anchor holes are filled with parylene as well. The parylene layer is subsequently etched with oxygen plasma, and the sacrificial layers are dissolved, leaving behind the final structure.

Again, note that all of these steps are accomplished using standard surface lithography methods. This virtues of this fabrication process are:

1. Neurocages are straightforward to fabricate; no bulk micromachining or backside alignment is necessary.

2. An arbitrary number of cages can be fabricated in arbitrary positions (in a hexagonal array, for instance).

3. The geometry of the neurocages can be easily adapted to accommodate working with different cell types.

2.6 Neurochip Yield

One neurochip occupies a 2 cm x 1 cm region diced from a 4-inch-round silicon wafer. The business end of the neurochip—the cages and electrical connections—actually fits on a 1-cm-square region. The other 1 cm square of the 2 x 1 cm diced chip is intentionally left blank for compatibility with manipulating individual neurons into cages. (See Section 4.2.2 for more on "loading" neurons into cages.) One wafer yields about twenty-two neurochips. Up to six wafers can be processed at once in the Caltech micromachining facilities, so one run generates 132 neurochips. A brisk pace of culturing neurochips requires that about 20 neurochips stay in circulation at any one time. That gives a rotation of 5 chips to be used spread out over 4 weeks. Chips can stay in circulation for at least 5 rounds of culturing, so about 20 chips are used per 6 months, or 40 per year. One run through the microfabrication procedure takes about 2 weeks, and generates a supply of neurochips that should last about 3 years.

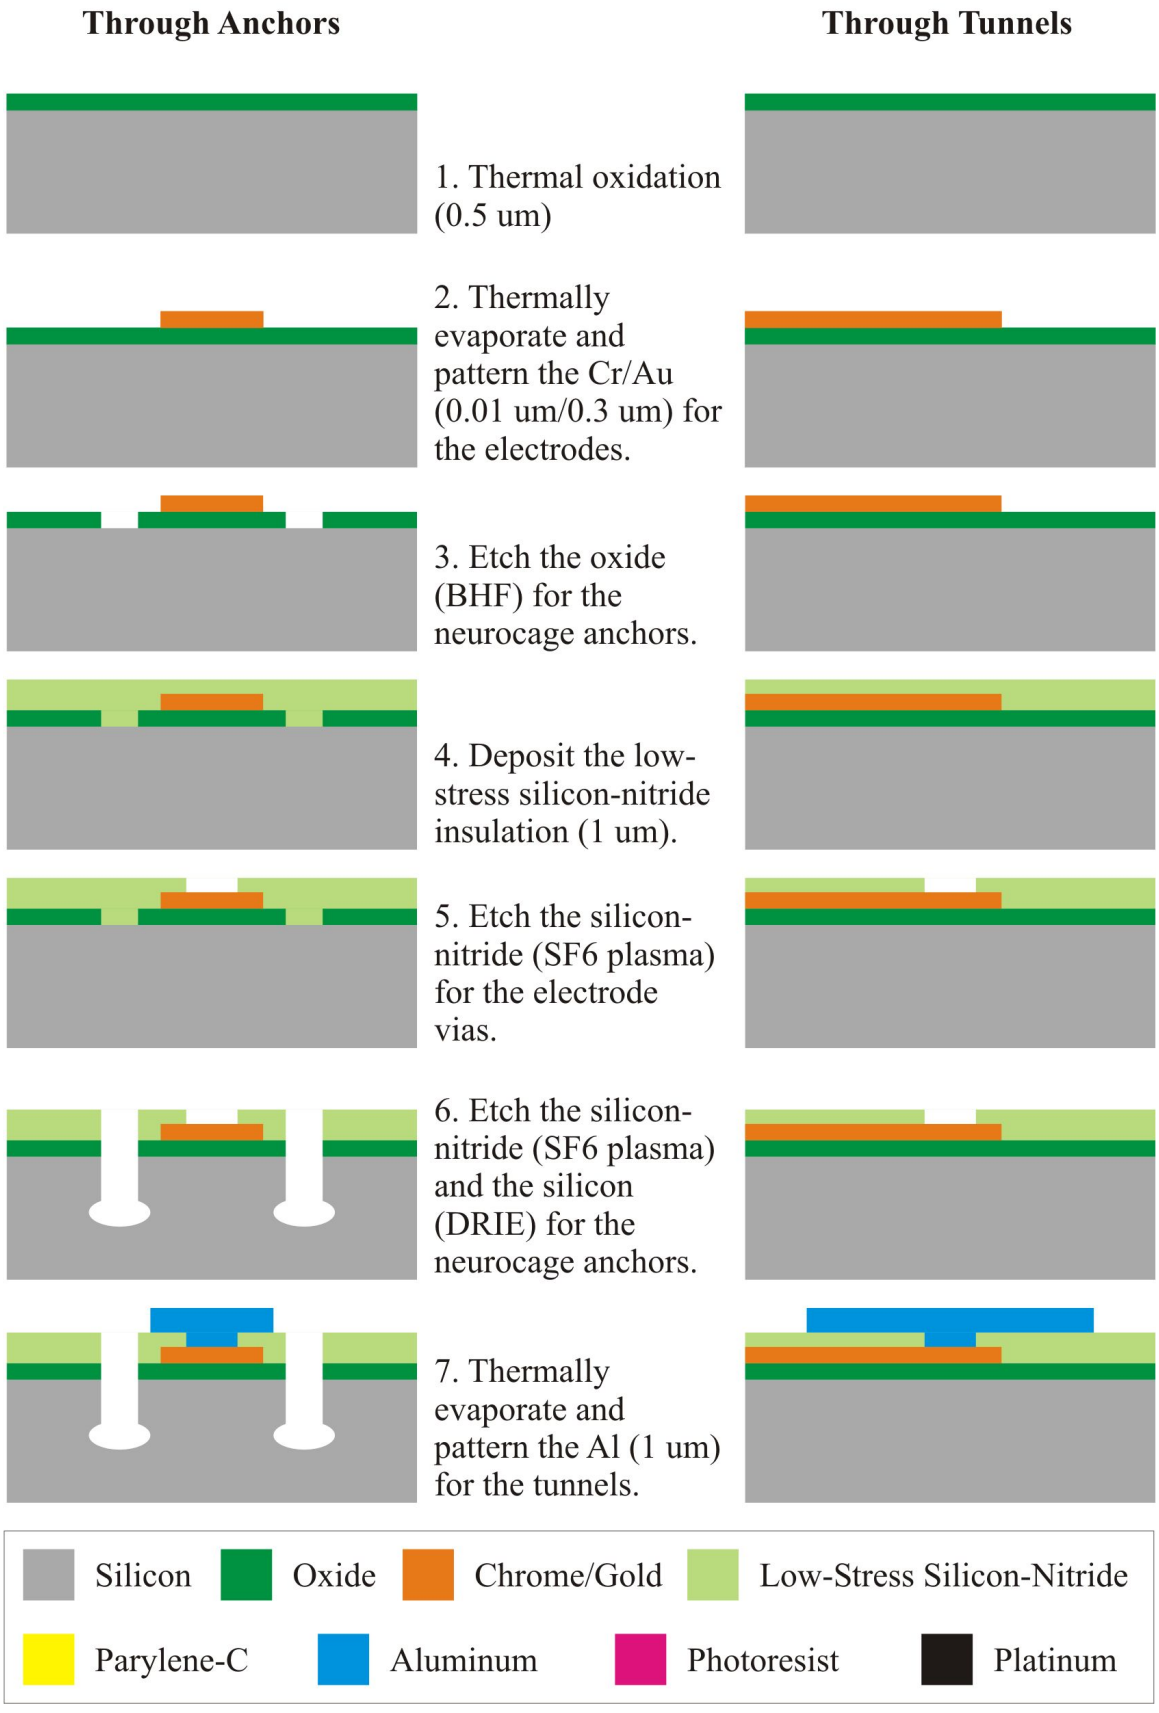

Figure 2.6: First 7 steps of the neurochip fabrication process

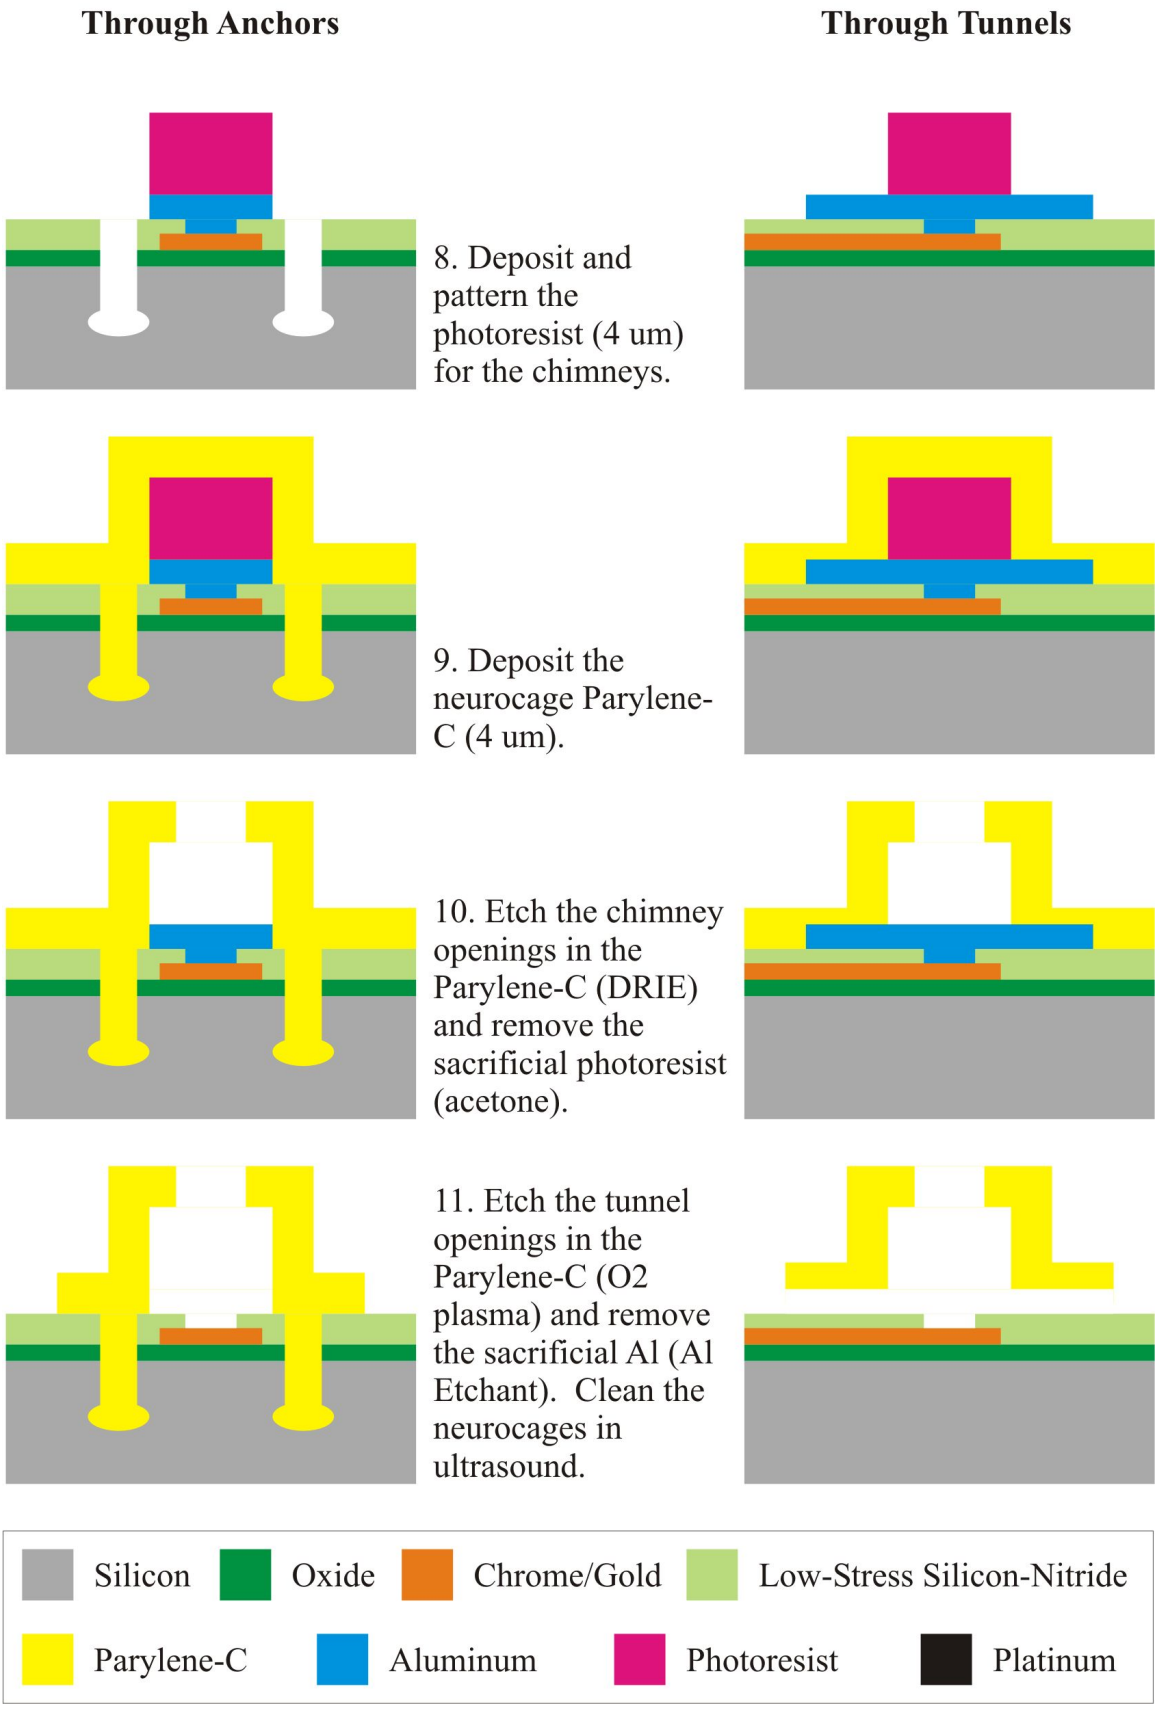

Figure 2.7: Steps 8–11 of the fabrication process

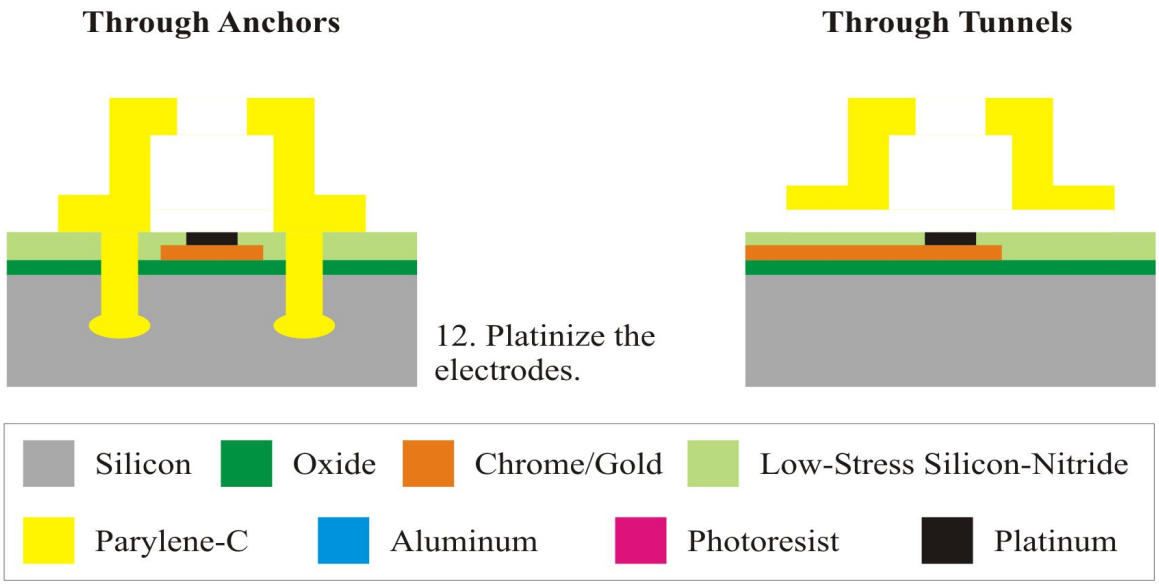

Figure 2.8: Final step of the fabrication process

2.7 Final Assembly

The end result of micromachining produces a large number of 2 cm x 1 cm neurochips on silicon wafers. Further assembly is required to interface the neurochip to external electronics and to prepare it for cell culture. Detailed, step-by-step assembly instructions can be found in Appendix A. Briefly, a single neurochip is glued into a custom-designed PC carrier board. The board is machined so that the neurochip sits flush in a milled pocket. Ultrasonic gold wire-bonds are made between the neurochip lead-bonding pads and carrier-side bonding pads. These bonds are covered in silicon epoxy. The epoxy serves to electrically insulate the wires, as well as add protection for the relatively fragile bonds. A 28-pin carrier is then soldered into the PC board. A platinum wire—which doubles as both the reference and current return electrode—is wrapped around the ground pin and electrically connected with silver epoxy (platinum does not solder). Finally, a 35 mm cell culture dish with a custom-milled window is then sealed over the top of the neurochip with silicone elastomer. The fully assembled chip is pictured in Figure 2.9. Upon close inspection the neurochip electrical leads are visible, as is the neurocage array. The leads appear as a bundle of parallel lines running roughly vertically. The neurocage array is the square dot in the middle of the left side of the chip.

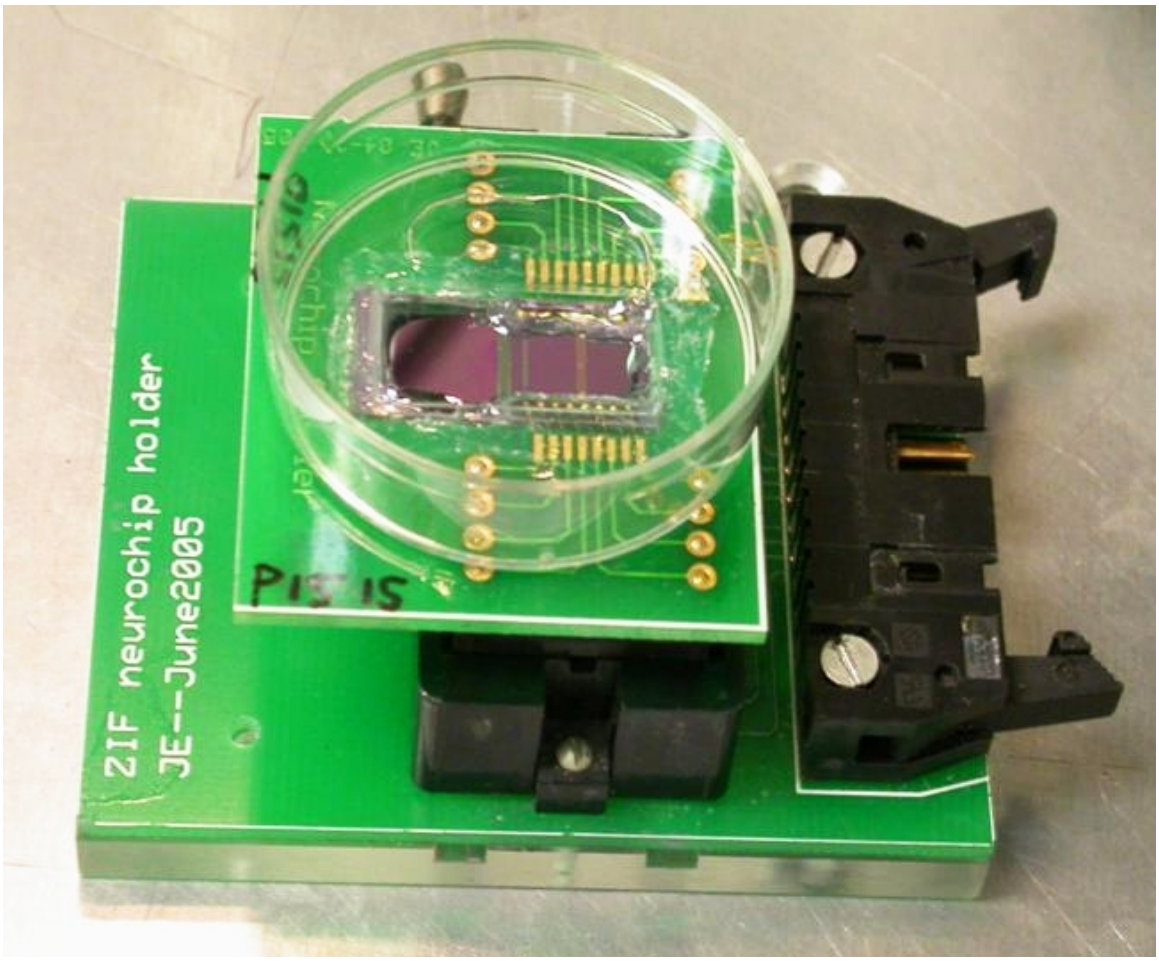

Figure 2.9: Fully assembled neurochip. The 2 cm x 1 cm neurochip diced from the silicon wafer is glued into a PC board. Gold wire bonds connect the neurochip to the PC board terminals. A 28-pin carrier is soldered into the board for connecting to a zero insertion force (ZIF) socket. A 35-mm dish with a custom-machined window is sealed over the top of the neurochip with Sylgard. A platinum wire is glued to the bottom of the culture dish. Electrical leads on the neurochip are visible as parallel bundles of lines running roughly from 1 o'clock to 7 o'clock. The neurocage array is also visible upon close inspection: it is the dot in the middle of the right half of the neurochip.

Chapter 3

Neurochip Electrodes and Hardware Interface

3.1 Introduction

This chapter describes the electronics and hardware for stimulus generation and low-noise recording of action potentials. We'll start with neurochip electrodes and work backwards through pre-amplifiers, buffer amplifiers, and finally to the computer controlling the whole system.

3.2 Simple Model for Neurochip Electrodes

The neurochip electrodes can be represented by a simple model as shown in Figure 3.1 below. A quick survey of the components: The electrode at the base of the cage is modeled as a lumped capacitor. The path in the resistive medium from the electrode to the top of the cage is modeled as a resistor, as is the "spreading" resistance from the top of the cage out to ground. The electrode lead-insulation layer-conductive saline constitutes the shunt capacitance. The origin and justification for the model is presented in the rest of the chapter.

3.3 Electrode Impedance

A neurochip electrode—the element which senses voltage drops present in physiological saline due to neural events, i.e., action potentials—is made out of gold. (The signal to be measured is on the order of a few tens of microvolts; see Chapter 6.) This voltage signal

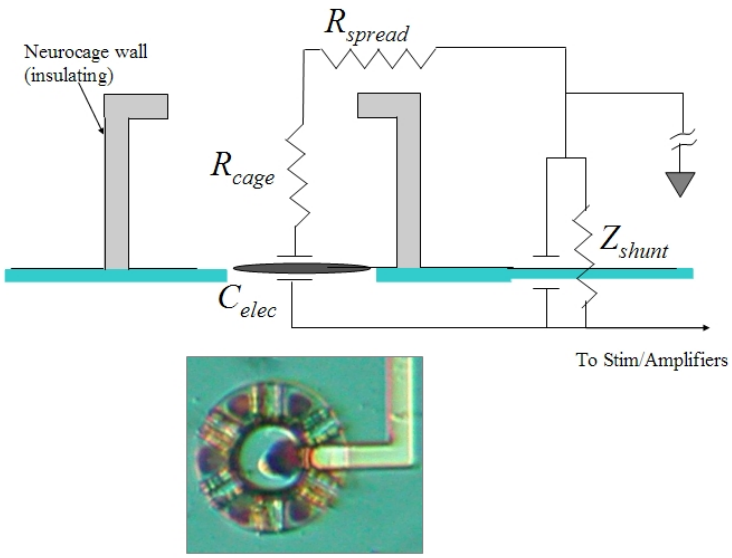

Figure 3.1: Cartoon neurocage model illustrating simplified electrical model of the neurocage electrical connections. An actual neurocage with platinized electrode and electrical lead extending out of the cage is shown below. Mean measured values, averaged over 48 electrodes, are: C_{elec} = 4300 pF; $R_{cage} + R_{spread}$ = 25 kΩ; C_{shunt} = 20 pF ; R_{shunt} = 5 MΩ. See text for discussion of each electrical element.

must, of course, be conducted from the saline onto the neurochip electrode. The interface between saline solution and a noble metal constitutes an electrolytic capacitor [56]. The capacitance property derives from electrochemical reactions which rapidly pass electrons into or out of the metal. Net electron flow due to Ohmic mechanisms is extremely small, and the DC resistivity of the gold-saline interface has a value of $\approx 10^6$ MΩ · μm^2. For a 10-μm-diameter electrode, the expected resistance is 10,000 MΩ. The capacitive reactance is much smaller, and therefore, the neurochip electrode–culture medium interface can be treated as a lumped capacitor. However, this capacitance, C_{elec}, is frequency dependent, and typically varies as $C_{elec} \propto 1/\sqrt{f}$, where f is the frequency of the signal to be measured [56]. For neurobiological signals, such as an action potential, the frequency of interest is about 1 kHz. (All values for impedance described in this section are determined at 1 kHz.) The method for the impedance measurement is described in Appendix B.

The capacitance of a noble metal is small, estimated to be about 0.2 pF/μm^2 at 1 kHz [56]. The corresponding theoretical capacitance of a 10-μm-diameter neurochip electrode would be about 15 pF. A measured value (N = 48 electrodes) was $C_{elec} = 22 \pm 9$ pF. The measured value might be higher than expected because the diameter of the gold electrode

is slightly greater than the target radius. Still, the unplatinized electrode capacitance is unacceptably small on two accounts. For one, the action potential signal must be transferred to amplifiers through several feet of ribbon cable. An estimate for the total stray capacitance of the ribbon cable and neurochip wiring is about 100 pF. So, with only about 20 pF capacitance for the electrode, the signal to be measured would be significantly attenuated by a factor of about 5. The expected size of the action potential signal is about 30 μV, so would be measured as a 6 μV signal. That's already too small for what can be cleanly recorded (see Section 3.7) Another reason 20 pF is too small has to do with the charging of the electrode when current stimulus is passed through it. The electrode will charge to a voltage inversely proportional to the electrode capacitance. Given the current stimulus parameters necessary to stimulate a neuron (see Chapter 5), a 20 pF electrode would charge to a voltage far beyond the one volt safety limit.

But never fear! It is possible to greatly increase the value of the capacitance by platinizing the electrode.

3.4 Platinization: Electrode Capacitance

Platinizing an electrode involves electroplating a spongy layer of platinum black onto the gold surface. Platinization increases the electrode capacitance (decreases the impedance) by increasing the surface area of the saline–metal interface. The capacitance is proportional to the surface area of the saline–metal interface. Platinizing has been found to increase the the electrode capacitance by a factor of 100 or more.

It is desirable to grow the platinum "bush" as large as possible to maximize the electrode capacitance, but it is important to keep it confined for several reasons. Firstly, platinum black is not compatible with reflection Nomarski optics of the Olympus BHMJ scope used to view neurons for all neurochip experiments. The platinum black is a visual black hole—nothing directly above it can be imaged. (This also makes loading neurons extremely challenging.) Furthermore it is extremely difficult, if not impossible, to assess viability of a neuron growing on top of platinum black. Second, if the platinum black expands too far out of the electrode region it can block the tunnel exits on the cage interior. This blockage denies an elongating neural process access to grow out through these tunnels. Finally, soma, axons, and dendrites avoid growing over the top of the platinum bush. In any case, a larger

platinum bush reduces the area over which a neuron can adhere. It seems prudent to leave the neuron ample space to grow.

Electrodes were platinized by applying a DC current density of 318 mA/cm^2 for 5 seconds. Electrodes were platinized once daily, then left to sit overnight in ddH20. Complete details on the platinizing protocol can be found in Appendix B. I repeatedly platinized, usually about 5–8 times, stopping when the platinum black bush started to impinge on the rest of the cage volume. Immediately after the final platinization the the electrode capacitance was measured to be $C_{elec} = 6600 \pm 900$ pF, (mean $\pm$ s.d.) measured over 48 electrodes. That represents a factor of 157 increase over the unplatinized value.

The electrode capacitance, however, is reported to decrease over time to 50% (or more) of the intial value when soaked in cell culture medium [43, 62]. This decrease might result from lipids and sugars packing into the spongy platinum black matrix which decreases the effective surface area. I soaked neurochips in cell culture medium, without neurons plated, for 2 weeks and measured the electrode capacitance one last time. The capacitance had indeed decreased to $C_{elec} = 4300 \pm 800$ pF, a 35% decrease. Even so, this value is still sufficiently high that signal attenuation is negligible, and stimulation is expected to be—and demonstrated to be—safe.

The platinization has been demonstrated to be robust: Through five (or more) rounds of experiments, the electrode capacitance does not degrade. The electrode capacitance was measured one final time on neurochips (N = 4) that had been previously used for 5 (or more) cell-culture sessions. A total of about 250 stimuli at a current density of 7000 mA/cm^2 were applied to each electrode during the course of these experiments. Chips were cleaned (according to the protocol in Section 4.3.4) after each round of a cell culture. Subsequently, the capacitance was measured to be 6400 $\pm$ 800 pF. Curiously, the measured real-part of the impedance increased to 36$\pm$2 kΩ. The cause for this increase is not clearly understood. (It could be due to a narrowing of the effective radius through which it conducts, as a result of cell debris which is not able to be throughly cleaned. This narrowing would have a minuscule effect on the total surface area of the sponge-like platinum bush; hence it would have very little effect on the capacitance.) Nonetheless, this increase in resistance does not significantly change the functioning of the electrode. That the capacitance can be restored to its initial value suggests the decrease noted after soaking for two weeks in culture medium is probably due to lipids and sugars packing into the platinum black matrix. Moreover, it

indicates that the stimuli did not damage the platinization.

The history of electrode capacitance is summarized in Table 3.1.

Table 3.1: Capacitance and corresponding reactance values for neurochip electrodes (mean ± s.d. for N = 48 electrodes). All values determined at 1 kHz. See text for details on the meanings of the "Electrode Status".

Electrode Status	C_{elec} (pF)	X_c (kΩ)
Gold Electrode (no platinization)	22 ± 9	7400 ± 800
After platinization, before soak	6600 ± 900	24 ± 3
Platinized, after 2 weeks soak	4300 ± 800	37 ± 7
After cell culture, application of stimuli, and cleaned	6400 ± 800	36 ± 2

3.4.1 Cage and Spread Resistance

The confined geometry of the neurocage renders the electrode-to-ground resistance non-negligible. Following the diagram above, the resistance can be broken into two stages: getting from the bottom to the top of the cage, R_{cage}; and spreading from the cage out to a distant ground, R_{spread}.

The cage resistance can be computed by considering the definition of resistance applied to the neurocage geometry:

$$R_{cage} = \int_a^b \frac{\rho(x)}{A(x)} dx \tag{3.1}$$

where x is the axial coordinate of the cylindrical cage, and ρ is the resistivity of the medium, taken to be a constant 70 Ω·cm.

The off-center geometry makes exact computation of R_{cage} difficult, but a reasonable approximation is given by considering that current flows roughly through a conical cross section toward the top of the cage. With this approximation, it is easy to show that:

$$\begin{aligned} R_{cage} &= \frac{\rho h_{cage}}{\pi\, r_{elec} r_{ah}} \\ &= \frac{(70\,\Omega \cdot cm)(10\,\mu m)}{\pi\,(7\,\mu m)(17\,\mu m)} \\ &\approx 19\,\mathrm{k}\Omega \end{aligned}$$

where h_{cage} represents the total height of the cage, from the silicon nitride insulation to the top lip of the access hole; r_{elec} is the radius of the electrode; and r_{ah} is the radius of the

access hole.

In the calculation above I used $r_{ah} = 17\ \mu$m, instead of 15 μm, because SEMs show that the access hole is over-etched by a few microns. Thus, it is reasonable to use a slightly larger value for the radius of the access hole. Also, I used $r_{elec} = 7\ \mu$m, instead of 5 μm, because the platinum layer extends a few microns outside of the patterned gold area, increasing the effective radius of the electrode.

The resistance from a small source spreading isotropically through a spherical geometry out to infinity (or at least a very distant ground) is computed as:

$$\begin{aligned} R_{spread} &= \int_{r_{ah}}^{\infty} \frac{\rho}{4\,\pi\,r^2} dr \\ &= \frac{\rho}{4\,\pi\,r_{ah}} \\ &= \frac{\rho}{4\,\pi\,(17\,\mu m)} \\ &\approx 5\,\mathrm{k}\Omega. \end{aligned}$$

Then the total theoretical resistance (real part of the electrode impedance) is 19 kΩ + 5 kΩ = 24 kΩ. effective radius of the electrode.

The measured real part of the impedance of a fully platinized electrode, averaged over 48 electrodes, is (mean $\pm$ s.d.): $R_{total} = R_{cage} + R_{spread} = 25 \pm 3$ kΩ. That value is in excellent agreement with the theoretical value.

3.4.2 Tunnel Resistance

The effect of the path from the inside to outside of the cage through the tunnels has been neglected to this point. For the long-tunnels version of neurocages, the resistance per tunnel should be about 2 MΩ. Six of them in parallel, make for a resistance of 333 kΩ. Since this value is much larger than $R_{cage} + R_{spread}$, the resistance of the tunnels is only a 5% correction. Furthermore, adding in the current path through the tunnels is only important if current actually flows there. Given that the cross section has a very low profile, it is reasonable to approximate that for current delivered through through the electrode during a stimulus, current primarily flows upwards, towards the top of the cage, not laterally through the tunnels. Evidently, this is actually the case: the measured resistance for the no-tunnels version of the neurocage is very nearly the same as the long-tunnels version,

indicating very little current flows through the low-profile (1-μm-high) slits at the bottom.

3.5 Shunt Impedance

It is important that the shunt impedances be as large as possible for high-fidelity recording of action potentials, and to avoid losses of the current delivered to the electrode. The neurochip system achieves the goal of having a shunt impedance much greater than the electrode–cage–ground impedance.

3.5.1 Shunt Capacitance of the Leads

Shunt capacitance, C_{shunt}, on the neurochip derives from two conducting materials being separated by a thin dielectric layer. The two conducting layers are the electrode lead and the saline solution. The thin dielectric separating them is the thin layer of insulation. This configuration essentially forms a parallel plate capacitor. Any uninsulated portion of the lead (direct gold–saline contact) also adds to the shunt capacitance (and shunt resistance). On older, reused chips this can be a concern [68].

Shunt impedance attenuates the neurobiological signal being measured. To avoid significant signal attenuation, C_{elec} must be much greater than C_{shunt}. A back-of-the-envelope calculation for the shunt capacitance is:

$$\begin{aligned} C_{shunt} &= \frac{\kappa \epsilon A}{d} \\ &\approx \frac{(3)(8.85\ \mathrm{pF/m})(5\ \mathrm{mm})}{(1\ \mu\mathrm{m})} \\ &\approx 1\ \mathrm{pF}. \end{aligned}$$

The value of the shunt capacitance was measured by placing a small drop of glue over the neurocage array to prevent conductive fluid from entering the neurocage and contacting the electrode, thus isolating the length of the leads as the sole conductive path. The measured value of the shunt capacitance on 1-μm-thick silicon nitride insulation neurochips was $C_{shunt} \approx 20$ pF. For 4-μm-thick parylene insulation, the measured shunt capacitance was 40 pF. While it is not clear what accounts for the difference between theoretical and measured values, indeed $C_{elec}/C_{shunt} \approx 100$. The signal loss, therefore, would only be on

the order of 1%, so the shunt capacitance can be considered negligible.

3.5.2 Shunt Resistance

Notice that electrode capacitance in combination with the shunt resistance to ground forms a high-pass filter, the time constant given by $\tau_{hpf} = C_{elec}R_{shunt}$. The timescale of recorded action potential signals is expected to be about a 0.5–2 msec, so it is important that $\tau_{hpf} \geq 5$ msec to avoid differentiating the signal significantly. Having measured the shunt impedance, the shunt resistance is calculated as the real part of the impedance. The measured value for 1-μm-thick nitride and 4-μm-thick parylene were both on the order of 5 MΩ. The time constant is calculated to be about 20 msec, a timescale very long compared to action potentials. Therefore, no signal differentiation is expected to occur. Furthermore 5 MΩ is large compared to the cage and spread resistances, so can be neglected as well.

One last note on parylene insulation: originally the parylene insulation layer was 2 μm. Experiments revealed that over time the insulation was failing as the shunt impedance levels dropped to about 500 kΩ in about two weeks, an unacceptably low level. Micrographs revealed that a 2-μm-thick parylene deposition can have pinhole defects. This allows fluid to very slowly seep underneath over time, and accounts for the large decrease in shunt impedance. 4-μm-thick depositions are immune from this problem.

3.6 External Hardware and Electronics

3.6.1 Recording: Pre-Amplifiers and Buffer Amplifiers

All signals were recorded using a 16-channel pre-amplifier and two 8-channel modules consisting of filters and buffer-amplifiers.

All neurochip electrodes were referenced to the same platinum ground wire, connected to power-supply ground. The fully assembled neurochip culture dish plugs into a ZIF socket. The electrode signal was transferred to the pre-amp board through about 8 inches of ribbon cable. Each electrode was AC-coupled into the pre-amp to guarantee no net charge is transferred through the electrode.

The 16-channel pre-amp board consists of simple x11 low-noise, non-inverting TL027 op-amps [55].

The signal output of the pre-amplifier drives a 2-pole low-pass filter set for 5 kHz. The low-pass filter output is AC-coupled into the buffer amplifiers. The buffer amplifiers have programmable gain. For a stimulus channel the gain is set to G = 1 so that saturation of the electronics will not occur until the electrode voltage is about 1V. For a recording channel the gain is set to G = 63.

3.6.2 Digitization/Computer Data Acquisition and Control

The output of the buffer amplifier is input to a 12-bit A/D card, National Instrument 6071-E. This A/D card has selectable gain between 1 and 100 and a dynamic range of 10V. A gain of 1 was used for a stimulus channel; saturation occurs at 1V (amplitude of neurochip signal). A gain of 50 was typically used for a recording channel.

The total gain for a recording channel was G_{record} = 11 x 63 x 50. The resolution of the recordings was 0.07 μV per bit. This high resolution allows action potential signals to be digitized and recorded cleanly.

3.6.3 60 Hz Sync Circuit

Neurochip recordings and stimulation were always triggered from a 60 Hz sync circuit which generated a trigger-pulse on the up-transitions in the AC line voltage. This technique time-locks every trial to occur at the same phase in the 60 Hz signal so that baseline 60 Hz signals can be subtracted off-line. To additionally reduce 60 Hz pick-up, all metal structures in the vicinity of the microscope and electrical apparatus were tied to the same ground potential via alligator-clip connections. The 60 Hz peak-to-peak noise was usually about 30 μ V peak-to-peak, and easily subtracted.

3.7 Johnson Noise

The theoretical RMS Johnson noise level for the electrode is given by:

$$V_{Johnson} = \sqrt{4k_B TRB} = 1.9\,\mu V \tag{3.2}$$

for a bandwidth B = 5 kHz, a temperature, T = 295 K, and $R = R_{total}$ = 25 kΩ.

In addition, the pre-amplifier specifications predict an RMS noise level of 1.4 μV. Adding the sources of noise in quadrature predicts a total noise level of $V_{rms} = 2.6$ μV.

The measured value of the RMS noise is typically about 2–3 μV.

3.8 Stimulus Generation

Extracellular stimuli were generated by 8-channel electrode interface modules. (See Appendix C.) All stimuli are generated as voltage pulses which drive the neurochip electrodes through 500 kΩ resistors on the preamp. So long as $Z_{elec} \ll 500$ kΩ, this pulse approximates a constant current source. Constant voltage stimuli could be produced by by-passing the 500 kΩ resistor (via manual switches on the preamp board) routing the voltage pulse directly to the neurochip electrode. The amplitude and duration of the pulses is controlled by custom LabView software. Current stimulus amplitudes were in the range of 0–20 μA, and the duration was between 200–400 μsec. The LabView software also generates a hardware sync trigger to control the relative timing of optical data acquisition.

3.9 Computer Interface Miscellany

In addition to the A/D card, an NI-6602 counter-timer card was used to generate timing pulses and the NI-6503 provided 8 additional lines of DIO needed to control switches on the drivers for stimulus generation. The entire data acquisition and stimulation system was controlled with drivers implemented in custom LabView programs. LabView was used for on-line display as well as automation of stimulus generation. Data were collected at 20 kHz, which permitted real-time display and saving of the data. The maximum possible rate of display plus saving on a 2.3 GHz Intel Pentium computer was about 25 kHz. The LabView software is described in more detail in Appendix D.

Chapter 4

Cell Culture: Survival and Trapping in Neurocages

4.1 Introduction

This chapter describes the cell culture methods used to generate and maintain neurochip cultures start-to-finish—from harvesting neurons from embryonic rats, to making and maintaining neurochip cultures.

I worked with hippocampal CA1 and CA3 pyramidal cells harvested from embryonic day 18 (E18) Wistar rat embryos. (Dentate gyrus cells are not yet present at this stage in development.) The hippocampus is an interesting part of the brain to study because it is known *in vivo* to serve as the center for consolidation of new memories. CA3/CA1 pyramidal cells were chosen because they are known to exhibit interesting properties, such as LTP, *in vivo* as well as *in vitro*. In addition, the pathway for information flow—that is, the connectivity—in the hippocampus is very well defined. *In vivo*, CA1 receives input from the CA3 cells via the Schaffer collaterals. CA3 makes dense connections within itself, but CA1 does not. (Dentate gyrus cells input onto CA3 *in vivo*, but, as previously mentioned, they are not present in my cell culture system.) Day 18 embryos were chosen because the CA1 and CA3 regions of the hippocampus have differentiated by this age, and because these relatively young neurons are resilient—i.e., they survive the "reverse brain surgery" procedure described below.

4.2 Cell Culture

4.2.1 Dissociation

All cell dissociation was graciously performed by Sheri McKinney. A description of the process follows. At 18 days gestation, embryos are removed by Cesarean section from a pregnant, CO_2-asphyxiated Wistar rat. Hippocampi are dissected from the embryonic brains and stored in ice-cold, oxygenated HBSS. The extracellular matrix of the tissue is weakened by incubation at 37 °C in 0.25% trypsin, followed by dilution in tissue culture medium supplemented with 5% equine serum to neutralize the trypsin. The partially digested tissue is centrifuged and re-suspended in tissue culture medium. The cells are then fully dissociated by gentle trituration with a sterile plastic 1 mL disposable pipette tip. This method gives a 90% yield of viable cells after 1 day in culture, with neurons composing about 95% of the population (determined by cell morphology after 1 day in culture). The other 5% are glial cells.

4.2.2 Neurochip Cultures: Chip Preparation and Loading Neurons

Neurochips are prepared for culture as described in Appendix E. Briefly, the neurochip culture dish is sterilized with a combination of ethanol and UV light. Alternating layers of polyethylene-imine (PEI) and laminin are applied to the surface of the neurochip to render the surface cytophilic and to promote neurite outgrowth.

Making a neurochip culture starts with plating a background mass culture on a dry surface. The mass culture consists of 30,000 total neurons split between two 15 μL drops on either side of the neurocage array (see Figure 4.1, top).

The mass culture is necessary because neurons won't grow without friends; a critical mass is necessary for survival. The drops are positioned so that they are about 2 mm away from the neurochip, a large enough distance that only some stray long axons impinge on the array. This isolation is necessary so that cells in neurocages can be unambiguously identified, and so that results of stimulus-response experiments are not confounded. The isolation is excellent until the culture is about 2 weeks old. Mature neurons can grow axons several millimeters long toward the array. Sometimes the mass culture migrates toward the array as glia are allowed to replicate. An investigator desiring perfect isolation may consider "painting" an agarose line to create a non-adhesive barrier for neuron growth into

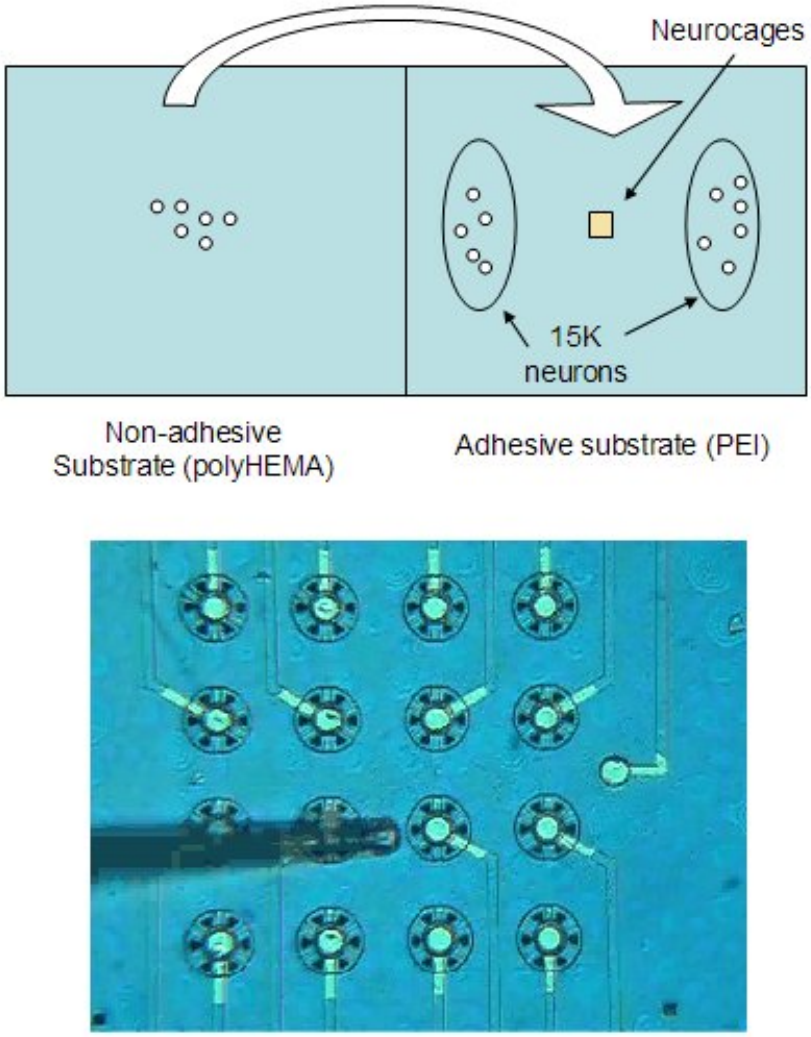

Figure 4.1: Neurochip cell culture. A "mass" culture is plated on either side of the neurocages. Individual neurons are loaded into each neurocage, one at a time. A neuron is selected from the non-adhesive substrate and carried over to the array with a custom micropipette. The micropipette is visible at the left of the bottom image.

the array.

The mass culture is incubated for one hour to allow the cells to anchor. Subsequently, the dish is flooded with about 3 mL of neurobasal medium (recipe in Appendix H). A piece of No. 1 coverslip coated in a non-adhesive substrate, polyHEMA (poly(2-hydroxyethyl methacrylate)), is placed in the dish on the blank side of the neurochip (see Appendix G). The culture dish is positioned under the microscope and a few thousand neurons are gently pipetted onto the polyHEMA non-stick surface. Individual neurons are manipulated with a custom-cut pipette tip (see Appendix G) with a manual pressure-driven assembly

to capture, push, pull, and position a neuron into a cage. This process is depicted in Figure 4.1, bottom. Neurons to be loaded into cages were rejected or selected based on the following criteria: Cells which were relatively very small or very large were avoided, as were irregularly shaped cells. Otherwise, cells of round and mid-sized anatomy selected for loading were were chosen at random.

Loading neurons continues one-by-one for about 30 minutes, until a cell has been loaded into each of 16 neurocages. At the end of the loading session, the dish is allowed to sit on the microscope stage for 5 minutes to allow the just-loaded neurons to anchor before transferring to the incubator. By this time, the CO_2-buffered medium has changed pH noticeably, from 7.4 up to about 7.7. Short-term exposure to alkaline pH evidently does not harm freshly plated neurons.

4.3 Neurochip Culture Maintenance

4.3.1 BDNF

To promote neuron survival and synaptogenesis, brain-derived neurotrophic factor (BDNF, Sigma B3795) is added to the culture dish at a concentration of 20 ng/mL on day 0 (just cultured) and again on day 3 [34, 58, 67, 32, 41, 40]. After about one week, neurons have matured morphologically to include extensive dendritic arborization. By this age, endogenous BDNF may be synthesized by hippocampal cells in the soma, axon and dendrite terminals [53]. A full study of the effect of BDNF on neurochip culture electrical activity and connectivity was not conducted. I did observe, however, that neurochip cultures not treated with BDNF during the initial stages of cell culture almost always failed to form suprathreshold connectivity.

4.3.2 Feeding

Starting 24 hours *in vitro*, 1/5 of the neurobasal-based culture medium (NB) is exchanged daily for DMEM-based medium (recipe in Appendix H). After a week of exchanging, the culture medium is actually a mix of 75% DMEM and 25% NB solutions. Subsequent feeding was done weekly with DMEM-based medium. The main difference in the formulation between NB and DMEM is the sodium concentration. In NB the NaCl concentration is 55 mM; in DMEM it is 110 mM. (The total osmolarity difference is also 55 mM.) Correspond-

ingly, the lower sodium concentration decreases the driving force for inward sodium current by about 15%.

Cultures are started in NB-based media because it is superior for short-term survival in low-density hippocampal cultures. The low-sodium concentration may aid survival by reducing the excitability of neurons, thereby protecting them from excitotoxic effects. This initial blessing, however, turns into a curse. Cultures grown in NB never exhibited any spontaneous activity or suprathreshold synapses even after 3 weeks *in vitro*. I make this statement after having studied about 20 cultures maintained in NB up to three weeks old. These cultures displayed an anatomically rich network, but were completely silent—they did not exhibit spontaneous activity or driven responses (see Chapter 7). (They invariably caused me to lose a few of my own neurons banging my head on the wall!) Cultures maintained in DMEM did indeed display functional suprathreshold synaptic activity.

4.3.3 Osmolarity

The osmolarity of the culture medium must be maintained between 270–320 mM during the first week. The osmolarity increases as the culture medium contains an increasingly large proportion of DMEM-based media. For the remainder of the culture lifetime, it is critical for survival that the osmolarity be maintained at 320 mM. Even in the incubator, about 50 μL of water evaporates per day. Every time the culture is removed from the incubator for observation or experiment, the osmolarity change is accelerated. Osmolarity was regulated by adding sterile ddH_20 to the culture dish two times per week.

4.3.4 Cleaning and Reuse

Neurochips are cleaned by first gently rinsing the culture medium out of the dish with ddH20, then 10 minutes application of 3% BM solution (ALA Scientific. They no longer sell this product. I have been told that BM is essentially concentrated contact-lens-cleaning solution). For those hard-to-get neurons, a 5% bleach solution applied for 1 minute is used. Careful not to overdo the bleach, or it can strip away platinization.

I routinely reused neurochip culture dishes up to 5 times or more. No degradation in quality of any kind was noted over the repeated culture sessions.

4.4 Typical Culture Development

For illustration, photos over the lifetime of a culture are provided below (Figures 4.3–4.9), all imaged using Nomarski differential interference contrast (DIC) optics. A comment is in order about what can be seen in these photos. The Nomarksi DIC optical technique tends to highlight differences in height along the surface. It is good for viewing large neurites—about 1 μm in diameter—but fails to capture smaller ones. An SEM of a neural process emerging from a tunnel (Figure 4.2) reveals that there are also many fine processes that are invisible with standard 10X and 20X Nomarski optics. It is important to keep this in mind when contemplating the richness of the network when viewing Nomarksi photos alone.

Figure 4.2: A neural process emerging from a neurocage tunnel. The origin of the neurite is the trapped neuron on the interior of the cage. Note that there are many thin processes branching off from the main process. Under Nomarski optics, only the thickest processes are visible; numerous thin processes are not. Scale bar: 10 μm

One other note is that soma are easily imaged in the cages for about the first 2 weeks in culture. Past that time, it becomes more difficult to view them under the microscope and cleanly identify them in the corresponding photographs. I speculate that this phenomenon might occur because: a) replicating glial cells originally plated a few millimeters away from the cages eventually spread into the array, sometimes capping off cages and generally making unambiguous identification more difficult, or b) by the time the culture reaches a more mature stage, the tendency to migrate caused a neuron to wedge up against the cage wall, obscuring it beneath the lip of the cage, making the geometry of the substrate on which the neuron grows non-planar. Both effects make it more difficult to see the edges of the cell membrane. Nonetheless, with careful inspection by a trained eye, it is usually possible to ascertain the viability of a caged neuron. Additionally, the state of the neural processes is a very good indicator for the viability of the neuron. A dying neuron's processes become segmented and disintegrate, leaving behind cytoskeletal remnants, whereas a healthy neuron's processes remain intact and continue to develop. It is easy to view and differentiate these cases.

The pictures that follow show typical neurochip culture development over time. The same culture is pictured from 0 DIV (just cultured) to 18 DIV. Rapid anatomical development occurs during this period—note the dendrite and axon growth, and rich network formation. Electrophysiology experiments revealed that neurochip cultures exhibit initial suprathreshold synaptic connectivity at about 10 DIV (see Chapter 7), just about the time a rich "neuropil" forms.

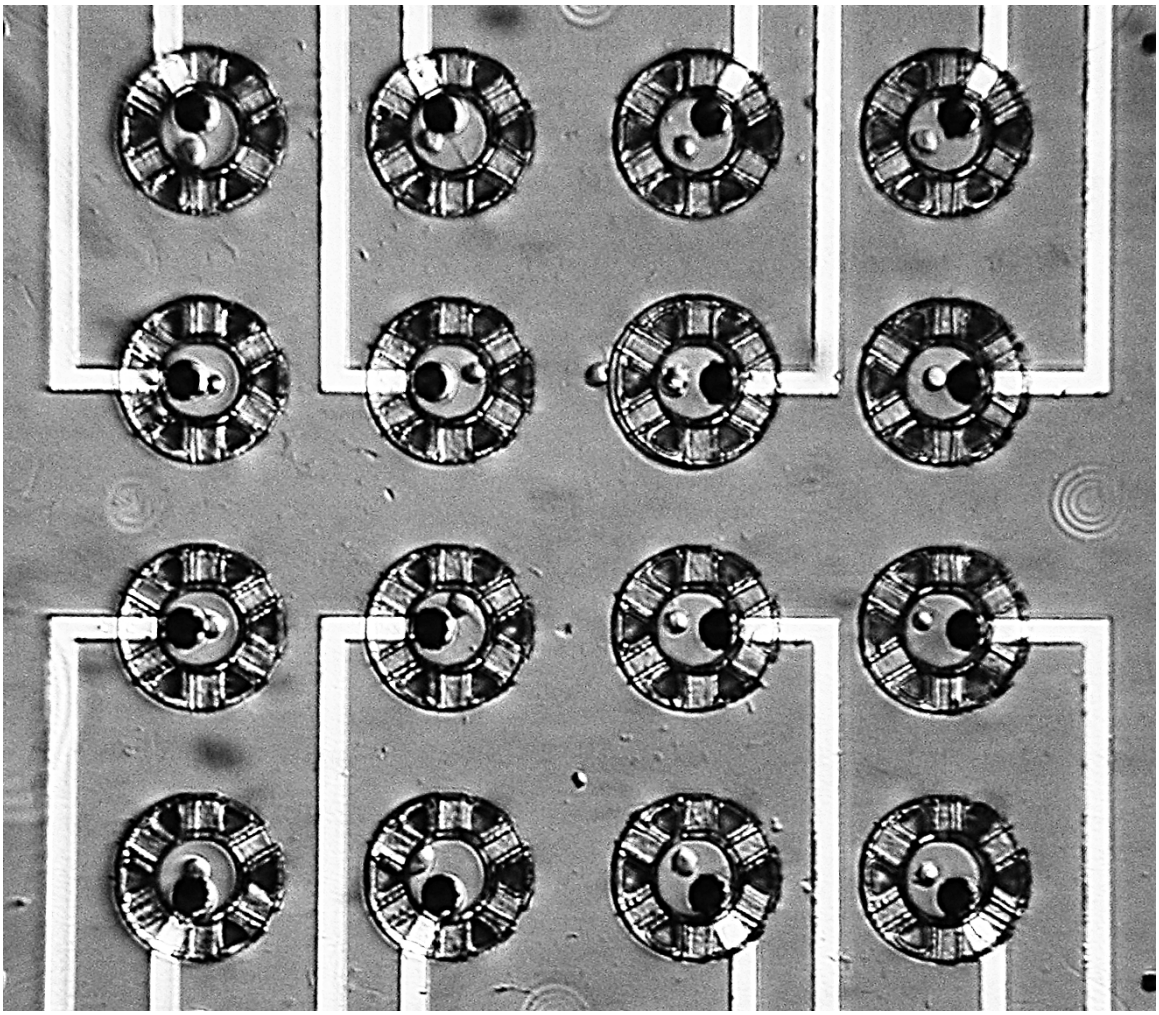

Figure 4.3: Neurochip culture: 0 days old (just cultured). Neurons are approximately spherical with a radius of about 5–10 μm, and have no axons or dendrites.

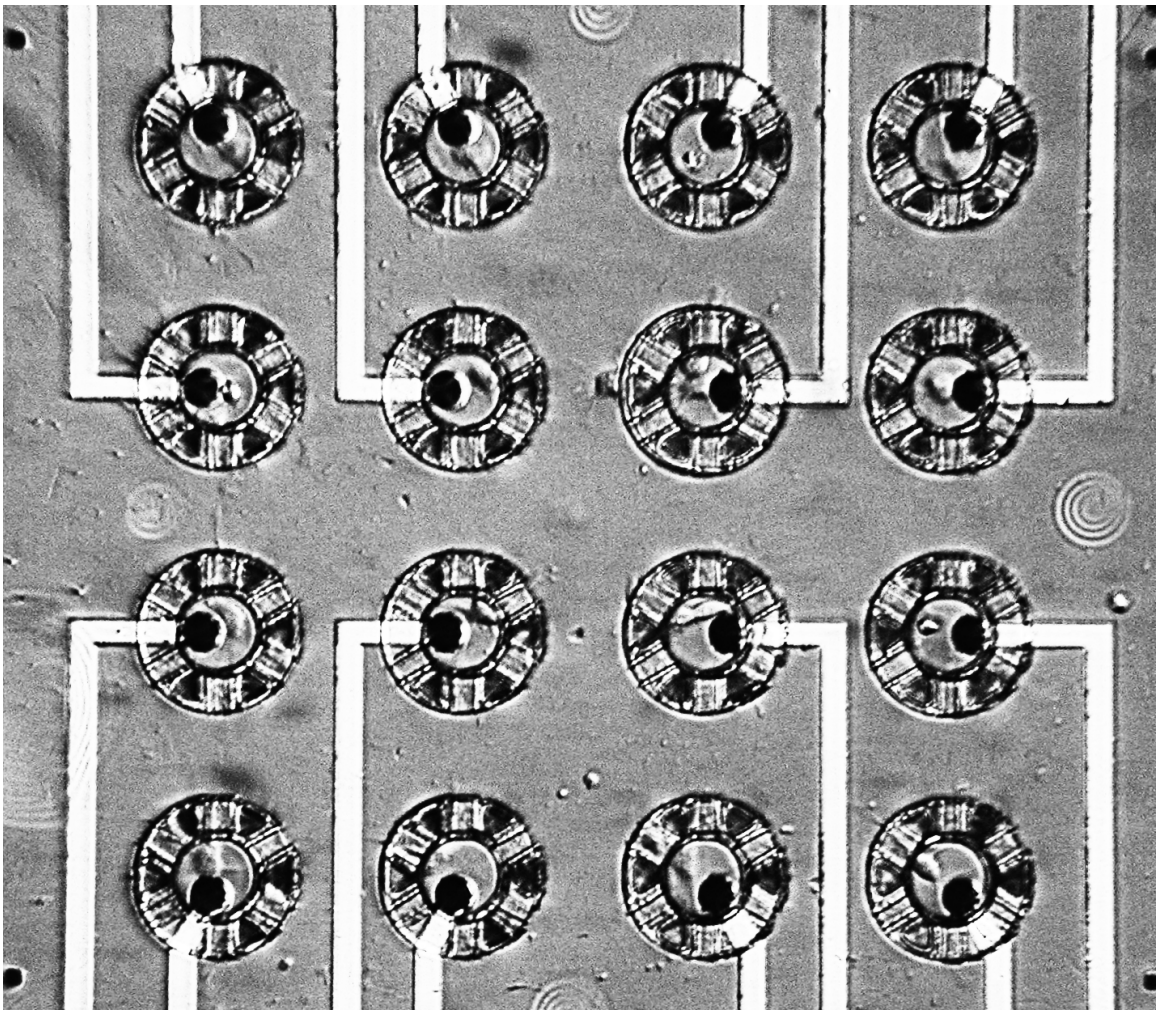

Figure 4.4: Neurochip culture: 1 day old. Neuron cell bodies have anchored to the surface and are beginning to flatten out and take on a pyramidal shape. They have begun to sprout neurites.

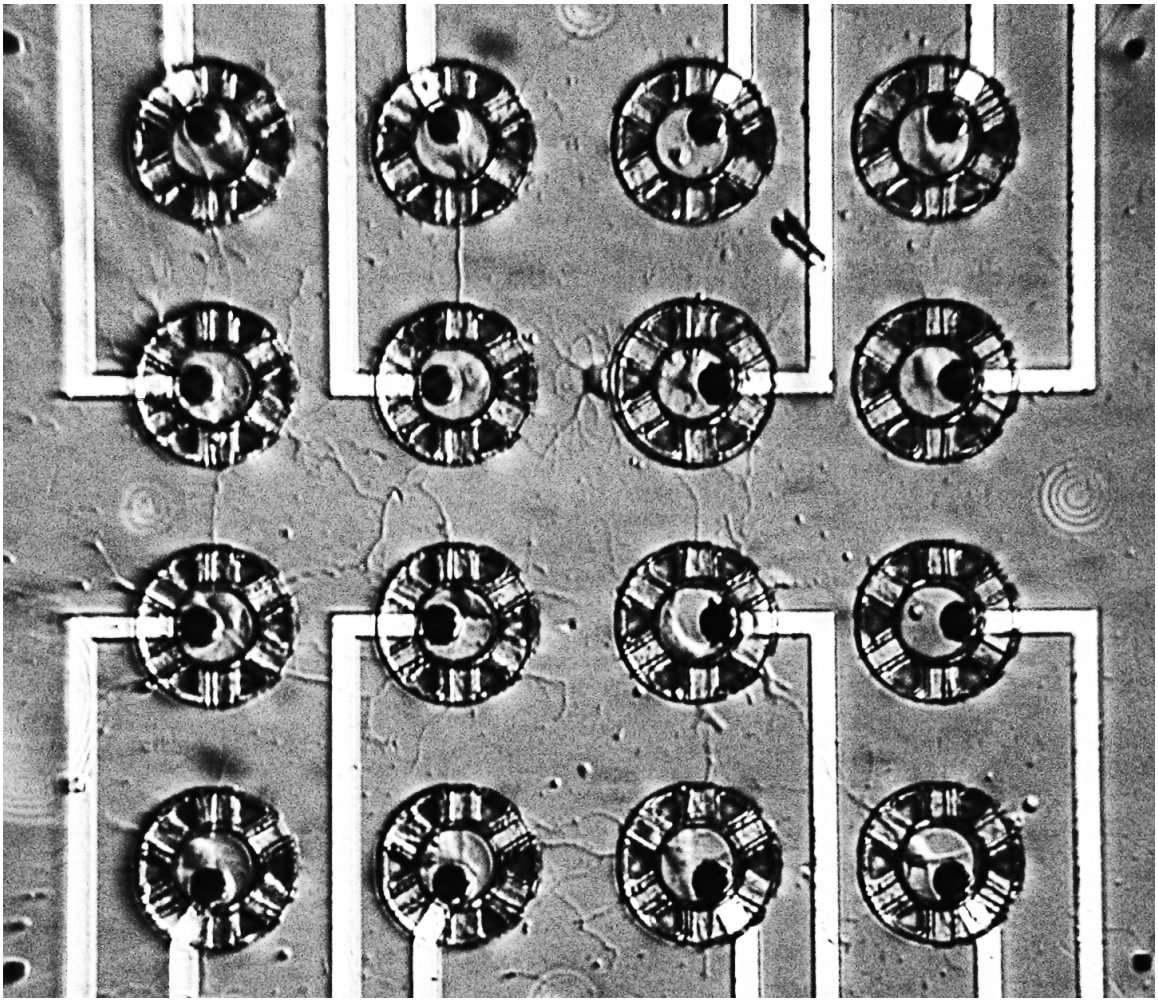

Figure 4.5: Neurochip culture: 3 days old. Neurites have grown significantly up to about 100 μm. Neurites have emerged through the tunnels out in the networking region, where they grow unconstrained.

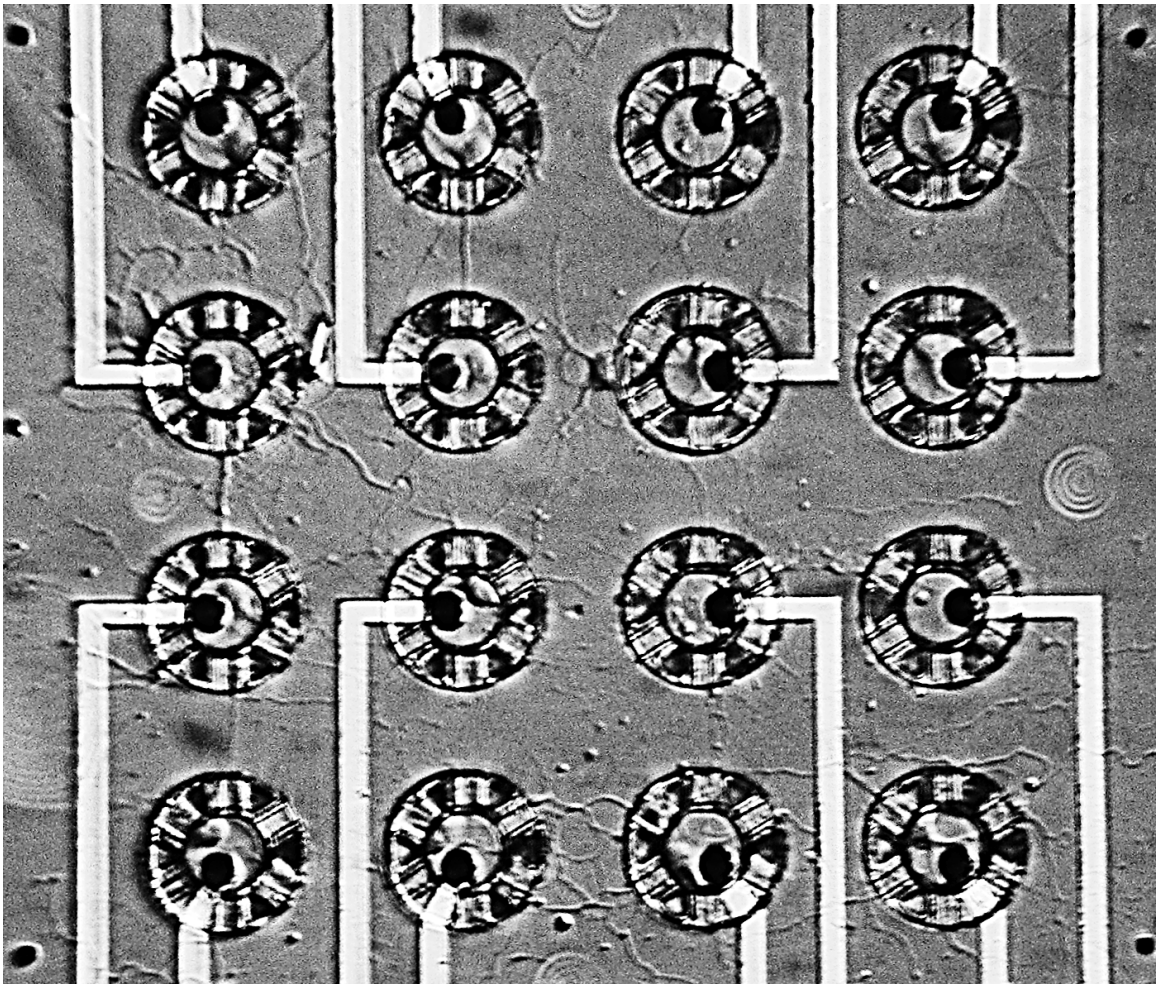

Figure 4.6: Neurochip culture: 5 days old. Neural process growth continues at a rapid rate. Initial networking is seen.

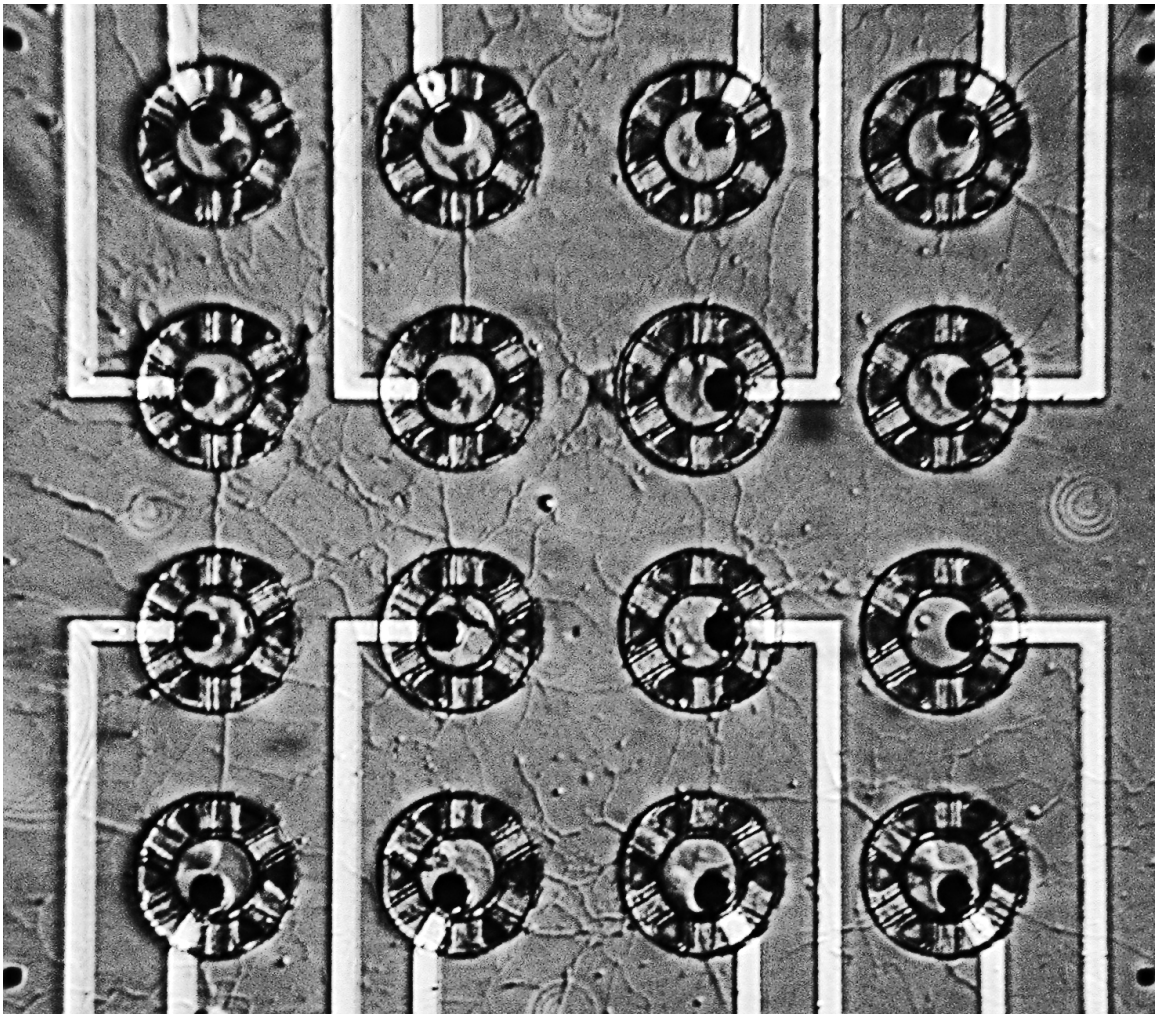

Figure 4.7: Neurochip culture: 7 days old. Neural process growth has continued at a rapid rate. The network has become richer.

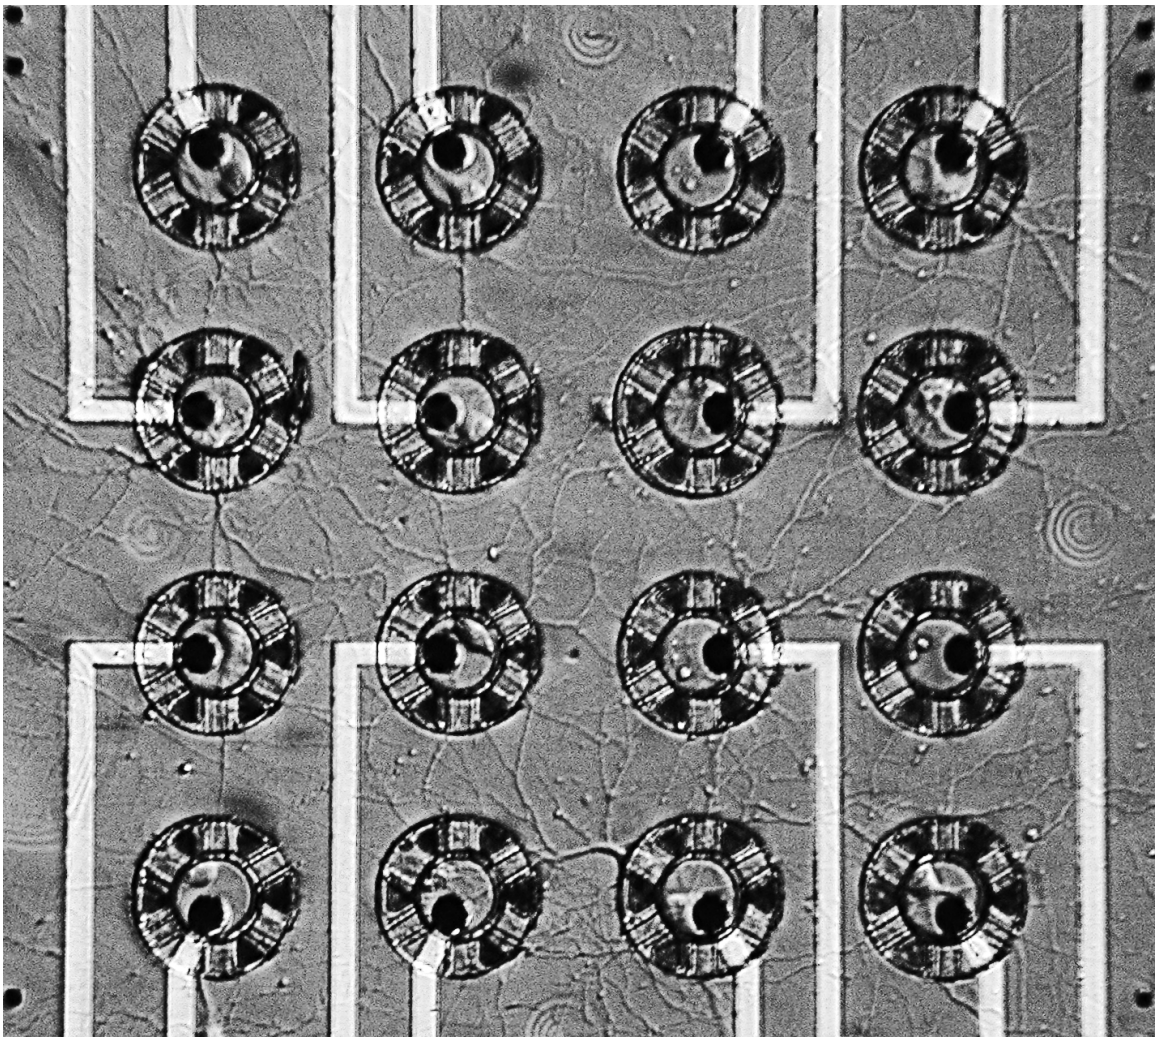

Figure 4.8: Neurochip culture: 10 days old. Soma are still trapped in cages. Process outgrowth through the tunnels is evident. A rich network has formed.

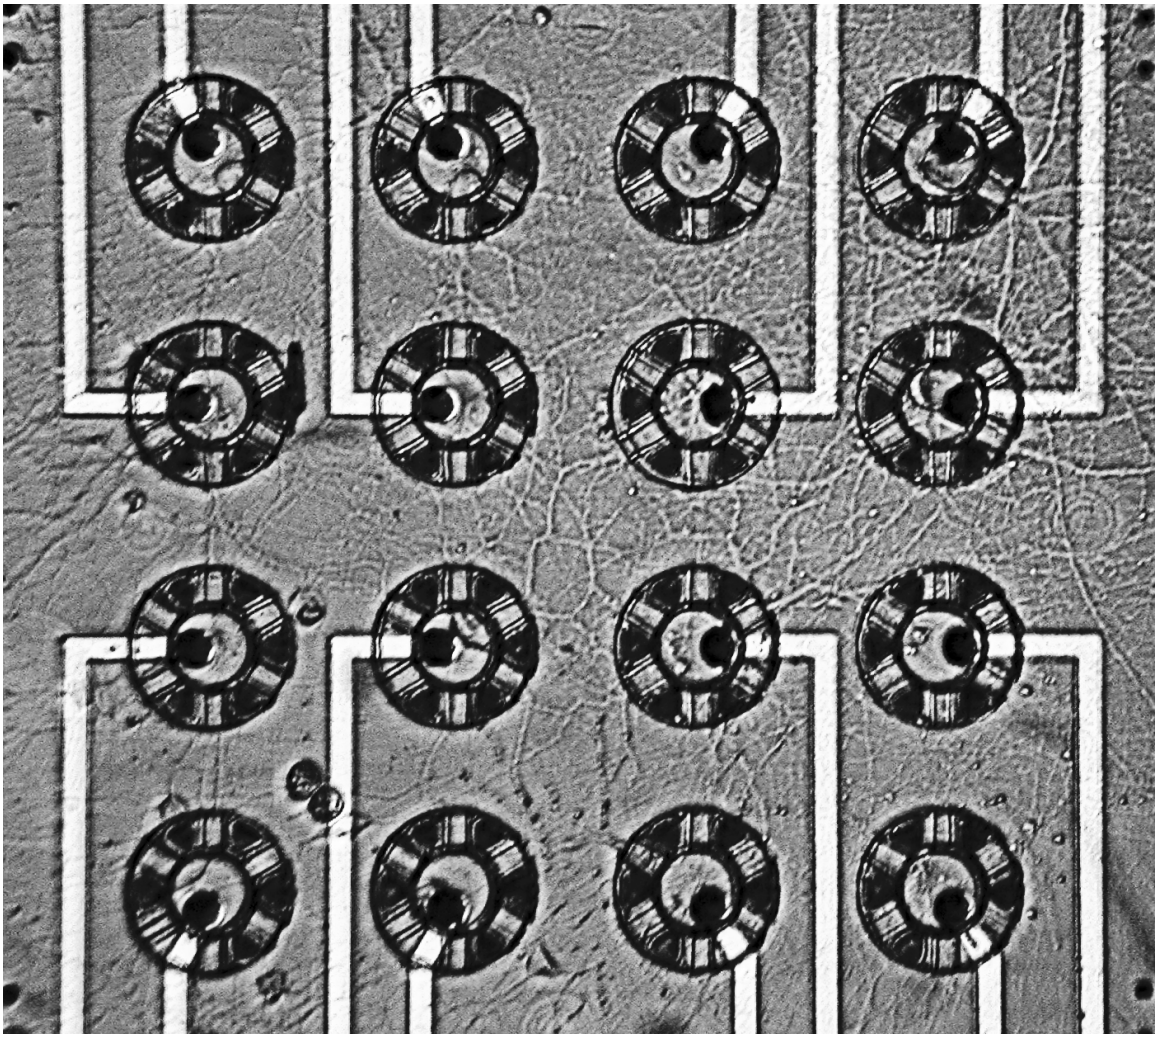

Figure 4.9: Neurochip culture: 18 days old. A few neurons have died off (bottom, left region), while the networking has become extremely dense elsewhere. In some cases the soma are difficult to cleanly identify, but the tangle of neural processes—continued network development—identifies most neurons as being healthy.

4.5 Survival in Neurocages

One obvious prerequisite for using the neurochip to study cultured neural networks over time is, of course, that the neurons survive and develop "long-term" in the neurocages. "Long-term" in this context means 3 weeks or more, for neurons begin to exhibit suprathreshold synaptic connectivity starting about 2 weeks *in vitro*. And, of course, the more neurons the better. The realistic expectation for the neurochip cultures is that survival will be sustained at approximately the 50% level. That goal is based on the fact that low-density ($300/\text{mm}^2$) control cultures survive only at about the 50% level over a few weeks (data not shown). Considering all the extra manipulation and isolation of the culture in the array, it is a small miracle that neurons survive at all. The goal, remember, is to investigate small cultures of about 10 neurons. Even studying cultures with 6 neurons is a 2-fold increase over the maximum number achievable with patch-clamp (or, one could consider the complexity to perhaps increase by a factor of 2^3).

Survival attained in the isolated neurocage array culture is indeed satisfactory. Figure 4.10 shows survival (mean ± s.e.m.) over time in the caged culture. Data was compiled from 41 "good" cultures, where "good" simply means that the small, isolated, culture in the array actually grew. About half of all neurochip cultures fell into this category. The other half—the "bad" cultures—simply did not start growing and died within a day. The factors determining the outcome of this binary event are unknown.

Not every culture was observed at the same age, which explains the small statistical fluctuations seen in the graph. (For example, one set of cultures was observed on, say, day 9, and another set—cultured the next week—was observed on day 12. If one set of cultures examined on day 12 had a particularly high survival rate, the mean survival on day 12 is higher than day 9. This effect is averaged over several sets of cultures.) Neurons were judged to be alive or dead based on anatomy. This judgment is easy during the first two weeks, as a viable neuron is seen to first flatten out, and then sprout axons and dendrites. The neural processes can be seen to elongate and branch extensively up until about two weeks old. A dead neuron, on the other hand, exhibits segmented processes and the soma shrinks into tiny clump a few microns in diameter. Determination of alive versus dead becomes more difficult from days 14–21 because the web of processes made it impossible to trace one back to the neuron of origin. However, optical and electrophysiology experiments

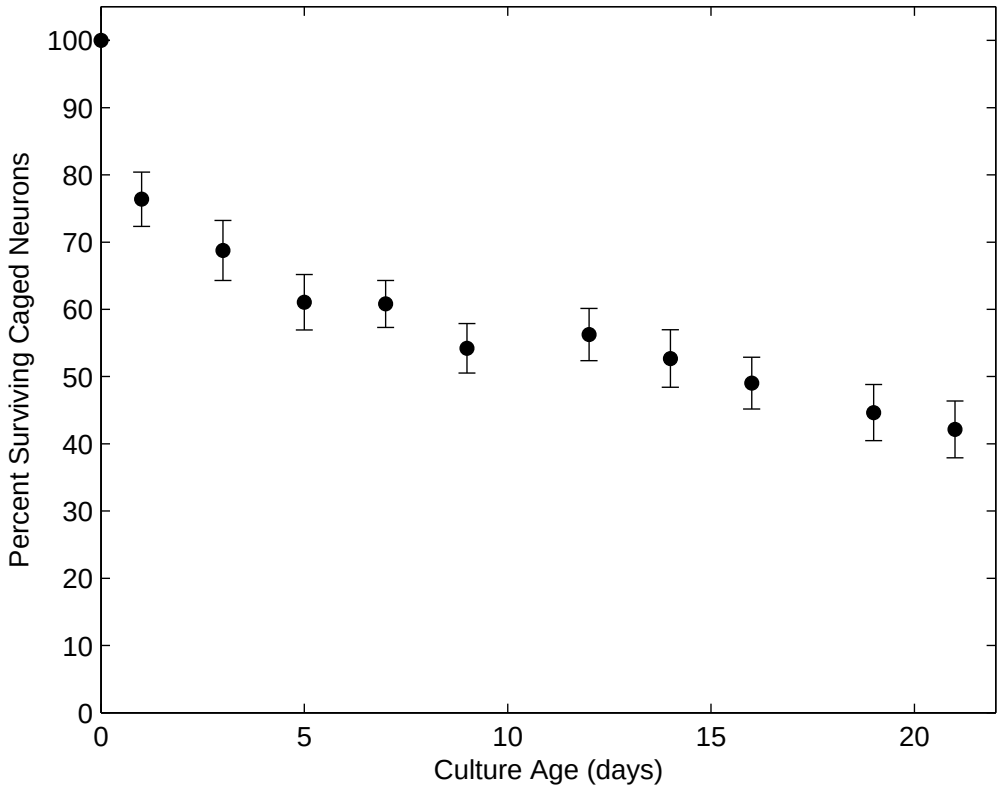

Figure 4.10: Survival of neurons (mean ± s.e.m.) in the isolated neurocage array culture over time. The large initial drop from day 0 to 1 indicates that not all neurons loaded were viable to begin with. The survival rate at 1 week is about 60 %; at 2 weeks about 52%; and at 3 weeks about 42%. Survival statistics in low-density (300/mm^2) control cultures are comparable.

with older cultures generally confirmed that visual inspection was accurate.

The initial precipitous drop between days 0 and 1 reflect the fact that some neurons are harmed during the dissociation process; not all plated and loaded cells are viable. Viability at the 75% level is also observed at 1 day in low-density control cultures. Over the next three weeks neurons survive well, with slow attrition rate. One neuron is lost per week, on average. At one week the survival rate is about 60%—about 10 out of 16 neurons surviving. At two weeks the survival rate drops slightly to about 52%. At week three about 42% of neurons are surviving. The second to third week is the critical time for investigating connectivity in the culture. Thus, I typically worked with cultures consisting of about 7–8 neurons.

4.6 Trapping Efficacy

Neurons must stay trapped, of course, or the whole purpose is defeated. The cages are in fact effective at trapping neurons at the $> 99\%$ level. Out of the 41 cultures followed here, essentially no escapes were noted. It turns out that the critical parameter for successful trapping is the tunnel height. In cages with a tunnel height of 1.7 μm, about 13% of neurons were observed to escape, on average, by migrating along their axon through the tunnel. The current and final design of 1.0-μm-high tunnels makes the escape rate essentially nil.

Chapter 5

Evoking Action Potentials with Extracellular Stimuli

5.1 Introduction

In order to probe neural network connectivity at the single-cell level over a timescale of weeks, the electrode within each neurocage must be capable of producing extracellular stimuli to reliably and safely evoke action potential (AP) responses from the nearby neuron.

Optical recording using a voltage-sensitive dye (VSD) was used to examine the response of a neuron to a particular stimulus. Electrical recordings cannot be used to assess whether a neuron was stimulated because the stimulus artifact (of the order of 1 Volt) is much too large for too long (by tens of milliseconds), swamping out any extracellular AP recording. The purpose of these experiments was to examine the change in fluorescence intensity in response to a stimulus and identify and characterize those stimuli which effectively and reliably evoke APs.

This chapter will begin with a quick theoretical overview of extracellular stimulation and present a very simplified model applied to a caged neuron. Returning to the experimental world, the physical setup for stimulation experiments—the optical and electronic components of the system—are detailed. Continuing along that path, an experiment will be described from start to finish to illustrate how a stimulation experiment is conducted and how AP responses are identified. The results of many (N = 66) such experiments are presented, to demonstrate that neurochip electrodes can indeed reliably and safely stimulate the nearby caged neuron. These results for single-cell stimulation are compared to MEA mass culture stimulation results. The chapter concludes by commenting on (debunking) a

couple of misconceptions regarding extracellular stimulation that remain pervasive in the neurobiology folklore.

5.2 Theoretical Considerations

The purpose of this section is not to make an extremely detailed model of extracellular stimulation. Rather, the point is to provide a simple theoretical framework that leads to an estimate of the required stimulation parameters. It will later be seen that this estimate is in reasonable agreement with the experimental results.

To elicit an AP from a neuron, a patch of membrane must be depolarized from its resting potential by an amount ΔV_m for a sufficiently long time to open the voltage-gated sodium channels. This initial influx of Na^+ must be large enough to generate an AP. Generally speaking, for rat hippocampal neurons, the transmembrane voltage must increase (from a resting potential of -70 mV) by about $\Delta V_m \approx 15$ mV. Voltage-gated Na^+ channel kinetics are on the timescale of about 0.2–0.8 msec [30], so $\Delta t \approx 0.5$ msec. The axon hillock is reported to have a locally very high density of Na+ channels and is, therefore, probably the region at which an AP initiates [15, 66]. Thus, this may be the critical patch of membrane which must be depolarized.

How might an *extracellular* current (or voltage) stimulus produce a change in *transmembrane* potential? After all, the stimuli is presented outside of the cell. Some background material is reviewed before presenting the actual stimulation model. A well-versed reader may wish to skip ahead to Section 5.2.3.

5.2.1 Model Neuron: Passive Electrical Properties

Neurons are typically modeled as having two passive electrical properties: capacitance and resistance. The capacitance, C_m, arises from the cell membrane acting as a thin dielectric separating charge (Na^+, K^+, Cl^- in solution). The resistance, R_m, is due to "leak" channels that span the membrane. A neuron, therefore, passively functions as a distributed parallel RC circuit. An estimate of the total capacitance (due to the entire cell membrane—soma and neurites) is $C_m \approx 30$–80 pF. Typical values for R_m measured by others are about 100 MΩ [61, 50, 8]. Therefore, the membrane time constant of the cell is $\tau_m = R_m C_m \approx 5$ msec. Finally, as a rough approximation, the surface area of the soma accounts for 1/4 the

total membrane surface area, thus 1/4 of the total capacitance: $C_{soma} \approx C_m/4 \approx 15$ pF. The number will become important later.

Recall for a parallel RC circuit that when current begins to flow, it is primarily capacitive current for times much shorter than the membrane time constant. That is to say, current through the resistor can be neglected for short times. Correspondingly in a neuron, for $\frac{t}{\tau_m} \ll 1$, a change in membrane voltage is associated with a capacitive current that dominates to the extent that the model of the cell membrane may be reduced to a capacitor alone. This result will become important when we consider possible mechanisms of stimulation.

5.2.2 Ohm's Law in Physiological Saline

When a current flows in saline from the neurochip electrode to a (distant) ground (or "return current wire"), Ohm's law holds. $\Delta V_{ab} = IR_{ab}$.

What do ΔV_{ab}, I, and R_{ab} correspond to in the neurocage arrangement? I is easy: it is just the total amount of current being passed through the neurochip electrode. Referring to Figure 5.1 below, R_{ab} is the resistance between points a and b and depends on the geometry of the region through which charge flows.

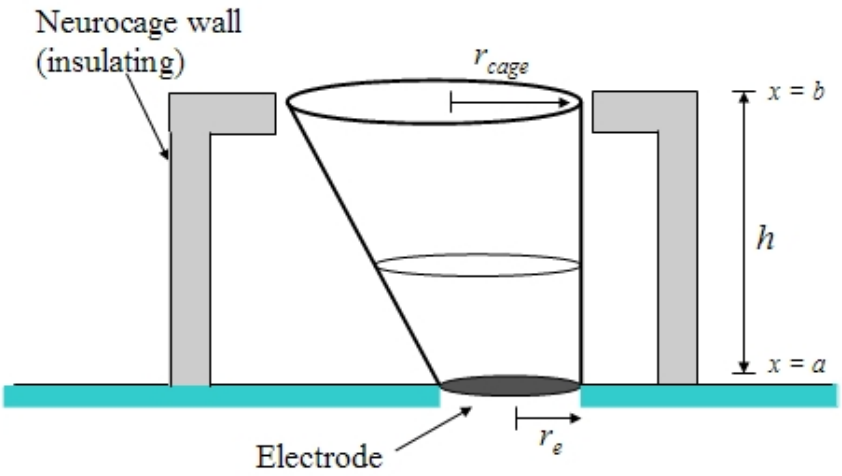

Figure 5.1: Geometry of the truncated conical cross section through which current flows from the electrode to the access hole at the top of the cage. The x coordinate points in the vertical direction. The radii of the electrode and access hole at the top of the cage are labeled r_e and r_{cage}, respectively. The total resistance of this geometry is $R_{ab} = \rho h/\pi r_e r_{cage}$.

Specifically,

$$R_{ab} = \int_a^b \frac{\rho(x)dx}{A(x)} \, . \tag{5.1}$$

where ρ is the resistivity of the medium. For physiological saline, resistivity is isotropic—ρ is a constant, and taken to be 70 Ω·cm. For constant current, the voltage difference between two points in the medium is given by:

$$\Delta V_{ab} = \rho I \int_a^b \frac{dx}{A(x)} \tag{5.2}$$

The geometry of the path of current flow can be crudely modeled in a neurocage as a truncated cone. The narrow end of this cone corresponds to the electrode, and the wider end-cap corresponds to the access hole at the top of the neurocage. Doing the integral in Equation 5.2 for the conical geometry yields the result:

$$\Delta V_{ab} = \rho I(b-a)/\pi r_a r_b = \rho I h/\pi r_e r_{cage} \tag{5.3}$$

where r_e and r_{cage} are the radii of the electrode and neurocage, respectively, and h is the distance height of the cone, in this case given by the height of the neuron. This approximation should be qualitatively correct. For an exact solution of voltage generated by current flow passed through a disk-shaped electrode into an isotropic infinite half-geometry (hemispherical, isotropic) see [71]. (However, note that their solution does not directly apply to the neurocage geometry since the neurocage acts as an insulator, thus violating the assumption of hemispherical isotropic geometry.) At any rate, seeking an exact solution for R_{ab} as a function of position is not particularly fruitful. Again, the main point here is that a reasonable expression for Ohm's law in the neurocage geometry has been obtained.

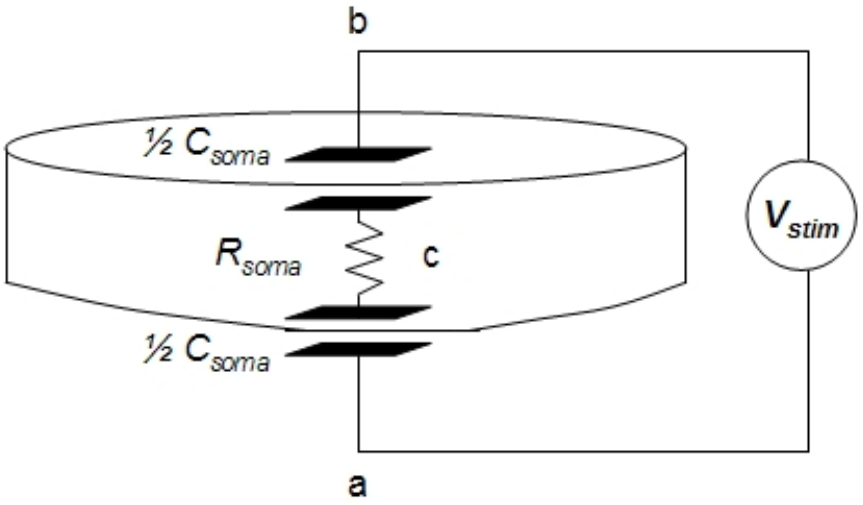

Figure 5.2: Electrical model of a hippocampal neuron for stimulation by an extracellular current pulse. Because the time scale of stimulation is short compared to the passive membrane time constant, the top and bottom membranes are modeled as capacitors. The intracellular solution, essentially saline solution, is modeled as a resistor R_{soma}. The extracellular current pulse generates a voltage drop in the medium between the top and bottom membranes, V_{stim}. The capacitance of the top and bottom membrane together is $C_{soma} =$ 15 pF. The resistance is $R_{soma} = 20$ kΩ. The RC time constant of this arrangement is 0.3 μsec.

5.2.3 How to Stimulate a Caged Neuron in Theory

With the requisite background at hand, one can now begin to think about how a neuron is actually stimulated extracellularly in the neurocage geometry. Keep in mind that the goal is to depolarize a localized region of the neuron, either by raising the intracellular voltage relative to the outside, or by decreasing the extracellular voltage relative to the inside. To begin, imagine a model of a neuron that includes only the electrical compartment of the soma, no dendrites or axons. (The effect of adding neurites into the model will be considered later.) As shown in Figure 5.2, the geometry of a neuron is crudely modeled to be a cylinder with dimensions corresponding to that of a hippocampal neuron. Let the radius be $r_{soma} \approx$ 10 μm. Let the height be $h_{soma} \approx 5$ μm.

Approximating the soma membrane as solely capacitive (see Section 5.2.1), the top and bottom halves are treated as a lumped capacitor. The interior of the cell is conductive

saline solution, so can be treated as an Ohmic lumped resistor. The *intracellular* resistance, R_{soma}, may be calculated as $R_{soma} = \rho h_{soma}/(\pi r_{soma}^2) \approx (70\,\Omega \cdot cm)(5\,\mu\text{m})/\pi(10\,\mu\text{m})^2 \approx 20\,\text{k}\Omega$. The series RC time constant of this arrangement is given by $\tau_{soma} = R_{soma}C_{soma} \approx (20\,\text{k}\Omega)(15\,\text{p}F) \approx 0.3\,\mu s$, extremely short on the timescale of neurobiological interest.

Next, imagine an electrode near the neuron through which a stimulus, I_{stim}, is applied to the electrode. The stimulus current in this crude model will emanate roughly vertically in the region near the neuron. Since the membrane capacitance is much smaller than the electrode capacitance, only a very small proportion of the flowing charge will be "absorbed" by the cell membrane. In other words, the current supplied by the electrode defines the extracellular potential change relative to a distant ground, ΔV_{extrac}, even in the region of the neuron. Thus, a potential difference is created between the top and bottom membranes.

When the stimulus is applied, the series RC circuit reaches a steady-state value within a time τ_{soma}—essentially instantly—and remains there for the duration of the stimulus, τ_{stim}. Due to the symmetric configuration (pictured in Figure 5.2), the cell's interior is an isopotential which approximately follows the average of the exterior potentials defined at points a and b. Therefore, the differences in potential across the top and bottom membranes, respectively, are given by:

$$\Delta V_m \approx \frac{1}{2}(V_b - V_a)$$

$$\Delta V_{ca} = V_c - V_a,\ \Delta V_{cb} = V_c - V_b$$

$$\Delta V_{ba} = V_b - V_a = -2\Delta V_{ca} = 2\Delta V_{cb}$$

where the last relation follows from the spatial symmetry.

This model predicts that no matter what the polarity of the applied stimulus (positive or negative) one side of the membrane will be depolarized while the other side hyperpolarized. For instance, in the case of positive current, $V_a > V_c > V_b$, so that the top side is depolarized while the bottom membrane is hyperpolarized. Also it should be noted that since the extracellular potential varies with position, it should be obvious that the relative position between neuron and electrode affects the amount of stimulation current required to elicit an AP.

With this model, we can estimate the amplitude of stimulation current required to elicit

an AP.

For successful stimulation we require that $\Delta V_m = \frac{1}{2}\Delta V_{extrac} \approx 15$ mV. Hence,

$$\Delta V_{extrac} = I_{stim} R = I_{stim} \frac{\rho h_{soma}}{\pi r_{cage} r_{electrode}}$$

$$I_{stim} = \Delta V_{extrac} \frac{\pi r_{cage} r_{electrode}}{\rho h_{soma}} = (30\,\mathrm{m}V) \frac{\pi\,(20\,\mu\mathrm{m})\,(5\,\mu\mathrm{m})}{(70\,\Omega \cdot \mathrm{c}m)\,(2.5\,\mu\mathrm{m})}$$

$$I_{stim} \approx 6\,\mu\mathrm{A}.$$

The upshot is that, based on this very simplistic model, a stimulus current of several μA in strength, lasting for several tenths of milliseconds, of either polarity should be sufficient to evoke an AP from a nearby neuron.

5.2.4 The Effect of Neurites

Now consider the effect of neurites added into the electrical model.

Neurites are considered to extend "far" from the cage and electrode. They are, therefore, in electrical contact—capacitively coupled to—the culture medium at reference potential, i.e., connected to ground. When the intracellular potential begins to change at the soma (but has not yet propagated through the dendritic arborization), a voltage difference exists between the soma and distant neurites. This causes current to be electrotonically conducted down the length of the neurites to hold the soma voltage closer to the resting potential. Hence, the effect of the neurites is to resist any intracellular change in potential due the the stimulus current.

This model predicts that, with intracellular potential held constant, the neuron is depolarized by decreasing the extracellular potential. In the framework of this model, therefore, a negative current pulse should successfully stimulate the neuron. The required minimum stimulus strength, furthermore, is predicted to be less than that without the neurites, by a factor of 2, give or take.

5.3 Dye Recording Apparatus

This section describes the electrical and optical instrumentation used to perform optical response experiments.

The heart of the system is an advanced CCD camera, the "NeuroCCD" from RedShirt Imaging LLC, that records changes in fluorescence intensity. It has 80x80 pixels, and a full frame rate of 2 kHz—fast enough to capture most neurobiologically important signals. If a higher rate is desired, it can record at 5 kHz, binning 3x3 pixels; or at 10 kHz from a swath of 80x12 pixels. The amplifiers have a selectable gain in the range of 1–30X. The camera has 90% quantum efficiency and low dark noise, which provides the possibility of measuring intensity changes with high precision, on the order of 1 part in 1000. An AP in the soma typically changes its fluorescence by 1–3% [17], so it should be cleanly recordable.

The system is built around an epi-illumination Olympus BHMJ microscope with the RedShirt camera mounted to the top-side trinocular port via a 3:1 demagnifying coupler (Thales-Optem) and custom-machined adapter sleeve. With this coupler, each pixel images an area approximately 1.5 μm squared, when using a 40X lens. The camera is supported by a counter-weight pulley system to prevent it from bearing weight on the microscope.

The fluorescence emitted from voltage-sensitive dyed cells (see next section) are imaged through a 40X water-immersion lens with NA = 0.8 (Nikon). Stained neurons are illuminated by a mercury arc mounted through the normal illumination port of the microscope. An optical feedback shunt regulator stabilizes the incident illumination nearly to the shot-noise limit [11, 10]. A standard Olympus "G" cube contains filters and a dichroic mirror to steer excitation light (Hg green, $\lambda = 546$ nm) to the neuron, and to capture fluorescence signals from it ($\lambda > 590$ nm). The intensity of the incident light must be carefully optimized. On one hand, one would like as many photons per area as possible to increase the fluorescence signal-to-noise ratio (SNR). On the other hand, however, even with improved VSDs, the illumination intensity must be minimized to reduce the effects of photobleaching and phototoxicity [38, 23]. The desired intensity is achieved by inserting an appropriate combination of neutral density filters into the light path, opening (or closing) the microscope aperture, and adjusting the hardware gain settings for the RedShirt system. No matter the intensity, the illumination time must always be minimized, as photo-bleaching and -toxicity are cumulative effects which occur after a mere few seconds of total illumination time. A

computer-controlled electromechanical shutter is inserted in the light path; it is normally closed, except for the short (100 msec) time interval during which a fluorescence signal is measured. Since both mounting ports must be occupied by the CCD camera and Hg-arc, respectively, no port is available for incandescent illumination for Nomarski optics. Instead, a fiber-optic dissection light is positioned to illuminate the culture dish in a bright field mode. Neurons are poorly visible in bright field mode, but it is still easy to navigate to the desired neurocage by following other well-imaged landmarks (mainly electrode leads which are highly reflective) to the neurocage.

The Neuroplex software suite from RedShirt was used to control all camera settings and optical data acquisition. Additionally, the Neuroplex software allows the optical traces ($\Delta F/F$ versus t) to be displayed and investigated immediately following the termination of data acquisition.

5.3.1 Stimulus Generation

Extracellular stimuli were generated by the eight-channel electrode interface modules in the Pine lab. All stimuli are generated as voltage pulses which drive the neurochip electrodes through 500 kΩ resistors on the preamp. So long as $Z_{elec} \ll 500$ kΩ, this pulse approximates a constant current source. Constant voltage stimuli are produced by routing the voltage pulse from the eight-channel module directly to the neurochip electrode, by-passing the 500 kΩ resistor. The amplitude and duration of the pulses is controlled by custom LabView software. Current stimuli amplitudes were in the range of 0–20 μA, and the duration was between 200–400 μsec. The LabView software also generates a hardware sync trigger to control the relative timing of optical data acquisition.

5.4 Voltage-Sensitive Dye Staining

5.4.1 General Properties of Dyes

Voltage-sensitive dyes (VSDs) are membrane-bound fast-response dyes. A membrane-bound dye molecule generally consists of a hydrophobic tail and a hydrophilic chromophore head. The dye orients itself in the cell membrane, but does not cross it. The hydrophobic tail inserts into the lipid bilayer, the chromophore head remains on the external side of the membrane. Incident excitation light excites the chromophore, which subsequently emits a

photon as it relaxes back into a lower energy state. The defining characteristic of these dyes is that the intensity of fluorescence changes when the membrane potential changes. (There are several possible physical mechanisms which might explain this amazing property—re-orientation, polarization, Stark shift, etc. We will not concern ourselves with how the dye works, and merely note that it does indeed work as claimed. Note that free dye molecules in solution (not oriented in the membrane) exhibit negligible fluorescence.) Thus, by tracking changes in fluorescence intensity, ΔF, one can track changes in membrane potential. In response to changing membrane voltage, fast-responding dyes change intensity within a few microseconds. Since neurobiological events—APs in particular—take place on a millisecond timescale, the fluorescence signal closely mimics the membrane potential change. Typically, one is interested in the fractional change of fluorescence intensity relative to the resting light intensity, $\frac{\Delta F}{F}$. The calibration of these dyes is typically $\frac{\Delta F}{F}/\Delta V_m \approx$ -1%/100 mV.

5.4.2 Dye Selection

Two fast-response, membrane-bound dyes were used for these studies, RH237 [26] and di-4-ANEPPDHQ ("di-4") [49]. RH237, developed by Grinvald some 20-plus years ago, was initially selected for (approximately the first half of the) experiments because of its history in the Pine lab (for example [43]). For the latter half of the experiments di-4 was utilized, as it was found to be equal, or superior, to RH237 in all regards. RH237 internalized within about 1 hr; di-4 did not internalize (at least over the time course of several hours). RH237 phototoxicity manifested after approximately 3 seconds of total illumination time; di-4 did not exhibit any notable phtoxicity even after as much as 10 seconds of total exposure time. The change in fluorescence intensity versus the change in membrane potential was approximately equal, about 1% per 100 mV for both dyes. xicity even after as much as 10 seconds of total exposure time. The change in fluorescence intensity versus the change in membrane potential was approximately equal, about 1% per 100 mV for both dyes.

5.4.3 Staining Protocol

The stock staining solutions are prepared fresh at a concentration of 1 mg/mL, dissolved in 95% ethanol. For staining, starting with a culture growing in normal physiological saline, the culture is gently rinsed 3 times with "Poo bath" saline (see Appendix H for recipe). After the third rinse, 7.5 μL are added to a culture in about 3 mL of Poo bath. The

staining time is 15 minutes, during which time the culture is placed back in the incubator. Following that, the culture is rinsed 3 times, as before, leaving the cells in saline. Note that this staining procedure is cytotoxic; 24 hours after the staining all that remained of neurons in the culture, even those not exposed to excitation light, was badly disintegrated chunks—probably cytoskeletal remnants.

Figure 5.3 shows an example of a neuron stained with VSD (di-4-ANNEPDHQ). The image is in three parts: on the left is a Nomarski phase-contrast image showing a neuron growing inside of a neurocage beside a platinized electrode (the black circle). Immediately to the right is the corresponding (overexposed) image of the same stained neuron illuminated with the mercury arc. The neuron is relatively intense, the edge effect of membrane staining is not pronounced due to the overexposure. The outline of the cage geometry is visible, and the electrode lead is also relatively bright, probably due to reflection from the underlying gold. The image on the far right is the corresponding CCD image—this is what the Red Shirt camera "sees". In Figure 5.3, the neuron is projected onto approximately 25 pixels. Cell bodies and processes traveling through tunnels are clearly visible in all three images. The pixels corresponding to the cell body are averaged to obtain the neuron fluorescence signal, F.

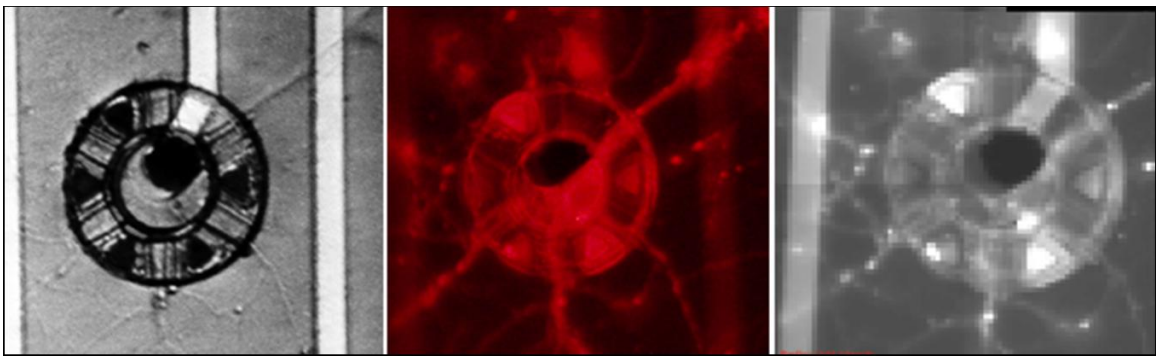

Figure 5.3: Corresponding images of Nomarski phase-contrast, stained neuron, and CCD image of fluorescence intensity (left to right). The signals from the CCD pixels corresponding to the area covered by cell body are averaged to obtain the neuron fluorescence.

5.5 Identification of APs from Fluorescence Traces

An optical measurement is an indirect measurement of cell membrane potential. Several criteria were used to identify APs based on the optical trace, $\Delta F/F$. If an optical response

met all the following criteria, an AP response was declared to have occurred:

1. Time Delay of Response: An AP is elicited within a few tenths of a millisecond following the end of the stimulus.

2. Amplitude of $\Delta F/F$: Given the calibration of the dyes, and that during an AP a cell changes membrane potential by an order of 100 mV, we expect a peak signal of $\Delta F/F \approx$ -1%.

3. Sign of $\Delta F/F$: A property of the VSDs is that fluorescence intensity actually decreases with a rise in membrane voltage (as the negative sign indicates above). Therefore, the sign of $\Delta F/F$ should be negative.

4. Width $\Delta F/F$: We know from patch clamp recordings of hippocampal neurons in a control culture that the width of the APs is about 3–5 msec wide. Therefore, the optical trace should also have a similar width.

5. All-or-nothing response: APs are intrinsically all-or-nothing events. Sweeping through stimulus strengths, therefore, a discontinuity in the magnitude of peak $\Delta F/F$ is expected to occur at the threshold current. Furthermore, for stimuli strengths greater than threshold, the magnitude of the response is expected to remain constant.

5.6 The Actual Experiment: Optical Data Acquisition

A typical experiment progressed as follows: Candidate neurons for optical response experiments were selected on the basis of prior visual inspection. Cells which appeared "healthy" were included in the study. Cultures were stained with VSD as described above. Bi-phasic current stimuli (both polarities) were delivered through the cage electrode with amplitude starting at 0 μA, incrementing by 2 μA, up to a maximum value of 20 μA (usually stopping at 16 μA). Since the response is "all-or-nothing" an obvious difference is observed when an AP is finally evoked. The amplitude at which an "all-or-nothing" response was determined was deemed to be the threshold current required to evoke an AP. Note that all stimuli in this section refer exclusively to *constant current* stimuli, unless explicitly noted otherwise. Voltage pulses were also examined and are briefly discussed in Section 5.8.3.

Full frames (80x80 pixels) were acquired at 2 kHz. Optical data were acquired usually for 80 msec, with the stimulus being delivered approximately in the middle of the interval. An "optical trace" was computed as the change in fluorescence over time divided by the resting light intensity (RLI) *spatially averaged over all CCD pixels onto which the neuron cell body is projected*. The spatial averaging increases the SNR in proportion to the square root of the number of pixels being analyzed. With a neuron typically being projected onto 25 pixels, spatially averaging decreased the RMS noise by a factor of 5 to about 0.1 %. For an AP response, then, the SNR was about 5. (The smallest response that could be cleanly noted was limited by the peak-to-peak noise, which was typically about 0.2%.) Typically the SNR of the optical system was large enough that only one trial for each stimulus was necessary. In instances when this was not the case, or when even high SNR were desired, multiple trials—usually between 3 and 5—were averaged.

5.7 Results

Figure 5.4 below shows sample optical traces. The red dashed line marks the onset of the stimulus, in this case a bipolar current stimulus, negative phase first, 0.4 msec per phase. The dotted black trace is a control trace, stimulus of 0 μA. The blue trace shows the response to a bi-phasic current 12 μA in amplitude. This trace strongly suggests an AP response: the peak response -1.3%; the (full) width is about 5 msec; and the response time is about 1 msec and clearly coincides with the onset of the stimulus. Sweeping through a range of stimulus strengths and noting the peak response generates a graph such as Figure 5.5. This shows discontinuities in response versus stimuli strength, as we'd expect to see if APs are evoked. Positive stimulus values indicate a bipolar stimulus delivered with positive phase first (then negative phase). Negative stimulus values indicate negative-phase first. Note the discontinuities at +6 and -8 μA. For current amplitudes larger than these, the peak response remains relatively constant. Taken together with the individual optical traces, therefore, -6 and +8 μA are noted as being the threshold stimuli to evoke APs in this case.

Such an experiment was repeated for N = 66 neurons, identifying thresholds for both positive- and negative-first bipolar current stimuli. Figure 5.6 shows the distribution of threshold currents for all neurons. Cells tested were in the age range of 6–31 days *in vitro*. There was no correlation noted between the age of the neuron and the stimulus threshold.

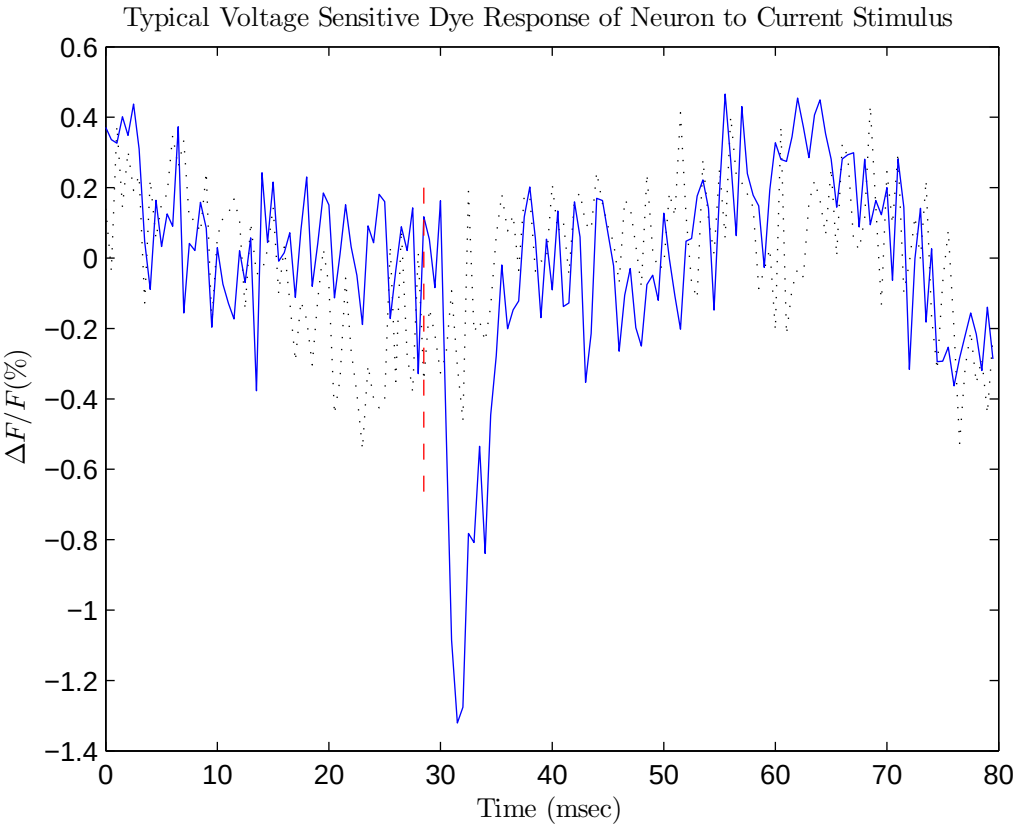

Figure 5.4: Optical traces in response to current stimuli. The dotted black line is a control optical trace (no stimulus delivered). The solid blue line is the optical response of a neuron to a 12 μA stimulus. The red dashed line marks the onset of the stimlus.

59 of the 66 (89%) of neurons tested exhibited action potential responses to negative-first stimuli, while 50 of 66 (76%) responded to positive-first stimuli. All current pulses were 0.4 msec per phase (0.8 msec total duration). The$\approx$ 10% of neurons which did not respond may have either been a) dead at the time of cell selection, b) did not survive the staining process, or c) simply did not respond. There is no way to know for sure which, but suspicion would tend toward option b. This conjecture is supported by the observation that in later network connectivity experiments (see Chapter 7) nearly all neurons tested were observed to drive AP responses in other neurons, meaning that the current stimulus is effective at nearly the 100% level.

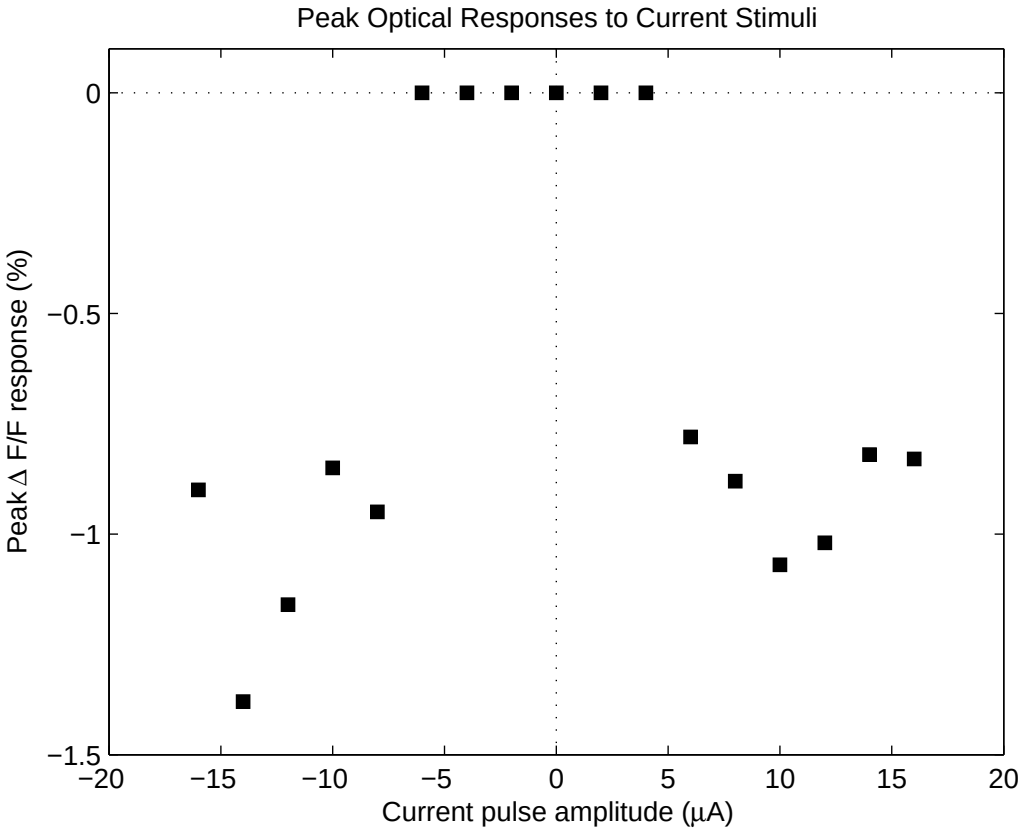

Figure 5.5: Peak $\Delta F/F$ responses versus stimulus strength. The sharp discontinuities occurring at +6 and -8 μA are identified as threshold stimulation currents.

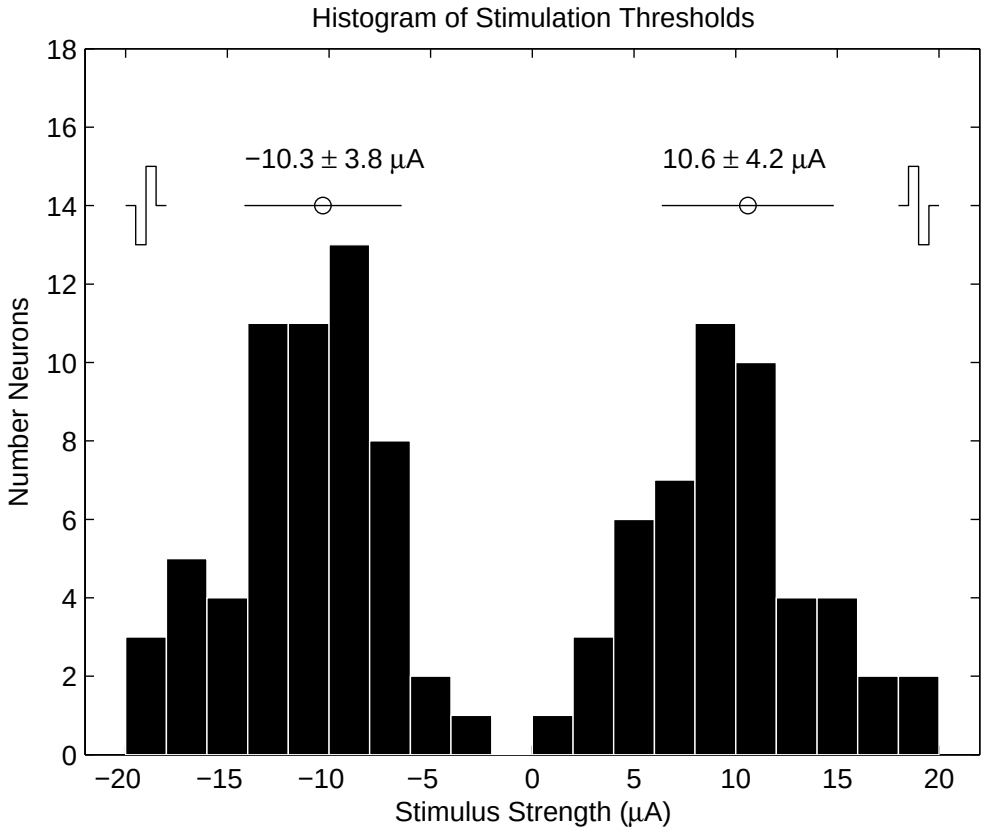

Figure 5.6: Histogram of bipolar current thresholds for N = 66 neurons. The average threshold values are about 10 μA for both positive and negative first stimuli.

Contrary to Maher [43], I found that both polarities are effective at eliciting action potentials. Wagenaar [68] found that positive-first stimuli were more effective than negative-first stimuli. However, in this study in the neurocage arrangement, I have found that negative-first stimuli are more effective.

It should be noted that the 66 neurons tested were tested under several combinations of: on/off-center electrode geometry; long/short tunnels. No matter the cage design, the threshold current distribution was comparable for all such combinations, with slightly more variation for the off-center electrode cage geometry. This can be explained by the fact that the neuron was sometimes far and sometimes close to the electrode. For centered electrodes, the symmetry dictated that the neuron-cage geometry was essentially invariant. Hence, with greater variance in relative position, it should come as no surprise that the threshold currents are slightly more variable as well.

A subset of N = 28 cells which were found to respond with APs were selected for further testing. The additional experiments examined a) responses to mono-polar stimuli to help elucidate biophysical origin of stimulation mechanisms, and b) responses to stimuli which were 0.2 and 0.3 msec per phase to find the appropriate stimulus width that should/must be delivered in practice.

The result of the monopolar stimulus experiment was that both positive and negative stimuli were capable of eliciting APs, but neither was as effective as the bipolar stimuli. Figure 5.7 represents the relative efficacy of each type of stimulus. The stimulus type (waveform) is shown at the bottom. Wagenaar obtained similar results [68], although it remains to be clearly understood why bipolar are more effective than monopolar pulses. That either polarity of current can elicit APs is consistent with the theory outlined in Section 5.2.

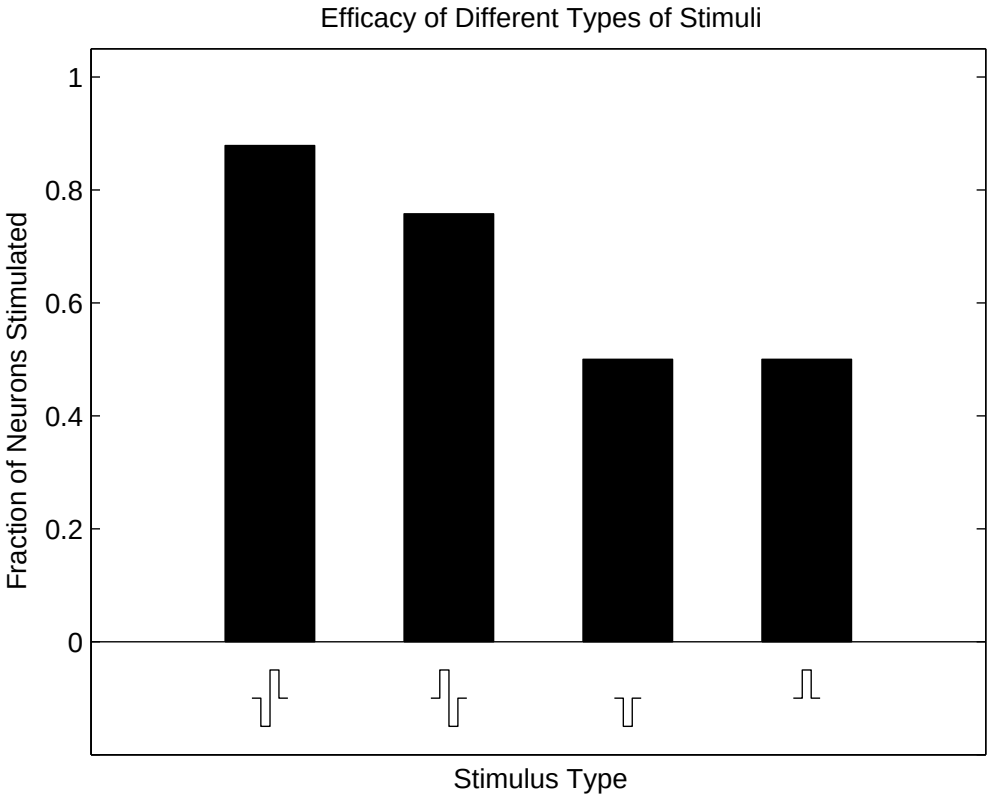

Figure 5.7: Relative efficacy of different pulse waveforms. N = 66 neurons were tested for bipolar stimuli. N = 28 were tested for monopolar stimuli.

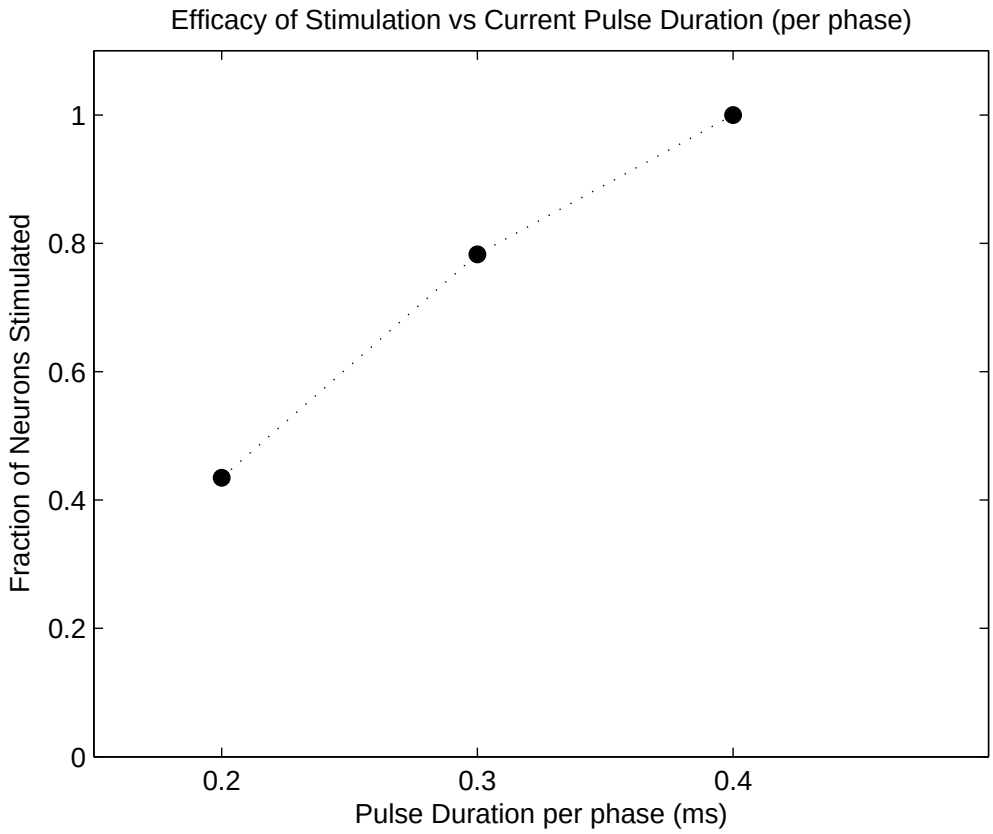

Figure 5.8: Dependence of stimulation efficacy on current pulse width

The result of the variable pulse-width experiment, displayed in Figure 5.8, revealed that efficacy of stimuli indeed strongly depends on the pulse width. N = 23 neurons were examined for this experiment. 0.2 msec stimuli were only about half as effective as 0.4 msec stimuli. While it is desirable to minimize the stimulus duration to minimize charging of the electrode (see Section 5.8), it is clear that stimuli must be at least 0.4 msec per phase. (Stimuli lasting for longer durations were not rigorously tested, but data did show that stimuli longer than 0.4 msec were no more effective.) The results suggest that sodium channel kinetics operate on a 0.4 msec timescale. That is to say, the cell must be depolarized for at least 0.4 msec to allow for the initial influx of sodium to generate the AP.

5.8 Safe Stimulation

While results presented in the previous section show that stimulation with current pulses is reliable, it must also be demonstrated that these pulse are safe—i.e., the stimulus must not harm (kill) the cell.

5.8.1 Current Pulses

Bipolar current pulses in the range of 0–20 μA, 0.4 msec per phase, were found to be non-harmful to the cell based on the following results:

1. Current stimuli were presented to a culture that was not stained. Stimuli were presented to the culture at a 1 Hz for 1 minute (a total of 60 current pulses for each neuron). Viewing the culture in Nomarksi optics before and after the experiment, no difference in visual appearance was noted. The cells continued to have a healthy appearance—plump soma, no decaying or fragmented neurites—even days later. Three cultures tested in this way all yielded the same result.

2. For more rigorous evidence that current stimuli are not harmful, a cell which was noted to be successfully stimulated by current pulses was tested again 2 hrs after the original experiment. Even 2 hours later the cell responded with AP responses in the exact same way. This experiment was repeated for N = 3 cells.

3. In later network connectivity experiments (see Chapter 7) stimuli were presented across the neurochip culture every couple of days while the synaptic responses were noted. Cultures were followed for weeks at a time with very little cell death noted, based on appearance. Moreover, as evidenced by enriching network connectivity over time, the neurons clearly must not have been harmed by the stimuli. Long-term survival rates in stimulated cultures were not in any systematic way different from that in non-stimulated cultures.

4. Only in one rare instance was a current stimulus noted to immediately harm (kill) a cell. During one of the optical dye experiments, a 16 μA current pulse killed a cell instantly. Previous to the stimulus the cell was noted to have nice edge-effect staining. Immediately after the stimulus the cell was seen to have a large amount of internalized

dye. This can only happen if the membrane is compromised. Following that stimulus, the cell would no longer respond to any stimuli—most likely it was dead. Again, this happened only once in 66 trials.

Based on these findings, bi-polar current pulses were deemed to be safe at the 98% level.

5.8.2 The 1 Volt Limit Relaxed

A common (mis)conception among electrophysiologists is that "going over 1 Volt" is dangerous.

Recall that the extracellular electrode-to-bath connection is modeled as a capacitor. A constant current stimulus charges this capacitor, changing the voltage across it according to:

$$\Delta V_c = \frac{I_{stim}\Delta t}{C_{elec}}. \tag{5.4}$$

Based on a report from 1972 [46] involving large-scale ($\approx$ 1 mm) electrodes, it was determined that water was electrolyzed when the capacitor charged to about 1 volt, producing oxygen and hydrogen gas bubbles. These gases are deleterious to cell membranes. J. Pine (personal communication) also reported having seen (N = 2) neurons be harmed immediately following current stimuli charging the electrode up to 1 volt. Through the years, this rumor has propagated so that it is taken as cold, hard fact.

However, from my data this notion can be relaxed. The capacitance for the neurocage electrodes was measured to be $C_{elec} \approx 4000$ pF. Bipolar current (charge-balanced) stimuli at 16 μA lasting 0.4 msec per phase were routinely used without any adverse effects to neurons noted.

This would correspond to charging the electrode-bath capacitance to:[1]

$$\Delta V_c = \frac{(16\,\mu\mathrm{A})(0.4\,\mathrm{m}sec)}{(4000\,\mathrm{p}F)} = 1.6\,\mathrm{V}. \tag{5.5}$$

Therefore, it is posited that 1 V is not a hard limit and stimulation is safe, up to about 1.5 V.

[1] Note that the pre-amplifiers (see Appendix C.1) saturate at about 1.1 V, limiting the electrode voltage that can be directly measured to this value.

5.8.3 Voltage Pulses

A "voltage pulse" consists of applying a constant voltage pulse to the electrode. In effect, this creates a brief, very large current transient which produces a voltage drop in the culture medium. Note that a voltage pulse does not mean the potential of the bath is moved directly. Using the electrode model discussed in Section 3.1, the electrode RC time constant is found to be $\tau_{elec} = R_{elec}C_{elec} = (25\text{k}\Omega)(4\text{n}F) = 100\mu sec$. This time constant dictates the duration of the current produced by the voltage pulse. The maximum amplitude of the current flowing due to the voltage pulse is computed as $I_{vp} = V_{stim}/R_{elec} \approx (1\text{ V})/(20\text{ k}\Omega) = 50\ \mu\text{A}$.

Wagenaar examined the efficacy of voltage pulses in mass cortical neuron MEA cultures [68]. His conclusion was that voltage pulses were more effective than current pulses, but no attempt was ever made to verify they are indeed safe. Indeed such an experiment would be impractical in that system due to the dense population of cells surrounding the electrodes.

I also examined voltage pulse stimulation, with the advantage of studying one cell at a time. Voltage pulses were found to be lethal in many cases. The optical response to stimulation was typically an extremely short-duration spike (on the order of 100 μsec). The time resolution of the RedShirt CCDcamera did not allow precise measurement of the width of the response, but it is safe to say that it is much less than 0.5 msec, probably closer to 0.1 msec—about the width of the expected large transient current. Sweeping through a range of voltage pulse amplitudes and noting the peak $\Delta F/F$ response, cells exhibited a curious "linear threshold" response, as shown in Figure 5.9—the amplitude of the peak $\Delta F/F$ response scaled linearly with the voltage amplitude. Note also, that the sign of the voltage pulse and $\Delta F/F$ were the same. Taken together, these data strongly suggest that the voltage drop generated by the current directly drove the potential across the membrane. This very likely means that voltage pulses severely compromised the cell membrane—probably irreversibly electroporating the neuron. In other words, the neuron was killed.

Of 25 cells tested, 19 exhibited this linear threshold response. 6 neurons were apparently stimulated without harm. Nonetheless, the main conclusion is that voltage pulses are lethal and should not be used in the neurocage regime.

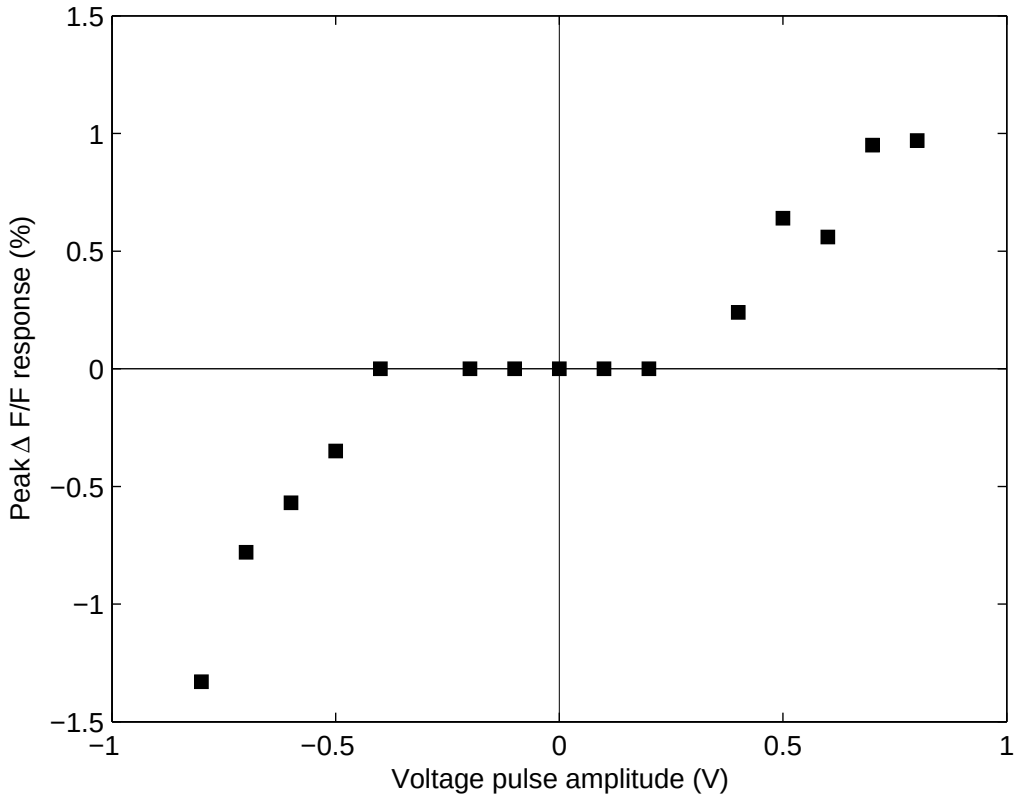

Figure 5.9: Peak optical responses $\Delta F/F$ to voltage pulse stimuli. The linear trend in the curve strongly suggests that the neuron cell membrane was severely, irreversibly damaged.

5.9 Conclusions

The results presented in this chapter show conclusive evidence that indeed caged neurons can be safely and reliably stimulated by the nearby electrode. The theory outlined on extracellular stimulation, moreover, is consistent with the results. In practice, for subsequent network mapping studies (see Chapter 7) stimuli 16 μA in strength, bi-phasic, negative-phase first, and 0.4 msec per phase in duration were used to probe the culture. This value was chosen as a compromise between stimulating almost all cells and avoiding any potential harm due to charging the electrode any higher than necessary.

Chapter 6

Extracellular Recording of Action Potentials with Neurochip Electrodes

6.1 Introduction

The purpose of this chapter is to demonstrate that the neurochip electrodes and associated hardware can record action potentials from the nearby neuron. This chapter is organized as follows: First the relevant theory will be briefly reviewed in order to understand the shape and size of expected extracellularly recorded AP signals. Then results of extracellular recording will be presented. It will be shown that the theory and experiment are in good agreement. This should not come as a big surprise, as extracellular recording of action potentials has been a well-understood and well-documented phenomena for at least 25 years [52, 28, 27].

6.2 Origin of Extracellular Signals

In the simplest terms, extracellular signals follow from Ohm's law. When a neuron fires an action potential it is, roughly speaking, a source/sink of current. Recall from Section 5.2.2 that physiological saline is an isotropic resistive medium—current flow within it generates a voltage drop. Putting the two together explains the origin of extracellular signals. The above statements will now be examined in more detail.

6.2.1 The Neuron as Current Source/Sink

When a neuron brought to threshold fires an AP there are two distinct phases to current flow:

1. Initially the voltage-gated Na^+ channels are activated. Sodium ions flow into the cell raising the intracellular membrane voltage, $V_m(t)$. The rise in membrane voltage is of order $\Delta V_m(t) \approx 100$ mV, and this change takes place in a time $\Delta t \approx 1$ msec. As the membrane voltage rises, the Na^+ channels inactivate, shutting off the active sodium current.

2. While the membrane voltage is still high, the active K^+ channels open. Potassium ions flow out of the neuron back into the surrounding medium, restoring the membrane voltage back toward the resting voltage. This decrease in $V_m(t)$ must also be of order 100 mV, and the kinetics of active potassium channels dictate that K^+ current flows for a relatively longer time, about 3 msec.

The values cited above are well known, and also match patch clamp recordings from hippocampal cells in control cultures in the Pine lab.

Estimates for the magnitude of the Na^+ and K^+ current can be obtained by considering the passive membrane properties of a neuron (see Section 5.2.1). A neuron at resting potential is essentially a charged capacitor—sodium and potassium ions are separated by the cell membrane. When the voltage-gated Na^+ channels open, the capacitor discharges as sodium ions flow into the soma. Similarly, for potassium current, the activated K^+ channels recharge the membrane capacitance. Recall that for a parallel RC circuit (R represents the passive membrane "leak" resistance here), the sodium and potassium currents can be approximated as entirely capacitive—current through the resistor can be neglected—for times $t << \tau_m$. Hence, the total current is approximately equal to the capacitive current:

$$I_{total} \approx I_c = C_{lm} \frac{\Delta V_m}{\Delta t} \tag{6.1}$$

where C_{lm} is the membrane capacitance local to the region of the membrane voltage change. For example, if the soma is brought to threshold (but the AP has not yet propagated to the distal dendrites), then C_{lm} represents the sum of the capacitances from the soma and nearby dendrites only. Recall from Section 5.2.1 that an estimate of the whole-cell capacitance is

60 pF and the capacitance due to the soma alone is approximated as 15 pF. Also recall that the membrane resistance of hippocampal neurons has been measured by others to be $R_m \approx 100\text{M}\Omega$.

Finally, then, an estimate for the sodium-current amplitude for an AP initiated in the perisomatic region is given by:

$$I_{Na} = -(15\text{ pF})\frac{(100\text{ mV})}{(1\text{ ms})} = -1.5\text{ nA} \tag{6.2}$$

and an estimate for the potassium current amplitude is:

$$I_K = -(15\text{ pF})\frac{(-100\text{ mV})}{(3\text{ ms})} = +0.5\text{ nA}. \tag{6.3}$$

Note the - sign on I_{Na} indicates an inward current, while the + sign on I_K indicates an outward current of positive ions. I'll make use of these current estimates below (Section 6.3) to get an estimate of the size of the expected extracellular voltage signal.

6.2.2 Critical Importance of Distributed Current Loops

Of particular importance for extracellular recording is that the ions flowing into (out of) one location—the soma and axon hillock, for instance—are electrotonically conducted down the cell interior, back out across the capacitance of the cell membrane, and finally back toward the active region to complete the current loop.

Another way to think of this process is the following: Na^+ ions entering the soma through voltage-gated channels create excess positive charge on the inside which, due to Coulomb forces, "push" interior ions away from the region of active Na^+ channels. Ions are passively (electrotonically) conducted down the still-passive segments (dendrites) where they cause local changes in the membrane voltage, thereby generating a capacitive current back out across the cell membrane. This outward capacitive current, in turn, returns to the exterior of the soma. A similar loop—in the opposite direction—is generated by the active potassium current.

The key feature to note is that an AP initiated in in the perisomatic region generates extracellular current flow along the length of the dendrites. It is this *distributed* current return in the saline solution that gives rise to the voltage difference—Ohmic drop—sensed

by the neurochip electrode. If the cell had no dendrites or passive processes, no external current would flow in the medium. For example, in the case of a spherically symmetric neuron with no dendrites, an AP essentially shorts the two sides of a capacitor together. Outside of the cell, the total enclosed charge is zero, so (by Gauss's Law) the electric field is zero everywhere; therefore the voltage difference between two points chosen anywhere outside the cell must also equal zero.

6.3 Expected Extracellular Signals Size and Shape

A neurochip electrode measures the voltage difference relative to a distant ground electrode (platinum wire in dish; all channels are referenced to the same ground electrode).

To estimate the size of this signal, I need to make 4 more approximations:

1. APs are initiated in the perisomatic region and the soma is electrically excitable. Changes in membrane potential at the soma are due to locally activated Na^+ channels. This means local inward sodium current dominates over any capacitive outward current spreading to the dendrites.

2. For a well-developed, highly-branched dendritic tree the extracellular current flows approximately isotropically. Hippocampal neurons older than 1 week *in vitro* fit this bill. With this in mind we can view a neuron as a symmetric current sink (source). An appropriate geometry can be chosen to match the neuron—either a hemisphere or a cylinder.

3. The length scale $\lambda = \sqrt{R_m/R_i}$ of a typical hippocampal neuron is large compared to the length scale of the neurocage tunnels. This means the majority of the outward capacitive current occurs outside of the cage region. This also means that the majority of the return current flows along a path from distant points exterior to the cage (i.e. "ground") toward the access hole and finally into the bottom region of the cage where the neuron is located. The resistance of the tunnels, therefore, can be neglected in this simple model.

4. The neuronal compartment closest to the electrode is the soma. In other words, the neurochip electrode is measuring a signal due primarily to the perisomatic current.

The two phases of the current for a prototypical AP discussed in the previous section generate potential differences which are spatially dependent. For the moment, consider what an electrode would measure if it were a infinitely small point electrode directly atop the neuron. This regime yields a estimate for the maximum expected signal to be recorded by neurochip electrodes.

With the above assumptions, the extracellular resistance—the R we need to apply Ohm's Law—is given by:

$$R_{exc} = R_{cage} + R_{spread}.$$

For the present computation, R_{cage} is the resistance from the cell body to the top of the cage (instead of from the electrode to the top of the cage). Approximating the neuron as a thin cylinder at the base of the cage, R_{cage} can be computed as was previously described in Section 5.2.2, but this time replacing the radius of the electrode with the radius of the neuron cell body, r_{soma}. R_{spread} represents the "spreading" resistance from the top of the neurocage access hole to infinity. Thus, the total neuron-to-ground resistance, R_{exc} can be computed as:

$$R_{exc} = R_{cage} + R_{spread} = \rho \left(\frac{h}{\pi r_{soma} r_{cage}} + \frac{1}{4\pi r_{cage}} \right) \tag{6.4}$$

where each of the terms in 6.4 correspond to those discussed in Section 3.4.1. Plugging in numbers ($r_{soma} \approx 10$ μm) yields a theoretical value of $R_{exc} = 24$ kΩ.

Finally, an estimate for the extracellular signal due to the inward Na^+ current is:

$$V = I_{Na} R_{exc} = (-1.5 \text{ n}A)(24 \text{ k}\Omega) = -36 \, \mu V.$$

where an estimate for I_{Na} was derived in equation 6.2.

This signal should last approximately 1 msec, i.e., the signal due to sodium current should have a full-width of about 1 msec. The biophysical origin of the sodium spike is depicted in Figure 6.1(a).

An estimate for the voltage generated by the outward K^+ current is:

$$V = I_K R_{exc} = (+0.5 \text{ n}A)(24 \text{ k}\Omega) = +12 \, \mu V$$

where an estimate for I_K was derived in equation 6.3.

This signal should last approximately 3 msec, i.e., the signal due to sodium current should have a width of about 3 msec. The origin of this potassium signal is depicted in Figure 6.1(b).

So, putting it all together, the expected prototypical extracellular signal measured by a neurochip electrode due to an AP in the soma should contain an approximately 1 msec downward spike 36 μV in amplitude, followed by 3 msec wide deflection with a peak amplitude of about 12 μV. This prototype waveform is indeed observed in neurochip recordings.

6.4 Recorded Action Potential Signals

6.4.1 Typical AP Signal

Figure 6.2 shows a prototypical recording of an AP. Three successive APs are aligned, indicating the spikes recorded on this channel are stereotyped, as one would expect. A control trace from a neurochip electrode (dotted) is also shown. A trace of the intracellular voltage during an AP obtained by patching onto a neuron in a control culture (dashed) is overlaid to help illustrate the relationship between the intracellular membrane potential and the extracellularly recorded signal. The large, fast (-45 μV, 0.5 msec) signal from the sodium spike is evident, as is the smaller, slower (+15 μV, 1.5 msec) potassium signal.

6.4.2 Departures from Prototypical Recordings

Surely we expect to also expect to see–and in fact have observed—a range of signal sizes dictated by the range of cell capacitances and variable relative position of neuron and electrode.

In the real situation, the electrode is neither infinitely small, nor is it positioned directly atop the neuron. A finite-sized electrode will measure the potential spatially averaged over its area. Of course, as the position of the electrode moves farther away from the neuron, the signal size will decrease. If we consider the neuron to be a thin disk (cylinder) of uniform current, then following the work of Wiley and Webster [71] yields an approximation for the expected losses in the recorded signal due to increasing neuron-electrode distance. They consider a model of a thin disk in an infinite, isotropic medium, but for the region not too far from the electrode the solution is still qualitatively correct. The maximum neuron-electrode

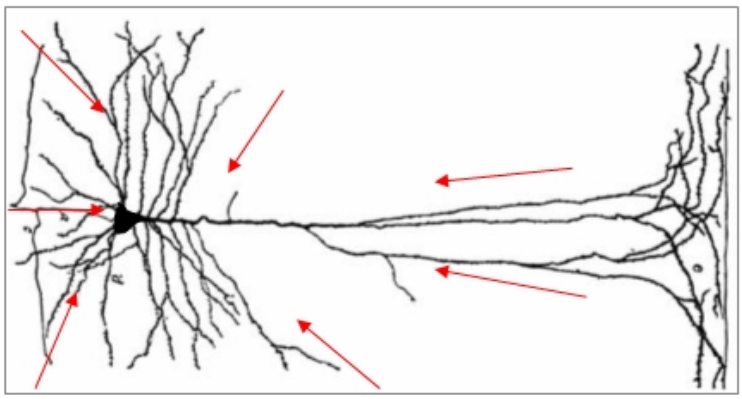

(a) Sodium inward current flow during the initial phase of the action potential. A local inward current at the soma causes a redistribution of charges along the neural processes. This inward current will be recorded as a negative potential by an electrode situated adjacent to the soma. The signal component is expected to last for about 1 msec.

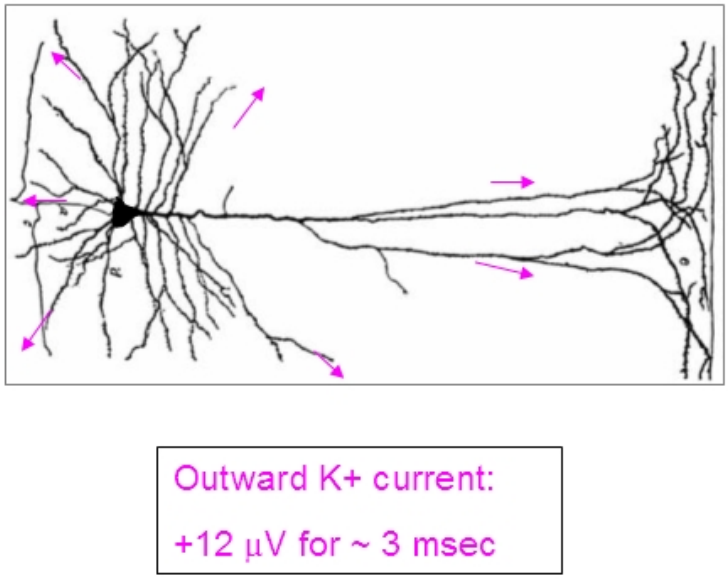

(b) Potassium outward current flow during the second phase of the action potential. The soma is sourcing outward current which redistributes charges along the length of the neural processes. This outward current will be recorded as a positive potential by an electrode situated adjacent to the soma. This signal component is expected to last about 3 msec.

Figure 6.1: The two phases of the expected neurochip signals recorded from an AP are generated by the active sodium and potassium currents.

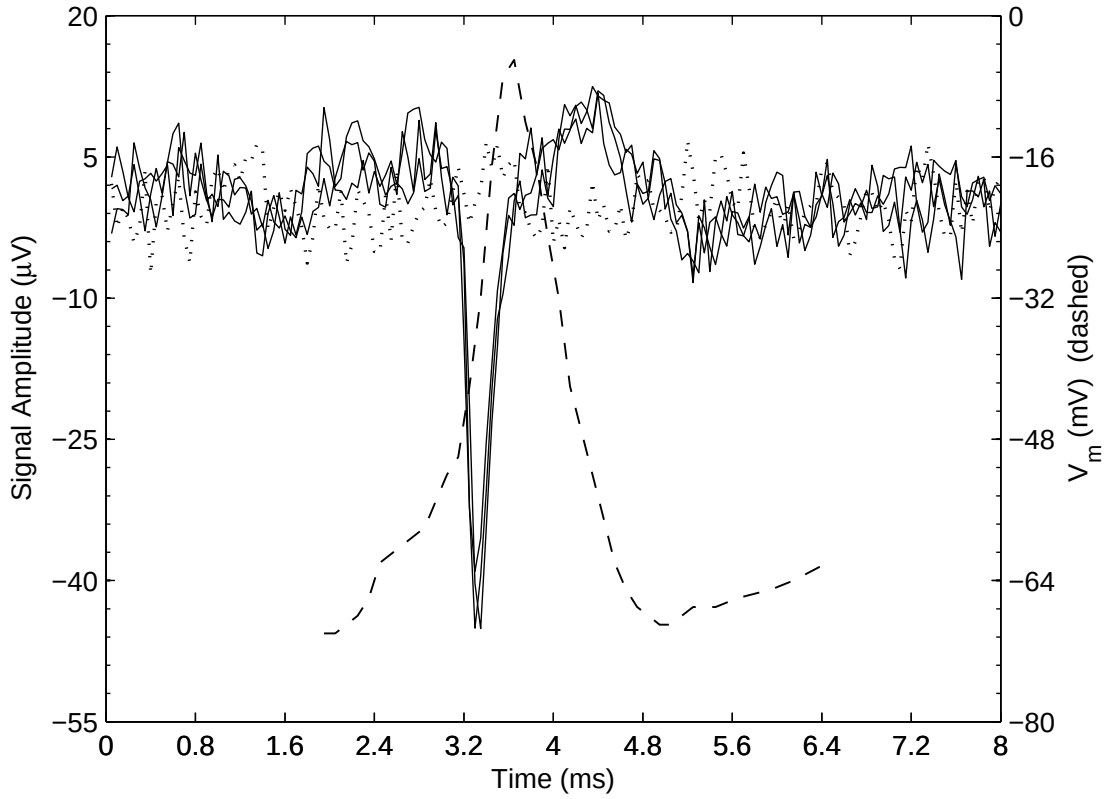

Figure 6.2: Prototypical neurochip extracellular recordings of APs. 3 successive APs are aligned at their maximum deflection to demonstrate the stereotyped shape of the signal (solid, black). Typical intracellular recording of membrane potential during an AP is overlaid (dashed) to illustrate the origin of the components of the extracellular signal. A control extracellular trace (dotted) is also shown for reference.

distance is constrained by the cage geometry to be no more than about 20 μm center to center. In many instances the electrode is very close—adjacent to—the neuron. Therefore, the average distance between neuron and electrode centers is on the order of 10 μm, or one electrode diameter. In this regime, the electrode would record about half the maximal signal.

6.5 Zoology of Recorded Signals

The previous section considered only a prototypical extracellular recording. As detailed above, the prototypical signal is bi-phasic, consisting of a large downward spike due to sodium current, followed by a broader, lower amplitude hump due to potassium current.

Other signal shapes are possible, however. Indeed they are recorded and shown in Figure 6.3. This section reviews the origin of other signal shapes drawing on the work of [14, 25, 54, 31].

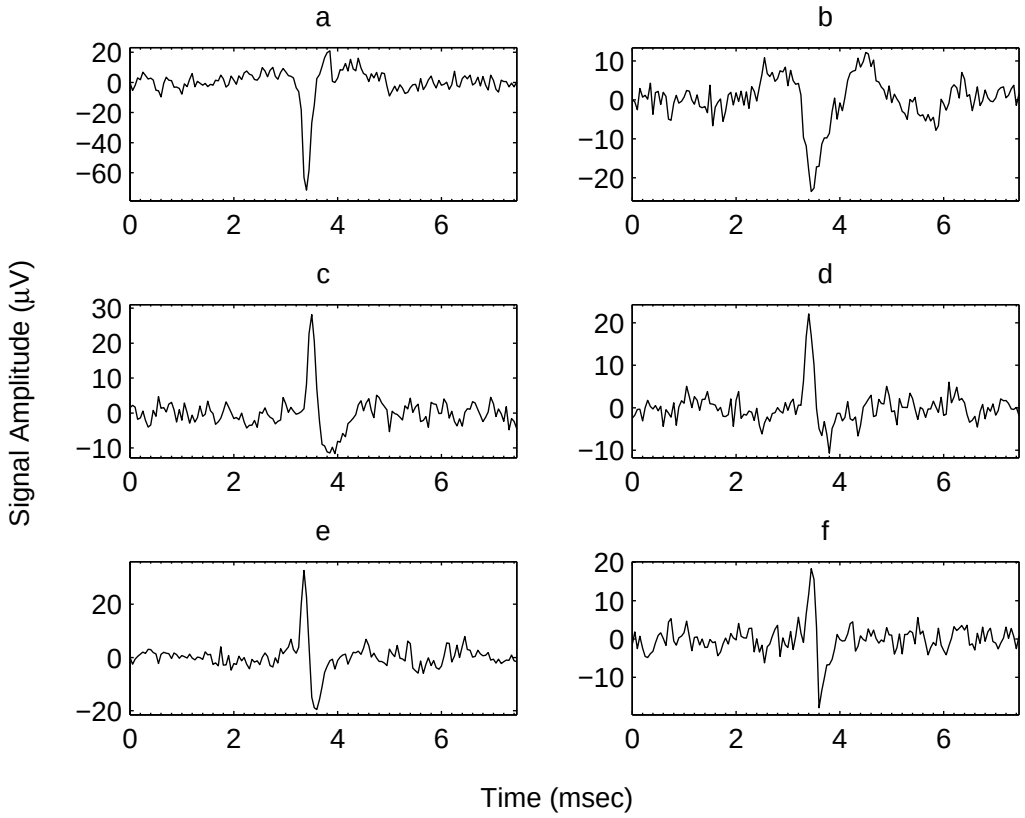

Figure 6.3: Various types of extracellular recordings of APs acquired with the neurochip. Waveforms can be broadly categorized into 3 types: negative spikes (a and b), positive spikes (c and d), and tri-phasic waves (e and f). Neurochip signal amplitudes generally were in the range of 20–70 μV. Note the vertical scale bars are different for each waveform plotted. See text for explanations of the origins of each type of signal. Waveforms (a) and (b) are both downward spikes. For waveform (b), positive-going components can be seen before and after the main downward sodium spike. The initial upward deflection is probably due to outward capacitive current due to slow dendritic stimulation, whereas the latter upward deflection is most likely due to the rectifying potassium current. Waveform (c), but not (d), exhibits a negative deflection after the initial upward spike. Waveform (e) contains positive- and negative-going phases which are asymmetric in amplitude and duration, while waveform (f) is symmetric in both respects. Note that features as small as 10 μV in amplitude can be cleanly identified as spikes.

6.5.1 The Positive Peak

Some recorded signals appear to be flipped on their head—the main component of the signal is an upward spike, or positive peak. An example of this type of signal is shown in Figure 6.3, waveforms c and d. The origin of the positive peak can be explained in two possible ways.

One explanation for these positive peaks is that the membrane excitability of the soma is low compared to the axon. In this case the soma acts primarily as a passive compartment. (Assumption number 1 in Section 6.3 is violated.) When the axon fires an AP, it sinks current from the extracellular medium, a large fraction of which flows into the soma. The current loop is completed by the capacitive current flowing back out of the cell body toward the axon. When the soma acts as a current *source* instead of a current sink, the sign of the recorded perisomatic extracellular signal is flipped. The extracellular signal "sees" sodium current (positive ions) flowing away from it, so records a positive peak. The positive peak magnitude should be comparable in size to the negative-going sodium spike (e.g., the signal in Figure 6.2); provided dV/dt is approximately the same in both the axonal and somatic compartments, the computation for the outward capacitive current at the soma is the same as before ($I = C_{soma}dV/dt$), with only the sign of I flipped.

Another explanation for positive peaks is as follows: The electrode is recording primarily from the apical dendrite trunk. (Assumption 4 in Section 6.3 does not hold.) This could be the case if the neuron is positioned opposite the electrode and the dendrite passes very close to, or over it. In this regime an initial positive spike is recorded because the part of the dominant current flowing into the axon hillock (perisomatic) region flows capacitively out of the apical dendrite. No negative waveform is recorded here because both the dendritic capacitance and sodium current conductance is relatively small. Thus, only a positive spike is recorded.

6.5.2 Tri-Phasic Waveforms

Tri-phasic waveforms result from a superposition of currents in space and time. The tri-phasic signal consists of an initial positive peak, followed immediately by a negative phase (see Figure 6.3, waveforms e and f). Assume the electrode is recording at a point close to the soma. Then the simplest explanation for tri-phasic shape is that at first the soma

is passively sourcing capacitive current to a distant, already active region of the neuron before being brought to threshold. This could happen, for instance, if the dendrites become de-polarized/excited to sink current locally before the AP is generated in the soma, and it begins sinking sodium current itself.

The delay between these two events dictates the temporal width of the upward component in the signal. Usually the delay is about 0.5 msec. The size of the initial upward component in the signal will be proportional to the amount of total current flowing into distant dendrites or axons. Typically, the positive phase has a magnitude less than the negative phase.

6.6 Can Subthreshold Signals Be Recorded?

Given that the typical noise level of the neurochip recordings is about 3 microvolts RMS, is it reasonable to expect that *subthreshold* signals can be cleanly recorded?

The answer: almost certainly not.

Consider that a subthreshold signal changes the membrane potential about 1/10th as much as an AP. Then the current flow into (or out of) the cell will also be 1/10th the estimate for AP currents. The expected signal is proportional to this current, so is also reduced by a factor of ten. Considering the estimates of signal magnitude derived in Section 6.3, this means we would hope to see signals on the order of 2–7 μV. That calculation is actually an over-estimate because the location of the EPSC occurs at distributed sites not necessarily close to the soma. In other words, for a subthreshold excitation, the current source/sink is farther away from the electrode than the soma. The maximum size of a subthreshold signal is about 5 μV, which is lost in the noise. Therefore, the neurochip is essentially limited to recording action potentials.

6.7 Specificity of Recorded Signals—Can Passing Axons Be Recorded?

The signals captured by an electrode are believed to be specific (solely due) to the neuron in the same cage based on the following argument. A segment of a passing axon originating

from a neuron in another cage is expected to have a capacitance of:

$$C_{pa} = c_m 2\pi r_{axon} L$$

where c_m is the capacitance per unit area of membrane, taken to be 1 μF/cm^2; r_{axon} is the radius of the axon ($\approx$ 1 μm); and L is the length of the segment passing by an electrode ($\approx$ 10 μm). Plugging in numbers yields a value of $C_{pa} \approx 0.5$ pF. The electrode area is much greater than that of the passing axon, so the measured signal is the spatially averaged potential generated by an AP in the axon. In the case that an axon is adjacent to the electrode, the average axon-electrode distance, d_{avg}, is about 5 μm. Therefore, the amplitude of the signal component due to the inward sodium current is given by:

$$V = \frac{\rho I}{2\,\pi d_{avg}} = \frac{(70\,\Omega \cdot cm)(0.5\,\mathrm{p}F)(100\,\mathrm{m}V/\mathrm{m}s)}{(2\,\pi)(5\,\mu\mathrm{m})} \approx 2\,\mu\mathrm{V}.$$

Clearly the signal due to the passing axon is too small to cleanly record at a Johnson noise level of 3 μV. A similar argument holds for recording from the dendrites of another neuron. For recording, therefore, a 1:1 correspondence between electrodes and neurons is believed to be firmly established.

6.8 Conclusion

The neurochip electrodes and associated electronics are capable of recording AP signals. The extracellular recordings have high SNR, typically in the range of 5–20. Simple models were proposed which yielded estimates for the size and magnitude of the recorded signals. Theory and experiment are in good agreement. Several waveform types have been recorded, all corresponding to APs. The origin of the shape and size of the signals observed is explained by considering the passive and active properties of neuron membranes, as well as the position at which the electrode records.

Chapter 7

Probing Connectivity in Small Cultured Neural Networks

7.1 Introduction

The previous chapters have demonstrated that:

1. Neurocages trap neurons with near 100% efficiency whilst still allowing dendrites and axons to grow normally.

2. The extracellular electrode reliably elicits AP responses from the adjacent neuron.

3. The extracellular electrode records AP signals from the adjacent neuron with high SNR.

Finally, it is time to put it all together. This chapter describes and shows the results of proof-of-concept experiments to demonstrate the capability of the neurochip in probing network connectivity. These experiments demonstrate that suprathreshold connectivity in small *in vitro* neural networks can be mapped at the single-cell level over three weeks (or longer). Herein lies the real power of the neurochip: the complete connectivity of culture can be mapped noninvasively, with single-cell resolution, at any time point, so that the culture development can be interrogated over the lifetime of the culture. This domain has been previously inaccessible to neurobiologists.

This chapter starts by providing an overview of the experiments, and then covers the specifics of data collection, processing, and analysis.

7.2 Why Connectivity Matters

One of the fundamental properties—if not *the* single most important property—of a circuit is its connectivity, i.e., how it is wired. It is known that the neural network wiring diagram has significant implications for neural information processing [12] and for the kinds of computations different circuit structures achieve [16]. In many *in vivo* sensory systems, use-dependent and experience-driven activity regulate the development of neural circuits, modulating neurite branch stability and synaptogenesis [33, 74]. Even cultures of dissociated neurons exhibit activity-dependent plasticity. Synapse strength and pathways are modified based on the type and timing of stimuli presented to the culture [5, 35, 57, 45]. Understanding the normal connections that form *in vitro* and how they develop and are modified will provide valuable insight into how *in vivo* networks develop and function.

7.3 Simple Network Connectivity Experiment—Overview

The basic idea to probe or "map" connectivity in a culture is straightforward: stimulate one neuron and record/examine the responses of all others. An AP evoked in a neuron other than the one being stimulated with the current pulse is termed a "driven network response". Repeating this for all cells in the culture probes all stimulus-response pairs, or all pre- and post-synaptic pairs. This experiment also yields data to define the synaptic delays. The procedure described above was repeated N = 10 times for better statistics, to investigate the "reliability" (probability of occurrence of a driven network AP response), and to observe the jitter in the response times.

Cultures investigated for these experiments were probed starting about 1 week *in vitro*, continuing every 2 or 3 days for the lifetime of the culture. Most cultures lasted about 3 weeks, some longer-lived ones stretched to 4 weeks and beyond. For an age range of 2–4 weeks *in vitro* cultures typically consist of six to ten neurons (see Figure 4.10). Neurochip cultures typically exhibited suprathreshold activity—spontaneous and/or driven—starting around 10–14 days *in vitro*. This time range for formation of functional synapses and network maturation is consistent with previous reports (for example, [3, 2, 42, 69]). As a general rule, cultures exhibiting spontaneous activity also exhibited driven network responses. The converse, however, was not always true: cultures not exhibiting spontaneous activity sometimes exhibited driven network responses.

For unknown reasons, not all "good looking" cultures developed suprathreshold synaptic connectivity. Out of 41 cultures grown and tested from March to July 2007, 17 did. In two cases I optically tested cultures which did not exhibit suprathreshold connectivity—driven or spontaneous—during electrophysiology experiments. In both cases, all cells in the culture were found to be excitable. This is an important bit of information: growing a small culture of electrically excitable cells does not guarantee they form synapses strong enough to drive APs in another cell.

7.4 Data Acquisition

7.4.1 Selection of Stimulation Electrodes

A network connectivity experiment begins by visually inspecting a neurochip culture to generate a list of electrodes to be stimulated. If a neurocage was judged to conain a healthy neuron, the corresponding electrode was stimulated. An inclusive approach was adopted so that a cage was still added to the list even when the neuron's appearance left the status of its health ambiguous. This visual approach was reliable for most cultures for ages up to about 3 weeks *in vitro*. Past that, glia would sometimes grow into the array, rendering visual inspection very difficult. In this case, the electrode list defaulted back to the latest one that was ascertained with a clear view. Thus, cultures older than about 3 weeks I could not ascertain which of the stimulated electrodes had a healthy neuron in its cage.

7.4.2 Computer-Automated Electrophysiology Sequence

Electrophysiology data from a neurochip culture was acquired in the following way: Each electrode in a list (see Section 7.4.1 above) was probed consecutively. Stimulation of the electrode sequence was repeated ten times. For example, if the list consists of electrode numbers x, y, and z, the stimulation sequence was [x, y, z; x, y, z; ... ; x, y, z] for a total of ten times. The inter-electrode delay—the time between stimulating one electrode and the next—was chosen to be 2 seconds, a sufficiently long time to avoid any paired-pulse stimuli plasticity effects [4].

To probe responses to a single electrode, two stimuli were presented consecutively: a control stimulus of 0 μV amplitude, followed immediately by the actual probe stimulus. The control stimulus data were used to measure the spontaneous firing rate immediately prior

to presentation of the probe stimulus. In this context, it is used to verify the absence of spontaneous APs. The probe stimulus was chosen to be -16 μA, bipolar, negative phase first, 0.4 msec per phase. This stimulus was chosen based on results presented in Chapter 5. This stimulus is expected to generate APs in $> 90\%$ of neurons, and is known to be harmless to nearly all cells tested with optical recording. Furthermore, throughout the stimulus-response experiment described here, visual examination before and after presentation of the stimulus revealed no damage or cell death.

Each trial (defined as the recorded responses to presentation of one stimulus to an electrode) was chosen to be 200 msec long, with the stimulus presented at 40 msec. This timing was chosen to have a post-stimulus window long enough to capture any resulting APs. 160 msec is a generous window because delay times were expected to fall in the range of 2–40 msec [24, 4].

After all electrodes in the list were probed 10 times, a baseline template file was generated by averaging 20 consecutive trials, each 200 msec long, with no presentation of a stimulus. A 60 Hz baseline template is generated for each electrode. These templates are utilized for off-line digital subtraction of baseline 60 Hz signals.

All recordings were initiated within a few minutes of extracting the culture from the incubator; data collection lasted about 10–15 minutes total. The temperature during recording was not precisely known, but decreased over time from 35° C to probably about 22° C.

All data were acquired with custom-built LabView programs (see Appendix D). The main program controlled data display and saving to hard disk. Data for each channel were digitized at 20 kHz (a total of 16 x 20 kSamples/sec = 320 kSamples/sec). Each trial was synced to a 60 Hz trigger pulse so that baseline 60 Hz signals could be subtracted offline. Another helper program automated the generation of stimuli and looping over the electrode list. The final product was a series of data files, each containing raw voltage versus time traces for sixteen channels by ten trials. The total number of data files generated was equal to the number of electrodes in the stimulation list times two. Each data file had a corresponding "description file" containing critical parameters of the data acquisition, such as which electrode number was stimulated, the stimulus settings, etc.

7.5 Pre-Processing—Artifact Removal

Raw data files containing voltage versus time traces need to be pre-processed, "cleaned" for further analysis. Pre-processing consists of two stages: 1) removal of the 60 Hz baseline and 2) removal of the stimulus artifact.

The first is accomplished by off-line digital subtraction of a baseline template. The template is computed as the average of N = 20 consecutive 60 Hz synced traces. 20 averaged traces are used to decrease the white noise present in the template. Averaging N traces reduces the noise by a factor of $\sqrt{N}$, a factor of about 4.3. Thus, when the averaged template is digitally subtracted from a single data trace, the noise level is increased by a factor of $\sqrt{1 + 1/(4.3)^2} = 1.03$. This increases the noise level only by 3%, which still permits high-SNR recordings of APs.

Stimulus artifacts are removed using the SALPA algorithm [70] implemented in MATLAB. (Cleaning 20 trials x 16 electrodes = 320 traces takes about 1 minute of off-line processing time.) Briefly, the SALPA algorithm locally fits a 3rd-degree polynomial to points in the trace following hardware signal saturation. This order of polynomial is of high enough order to make a good fit to the stimulus artifact shape, but low enough order that it does not over-fit and lose putative AP signal "spike" waveforms. As the signal comes out of saturation, it is possible that a good fit is not obtained. The voltage trace is zeroed out ("blanked") at these points, meaning that I am blind to recording post-synaptic responses during this time. Figure 7.5 shows the distribution of the observed blanking times for all trials from experiments conducted with 17 neurochip cultures. Both the probability- and cumulative-distribution functions are shown. The SALPA algorithm typically blanked out the traces for about 1–2 msec following the end of the stimulus. Blanking times > 2.5 msec were observed in about 10% of the trials. Blanking times exceeding 3 msec were observed in only about 2% of the trials. Therefore, we can expect to detect nearly all responses occuring with a delay time (see Section 8.4 for definition) of ≥ 3 msec. We can expect to detect about 75% of all responses occuring with a delay of ≥ 2 msec. Detection of responses occuring $\leq$ 1.5 msec delay will be successful at about the 50% level. Response times occuring with less than 1 msec delay will be missed quite frequently, about 75% of the time. However, I do not expect to see many responses occurring as fast (or faster than) 1.5 msec [24].

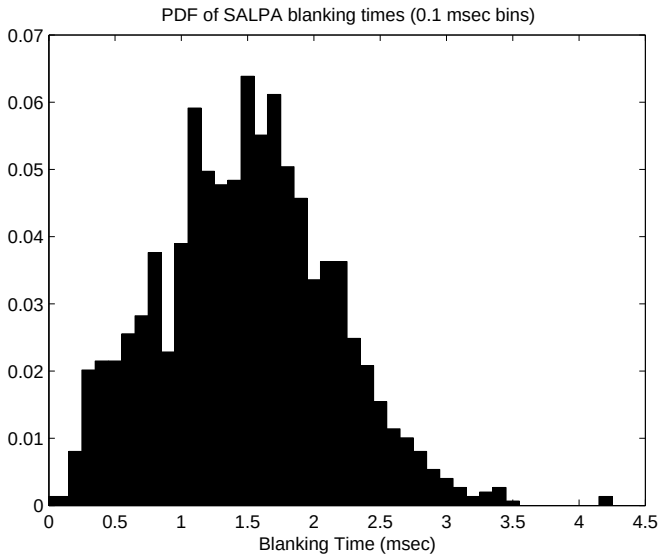

(a) Probability distribution of observed SALPA blanking times.

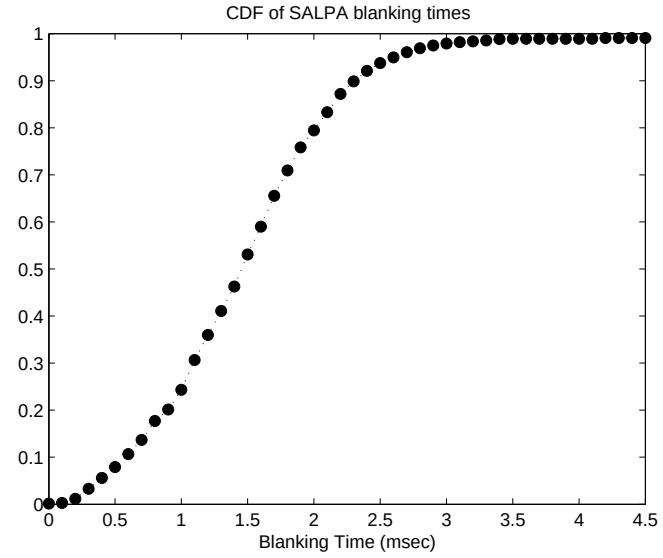

(b) Cumulative distribution of observed SALPA blanking times.

Figure 7.1: Distribution of the "blanking times" due to the SALPA algorithm. The blanking time is defined as the length of time following the stimulus for which the SALPA algorithm zeros out the trace. Most blanking times fell in the range of 1–2 msec. About 98% of observed blanking times were ≤ 3 msec.

7.6 "Cleaned" Traces

Figure 7.2 shows an example of data traces which have been fully pre-processed. One (of ten total) traces with probe stimulus presented to electrode 1 is shown for clarity. The trace on channel 1 at low gain shows the stimulus. The 60 Hz signal has been reduced to undetectable levels. The RMS noise level is about 2–3 μV on most electrodes. Action potentials responses are seen on electrodes (labeled "Chan", in Figure 7.2) 2, 3, and 14.

For recording electrodes, the data points adjacent to the time at which the stimulus was presented have been blanked out (set to zero) by SALPA for only about 2 msec, with blanking time slightly variable among electrodes. Spikes can be detected, therefore, starting about 2 msec after the stimulus. It is possible that some spikes resulting from very fast transmissions go undetected during this time. That is a limitation inherent in the system. The trade-off is that spikes riding on the stimulus artifact waveform can now be cleanly detected.

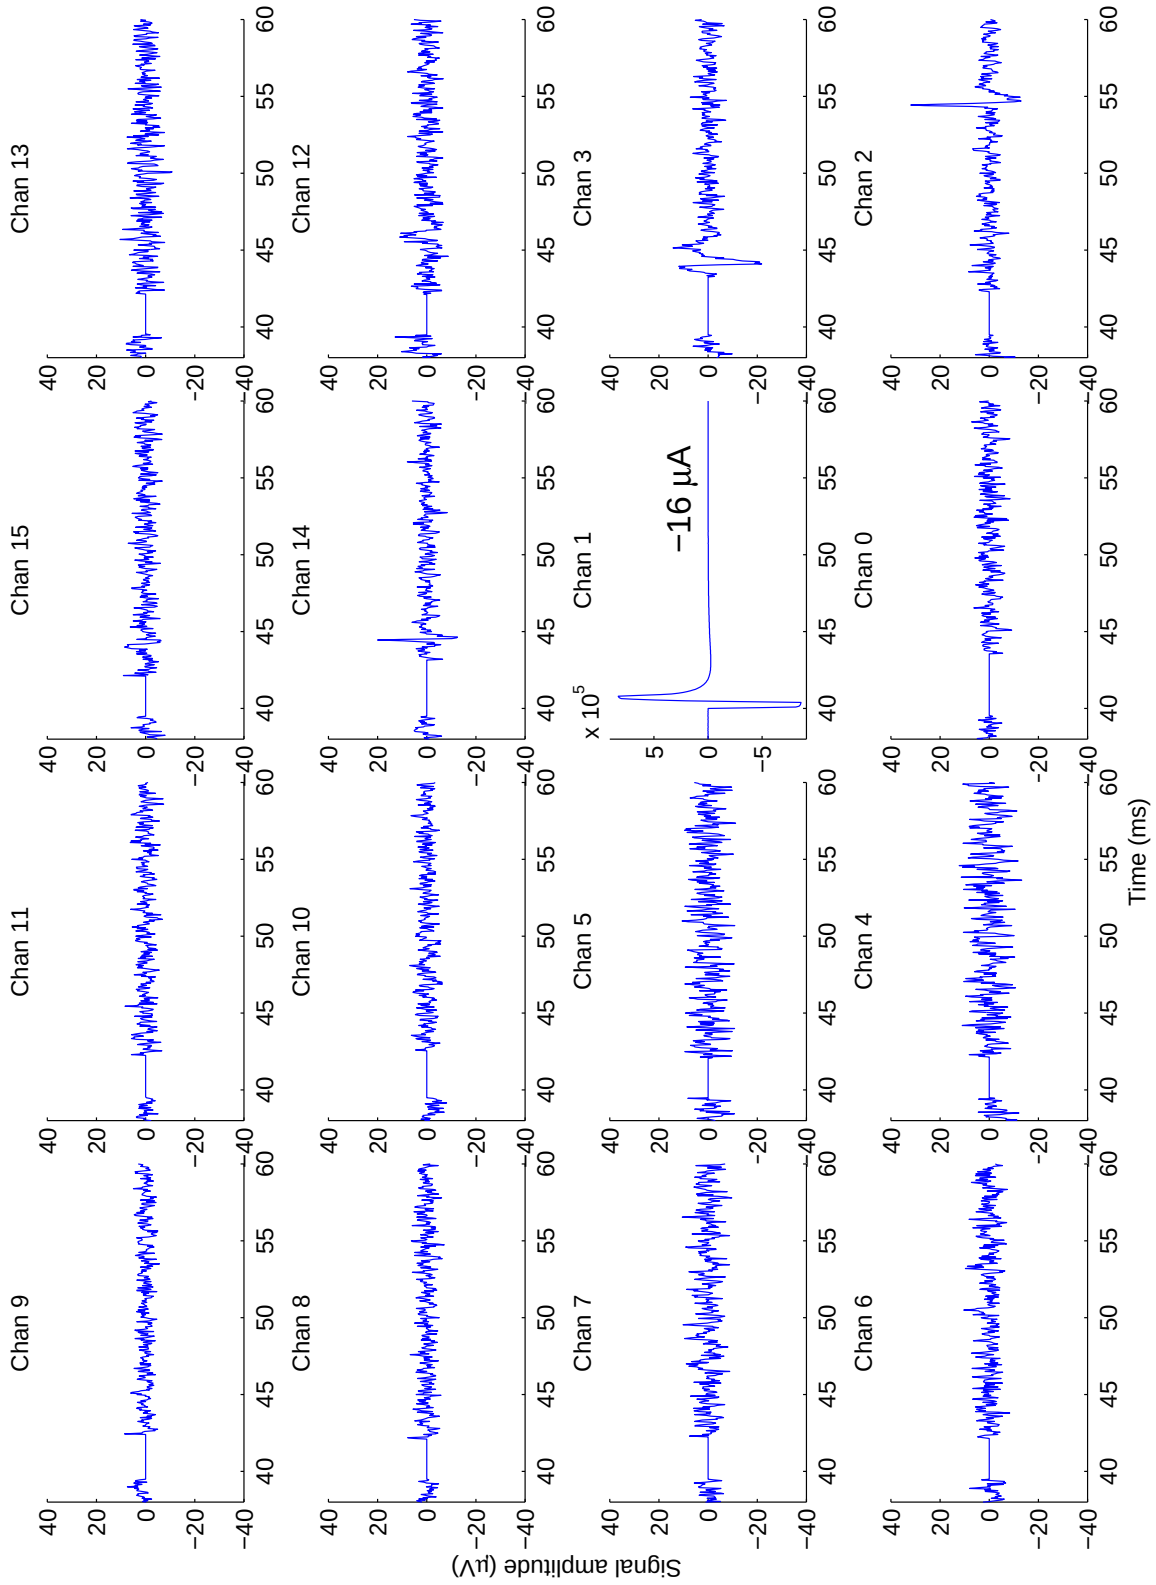

Figure 7.2: Pre-processed "cleaned" traces. The 4x4 layout matches the physical orientation of the neurocages. Only 1 of 10 trials is shown to aid clarity. A stimulus was presented on electrode 1 at time $t = 40$ msec, resulting in spikes on electrodes 2, 3, and 14 at times $t \approx 55$, 44, and 45 msec, respectively. The total trial lasted for 200 msec, but only the time range of 38–60 msec is shown for visual clarity.

7.7 Spike Detection and Raster Plot

Having artifact-reduced raw data traces, spike detection is much more robust.

I implemented a simple threshold spike-detection algorithm in MATLAB. First, each trace was searched for points at which a threshold voltage was exceeded. Points on a putative spike(s) were defined as the set of all points such that $|V(k)| \geq V_{thresh}$, where k is an integer time index. I chose an empirically tuned value of V_{thresh} = 5*RMS noise level. The algorithm with this threshold was found empirically to perform very well—no false positives, no missed spikes. That is no surprise, given the high SNR of the recordings.

Multiple spikes occurring on a single electrode were separated based on a refractory period criterion. When a difference in consecutive time indices exceeded a refractory time, a new spike was declared. I chose a refractory time of $T_{refractory}$ = 1 msec based on neurobiological grounds. This value was also long enough to be robust to noise levels present in the pre-processed traces. Having detected spikes, a raster plot is generated as a visual tool to aid in further analysis. Figure 7.3 shows an example raster plot for 10 trials, with the stimulus presented on electrode 1. The raster is the companion to Figure 7.2. The stimulus electrode is marked by a red bar denoting the timing of the stimulus. For recording electrodes, the stimulus timing is marked as a cyan bar. Each vertical tick mark represents a detected spike. Consecutive trials are stacked in rows along the vertical axis, while the horizontal axis marks the time at which a spike was detected. The raster serves as an excellent visual aid to investigate the reliability and timing of stimulus-response pairs. In this case, viewing the raster, it is easy to see that the AP responses on electrode 2 ("Chan 2", in figure 7.3) were diffusely spread in time, while the jitter of about $\pm$ 1 msec in the responses on electrodes 3 and 14 were much smaller. Electrode 3 may have a weak secondary response at time $t \approx 55$ msec. Electrode 14 shows a weak, temporally diffuse response about the same time.

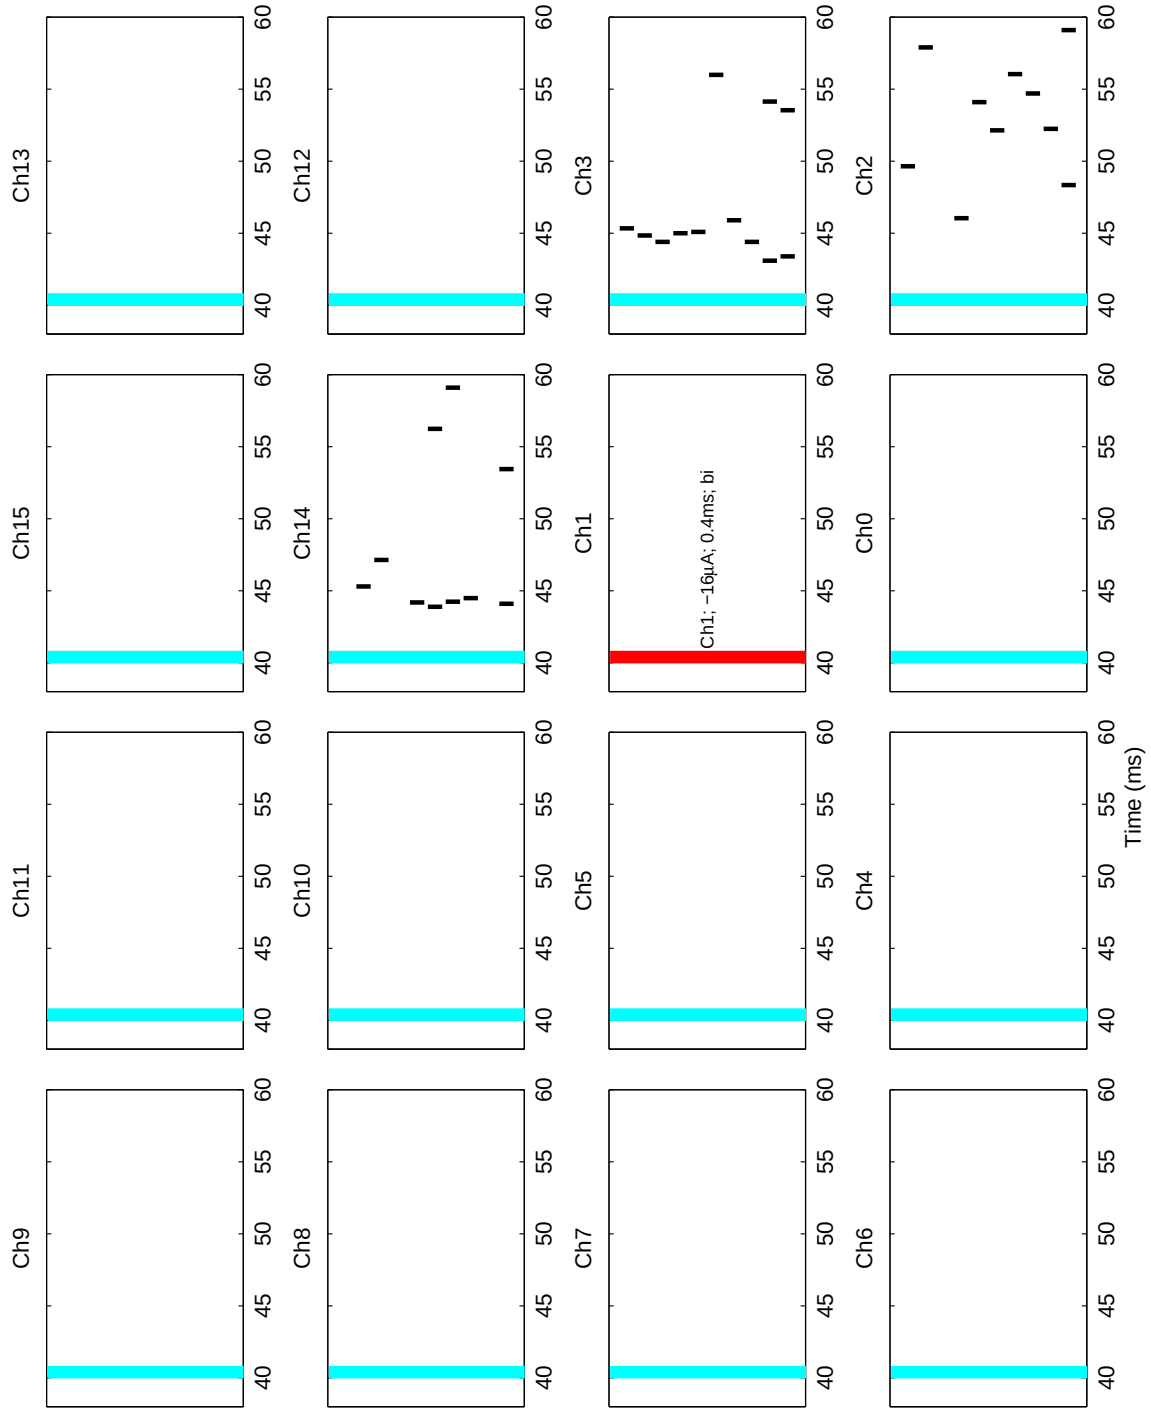

Figure 7.3: Raster plot corresponding to pre-processed traces of Figure 7.2. The resulting spike times from all 10 trials are displayed. Consecutive trials are stacked vertically with trial 1 at the bottom, and trial 10 at the top of each electrode's display. Again, the 4x4 layout matches the physical orientation of the neurocages. A stimulus was presented on electrode 1, resulting in spikes on electrodes 2, 3, and 14. The horizontal axis is time in msec. The time range displayed is from 38–60 msec, the same time scale as in Figure 7.2.

7.8 Synaptic Response: Connection Detection

The goal of this section is to describe how the spike times—with the aid of the raster plot—can be used to generate a network connectivity map. Two neurons, A and B, are defined as "connected" if an AP in neuron A evokes an AP response in neuron B with some degree of regularity (to be defined in detail shortly). In other words, A and B are the pre- and post-synaptic cells, respectively. Connectivity maps can be generated by considering the spike timing information acquired from multiple trials. The raster plot makes it easy to see connectivity.

For example, in Figure 7.3 there is clearly a "significant response" (to be defined rigorously below) on electrodes 14 and 3 at times of about $t \approx 44$ and 45 msec, 4 and 5 msec delays after the stimulus, respectively. The jitter noted over ten trials is about 1–2 msec for both electrodes. There is also diffuse response on electrode 2. The average response time appears to occur at about $t \approx 55$ msec, with a large jitter of about 5 msec. When a multiplicity of distinct responses is noted they are referred to as "first-", "second-", and "third- generation" responses. For the raster presented in Figure 7.3, electrode 3 has a robust (10/10 responses) first-generation response at about $t = 45$ msec (5 msec delay), and apparently a less robust (3/10 responses) second-generation response at $t = 55$ msec (10 msec delay).

The first-generation responses on electrodes 3 and 14 are clearly "significant" responses. Is an evoked AP response in 3 out of 10 trials—such as the secondary response on electrode 3—a "significant" response? And what about the diffuse response on electrode 2, with evoked action potentials spread out over a 15 msec window. The answer depends on comparison to levels of spontaneous activity.

For the small, isolated cultures on the order of 10 cells that I investigated, the levels of spontaneous activity were very low. By very low I mean spontaneous spikes essentially never occurred (were never detected) within a 200 msec control window. They were so rare that it is hard to accurately estimate the spontaneous activity firing rate. In this limit, I considered the spontaneous firing rate to be approximately zero spikes per second.

Therefore, I deemed it sensible to consider a response "significant" whenever a total of 3 or more spikes (out of 10 trials) occurred within a time jitter window of $T_{jitter} = 2$ msec. (Appendix I presents a mathematically more rigorous justification.) I tuned the width for

the jitter window based on the empirical observation that counting spikes over this time width detected several different classes of response types: strong and tightly timed (such as the first-generation responses on electrodes 3 and 14); weak and relatively tightly timed (3 spikes/10 trials responses on electrode 3, but not on 14); and strong but relatively diffuse in time (such as the response on electrode 2). Notice that responses which were weak and diffusely timed are excluded. In other words, the width of the jitter window was tuned so that the result of the quantitative detection scheme matched the result of visual inspection.

7.8.1 Connection Detection "Clustering" Algorithm

To formalize the process I use for detection significant responses, I present the algorithm in more detail here.

I start by counting the number of APs (spikes) on neuron B in response to neuron A, occurring within a 2-msec-wide window,[1] centered at time t, and divide by the number of trials:

$$SC_{A,B}(t) = \frac{1}{N_{trials}} \sum_{k=1}^{N_{trials}} \sum_{i=1}^{N_{spikes}} u(t^k_{A,B}(i) - t) \tag{7.1}$$

where $u(t)$ is defined by: $u(t) = 1$ for $t \in$ [-1, 1] msec; $u(t) = 0$ otherwise. Also, in 7.1, $t^k_{A,B}(i)$ is the time at which the ith spike during the kth trial occurs in neuron B when neuron A is stimulated. If no spike occurred during trial k, then the contribution of the kth element in the sum is set to zero.

Next I make a list of all times t^* such that $SC_{A,B}(t^*) \geq 0.3$. I set a criterion for a minimum difference of consecutive times in the list, T_{sep}, to distinguish between first- and higher-generation responses. When (if) two consecutive times in the list differ by more than T_{sep}, i.e., $t^*_{j+1} - t^*_j > T_{sep}$, the list of spike times, $t^k_{A,B}(i)$ is partitioned into two separate groups: One group contains spike times $t^k_{A,B}(i) < t^*_{j+1}$ and the other contains spike times $t^k_{A,B}(i) \geq t^*_{j+1}$. Each group of points represents a distinct response—i.e., the first group of points represents the first-generation response, the second group represents the second-generation response, etc.

I chose a heuristic value of $T_{sep} = 1$ msec based off of the typical refractory period

[1] Note that choosing a 2-msec-wide window limits the detection of significant responses to those which contain some spikes that are fairly precisely timed—i.e., three spikes with a response time jitter of ≤ 2 msec.

of hippocampal cell, and also based of visual inspection of raster plots. In essence, this procedure functions as a primitive clustering algorithm to distinguish between first- and higher-generation responses. (Note, however, that this algorithm is preferable to a true clustering algorithm because it builds in *a priori* knowledge of the underlying neurobiology.)

Finally, the time at which the mth generation responses occurred, $T_{A,B}^{(m)}$ was computed by averaging the spike times in the mth group of points.

7.9 Limiting the Complexity of the Connectivity Map

The raster in Figure 7.3 hinted at the complexity in the network manifested as second- and third-generation responses. These later responses are presumably due to polysynaptic connections. It may well be possible to more fully de-construct polysynaptic pathways by finding a chain of monosynaptic connections with the sum of delay times that matches the polysynaptic delay. That tack has not been followed here. Instead, I seek only to *define a connectivity map based on first-generation responses only.* This approach simplifies the analysis while still demonstrating the capabilities of the neurochip. Additionally, it will be shown shortly that connectivity maps including only first-generation responses illustrate very interesting dynamics during development in cultured neural networks.

It should be noted here that autapses can not be detected in the current system because a electrode being stimulated carries an artifact that is 1000 times larger than the expected AP signal for tens of milliseconds. Also it should be noted that only excitatory responses can be observed with the neurochip. If an inhibitory neuron is stimulated, no APs will be fired in other neurons, so no post-synaptic responses will be noted. Any cell which is seen to be input-only, i.e., one which is driven by other cells but drives no others, could be inhibitory. Post-hoc immunostaining could potentially identify GABAergic inhibitory neurons. This topic is addressed more fully below.

7.10 Delay Times: Mono- or Poly-Synaptic?

One desirable trait of a connectivity map is that it should encode some information about delay times. Here I define the synaptic delay time of the first-generation response of neuron B to neuron A to be the difference in time between the end of the stimulus and the time

at which a significant response occurred:

$$\delta t_{A,B} = T_{A,B} - t_{endstim}$$

, where $T_{A,B}$ is defined in Section 7.8.1. This definition is reasonable because, as Figure 5.4 shows, the electrically stimulated neuron does does not fire at the beginning of the stimulus; in fact it fires somewhere in between, probably closer to the end of the stimulus.

Fitzsimonds and Poo [24] investigated onset latencies (delay times) of excitatory post-synaptic currents (EPSC) in isolated triplets of hippocampal neurons grown with glia in low-density cultures. They measured the onset latencies of monosynaptic EPSCs to be in the range of 1.5–2.6 msec. As a first guess, I would expect the monosynaptic delays observed in neurochip cultures to fall into a similar range. Therefore I have initially defined fast connections—putatively monosynaptic—to have ≤ 5 msec delay. I have initially defined a middle temporal range for polysynaptic connections to fall between 5–20 msec. This range was chosen because it includes putative bi-, tri-, and quadra-synaptic connections. Long delay times could result from high-order synaptic connections. Any delay longer than 20 msec falls into this category.

7.11 Connectivity Maps

The connectivity maps I generate are color-coded discretely, according to the three categories described above. Short delay (< 5 msec) are colored in red; middle-range delays (5–20 msec) are colored green; long delays (> 20 msec) are colored black. Off-line data processing and graphics rendering were implemented with custom-built MATLAB software.

Figure 7.4 shows a connectivity map generated from probing a 13-day-old culture consisting of 12 neurons. The culture connectivity map here is derived, in part, from the data shown above (Figures 7.2 and 7.3), with electrode 1 being stimulated. An arrow is drawn from the pre-synaptic neuron, terminating in an arrowhead at the post-synaptic neuron to denote a functional connection. The raster in Figure 7.3 can be utilized to understand the map. For example, a stimulus on electrode 1 resulted in an AP with 4 msec delay from the neuron adjacent to electrode 14. Thus, a red arrow is drawn from electrode 1 to electrode 14. The neuron adjacent to electrode 3 responded with a longer delay of slightly more than

5 msec. Therefore, a green arrow is drawn from electrode 1 to electrode 3. Finally, the response from a the neuron adjacent to electrode 2 occurs with an average delay of about 15 msec, also considered to be a mid-range delay. Therefore, a green arrow is drawn from electrode 1 to 2.

In the connectivity map, circles in black outline represent electrodes which were probed. Circles that are grayed out represent electrodes which were not probed. The connectivity map shown here is the same as the 13 DIV map shown in Figure 7.5. Note that this visualization of the network connectivity only shows the starting and ending points of a pathway—it alone does not reveal the presence of any interneurons in a pathway. For example, a pathway was detected from neuron 14 to neuron 2. There are also connections between neurons 14 and 3, and between neurons 3 and 2. It could be the case that a direct connection (monosynpatic pathway) from 14 to 2 does not exist, and that the pathway is actually di-synaptic: 14-3-2.[2]

[2]However, an alternative view of network connectivity which attempts to discern whether a pathway is mono- or di-synaptic, or both, is presented in Appendix K.

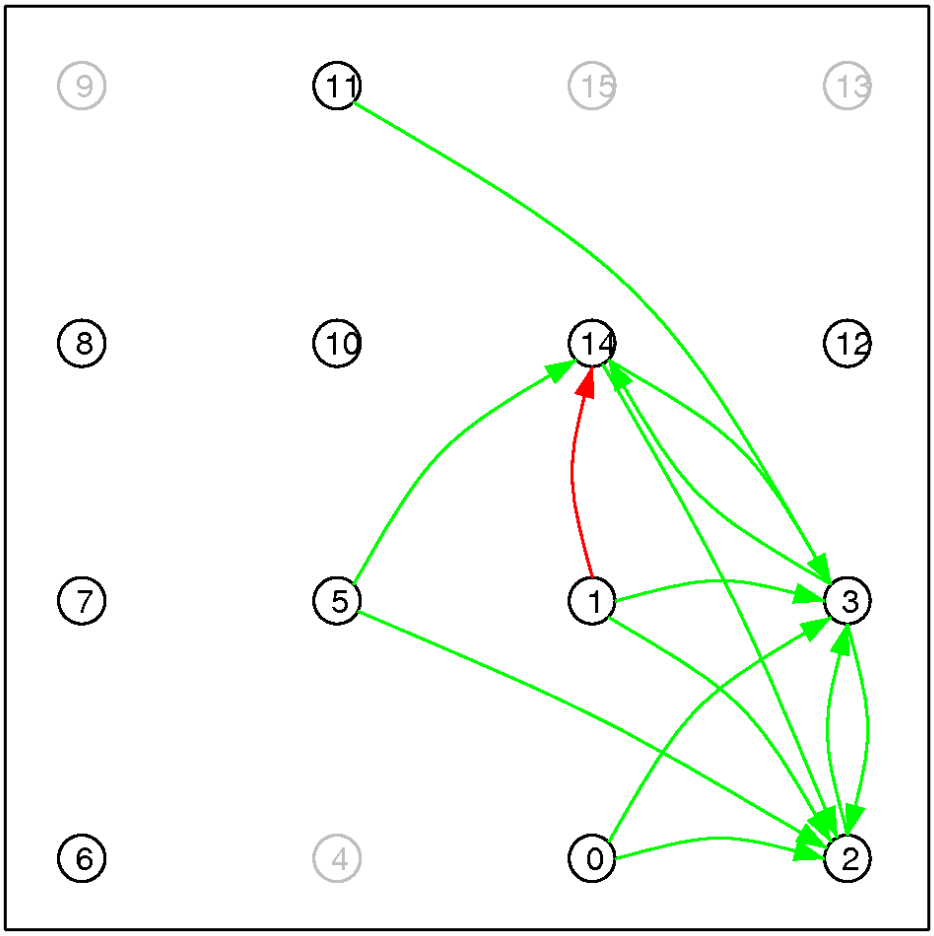

Figure 7.4: Color-coded connectivity map for a 13-day-old culture consisting of 12 neurons. The discrete colors code for response delays. Red < 5 msec; Green $= 5$–20 msec; Black > 20 msec. Not all neurons stimulated exhibit suprathreshold connectivity (electrode 12, for example). Only electrodes corresponding to circles outlined in black were stimulated. The connectivity map displayed here corresponds to map at 13 DIV in Figure 7.5.

Now, having the ability to noninvasively stimulate a culture and build a map from the data acquired at one time point means that I can come back in a day or two—or whenever—and do it all over again! That's the heart and soul of the neurochip system, the unprecedented capability that it offers.

As an example, the culture which has been discussed in detail in this chapter was investigated at 6 different time points: at ages of 9, 11, 13, 16, 20, and 22 days old. For each day the culture was probed, a connectivity map was generated. The results from each of the six days are displayed in Figure 7.5, referred to as a "culture evolution" figure. (The connectivity map displayed in Figure 7.4 is the 13 DIV map in Figure 7.5 below.) The color-code remains the same for all days. These data sets, with connectivity probed at the single-cell level over several weeks, are unprecedented.

More culture evolution figures for several other are provided in Appendix J for further viewing, contemplation, and amusement.

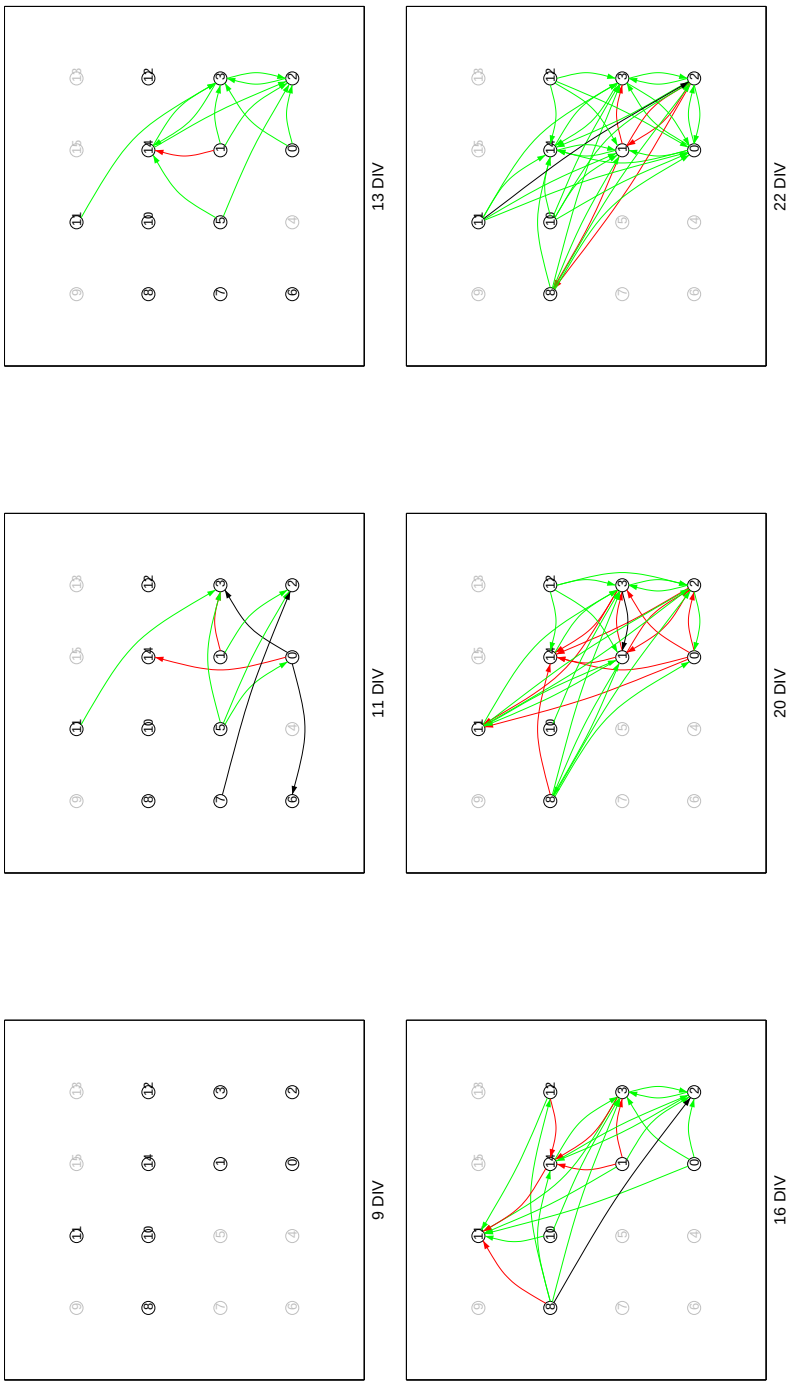

Figure 7.5: Evolution of a cultured neural network's connectivity over 3 weeks. The color code is the same as before. The culture undergoes a rapid maturation period. At 9 days old, no suprathreshold connectivity was observed. Within 4 days, many such connections existed. The culture continued to mature up until 22 days old, at which point the culture was terminated. The neurochip system is the first to provide such detailed data over time.

Chapter 8

Network Analysis: Results from Initial Studies

8.1 Network Analysis: Some Basic Statistics

With this type of data acquired with the neurochip at hand, one can begin to ponder many questions. How richly is the network connected? In other words, how many suprathreshold connections does each cell make? What are the typical synaptic transmission delay times? Do transmission times generally speed up? Do only the fast ones speed up while the slow ones slow down or disconnect? Do the connections become more reliable over time? That is, does the probability of suprathreshold transmission increase? How and why are new connections formed? Do the pre-post synaptic partners remain stable over time? How do neurons choose their targets? Is *in vivo* connectivity preserved *in vitro*? Do CA3 neurons drive CA1 neurons; CA3 drive other CA3; but CA1 neurons do not drive CA3?

To begin to answer some of these questions, I have compiled some basic statistics/metrics about the neurochip cultures. From the set of 17 cultures that exhibited suprathreshold connectivity, I chose a subset of 10 from which the following statistics are compiled. The criterion for including a culture for analysis was that its connectivity be measured at 3 or more time points during the first 21 days *in vitro* (DIV). This criterion prevents reporting spurious fluctuations due to a culture that was probed at only one or two points in time.

Throughout the rest of this chapter, I quantified the various network statistics/metrics at various developmental stages: between 7–10, 11–13, 14–16, 17–19, and 20–21 days. The choice for the five different time segments is heuristic, and based off of the results for average number of connections over time (see Section 8.2). Table 8.1 shows the total number of

neurons with which I experimented, as well as the total number of connections detected during each of these developmental periods. Note that if a network was probed on two days within the same developmental period, the number of neurons and connections is double-counted. For example, if a culture with 10 neurons was probed on days 11 and again on 13, then it contributed 20 neurons to the sum for N_{cells} during the second period.

Table 8.1: The total number of neurons tested($N_{neurons}$) and detected connections ($N_{connections}$) during each of five (heuristically chosen) developmental periods.

Time (DIV)	$N_{neurons}$	$N_{connections}$
7–10	32	11
11–13	66	66
14–16	90	158
17–19	78	283
20–21	34	81

8.2 Number of Connections versus Time

To better quantify how networks develop over time, I counted the number of suprathreshold connections per cell at various time points. Only first-generation responses were counted—i.e., second-generation (and higher-order) responses were not considered here. For each of 10 cultures, the average number of suprathreshold connections per stimulated cell versus time is reported in Figure 8.1.

This data can also be viewed a different way: Figure 8.2 shows the data in terms of the fraction of the maximum possible number of connections. Since autapses can not be detected, the maximum number of possible connections equals $n_{neurons}(n_{neurons}-1)$, where $n_{neurons}$ is the number of stimulated neurons in a culture. The "connectivity fraction" (CF) is defined as $CF = n_{connections}/n_{neurons}(n_{neurons}-1)$, where $n_{connections}$ represents the total number of connections in the culture. For example, a value of 0.2 means that 20% of the total possible connections are actually connected; a value of 1 would indicate a fully connected culture.

The data show a few obvious trends. Cultures generally first start to exhibit suprathreshold connectivity around 10–12 days old. A few cultures blossom late, at about 15 days old. This time frame is in accord with the development of spontaneous activity noted in small

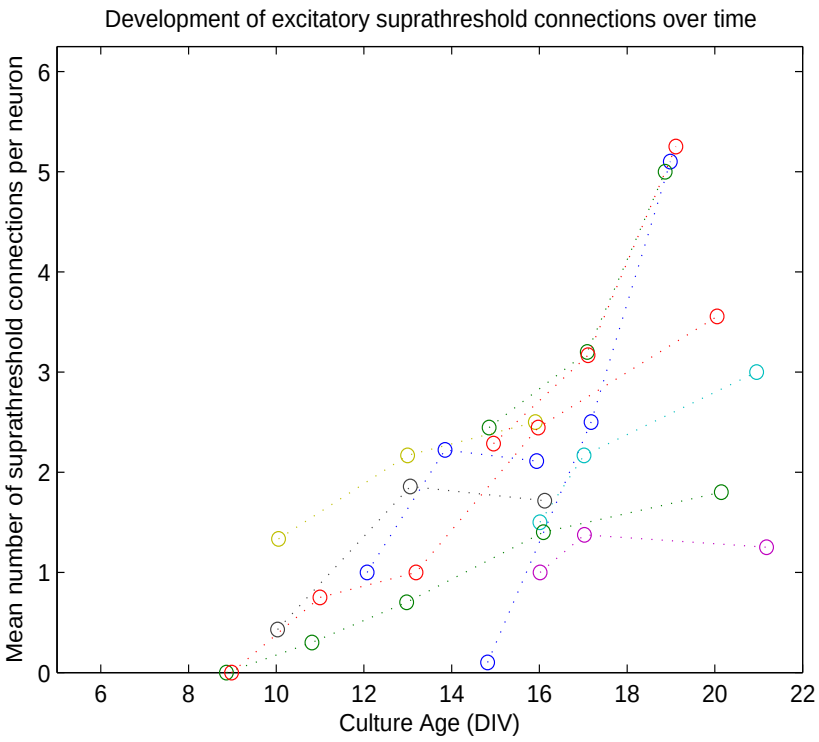

Figure 8.1: Average number of suprathreshold connections versus time for 10 neurochip cultures. Error bars are not shown to aid visual clarity. Suprathreshold synaptic connectivity first appears around days 10–12, with some late bloomers exhibiting initial suprathreshold connections at 15 days old. The number of connections is seen to increase over time and generally peaks at about 3 connections per cell, with some exceptions that become very richly networked with an average of about 5 connections per cell.

neurochip cultures, as well as in larger, denser MEA cultures [42, 69]. The number of connections increases over time, in a roughly linear fashion. At two and a half weeks old, each cell is connected, on average, to about 2 other cells. Inter-culture variance increases with age. Three cultures become densely connected with an average of 5 connections per cell at 18 days old. Most cultures develop a maximum of about 2–3 synapses per cell by three weeks.

In terms of the connectivity fraction, by the time the culture is relatively mature—older than about 14 DIV—the connectivity fraction falls within a range of about 0.2–0.5. In other words, about 20–50% of the total possible connections actually form. Other investigators studying low-density hippocampal cultures on microislands have measured connectivity fractions of about 20%. (G. Bi, personal communication). So, the neurochip cultures are similar or, in some cases, more richly connected.

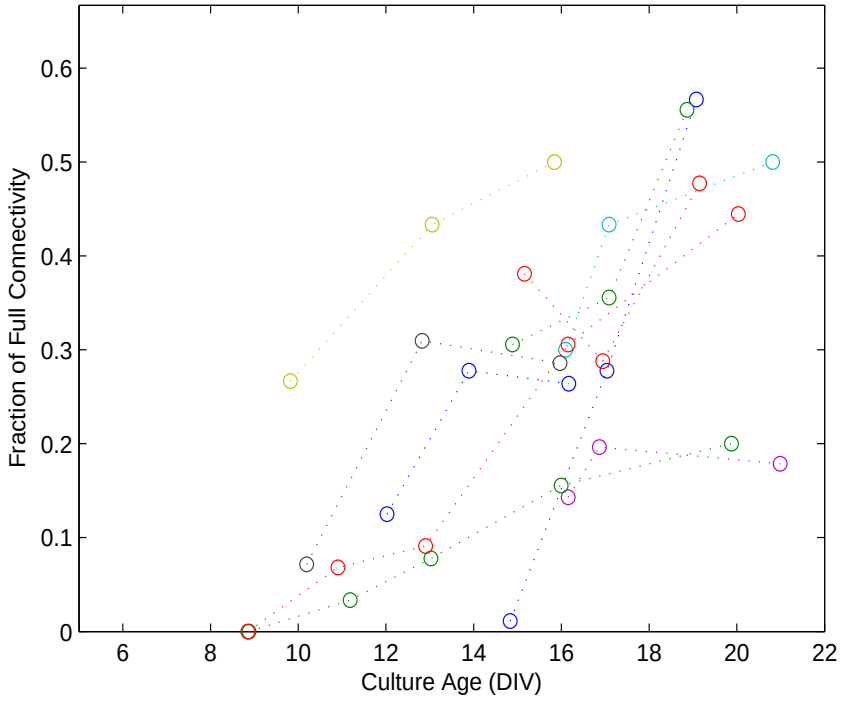

Figure 8.2: Connectivity fraction versus time. The connectivity fraction expresses the number of connections divided by the maximum possible number of connections. Error bars are not shown to aid visual clarity. The connectivity fraction for mature cultures ($t >$ 14 DIV) falls in the range of about 0.2–0.5, indicating that about 20–50% of the possible connections are actually formed.

8.3 Size versus Complexity

Another natural question to ask is "Are cultures with more neurons more fully connected?" In other words, does the complexity of the network, the richness of the networking, depend on the number of neurons? Figure 8.3 shows that the answer is "No". This graph shows the (average) number of stimulated cells versus the maximum achieved connectivity fraction for 10 neurochip cultures. For 6 out of 10 cultures, the number of stimulated neurons remained the same for all time. For the other 4 cultures, the stimulation list changed over time (to account for dying cells), so the average number of stimulated cells over time is used. Figure 8.3 shows that cultures with the smallest and largest number of neurons became the most richly networked, while cultures in the mid-range developed less connectivity. Hence, the richness of connectivity is not correlated to the number of neurons in the culture.

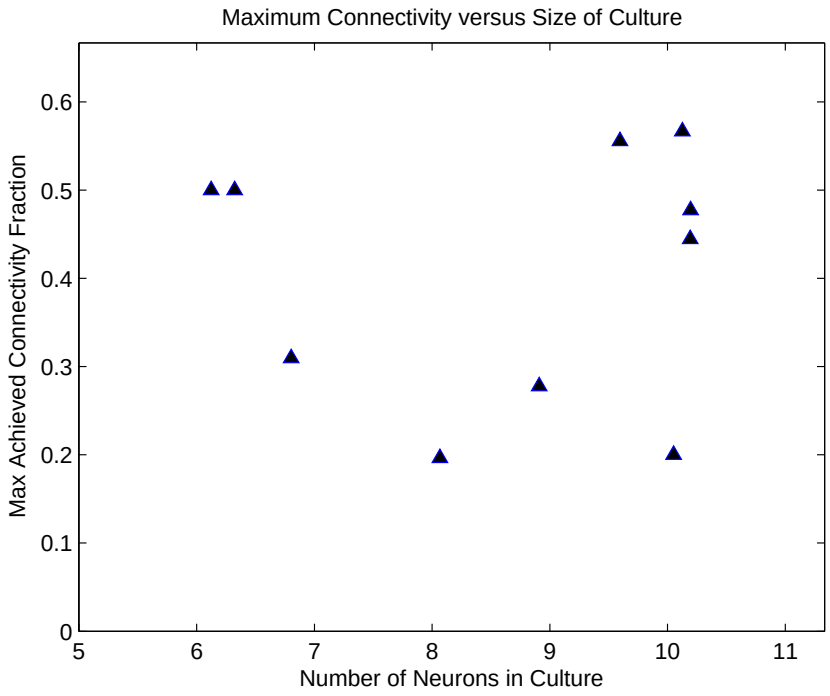

Figure 8.3: Maximum achieved connectivity fraction versus (average) number of stimulated neurons for 10 neurochip cultures. The richness of connectivity is not correlated to the number of neurons in the culture.

8.4 Suprathreshold Connection Delay Times

Another basic network metric to quantify is the synaptic transmission delay time. Recall that this quantity (abbreviated here as δt) is defined as the difference in the response time and the time at which the stimulus end (see Section 7.10). I investigated the delay times at various developmental stages listed in Table 8.1. If a culture was tested at two days within the same stage, say at day 11 and day 13, both were included in this analysis. Another reminder: only first-generation responses are counted here.

The terminology of "fast", "medium" and "slow" connections refer to synaptic transmission delays, respectively, of < 5 msec, 5–20 msec, and > 20 msec. Fast, medium, and slow connections correspond to "short", "medium" and "long" delay times, respectively. The color code is the same as for the connectivity figures shown above: red = fast; green = medium; black = slow.

Figure 8.4 shows the normalized distribution of delay times at each of these stages for the 10 neurochip cultures analyzed here. The companion figure, 8.5, shows the fraction of

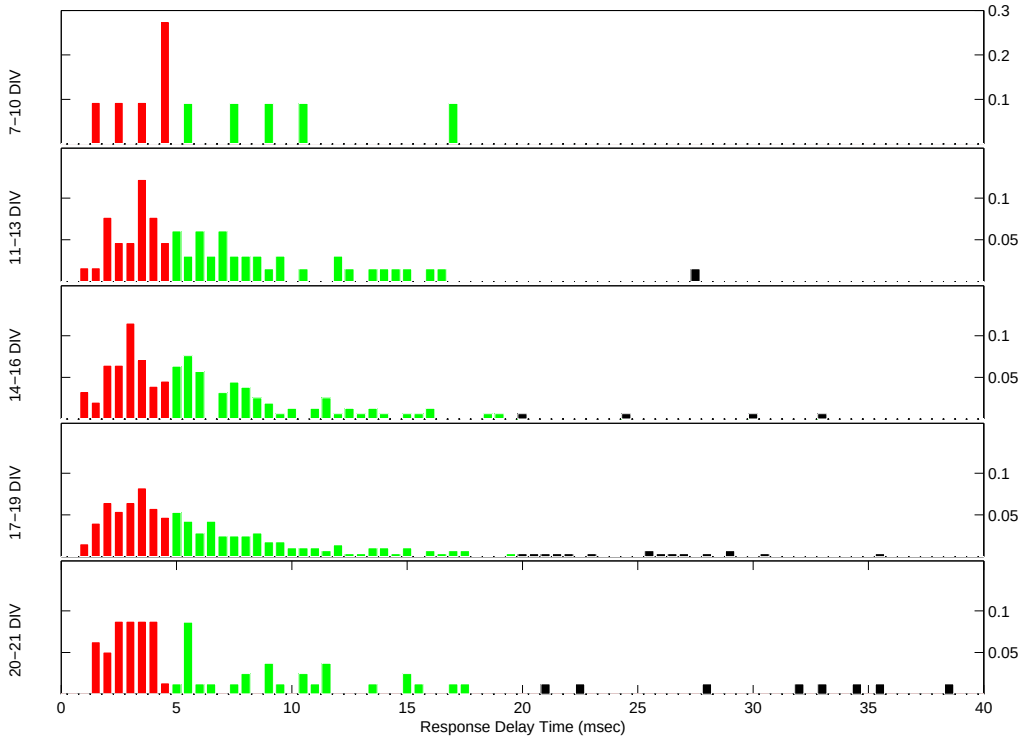

Figure 8.4: Normalized distribution of connection delay times at various developmental stages. The red–green–black color code corresponds to short ($<$ 5 msec), medium (5–20 msec), and long delay times ($>$ 20 msec). This is the same color code used in the connectivity figures in Chapter 7. Note the vertical scale of the top graph is different than the rest. Most delay times fall in the range of 3–10 msec. The distribution of delay times for $\delta t < 3$ msec shown here is an imperfect measure—some are lost in the stimulus artifact blanking by SALPA, which is typically about 1–2 msec long. The increasing frequency of long delays in older cultures is consistent with the hypothesis that they are generated by connections to the distant mass-culture. Overall, the distribution does not change much from 7–21 DIV.

each type of connection—fast, medium, and slow—over time.

Fitzsimonds and Poo measured delays for monosynaptic excitatory post-synaptic currents (EPSCs) in low-density hippocampal cultures in the range of 1.5–2.6 msec [24]. Few connections ($\approx$ 15%) with such short delay times are noted in the neurochip cultures, partially due to the limitation of the minimum time at which an AP can be detected after a stimulus. Taking into account the distribution of blanking times shown in Figure 7.5, I roughly estimate that I would fail to detect about 1/2 of all APs which occur with a delay of 1.5 msec; about 1/4 of APs with a delay time of 2 msec; and only about 1/20 APs with

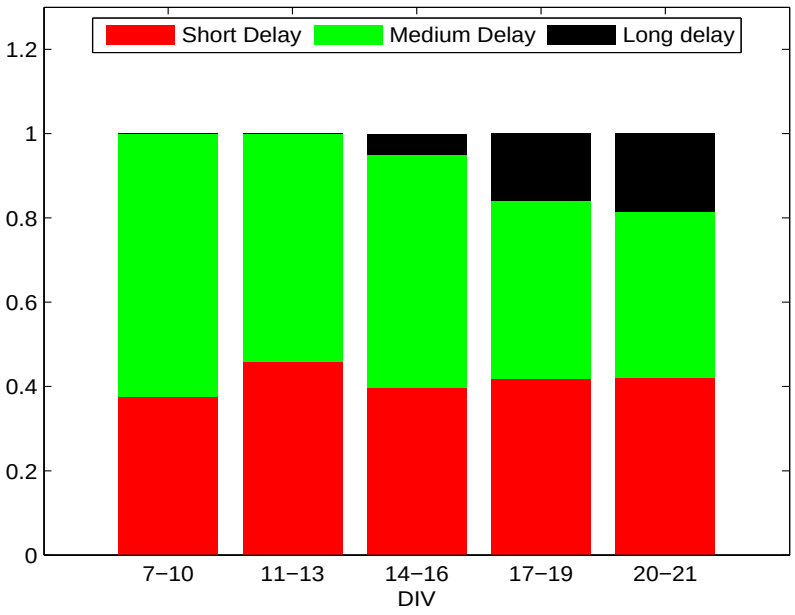

Figure 8.5: Fraction of each type of delay—fast, medium, and slow—observed at various developmental stages. The fraction of fast connections remains relatively constant. The majority of delays are in the medium range, 5–20 msec. Long delay connections are not present in younger cultures; in older cultures their fraction increases over time.

a delay of 2.5 msec. (Only about 2% of APs occuring with a delay time of 3 msec (or greater) would be blanked out by the SALPA algorithm. Thus, on average, I probably miss about 1/4 of all responses with delay times between 1.5–2.5 msec. This adjustment would make the distribution in Figure 8.4 flatter for very short delay times, but still would not move my results into agreement with theirs. This discrepancy is not clearly understood at this time; one possibility is that differences in pharmacology—they use high K^+ cell culture medium—produced different results.

The majority ($\approx 65\%$) of delay times fall in the range of of 3–10 msec. The question remains open as to whether short and medium delays correspond to monosynaptic connections and higher-order connections, respectively. [1]

[1] J. Pine and I later performed an initial analysis on the delay times and the connectivity maps to determine whether the the medium delay connections could have resulted from polysynaptic responses. The timing of each connection was examined as to whether it could reasonably result from a chain of two observed first-generation responses (see Appendix K). Putative disynaptic connections in young cultures are very rare ($< 10\%$). They are not particularly abundant in older cultures either. On average, about 20% of all detected connections were thought to be disynaptic although some cultures were noted to have as many as 40%. The implication was that the the majority of delays of ≤ 10 msec actually correspond to a monosynaptic connection. The difference between our measurements and Poo's may result from different pharmacology and network geometry.

The long delays ($\delta t > 20$ msec) are not present until days 11–13. The most plausible explanation for the long delays is that they result from imperfect isolation of the culture in the array and the mass culture. The stimulated neuron in the cage probably stimulates the mass culture via long axons which have grown in toward the array over several weeks. Through some high-order polysynaptic connections, the signal volleys back into the array, stimulating another caged neuron.

The proportion of fast connections does not change much over time. For young cultures (7–10 DIV), about 30% of the connections are fast. For older cultures about 40% are fast. The proportion of medium delay connections decreases over time from about 65% to 40%. This change can be attributed to a increasing number of long delay connections over time.

8.5 Response Time Precision

Spike-timing precision is a very important property. The temporal fidelity of a network is critical for its ability to properly process information [60]. Many *in vivo* systems, such as the rodent somatosensory (whisker barrel) cortex efficiently represent stimulus information using a spike-timing code [51]. Millisecond-precision timescales are also known to be important for the function of the visual system in fly, fish, and mammalian visual system [9].

How precisely timed are the network responses in neurochip cultures? And how does the temporal precision change as cultures mature?

To answer this question, I examined the response time jitter, defined as the standard deviation of the spike times that were clustered for the first-generation responses (according to the procedure described in Section 7.8.1). The distribution of jitter in response time for the (now familiar) five developmental periods is shown in Figure 8.6. Note that the algorithm used to define significant responses counts spikes in a 2-msec-wide window, essentially placing an upper bound on the maximum possible jitter of about 1 msec. The jitter times are broadly distributed over a 1 msec interval. Interestingly, similar to *in vivo* systems, neurochip networks become more temporally precise, on average, as they mature over the first 19 DIV. The corresponding mean jitter values are: 0.81, 0.68, 0.61, and 0.50 msec—an $\approx 38\%$ improvement in precision. The mean jitter time increases from the fourth to the fifth period, possibly because the increasing number of connections to the mass culture which

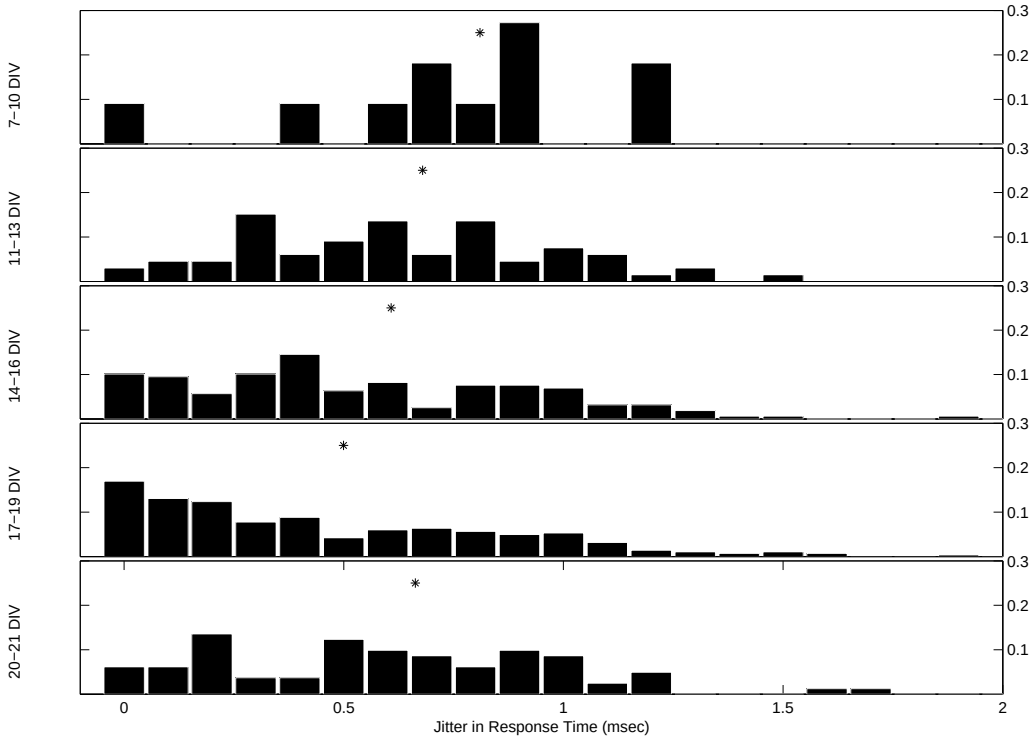

Figure 8.6: Normalized distribution of jitter in response times, broken down into five developmental periods. Data were culled from 10 neurochip cultures. The jitter in responses is broadly distributed over the interval of 0–1.25 msec. The average response jitter is marked by an asterisk for each period. The average jitter time decreases over the first four periods by about 38% (from 0.8 msec to 0.5 msec), but increases between the penultimate and final period.

tend to be less precisely timed.

There is another important piece of information hidden in Figure 8.6. The fact that the distribution is relatively smooth speaks to the issue of whether stimulation of specific—whether a current pulse stimulates only the neuron adjacent to the stimulus electrode, or if it might also directly evoke APs by stimulating a passing axon from another neuron. The latter is termed "direct response" by Wagenaar et al. [68]. He examined much larger cortical cultures ($> 50{,}000$ neurons) grown on MEAs and found that direct responses were identifiable based on their very high temporal precision (jitter in response times of ≤ 0.25 msec). If I were stimulating (a significant fraction of) passing axons, I might expect to see a bimodal distribution or discontinuity in the distribution of jitter times. This distribution

would have a (relatively) small mode at peaked at about 0.25 msec, and another larger mode at longer jitters. Since I measured a relatively uniform distribution (not bimodal), I conclude that either 1) that stimulation is highly specific; I am not stimulating passing axons from other neurons, or 2) a small neurochip culture with about 10 neurons is not comparable to much larger, denser cultures. A definitive study using synaptic blockers (AP5, CNQX) to prevent excitatory synaptic transmission should be done to distinguish between these possibilities.

8.6 Suprathreshold Connection Reliabililty

The fidelity, or probability, with which a an AP in neuron A drives another neuron B is another relevant network metric. How strong, or "reliable" are the suprathreshold connections in neurochip cultures? Does the synaptic transmission line transmit 100% of the time? Only some of the time? Do fast connections get stronger over time? Do slow connections become weaker/less reliable over time?

I define the "connection reliability" to be the number of times out of N_{trials} (= 10) trials that stimulating cell A drives an AP in cell B. Specifically, the connection reliability factor, $CR_{A,B}$, is computed as the sum of the number of spikes times in the cluster corresponding to the first-generation response divided by the number of trials (see Section 7.8.1).

In this context, a "strong" connection is one which has high reliability (> 0.7, a response to more than 7 out of 10 trials). A "weak" connection is one which has low reliability ($<$ 0.5, a response to fewer than 5 out of 10 trials). Note that, by definition, the minimum connection reliability to declare a significant response as having occurred is $CR_{A,B} = 0.3$. Also, note that it is possible to have $CR_{A,B} > 1$ if, on average, more than 1 spike per trial occurred in a single generation (cluster of points). This case is rarely realized; values for $CR_{A,B}$ are almost always ≤ 1.

Figure 8.7 shows the normalized distribution of reliabilities for the five developmental periods. The asterisk over each histogram marks the mean value of the distribution. The left-most column shows the distribution for all connection types—fast, medium and slow. The right side of the figure breaks down the connection reliability according to the delay times. The second column shows the reliability for the fast connections; the third column for the medium delay connections; and the right-most column for the slow connections.

Each row represents a different developmental stage.

The data are not particularly dramatic. Nonetheless, a few remarks are in order: It is interesting to note that the average reliability for fast and medium delay connection types tops-out at about 0.6–0.7. It is as if the culture is satiated; or perhaps it is not able to produce any more functional synapses. For all connection types, the reliability changes only slightly over time.

The reliability value averaged over all connection types may not be the best indicator; one should also consider the distributions for each of the subtypes (fast, medium, and slow connections).

Over a 2 week span, the fast connections strengthen from an average value of about 0.6 in young cultures to an average value of about 0.7 in more mature cultures. The medium delay connections also become stronger over the first four developmental periods from about 0.35–0.7, but curiously regress during the last period (20–21 DIV) to about 0.55.

Compared to the medium delay connections, the fast connections develop stronger connections more quickly. For example, for 11–13 DIV the fast connections have shifted to an average reliability of about 0.7, while the medium connections require about 5–7 more days to achieve such strength. The long delays rarely become strong. For the most part, the reliability of connections with long delays is in the range of 0.3–0.4. This result is compatible with the idea that long delay connections result from volleys from the mass culture. In this case, the pathway from a stimulus on cell A evoking an action potential in cell B involves many intermediate cells—the pathway is high-order polysynaptic, and therefore, less reliable. Also, each cell involved in the high-order synaptic pathway will cause the total response time to be more randomly distributed. According to the definition of $CR_{A,B}$ above, the reliability of a diffuse response (large jitter in the timing of responses) will be computed as a lower value.

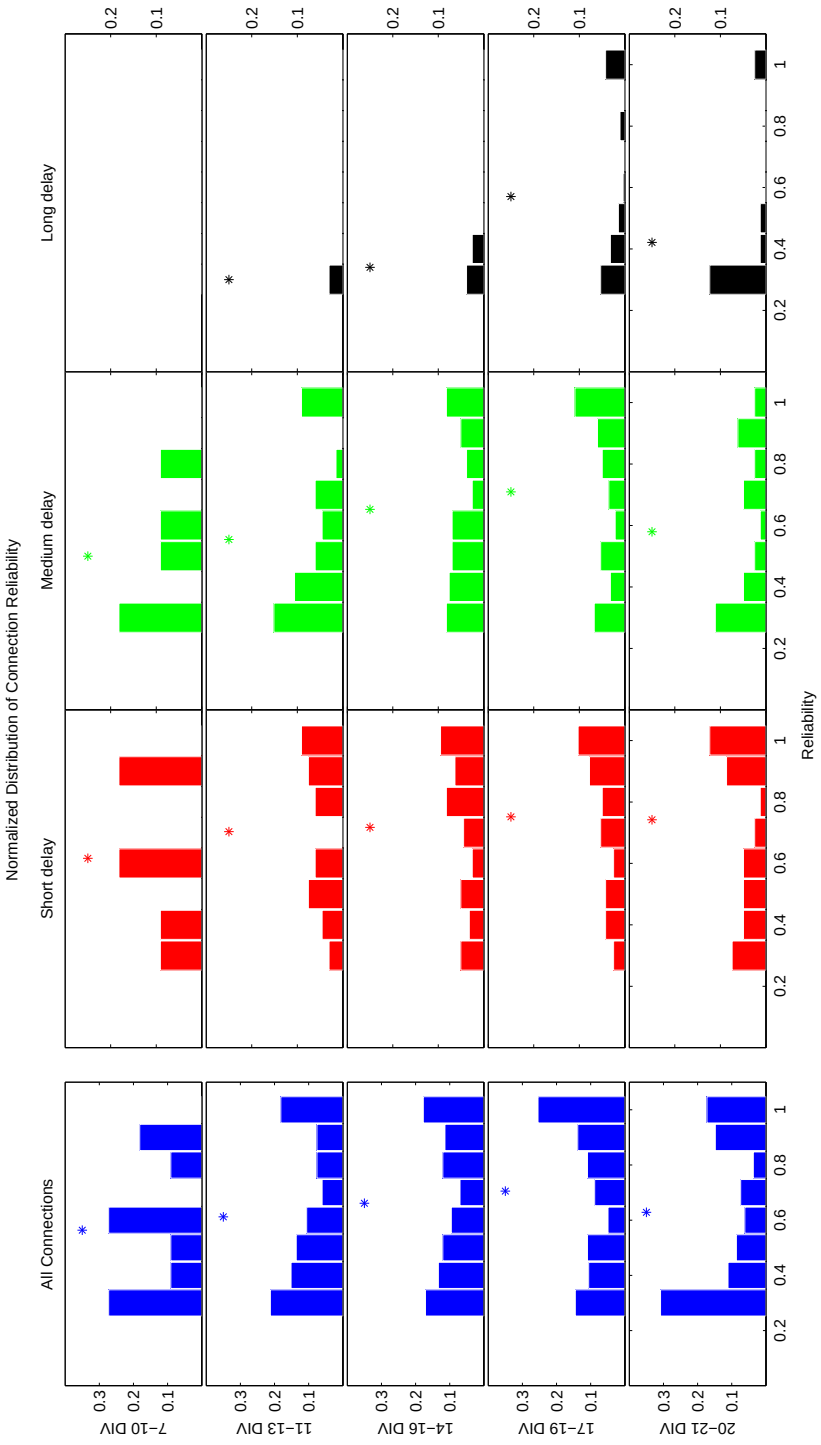

Figure 8.7: Normalized distribution of connection "reliability", broken down into five developmental periods. Data was culled from 10 neurochip cultures. The asterisk over each histogram marks the mean value of that distribution. The left most column shows the weights for all subtypes of connections—fast, medium, and slow. The right 3 columns break down the weights according to connection delay time. While there is no strong trend, the fast and medium delay connections become more reliable over time. The time-course of the strenghtening of the fast connections is accelerated relative to the medium delays. The reliability for connections with long delays remains weak.

8.7 Is *In vivo* Connectivity Preserved?

One wonders how a neuron chooses its post-synaptic targets. Are the connections made randomly without regard to cell type? Or do CA3 neurons modify their synapses over time so that their input drives CA1 targets whereas CA1 cells do not input back onto CA3? Is hippocampal connectivity preserved *in vitro*?

One way to approach this question is to quantify, for each cell, the number of inputs versus outputs. The number of inputs to a cell indicates how many other cells can drive an AP in it. The number of outputs from a cell indicates how many cells are driven to fire APs by it. In other words, the number of inputs is the number of pre-synaptic partners; and the number of outputs is the number of post-synaptic partners.

If hippocampal circuitry is preserved *in vitro*, one would expect that a CA1 cell has several cells which drive it, and no cells driven by it—several inputs, no outputs. CA3, on the other hand, would be expected to have 2 potential targets—CA1 and other CA3 cells—for input. CA3 cells could be driven by other CA3 cells. So one might expect to see CA3 driving APs in 2 cells, while being driven by 1. The ratio of outputs to inputs should be 2:1. Figure 8.8 displays, for each cell, the normalized number of inputs (pre-synaptic partners) versus the number of outputs (post-synaptic partners). The normalization is with respect to the number of stimulated neurons in a culture minus one—autaptic connections can not be detected, and therefore, are not considered here. So, the value of $Outputs = 1$ for a neuron means it drives all other neurons in the array. A neuron with value of $Inputs = 1$ indicates that it is driven by all other stimulated cells in the culture. Graphically, a cell which is input only ($Outputs = 0$, $Inputs \neq 0$) will lie along the horizontal axis. A cell which is output only ($Inputs = 0$, $Outputs \neq 0$) will lie along the vertical axis.

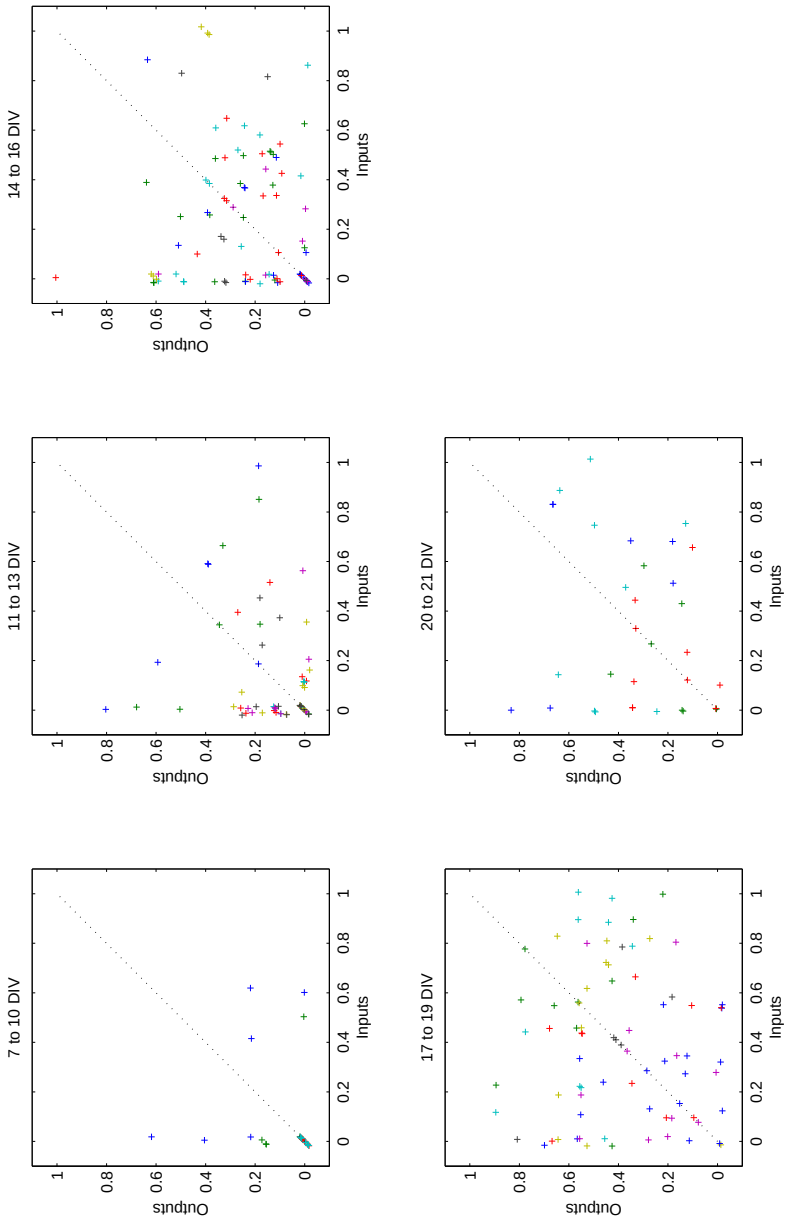

Figure 8.8: Number of input versus outputs for individual cells over different developmental periods. The number of inputs and outputs are expressed as a fraction of the total number of target cells minus one (autaptic connections are not considered, since they cannot be detected on a 20-msec timescale). The diagonal dotted line serves as a guide to the eye. Except for a noticeable population of cells that are output only (the column where $Inputs = 0$, but $Outputs \neq 0$), the ratio of inputs to outputs appears to be distributed randomly, suggesting that hippocampal connectivity is not preserved *in vitro*.

In neurochip cultures it is believed that the relative proportions of CA1 and CA3 cell types are roughly equal. Therefore, if indeed CA1 only receives inputs, there should be a cluster containing about 1/2 the total points lying along the horizontal axis. Some neurons, but nowhere near half, meet this criterion. The distribution of points is too uniform to support the conclusion that CA3/CA1 connectivity is preserved *in vitro*. Post-hoc immunostaining could be used to definitely answer the question.

The small cluster of output-only points is worth noting. As mentioned previously, these are represented by the points that lie along the vertical axis, but not at the origin. At all developmental stages, there are more output-only than input-only cells. For example, for cultures between 14–16 DIV, out of 90 total cells, 26 were output-only, whereas only 9 were input-only. The neurobiological significance of this finding is not clear, but one might conjecture that these output-only cells are of type CA3.

8.8 Upper Bound for the Fraction of Inhibitory Cells

Another interesting observation based on Figure 8.8 is the small number of neurons which were input only. These are represented by the points that lie along the horizontal axis, but not at the origin ($Outputs = 0$, $Inputs \neq 0$). This population of cells represents inhibitory cells and/or neurons with excitatory connections too weak to drive AP potential responses in other neurons. This places an upper bound on the proportion of inhibitory cells in the neurochip cultures. For the 3rd developmental period (14–16 DIV) 9 of 90 cells are input only; during the fourth developmental period (17–19 DIV), 5 out of 78 are. Therefore, the maximum percentage of inhibitory cells is about 10%. Post-hoc immunostaining could be also be used for a definitive measurement.

Chapter 9

Concluding Remarks and Future Direction

9.1 Summary of Major Results Achieved

The work presented here demonstrates a system which offers new and more powerful capabilities for studying cultured neural networks. I have shown that small *in vitro* neural networks can be grown in neurocages and that the electrodes can safely, noninvasively stimulate and record APs over the lifetime of the culture.

To review, the major advantages of the neurochip system are:

1. Specificity: long-term 1:1 neuron-to-electrode correspondence

2. Timescale: in electrical contact with all constituent neurons over the lifetime of the culture

3. Numbers: networks studied consist of about 6–10 neurons.

4. Unconstrained Process Growth: mechanical traps permit axons and dendrites to grow as they please, unconfined to any predefined geometry.

The stimulus-response experiments illustrate the potential power of the neurochip in probing cultured neural network connectivity at the single neuron level over 3 weeks or longer. Some future experiments with the neurochip are discussed below.

9.2 Future Direction: Scaling Up

While cultures of 6–10 neurons are interesting, cultures consisting of more neurons would surely be even more interesting. There are no obstacles to scaling the neurocage array up to 64 cages in an 8x8 array. In fact, significant progress has already been made on this front[63]. With a 64 cage array, one could expect to grow cultures consisting of 24–40 neurons, 10 times more neurons than are studied using conventional (patch clamp) techniques, and potentially with much richer connectivity.

Also in the works is porting the neurocages to a thin glass substrate. The reason for moving off silicon wafers is that a neurochip built upon a glass substrate is compatible with a computer-automated laser-tweezer system. [13] The laser-tweezers automates the process of loading single neurons into cages, a major improvement over the current manual loading method, particularly for larger numbers of cages.

9.3 Future Direction: Experiments

The potential uses and experiments with the neurochip are many. Several interesting experiments immediately come to mind, all within arms' reach. This list is by no means meant to be all-inclusive. I encourage others to adopt the neurochip system and invent experiments to aid their studies.

9.3.1 Immunostaining

Post-hoc immunostaining of neurochip cultures can be used to identify the type of cell present in each neurocage. Having collected and analyzed connectivity data, and having identified cell types, the issue of whether CA3/CA1 connectivity is preserved in cultured neural networks could be answered. The Ghosh lab at UC San Diego has been addressing this same question using patch-clamp electrophysiology. Using the neurochip instead could potentially provide richer electrophysiology data at a faster rate. (The Pine lab now greatly benefits from the Ghosh lab's expertise with immunostaining.)

9.3.2 Beyond Stimulating a Single Cell

So far, I have examined connectivity in terms of stimulating one cell and looking at the responses of all others. What about stimulating more than one cell at a time? Do the results sum linearly? Can an investigator find combinations of stimulated cells which produce an AP in another cell, whereas stimulating each of the cells alone does not evoke an AP? Can the stimuli from 2 (or more) cells be offset to cause an AP in another cell, yielding transmission delay information from individual subthreshold connections? In short, it seems possible to gain knowledge about networks by more fully probing them with more complex stimuli.

9.3.3 Plasticity: Patterned Electrical Stimulation

The next obvious experiment is to apply patterned or chronic electrical stimulation to the neurochip cultures.

Previous experiments in which a small number of neurons were manipulated have revealed that *in vitro* network function is readily modified in response to imposed electrical activity [5, 35, 22, 59]. Activity-dependent dynamics in the hippocampus are crucial for learning and storing new memories. In addition, many *in vivo* sensory systems demonstrate that use-dependent and experience-driven activity regulate the development of neural circuits, modulating neurite branch stability and synaptogenesis [33, 74].

The neurochip could be used to apply various patterns of stimulated activity on the network to study plasticity mechanisms, and "learned" responses on different timescales—ranging from minutes to days—could be investigated. Stimuli could be either acute or chronic. Comparing activity between neurochip networks which are subject to stimulation versus those that are not, or comparing a single network's activity before and after stimulation, will reveal how activity-dependent effects shape the network's input-output responses. These studies could help elucidate how normal circuit formation and plasticity mechanisms give rise to network function for both *in vitro* and *in vivo* systems.

9.3.4 Subthreshold Synapses: Integrating with Optical Data Acquisition

A limitation of the neurochip system is that subthreshold events can not be cleanly recorded. Augmenting the suprathreshold connectivity data with subthreshold data would yield a

fuller picture of the state of the network and more fully elucidate how the network evolves.

In theory it should be possible to cleanly record subthreshold responses with optical techniques. Averaging many trials (10 or more) will be required to reduce noise levels low enough to cleanly detect the 0.1% changes in fluorescence expected for subthreshold events. The main issue with optical data acquisition at the moment is phototoxicity. Every culture I stained with potentiometric dye perished within 24 hours time. That makes optical data acquisition a terminal experiment—highly unfavorable. Even neurons exposed to very low levels of light–never directly exposed to the mercury excitation light—died. The cause of widespread cell death is unknown, but may be due to exposure to ethanol, as all dyes are dissolved in it. If the ethanol proves to be the toxic ingredient, it may be possible to dissolve the dye in distilled water or other non-toxic solvent.

9.3.5 Pharmacology: Fucose

The neurochip system can be utilized as a platform to examine the effects of various pharmacological agents on network formation and function.

As an example, fucose α (1-2) galactose has been implicated as a molecule that influences neuronal morphology and synapse formation [48, 37]. Inhibiting normal fucose function in animals has been found to impair memory formation [48]. This molecule could be added to neurochip cultures and compared to others which do not have the molecule added. Perhaps those cultures with fucose added would develop more reliable synapses. Stimuli could be imposed on cultures with and without the fucose molecule. Do cultures with higher concentrations of fucose exhibit a higher degree of plasticity?

The effects of a wealth of other growth factors—such as BDNF, FGF_2, and NT-3, to name a few—could be studied.

Appendix A

Neurochip Assembly Instructions

A.1 Step-by-Step Assembly

Methods and materials for chip assembly are listed below, along with tips and cautions. There is a lot of doing, then waiting, and the total assembly time can be as long as 2 weeks.

1. The microfabrication wizards deliver several dozen diced neurochips, each 2 cm x 1 cm, on a 500-μm-thick silicon wafer.

2. Place chip in milled PC board carrier, and glue around the edges to secure the chip in the pocket. (The PC board was designed by Jon Erickson, the manufacturing done by Advanced Electronics (Aurora, CO. www.4pcb.com).) The PC board contacts are soft-plated gold. The glue is our favorite biocompatible silicone adhesive, DAP, made by Dow Corning, available from Ace Hardware (UPC number: 070798006980).

   Be careful to not get glue on the contacts and make sure to keep the glue below the carrier surface so that the wire-bonding tip does not have to travel up and over a glue glob.

3. Wire-bond neurochip to carrier (gold ultrasonic bonds). This is done in micromachining lab. Settings are listed in their logbook. First bond to neurochip, then to chip carrier.

4. Protect and insulate the lead bonds. Build a dam between the neurochip lead bonding pads and neurocage array. It should be built very close to the bonding pads using DAP silicone glue. The dam is used to prevent uncured Sylgard from flowing toward, or into, the neurocage array. The dam doesn't need to be very high, just above the

level of highest wire bonds. Best to keep the dam height to a minimum so that it is easier to attach the culture dish at later stage. Then put drops of Sylgard 184 (Dow Corning) over leads to mechanically protect and electrically insulate them. Sylgard is mixed at 5:1 (instead of 10:1) to make it cure faster and to prevent it from being too runny. Let freshly prepared Sylgard cure for 20 min at 60°C prior to applying. It should just be a bit sticky when applied to neurochip/carrier. Even so, Sylgard is runny, and as it cures tends to spread out. Hence the need for the dam, otherwise a large region of neurochip can be covered up (even the cages!). Do, however, allow Sylgard to run on the PC board to cover carrier-side leads. Cure for a minimum of 24 hours at 60° C, and apply more as needed to ensure lead bonds are covered. A 48 hr cure is best.

5. Solder dual in-line plug (DIP) to carrier. DIP electronics are from Aries electronics (www.arieselectronics.com, NJ, Part No: 28-6625-21). This is a 28-pin plug, and the neurochip has only 17 connections. Hence, pull out the 11 pins that aren't needed. Hold the DIP with regular pliers, and pull extraneous pins with needle-nose pliers.

   Soldering is done in a little wood sculpture I made that has a window for the neurochip, so that its surface is not harmed when flipped upside down. Do NOT solder the ground pin yet.

6. Add platinum (Pt) ground/reference wire. Pt is 32-gauge and slips nicely through the hole in the PC board carrier even with the DIP pin present. A tail of Pt should be wrapped around its DIP pin. First solder the pin. Solder does not wet Pt—even though the mechanical connection seems to be firm—so use silver epoxy to make a more robust connection. Conducting silver epoxy is Silver Epoxy Type 50 from Transene (www.transene.com). Add a small amount of silver epoxy to pin and to top side. This should cure in 1 hr at 55°C, but I let it cure overnight to be sure.

7. Sylgard the solder joints to protect them from humidity/corrosion in incubator. Sylgard is mixed 5:1 as before; let it cure 20 min in oven before working with it. Then brush it over the top side of the chip carrier. The bottom joints could be done, but I choose not too, since they are so large and don't seem to degrade even in heat and humidity. Let it cure at least 24 hrs (48 is preferred).

8. Attach culture dish to chip. Dishes are made-to-order machined by the wizards (Mike Roy or Steve Olson) in the machine shop downstairs. The substrate is a 35-mm Corning dish (Corning part number: 430165) with a machined 23 mm x 12 mm window in the center. First paint on a fairly thick layer of DAP glue to the bottom of dish around the perimeter of the window. Very carefully slip the Pt wire up through the window and place the dish over the top. It is best to leave spare room on the neurochip side of the 2 cm x 1 cm chip at the expense of shortening the exposed area of the blank side. Let dry at least a couple of hours, overnight is better. Then comes the trickiest step of all: sealing the dish with Sylgard. Again, mix at 5:1, then pre-cure at least 20 min. It is critical to have just the right consistency of Sylgard: not enough pre-cure, and it runs everywhere. Too much pre-cure, and it is too stringy and stiff to work with. Nonetheless, err on the side of too long. Very carefully, paint Sylgard around the inside edge of the window, over the DAP glue. Let cure for a couple of hours, then check to see if good seal was made. If not, repeat process. (I usually have to do 2 layers or more.) After the last Sylgard layer is painted, let cure 24–48 hours.

9. Platinize electrodes. See Appendix B for full details on platinization of gold electrodes. In brief, set current i = 5 V/20 MΩ for 5–10 seconds. Observe the progress of platinization with the BHMJ microscope. During the first platinization, a platinum black layer deposits within seconds. For the remaining platinizations, observe the growth of the platinum black spongy layer. At this time, I do not believe it is necessary to treat with H_20_2, as Maher did. In fact, H_20_2 is known to severely damage parylene surfaces and should never be used with parylene chips. It also can have a bad effect on low-stress silicon-nitride insulation.

   Soak overnight in ddH_20 between and replatinize more times. This may help release weak spines, so that a stronger spongy layer of platinum black builds to make higher capacitance. Visibility can be hampered, however, if the platinum black is allowed to continue to grow out of the hole, spilling into the free region.

10. Soak the dish. Chips need to be soaked for 1 week after platinization for two reasons. One is to soak out the lead in the platinizing solution, which is toxic to neurons. The other reason is that chips seem to have better neuron adhesion if they are soaked prior to first use. This may have to do with monolayers of junk/cytophobic chemical

left over from the microfabrication or assembly process. Whatever the cause, it is beneficial to soak the neurochip in saline solution, changing the solution daily.

11. Experiment and enjoy!

Appendix B

Platinization and Electrode Impedance Measurement

B.1 Platinization Solution Recipe

Platinizing solution, also known as Kohlrausch solution, consists of the following:

1. 200 mL ddH_20
2. 0.01% lead acetate (w/v) 20 mg
3. 1% chloroplatinic acid (w/v) 2 g
4. 0.0025% HCl 5 uL

It should be stored in the dark, and when done so the shelf-life is apparently infinite. I inherited a stock of platinizing solution nearly 10 years old and it worked perfectly fine.

B.2 Procedure for Platinizing Neurochip Electrodes

1. Fill neurochip culture dish with about 3 mL platinizing solution, making sure that the fluid penetrates into the neurocages to contact the electrodes. It is helpful to put a small drop of ethanol over the cage array before adding platinizing solution. The low surface tension of the ethanol lets it fill the neurocage.
2. Pass a DC current of 1/4 μA through a 10- μm-diameter cage electrode, a current density of $j = 318$ mA/cm^2, for 5 seconds. The electrode should be held negative relative to the voltage of the platinizing solution (= case ground).

3. Remove and save platinizing solution. 3 mL of platinizing solution can be reused many times.

4. Immediately rinse the dish three times with ddH_20 and let sit several hours (overnight) in ddH_20.

5. The next day, drain dish and dry with N_2 gas hose.

6. Start back at the first step, platinizing again for 5 seconds at the same current as before.

7. Repeat this procedure 5–8 times until the desired capacitance has been reached or until the platinum bush becomes too large and starts impinging on the rest of the cage volume.

Quick aside: Curiously, platinizing one time for 25 seconds produced a drastically different result than 5x5 second platinization. The former grew a much larger bush and had a lower capacitance.

B.3 Electrode Impedance Measurement

The circuit for platinization and impedance measurement is maintained in the Pine lab. The electrode impedance measurement portion is essentially a voltage divider. A sinusoidal current is fed into a large resistor, $R = (20\ \mathrm{M}\Omega)$ in series with the neurochip electrode such that:

$$V_{elec} = V_{in}\left(\frac{Z_{elec}}{Z_{elec} + R}\right). \tag{B.1}$$

Rearranging to solve for the electrode impedance, Z_{elec}:

$$Z_{elec} = R\left(\frac{V_{elec}}{V_{in} - V_{elec}}\right). \tag{B.2}$$

By measuring the voltage across the neurochip electrode to ground, the magnitude of the impedance of the electrode can be calculated. By measuring the phase difference between the input signal and the voltage across the electrode, the real and imaginary parts can be

calculated.

$$\begin{aligned} \Re(Z) &= Z_{elec} cos(\phi) \\ \Im(Z) &= Z_{elec} sin(\phi) = X_c \end{aligned}$$

When measuring electrode impedance, the imaginary part of the impedance—the capacitive reactance—is ascribed to the capacitance of the electrode-saline interface; and the real part is ascribed to the neurocage geometry-dependent Ohmic conduction from the base of the cage out to a distant ground.

The shunt impedance of the cage electrode leads was measured on two chips. To make this measurement the conduction path to the electrode is blocked with a drop of silicon adhesive placed on top of the neurocages. The only path for current flow is from the lead, across the insulation to the saline solution. The reactance is ascribed to the shunt capacitance, and the real part is ascribed to the shunt resistance along the length of the leads.

Appendix C

Data Acquisition Hardware

C.1 Pre-Amplifier

Figure C.1 shows the circuit diagram of the preamplifier used for data collection during experiments described in this thesis. One of sixteen identical channels is shown. V_{elec} is the measured neurocage electrode voltage. The input signal drives an AC coupling ($\tau = 2.2$ sec). *Preampout* is the voltage output of the gain of 11, non-inverting op-amp. *StimIn* is the stimulus input, a sequence of timed voltage-pulses driving a 500 kΩ resistor to approximate a constant current pulse. The units of resistors and capacitors are given in Ω and Farads, respectively.

C.2 Filters and Buffer Amplifier

Figure C.2 shows the circuit diagram for the filters and buffer amplifier of the data acquisition hardware system. The system consists of a filter stage, buffer amplifier which drives a low-pass filter. The 2-pole low-pass filter has a cutoff frequency of 5 kHz. The gain of the op-amp is programmable, and can be set to either 63 or 1. The LabView software automatically sets the switch. For a stimulus channel the switch is closed (G = 1), and for a recording-only channel the switch is left open (G = 63). The output of the op-amp drives a low-pass filter set for 15 kHz. The output, V_{out}, is routed to the computer data acquisition card, NI-PCI-6071E. The units of resistors and capacitors are given in Ω and Farads, respectively.

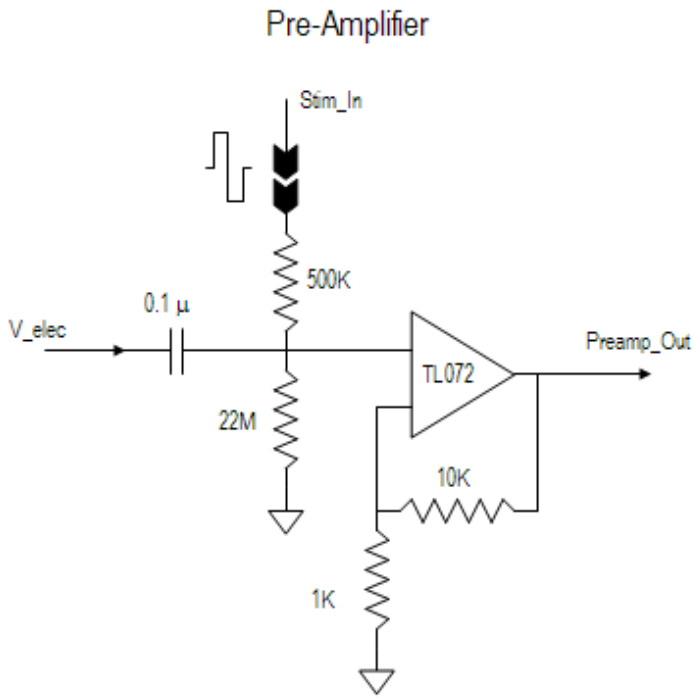

Figure C.1: Pre-amplifier circuit diagram for one of sixteen identical channels

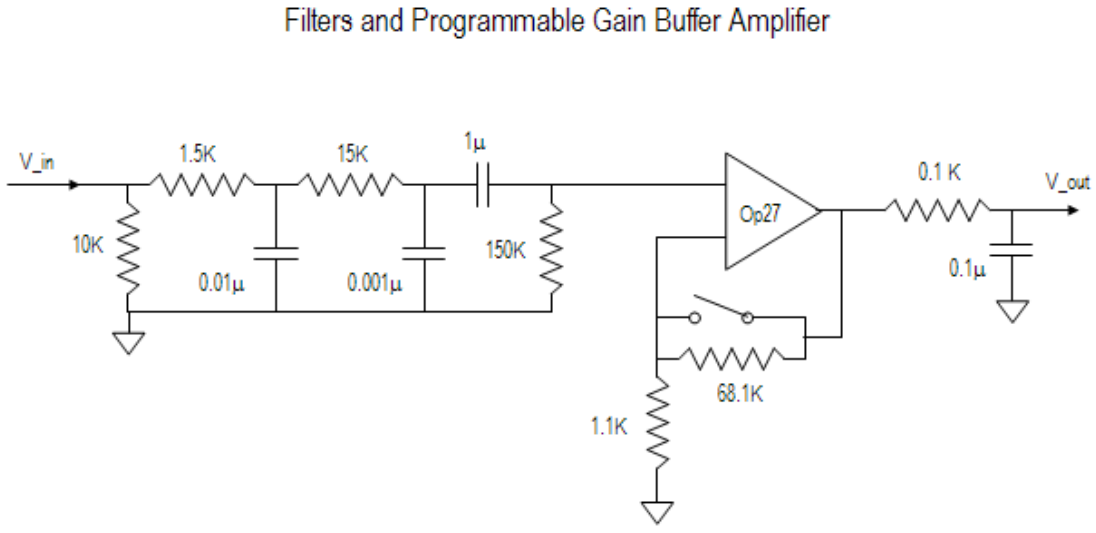

Figure C.2: Filters and buffer amplifier circuit digram for one of sixteen identical channels

C.3 Stimulus Generation

. The stimulus generation circuit produces timed voltage pulses. Both the timing and the amplitude of the voltage pulse is programmable with the LabView software (see Appendix D). 12 DIO bits (on the NI-PCI-6503 DIO board) set the voltage amplitude via a D/A converter. The output of the D/A drives buffer amplifiers to produce a voltage pulse in the range of $\pm$ 10 V. Channels are selected for stimulating by setting DIO bits which close (or open) a particular stimulus channel. The timing of the pulses is controlled via the NI-PCI-6602 counter/timer card, which outputs pulses at the desired stimulation time, relative to the start of data acquisition. The width of the pulses is also programmable, and set with LabView software.

Appendix D

LabView Data Acquisition Software

The basic functionality of the LabView (National Instruments; Austin, TX) data acquisition and stimulus generation progam— Neurochip-legacy.vi—is described below. Full documentation is available in the Pine lab, or by contacting me, Jon Erickson. Before executing the program, adjust all settings (described below) to taste. Execution is initiated by clicking Control + r or by clicking the white arrow at the top left of the control panel.

Figure D.1 is a screenshot of the front (control) panel. This appendix highlights its features.

Displays: The front panel is dominated by the 16-channel real-time display. This 4x4 array of graphs is laid in a pattern isomorphic to the neurochip electrodes. The corresponding electrode number is labeled in the upper-right-hand corner of each display. The scale of the y-axis (signal amplitude) is controlled by the control labeled "Y-axis-zoom factor" above the top-right corner of the 16-channel display.

Gain Settings: The gain for each channel can be independently set. Default gains are set with the controls labeled "stim gain default" and "record gain default" in the upper right hand corner. For a stimulation channel the default gain is G = 1/2, to maximize the dynamic range. For recording-only channels, the default gain is typically set to 50.

Data Acquisition Timing: Data can be acquired continuously by setting the slider to "continuous". In this mode, the display runs (and data is optionally acquired) until the "STOP" button is clicked. For finite data acquisition, set the slider to "finite". In this mode, the display and saving of data is stopped by the software after the desired number of samples has been taken. The time for which data is acquired in finite acquisition mode

is set in the control "trial length (ms)". Operating in finite mode also allows multiple trials to be programatically repeated—i.e., LabView controls a sequence of N trials, entered in the control labeled "Number trials".

Saving Data: Data for all 16 channels is saved when the boolean control "Save File(s)" is turned on. The file path to which data is saved is entered in the text control box labeled "File Path".

Stimulation Channels and Parameters: In the upper-left-hand corner are the controls to select the stimulus channels and parameters. The array of boolean controls is used to choose which, if any, channels are stimulated. The 4x4 array is isomorphic to the physical layout of the electrodes on the neurochip. The slider underneath provides a convenient control to turn all channels on or off. Typically it is left in a neutral position, meaning that stimulus channels are selected according to the boolean array above.

Immediately to the right are the controls to set the stimulus amplitude: Note that this sets the DAC voltage. So, for a current pulse, the amplitude of the current is computed as $I_{stim} = V_{DAC}/500\ \text{k}\Omega$.

The slider controls to the right set the polarity of the stimulus (positive or negative first) and whether the pulse is mono- or bi-polar. The timing of the stimulus is set in the digital controls to the right of the sliders. The time at which the stimulus is presented is set in the control "stim delay (ms)". The length of the A phase and B phase of the stimulation pulse are set in the fields "A length (ms)" and "B length (ms)", respectively.

Additional, more advanced features are available. Consult the full documentation in the Pine lab for further details.

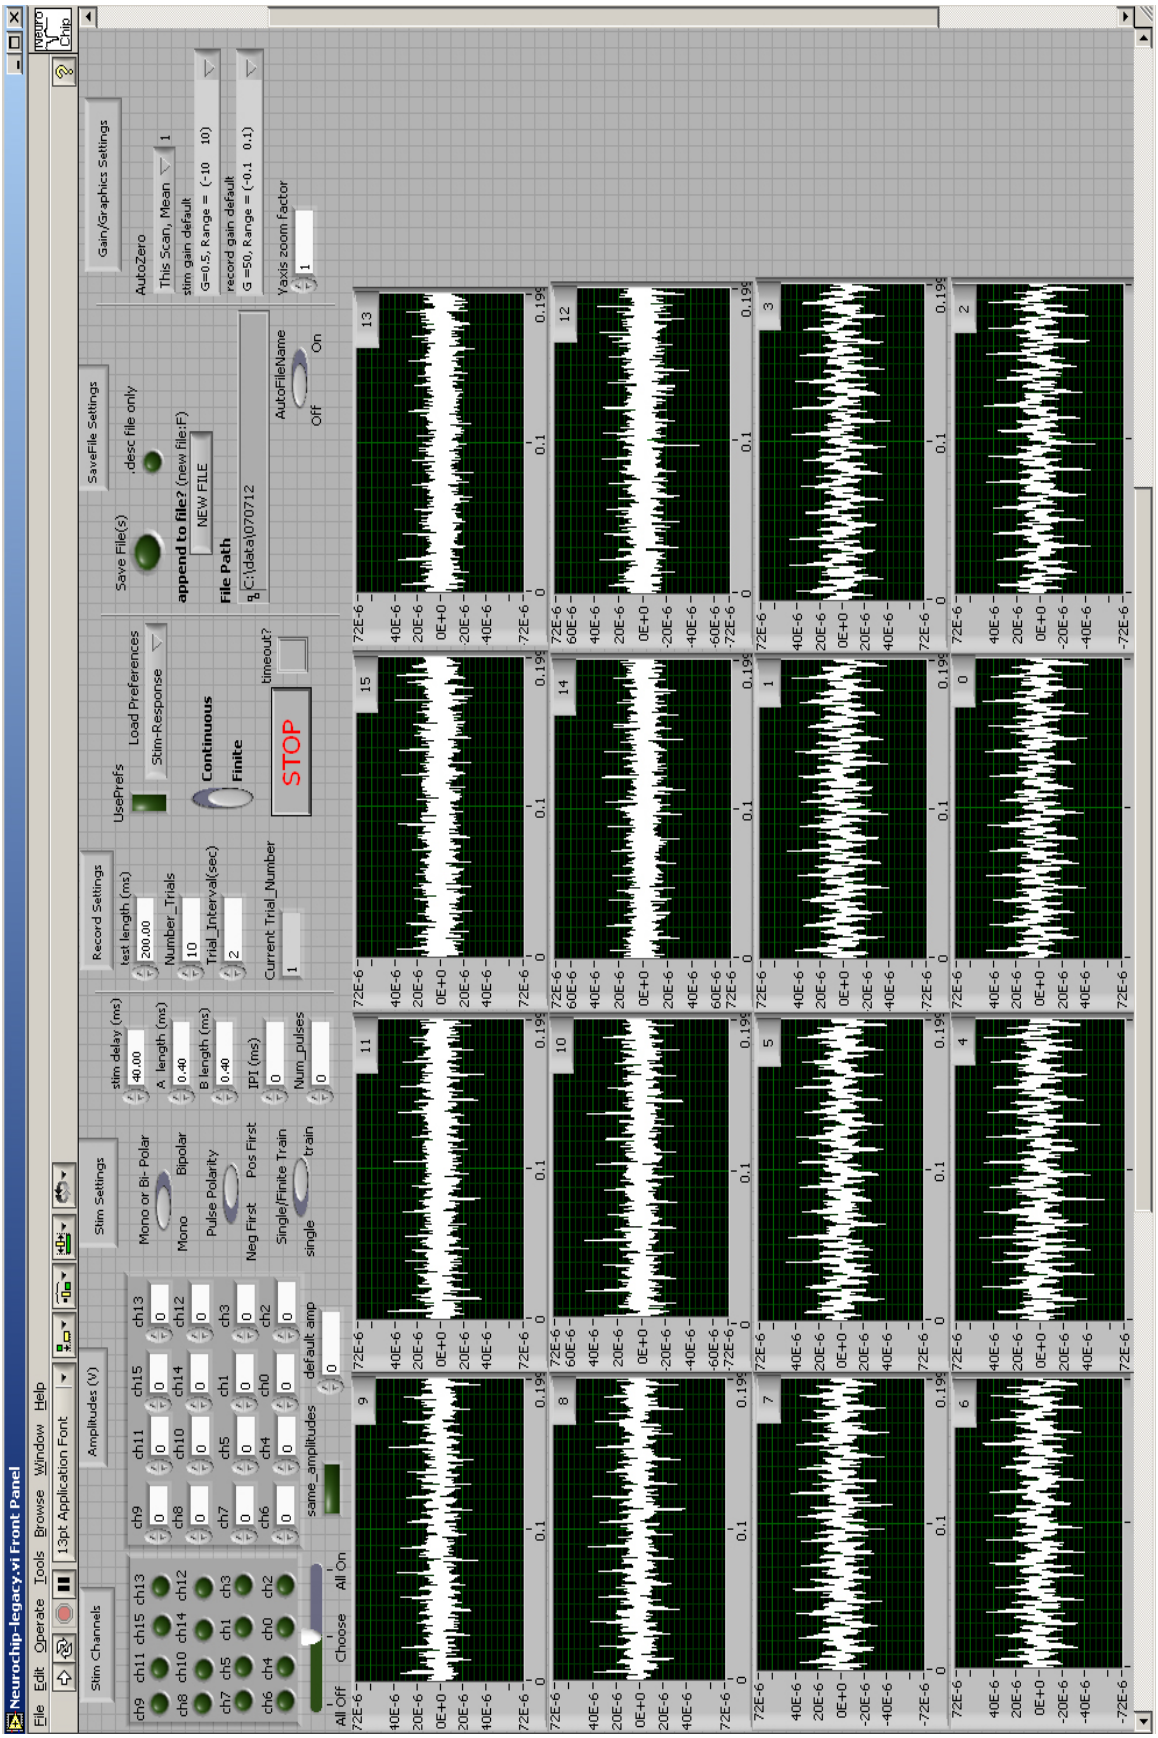

Figure D.1: Screenshot of LabView data acquisition software control panel

Appendix E

Protocol for Neurochip Cell Culture Preparation

E.1 Preparation of Neurochip Surface

1. Sterilize: 30 min of UV and/or soak dish in EtOH for 30 min and allow to evaporate.

2. Prepare surface by coating alternating layers of 0.05% (w/v) PEI (Sigma, P3143) and laminin (Sigma L2020,1mg/mL). One layer of each is sufficient to grow cells. If the culture is to be used in a dye-staining experiment, I deposit 3 layers of each known as the "layer-by-layer" (LbL) method $[PEI/lam]_3$. Or repeat the cycle as many times as you like (the references listed below deposit up to 8 layers of each).

   To start, make sure dish is completely dry. Add enough PEI, usually about 200 μL to cover the neurocage side. Incubate $>$ 4 hrs. I usually incubate overnight. (I am not sure it matters, but it is my contribution to cell-culture folklore.) Rinse PEI thoroughly with ddH20 3X, filling and aspirating dish with 2–3 mL at a time. Then do 2 rinses with HBSS (Hank's). After the final Hank's rinse, leave 3 mL in dish. To the Hank's solution in the dish add 3 μL of 1 mg/mL laminin. Incubate overnight. Next day drain HBSS so that just the trough is filled. Rinse 3X in ddH_20. (The reason for this step is that Hank's can leave behind nasty salt crystals when the dish is allowed to completely dry.) Then let completely dry. Repeat.

   The reason for LbL is to block dye absorption into neurocages. The parylene is a hydrophobic molecule. PEI and laminin are hydrophilic. Thus the dye "sees" a PEI/lam treated surface to have the same properties as a cell membrane and readily inserts. However, with a large enough hydrophilic domain built up, the dye can't

"see" the hydrophobic parylene domain, thus blocking dye absorption. In summary, LbL surface treatment is beneficial for blocking dye absorption into the parylene neurocages.

Helpful LBL references: [1, 39]

E.2 Neurochip Cell-Culture Technique: Loading

1. Plate "background" mass cultures around edges of dry neurochip, as depicted in Figure 4.1. Typically I plate 30 K per 40 μL, divided evenly between two 20 μL drops on either side of the cages. Make sure the chip has had ample time to completely dry, or these drops will tend to wet the chip, spreading large numbers of neurons into regions where they are not desired. Let sit 10 min in laminar flow hood, with the lid on to prevent evaporation. Then move chips to incubator. Incubate 1hr, then flood with 2.5 mL medium.

2. Inspect neurocages for bubble formation. Use 1 mL pipette to gently blow off any bubbles remaining on top of cages and on the blank chip side.

3. Place non-stick poly-HEMA chip firmly onto the blank side of the neurochip.

4. Squirt a few thousand neurons into the dish onto the polyHEMA chip. From here, neurons will be picked up with a custom cut micropipette (see Appendix G) and transported into the neurocages. Best to use neurons straight at maximum density (usually the density after spinning-down cells is a few million cells per milliliter) to minimize the volume that must be squirted. Minimizing the volume squirted helps to keep the cells localized on the polyHEMA chip

5. Place micropipette loading tip in dish just above polyHEMA surface and repeatedly place a neuron into a vacant neurocage. The maximum time the dish should be allowed to sit out is about 40 min. Past this time the pH of the medium has changed enough that cells might be damaged.

6. When finished loading, wait a few minutes, place the cap on the culture and gently transport to the incubator.

E.2.1 Comments on Manually Loading Neurons

Manually loading neurons into cages takes much practice and patience. The loader must ultimately develop his/her own technique. Having spent hundreds of hours performing this task, I will try to share a few tips here.

Loading neurons begins with squirting a large number (about 5,000) onto the non-stick polyHEMA chip. A small volume (2 μL) works best. Otherwise, neurons tend to be ejected far from the pipette tip, and spread too diffusely which makes finding a neuron annoying.

After locating a neuron, place the sucker (glass micropipette tip) about 20 μm away from the neuron and begin to gently suck to capture the neuron in the loading tip. By gently, I mean turn the dial on the control syringe about 1/4–1/2 a turn. Move slowly, wait a few seconds before applying more pressure. If you have to turn more than 2 turns, something is wrong.

After capturing the neuron, use the fine vertical adjust on the Leitz micromanipulator to raise the tip above the polyHEMA surface by 1/2 turn.

Use the coarse translation knob on the stage to position the neurocage array under the loading tip. The cages will be out of focus because they are below the plane of the polyHEMA chip. Refocus using the microscope focusing knob(s).

After refocusing, move the tip to the desired spot over the neurocage array then gently begin to lower the loading tip toward the array. Make sure it does not contact the neurochip surface because a) the chip surface, cages, and leads could be damaged and b) the sticky substrate on the neurochip will transfer onto the tip. This makes it very difficult to control sucking and pushing of neurons in the loading tip. I bring the tip down to within about 10–20 μm of the neurochip surface. Practice makes perfect.

Gently eject the neuron out of the loading tip. Again, 1/4 or 1/2 turns, should do. Wait a few seconds before turning the knobs, or the neuron will come flying out of the tip in an uncontrolled manner.

The neuron needs to be gently steered into the neurocage, until it finally contacts the chip surface inside the cage and anchors. This is the tricky part. I typically use 1/4 or 1/2 turns to gently push or pull the neuron as needed. I used the joystick on the Leitz micromanipulator to position the tip so that it was in line with the neuron and cage. It typically takes about 10 seconds for the neuron to settle toward the neurochip surface.

Translation: there is a window of less than 10 seconds to get the neuron in position over the cage. If I sensed the neuron was not going to get in position in time, I would more forcefully suck it back into the tip with one or two full turns.

Appendix F

Potentiometric Dye Staining Protocol

F.1 Staining Method

My potentiometric dye of choice is di4-ANEPPDHQ (Sigma, D36802). The dye stock solution is 1 mg/mL dissolved in EtOH. (Note: EtOH is highly neurotoxic, even at very low concentrations.) For staining 7.5 μL of stock solution are added to a culture in 3 mL of "Poo saline" (see Appendix H). The Poo saline bath is used for three rinses before staining, the last one stained, and three rinses after, leaving the cells in Poo saline. For each rinse cycle, gently pipette off most of the solution in the dish, always making sure to leave a small visible amount of liquid covering the neurons—never allow them to dry out! The staining time is 15 minutes, though 10 minutes seems to work well. I put the culture in the incubator for these 10–15 minutes. The dye seems to stay in the membrane quite well for several hours. No internal staining has been noted over time scales of 4 hrs. There is a rumor that the dye may be sufficiently soluble in water, which can be tested by drying down 7.5 μL in alcohol and then redissolving in 3 mL of saline.

Appendix G

Protocols for PEI, polyHEMA, and Making Neuron "Suckers"

G.1 PEI Solution Recipe and Protocol

PEI solution: 0.05% by weight in Borate Buffer (BBS) described below: PEI is very gooey and hard to work with. We usually make a 10X stock solution (0.5% PEI), diluted in BBS, and for final working PEI dilute by a factor of 10 again (also in BBS).

Borate Buffer: 3.1 g boric acid plus 4.75 g Borax in one liter of ddH20. Adjust pH to 8.4 (with NaOH). It is very important to adjust the pH before mixing solution; otherwise borax/boric acid crystals will not dissolve. Filter sterilize at 0.2 μm.

Reagents description:

1. PEI (polyethylene-imine): (50% w/v) Sigma, P-3143
2. Borax (Sodium tetraborate): Sigma, B-0127 (anhydrous 99%)
3. Boric acid: Fisher Scientific, A73-500.

G.2 How to Make Micropipette Loading Tips

1. Pull glass tubes (see below) using the following settings on puller:

Heat = 600, Pull = 170, Vel = 70, Del=125.

(This is currently stored as program number "07" on the puller.) P = 500 is another specified setting (not sure if this is relevant).

Glass capillary tubes come from: World Precision Instruments, Sarasota, FL, (813)-371-1003

Glass 1BBL W/Fil 1.0 mm 4IN 1B100F-4

2. Cut tips using diamond cutter. Use scale graticule. Under 10X one small division is 10 μm. I cut so that the inner diameter is about 40 μm (OD will be about 50-60 um).

Saw back and forth, slowly increasing pressure of tip into cutter. Continue until tip breaks. Examine tip for clean break. If it broke cleanly, keep it. If not, reject it.

G.3 How to Make polyHEMA Non-Stick Chips

Mix 40 mg Poly(2-hydroxyethylmethacrylate, "polyHEMA" (Sigma, P3932) in 2 mL EtOH. Let dissolve overnight with gentle stirring/rocking.

To coat surface: place a 15-mm-square No. 1 cover slip in a 35-mm petri dish. Deposit several small drops, each about 50 μL, to cover the entire cover slip, and cover the Petri dish. Keeping the dish covered is critical because the slow evaporation of EtOH leaves a very nice, optically smooth surface, whereas fast evaporation leaves behind a "cracked" polyHEMA surface. On a hot or dry day, it may even be necessary to put a covered culture dish inside of another covered, larger chamber.

When first coat is dry, repeat procedure at least 1 more time. A minimum of 2 coatings makes the non-adhesive surface much more uniform than a single coating. This protocol was given to me by Gary Chow.

I like to put small drops of silicone glue on the polyHEMA surface after it has dried. The glue drops are spaced approximately for 5x5 mm square pieces cut with a diamond scribe. The glue drops are used to mark the side of the coverslip with the deposited polyHEMA–so that I can tell which way is up. These glue drops are also very useful for handling the polyHEMA non-stick chip while transporting them into a neurochip culture dish.

More hints and tips: Our experience indicates that raw polyHEMA from Sigma goes bad within a couple of months. This usually manifests by crystals looking "wet" or clumping together. A good, fresh stock should be as dry as table salt.

Appendix H

Bath Solution Recipes

H.1 Recipe for Poo Saline

To make a 400-mL stock of Poo bath (the name derives from neurobiologist M. M. Poo, from whom the recipe comes), mix 400 mL double distilled water with the following:

Reagent	Concentration (mM)	Quantity
Hepes	10	4 mL of 1 M solution
KCl	3	0.090 g
Glucose	8	0.58 g
NaCl	145	3.4 g
CaCl2	3	0.18 g
MgCl2	2	0.16 g

After mixing all ingredients, sterile filter the solution using an 0.2 μm filter from Nalgene.

The final product should be 320 mM. Check with osmometer to make sure everything is rational. I made fresh Poo saline about every three months, just to be safe.

H.2 Recipe for Neurobasal Culturing Medium

Neurobasal (NB)-based medium is used for hippocampal cell culture. It is reported to be optimized for survival of neurons in low-density dissociated cell culture [7]. I have informally studied survival in it compared to DMEM-based medium and did see good evidence to

substantiate this claim. The main difference between NB- and DMEM-based media is the sodium concentration—NB has the half the normal sodium concentration (55 versus 110 mM). This may help to keep neurons from becoming electrically active early on, preventing them from excitotoxicity.

Mix the following in a 50 mL tube:

Recipe for Neurobasal Cell-Culture medium		
Reagent	Quantity (mL)	Vendor
Neurobasal Medium	46	Invitrogen 21103-049
B27	1	Invitrogen 17504-044
Glutamax	0.125	Invitrogen 35050-061
Equine Serum	2.5	Hyclone SH30074.03

H.3 Recipe for DMEM-Based Medium

DMEM medium is a standard cell-culture medium. I used it here for neurochip cultures, gradually changing the solution from NB- to DMEM-based medium. The reason is that cultures evidently were much more electrically active at a few weeks *in vitro* in DMEM based medium, compared to NB. DMEM is supplemented with 10% Ham's F12 (a vitamin mix, similar to B27), and 5% equine serum (necessary for glial cell replication). Mix the following ingredients in a 50 mL tube:

Recipe for DMEM Cell Culture Medium		
Reagent	Quantity (mL)	Vendor
High glucose DMEM	43	Invitrogen 10313021
Ham's F12	4.5	Invitrogen 31765-035
Equine Serum	2.5	Hyclone SH30074.03

Appendix I

Quantifying Probability of Driven Responses

To quantify and justify the idea that observing 3 spikes in a window of 2 msec constitutes a significant response, consider the following:

Spontaneous spikes are generated by a Poisson process. An (over-)estimate for the spontaneous spike rate, is 5 spikes per ten 200-msec trials. So $r_{spont} \leq 5$ spikes per 2000 msec. Now the question is: given this spontaneous firing rate, what is the probability of observing 3 or more spikes in a time window of 2 msec over 10 total trials? To start, consider the probability of one spike occurring during a 2 msec interval. Plugging into the Poisson distribution yields:

$$p(x,r) = p(x=1, r=1/1000) = \frac{e^{-r}r^x}{x!} = \frac{e^{-1/1000}\left(\frac{1}{1000}\right)^1}{1!} \approx \frac{1}{1000} \tag{I.1}$$

where x represents the number of observed responses in a jitter window of $\Delta t = T_{jitter} = 2$ msec, and r corresponds to the number of spikes observed due to spontaneous activity in the same time jitter window. Now, given the probability of one spike occurring in this window, the probability, that X ≥ 3 or more spikes occur within the jitter window is computed as follows:

$$P\left(X \geq 3\right) = \sum_{1}^{N_{windows}} \sum_{k=3}^{10} \left(\frac{1}{1000}\right)^k \binom{10}{k} \tag{I.2}$$

where $N_{windows}$ is the number of windows of width T_{jitter} in which spikes may fall. How many are there? Realistically—and empirically verified to be true—is the fact that a synaptic response almost always occurs within about 20 msec. For a sampling rate of 20 kHz,

there are about (20 msec) x (20 windows/msec) = 400 windows. Computing the above sum reveals that the probability of detecting 3 or more spikes in a single window of width 2 msec is about 0.00005. Therefore, if 3 or more spikes are detected within a jitter window of $T_{jitter} = 2$ msec for 10 consecutive trials, they are almost certainly generated in response to the stimulus (not by spontaneous activity). Seeing how 3 or more spikes are very unlikely to be generated by spontaneous activity, one might be tempted (and justified) to lower the threshold to 2 spikes to declare a significant response. The point is moot, however, since a large majority of synaptic connections resulted in more than 3 out of 10 action potential responses.

Appendix J

Supplemental Connectivity Maps

This appendix includes connectivity maps over time. I have culled these out of 16 cultures examined. The level of complexity in a culture is highly variable, but all are seen to mature very quickly between about 7–17 DIV.

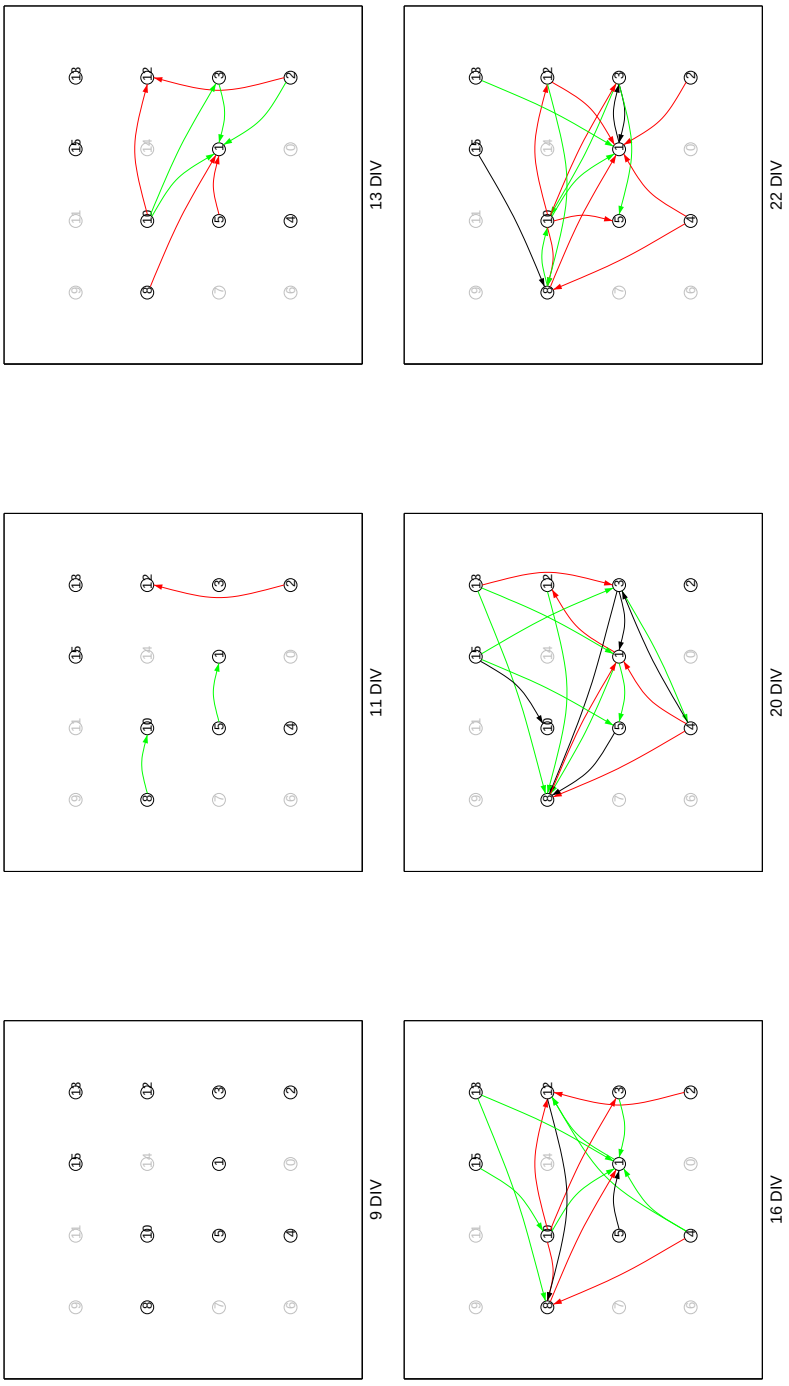

Figure J.1: Connectivity Maps: Network develops rapidly between 11–13 days.

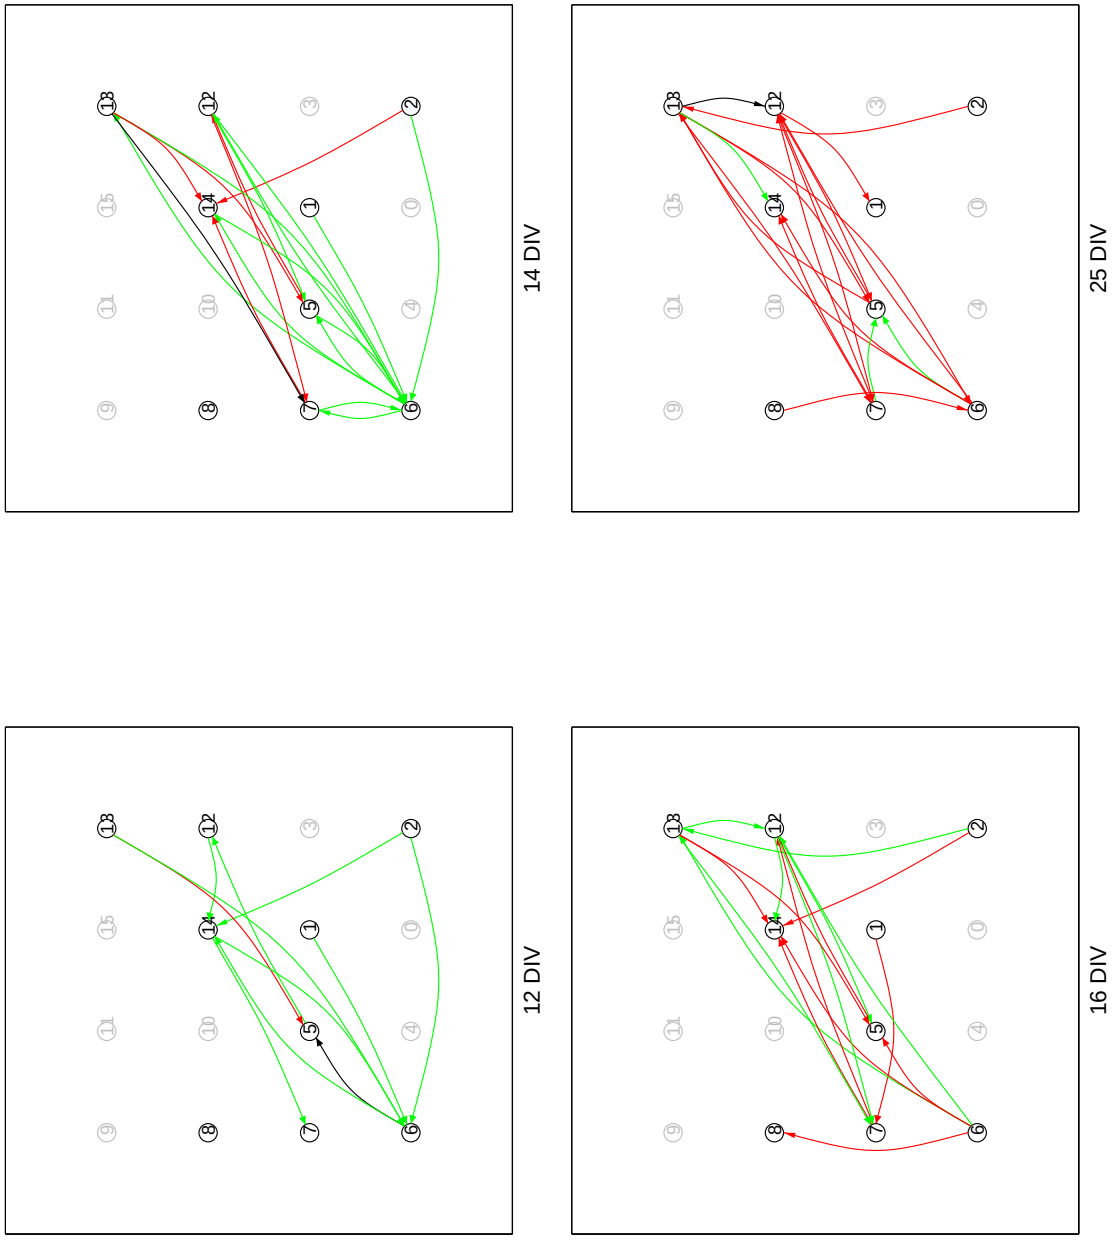

Figure J.2: Connectivity Maps: During the last days, the transmission delays are decreased so that most connections change from medium to fast.

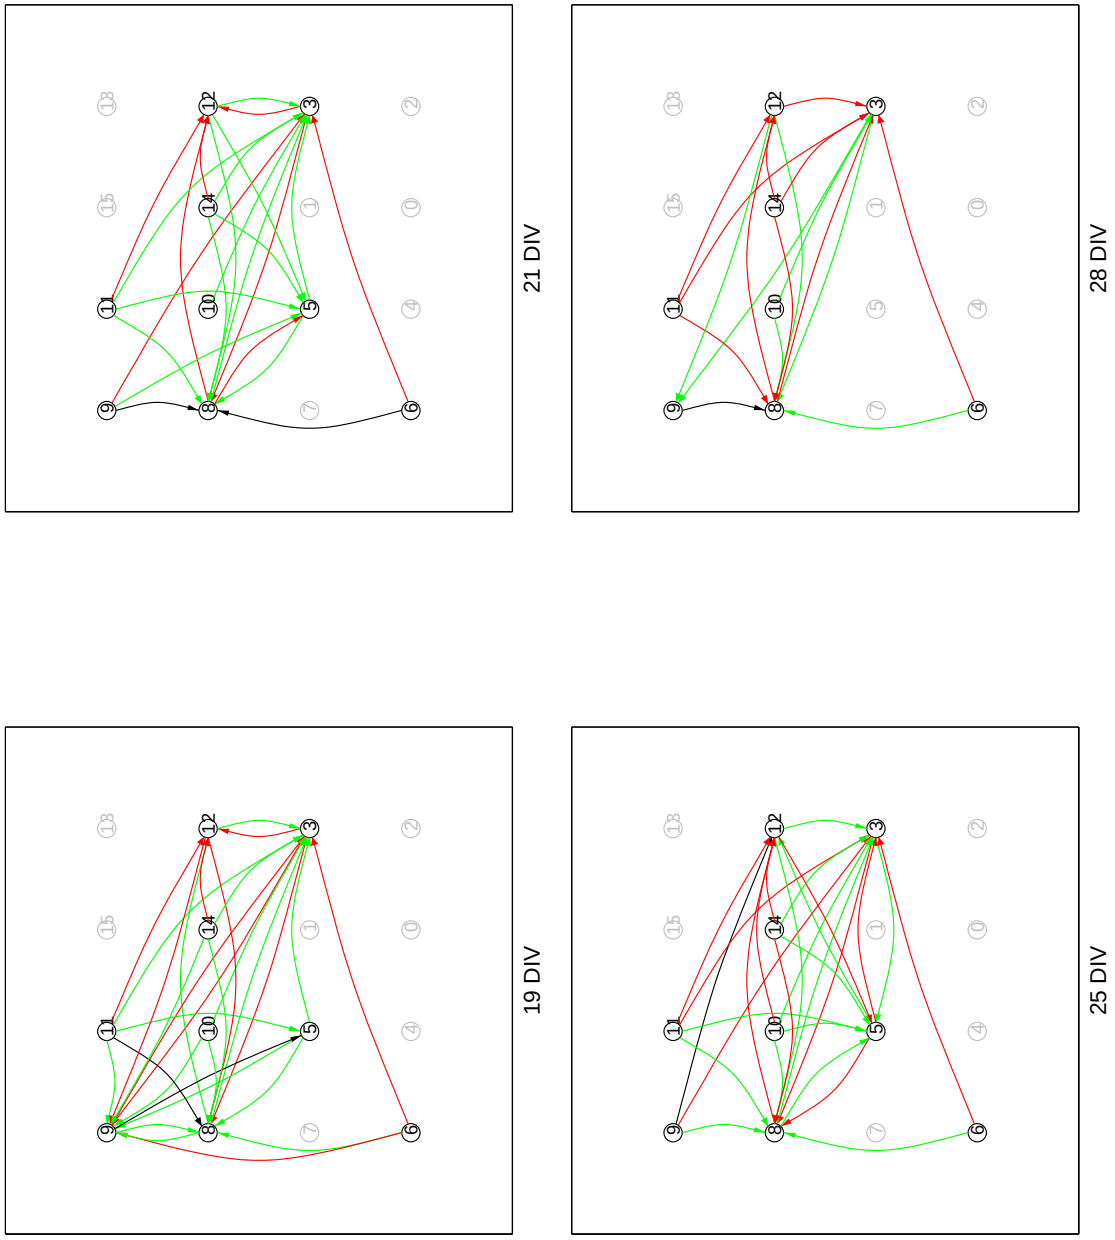

Figure J.3: Connectivity Maps: Many connections are stable from 19–28 days old. The culture is starting to die out at 28 days.

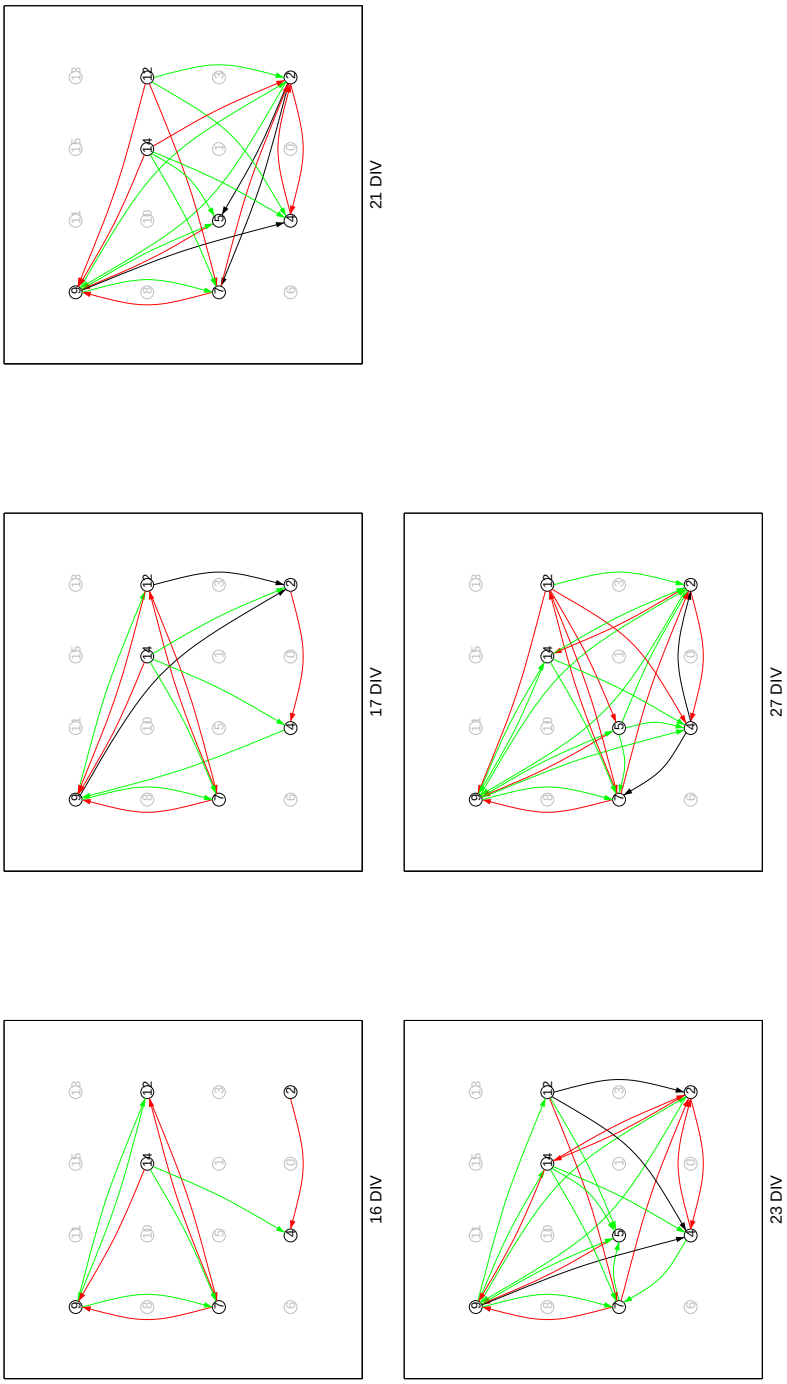

Figure J.4: Connectivity Maps: Array-wide connections are made. Culture is dynamic even at 4 weeks old.

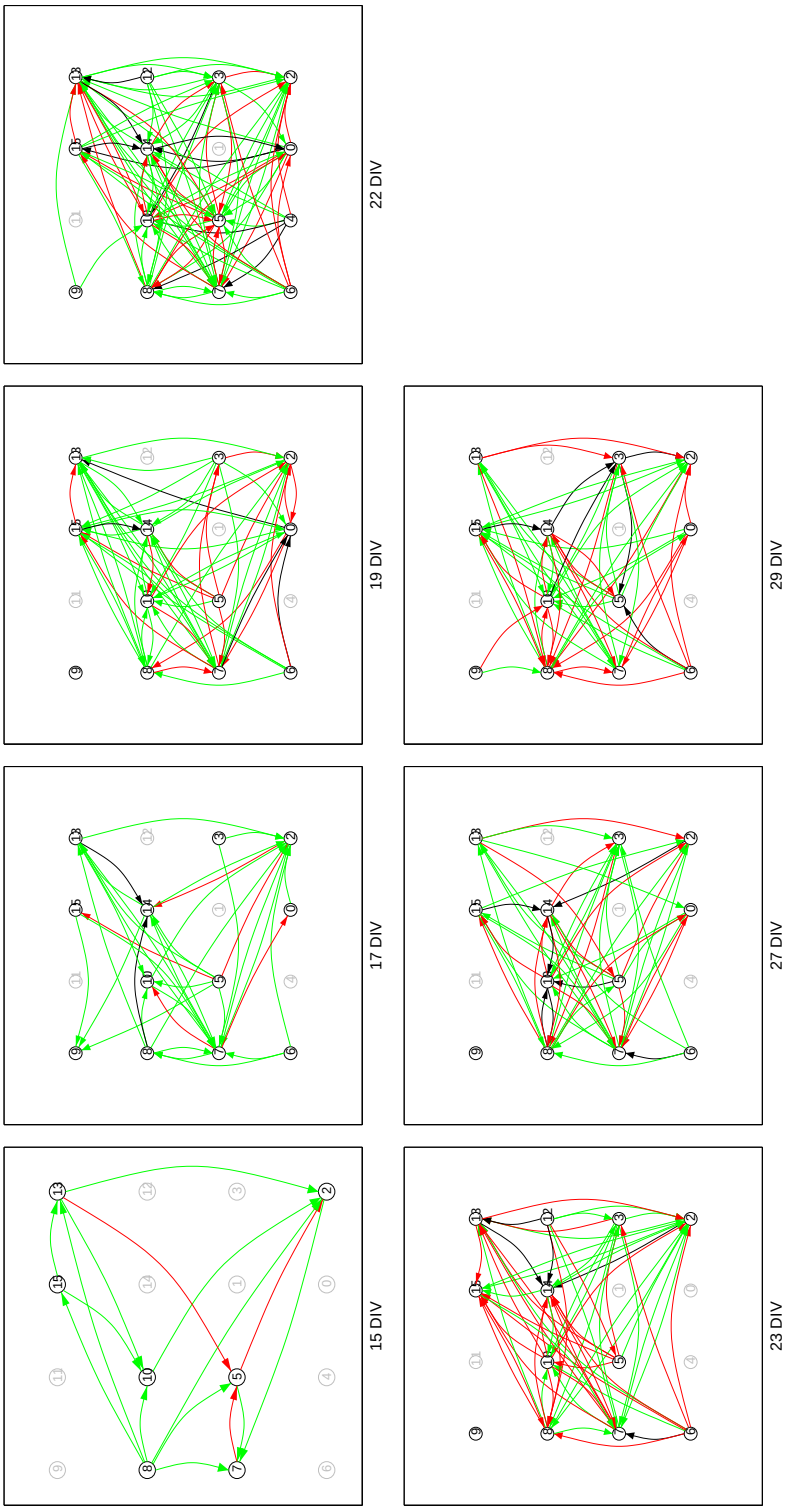

Figure J.5: Connectivity Maps: Extremely dense networking

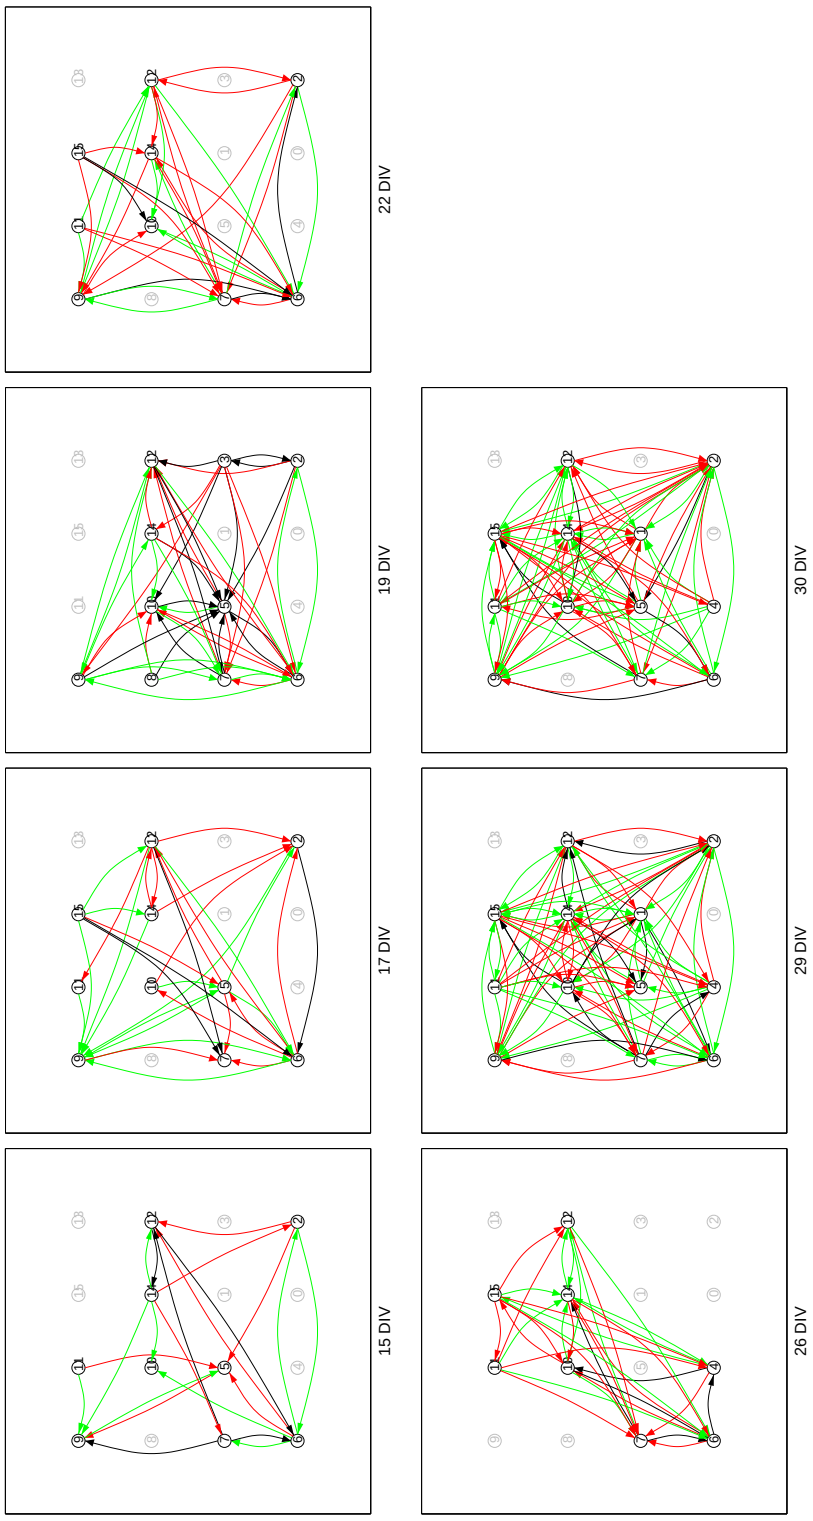

Figure J.6: Connectivity Maps: Another culture with extremely dense networking

Appendix K

CultureState: Another View of Connectivity

Initial work to further analyze connectivity—to determine whether detected connections arise from mono- or di-synaptic pathways, or both—has been completed. To discern a mono- versus a di-synaptic connection I imposed a simple set of rules:

1. For a connection from neuron A to neuron B, first-generation responses also exist from A to neuron C and also from C to B—i.e., C is the interneuron in the $A-C-B$ pathway.

2. The delay times for the responses from A to C and from C to B meet the criterion that they approximately sum to the delay time measured for the A to B response:

$$|\delta t_{AB} - (\delta t_{AC} + \delta t_{CB})| < \epsilon$$

   where ϵ is chosen (generously) here to be 2 msec.

If a pathway meets both the criteria above, then this connection was tagged as a disynaptic pathway. In the case that there were multiple generations of responses, if the disynaptic pathway delay time was equal to the second-generation time delay, then this connection was tagged as having both mono- and di-synaptic pathways. Note that, strictly speaking, the mono- and di- synaptic terminology has been misappropriated, since I only search a chain of connections 2-deep. Therefore, if either of the identified $A-C$ or $C-B$ pathways is itself disynaptic, then the $A-C-B$ pathway is actually higher-order, at least tri-synaptic. (A tree-search algorithm could be used to full de-construct all pathways, but that remains for future work.)

The results of the dissection of connections into mono-, di- synaptic pathways, or both, can be visualized as in the Figure K.1. Each sub-panel in the 4x4 display shows the responses to the cell being stimulated. Inside each sub-panel is a 4x4 grid of squares which is isomorphic to the physical geometry of the neurocages. Boxes outlined in black indicate the neuron was tested, while a grayed-out box indicates that no neuron was present/tested. The solid black square represents the stimulus electrode. The other boxes are colored according to the observed first- and (possibly) second- generation response times. In the case that the connection is putatively monosynaptic–i.e., no potential di-synaptic connection was found–the box is colored according to the monosynaptic response time. In the case that the response is putatively solely disynaptic–i.e., the first-generation response results from a chain of first-generation responses, the box is broken into sub-boxes: The top-half is the measured first-generation, di-synaptic response time. The colors of the two 1/4 boxes covering the bottom-half correspond to the response times of the two first-generation responses which are linked to form the di-synaptic response. Finally, in the case that both mono- and di-synaptic pathways were detected, the box is broken into a total of 4 sub-boxes as follows: The box covering the left-half is color-coded according to the first-generation, monosynaptic response. The right half is broken into thirds, with the top representing the total di-synaptic delay time, and the bottom two 1/3-boxes colored according to the two pathways which chain to form the disynaptic response.

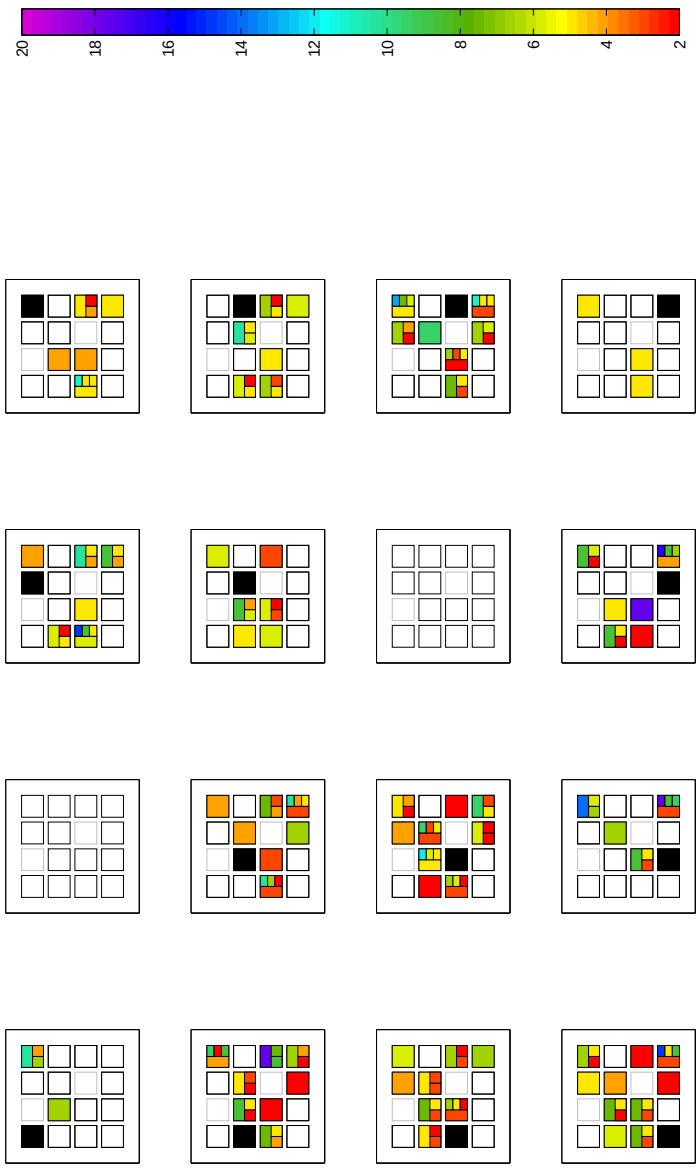

Figure K.1: CultureState provides a visualization of connectivity to discern between mono- and di-synaptic pathways. See text for full description. This particular figure corresponds to the connectivity map shown in J.5 at 22 DIV. The color-bar at right is the color-key for the delay times in units of msec. The example shown is particularly rich in the number of detected disynaptic connections.

References

[1] H Ai, H Meng, I Ichinose, S A Jones, D Mills, Y Lvov, and X Qiao. Biocompatibility of layer-by-layer self-assembled nanofilm on silicone rubber for neurons. *J Neurosci Methods*, 128(1–2):1–8, Sep 2003.

[2] F Arnold, F Hofmann, C Bengtson, M Wittmann, P Vanhoutte, and H Bading. Microelectrode array recordings of cultured hippocampal networks reveal a simple model for transcription and protein synthesis-dependent plasticity. *J. Physiol. (Lond.)*, 564:3–19, Apr 2005.

[3] H Bading, M Segal, N Sucher, H Dudek, S Lipton, and M Greenberg. N-methyl-D-aspartate receptors are critical for mediating the effects of glutamate on intracellular calcium concentration and immediate early gene expression in cultured hippocampal neurons. *Neuroscience*, 64:653–664, Feb 1995.

[4] G Bi and M Poo. Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature*, 401(6755):792–796, Oct 1999.

[5] G Bi and M M Poo. Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu Rev Neurosci*, 24:139–166, 2001.

[6] D Branch, B Wheeler, G Brewer, and D Leckband. Long-term maintenance of patterns of hippocampal pyramidal cells on substrates of polyethylene glycol and microstamped polylysine. *IEEE Trans. Biomed. Eng.*, 47(3):290–300, 2000.

[7] G J Brewer, J R Torricelli, E K Evege, and P J Price. Optimized survival of hippocampal neurons in B27-supplemented Neurobasal, a new serum-free medium combination. *J Neurosci Res*, 35(5):567–576, Aug 1993.

[8] J Buchhalter and M Dichter. Electrophysiological comparison of pyramidal and stellate

nonpyramidal neurons in dissociated cell culture of rat hippocampus. *Brain Res. Bull.*, 26:333–338, Mar 1991.

[9] D A Butts, C Weng, J Jin, C Yeh, N Lesica, J Alonso, and G B Stanley. Temporal precision in the neural code and the timescales of natural vision. *Nature*, 449:92–95, Sep 2007.

[10] C B Chien and J Pine. An apparatus for recording synaptic potentials from neuronal cultures using voltage-sensitive fluorescent dyes. *J Neurosci Methods*, 38:93–105, 1991.

[11] C B Chien and J Pine. Voltage-sensitive dye recording of action potentials and synaptic potentials from sympathetic microcultures. *Biophys J*, 60(3):697–711, 1991.

[12] DB Chklovskii, B W Mel, and K Svoboda. Cortical rewiring and information storage. *Nature*, 431:782–788, 2004.

[13] G Chow. *Laser tweezers for moving live, dissociated neurons.* PhD thesis, Caltech, 2007.

[14] E Claverol-Tinture and J Pine. Extracellular potentials in low-density dissociated neuronal cultures. *J Neurosci Methods*, 117(1):13–21, May 2002.

[15] C M Colbert and D Johnston. Axonal action-potential initiation and Na+ channel densities in the soma and axon initial segment of subicular pyramidal neurons. *J Neurosci*, 16(21):6676–6686, Nov 1996.

[16] A Destexhe and E Marder. Plasticity in single neuron and circuit computations. *Nature*, 431:789–795, 2004.

[17] M Djurisic, M Zochowski, M Wachowiak, C Falk, L Cohen, and D Zecevic. Optical monitoring of neural activity using voltage-sensitive dyes. *Methods Enzymol*, 361:423–451, 2003.

[18] J Erickson, A Tooker, Y-C Tai, and J Pine. Caged neuron multi-electrode array. In *Proc. 4th International Meeting on Substrate Integrated Multi-Electrode Arrays*, pages 52–53, Reutlingen, Germany, July 2004.

[19] J Erickson, A Tooker, Y-C Tai, and J Pine. Caged neuron multi-electrode array: a new generation of neurochip. In *Proc. Society for Neuroscience Annual Conference*, Washington, DC, November 2005.

[20] J Erickson, A Tooker, Y-C Tai, and J Pine. Progress on the caged-neuron MEA project. In *Proc. 5th International Meeting on Substrate Integrated Multi-Electrode Arrays*, pages 192–193, Reutlingen, Germany, July 2006.

[21] J Erickson, A Tooker, Y-C Tai, and J Pine. Caged neuron multi-electrode array: a new tool for probing connectivity in cultured neural networks. In *Proc. Society for Neuroscience Annual Conference*, San Diego, CA, November 2007.

[22] D Eytan, N Brenner, and S Marom. Selective adaptation in networks of cortical neurons. *J Neurosci*, 23(28):9349–9356, Oct 2003.

[23] J B Feix, B Kalyanaraman, C F Chignell, and R D Hall. Direct observation of singlet oxygen production by merocyanine 540 associated with phosphatidylcholine liposomes. *J Biol Chem*, 263(33):17247–17250, Nov 1988.

[24] R M Fitzsimonds, H J Song, and M M Poo. Propagation of activity-dependent synaptic depression in simple neural networks. *Nature*, 388(6641):439–448, Jul 1997.

[25] C Gold. *Biophysics of extracellular action potentials*. PhD thesis, Caltech, 2007.

[26] A Grinvald, R Hildesheim, I C Farber, and L Anglister. Improved fluorescent probes for the measurement of rapid changes in membrane potential. *Biophys J*, 39(3):301–308, Sep 1982.

[27] G W Gross. Simultaneous single unit recording in vitro with a photoetched laser deinsulated gold multimicroelectrode surface. *IEEE Trans Biomed Eng*, 26(5):273–279, May 1979.

[28] G W Gross, E Rieske, G W Kruetzbert, and Meyer A. A New fixed-array multi-microelectrode system designed for long-term monitoring of extracellular single unit neuronal activity in vitro. *Neuroscience Letters*, 6:101–105, 1977.

[29] Q He, E Meng, Y-C Tai, C Rutherglen, J Erickson, and J Pine. Parylene technology for mechanically robust neuro-cages. In *Proc. Transducers*, pages 995–998, Boston, MA, October 2003.

[30] B Hille. *Ionic Channels of Excitable Membranes*. Sunderland, MA: Sinauer Associates, 1992.

[31] G R Holt and C Koch. Electrical interactions via the extracellular potential near cell bodies. *J Comput Neurosci*, 6(2):169–184, Mar 1999.

[32] H W Horch and L C Katz. BDNF release from single cells elicits local dendritic growth in nearby neurons. *Nat. Neurosci.*, 5:1177–1184, Nov 2002.

[33] J Y Hua and S Smith. Neural activity and the dynamics of central nervous system development. *Nat. Neurosci.*, 7(4):327–332, 2004.

[34] N Y Ip, Y Li, G D Yancopoulos, and R M Lindsay. Cultured hippocampal neurons show responses to BDNF, NT-3, and NT-4, but not NGF. *J Neurosci*, 13(8):3394–3405, Aug 1993.

[35] Y Jimbo, T Tateno, and H P Robinson. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys J*, 76(2):670–678, Feb 1999.

[36] S B Jun, M R Hynd, N Dowell-Mesfin, K L Smith, J N Turner, W Shain, and S J Kim. Low-density neuronal networks cultured using patterned poly-l-lysine on micro-electrode arrays. *J Neurosci Methods*, 160(2):317–326, Mar 2007.

[37] S.A. Kalovidouris, C.I. Gama, L.W. Lee, and L.C. Hsieh-Wilson. A role for fucose alpha(1-2) galactose carbohydrates in neuronal growth. *J. Am. Chem. Soc.*, 127:1340–1341, Feb 2005.

[38] B Kalyanaraman, J B Feix, F Sieber, J P Thomas, and A W Girotti. Photodynamic action of merocyanine 540 on artificial and natural cell membranes: involvement of singlet molecular oxygen. *Proc Natl Acad Sci U S A*, 84(9):2999–3003, May 1987.

[39] P J Lein, G A Banker, and D Higgins. Laminin selectively enhances axonal growth and accelerates the development of polarity by hippocampal neurons in culture. *Dev Brain Res*, 69(2):191–197, Oct 1992.

[40] V Lessmann, K Gottmann, and R Heumann. BDNF and NT-4/5 enhance glutamatergic synaptic transmission in cultured hippocampal neurones. *Neuroreport*, 6:21–25, 1994.

[41] E Levine, C Dreyfus, I Black, and M Plummer. Brain-derived neurotrophic factor rapidly enhances synaptic transmission in hippocampal neurons via postsynaptic tyrosine kinase receptors. *Proc. Natl. Acad. Sci. U.S.A.*, 92:8074–8077, Aug 1995.

[42] E Maeda, H Robinson, and A Kawana. The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J. Neurosci.*, 15:6834–6845, Oct 1995.

[43] M P Maher, J Pine, J Wright, and Y-C Tai. The neurochip: a new multielectrode device for stimulating and recording from cultured neurons. *J Neurosci Methods*, 87(1):45–56, Feb 1999.

[44] M P Maher, J Wright, J Pine, and Y C Tai. A microstructrue for interfacing with neurons: the Neurochip. *Proc. IEEE Med. Bio*, 20(4):1698–1702, 1998.

[45] H Markram, J Lubke, M Frotscher, and B Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275:213–215, 1997.

[46] C Marrese. Preparation of strongly adherent platinum black coatings. *Analytical Chemistry*, 59:217–218, 1987.

[47] E Meng, Y-C Tai, J Erickson, and J Pine. Parylene technology for mechanically robust neuro-cages. In *Proc. Micro Transducers and Systems*, Squaw Valley, CA, October 2003.

[48] H Murrey, C Gama, S Kalovidouris, W Luo, E Driggers, B Porton, and L Hsieh-Wilson. Protein fucosylation regulates synapsin Ia/Ib expression and neuronal morphology in primary hippocampal neurons. *Proc. Natl. Acad. Sci. U.S.A.*, 103:21–26, Jan 2006.

[49] A L Obaid, L M Loew, J P Wuskell, and B M Salzberg. Novel naphthylstyryl-pyridium potentiometric dyes offer advantages for neural network analysis. *J Neurosci Methods*, 134(2):179–190, 2004.

[50] A Offenhausser, C Sprossler, M Matsuzawa, and W Knoll. Electrophysiological development of embryonic hippocampal neurons from the rat grown on synthetic thin films. *Neurosci Lett*, 223(1):9–12, Feb 1997.

[51] S Panzeri, R S Petersen, S R Schultz, M Lebedev, and M E Diamond. The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron*, 29:769–777, Mar 2001.

[52] J Pine. Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J Neurosci Methods*, 2(1):19–31, Feb 1980.

[53] M M Poo. Neurotrophins as Synaptic Modulators. *Nat Rev Neurosci*, 2:24–32, 2001.

[54] W Rall. Electrophysiology of a dendritic neuron model. *Biophys J*, 2:145–167, Mar 1962.

[55] W Regehr, J Pine, and D Rutledge. A long-term in vitro silicon-based microelectrode-neuron connection. *IEEE Trans. Biomed. Eng*, 35(12):1023–1032, 1988.

[56] D A Robinson. The electrical properties of metal microelectrodes. *Proc. IEEE*, 56(6):1065–1072, June 1968.

[57] M E Ruaro, P Bonifazi, and V Torre. Toward the neurocomputer: image processing and pattern recognition with neuronal cultures. *IEEE Trans Biomed Eng*, 52(3):371–383, Mar 2005.

[58] H E Scharfman. Hyperexcitability in combined entorhinal/hippocampal slices of adult rat after exposure to brain-derived neurotrophic factor. *J Neurophysiol*, 78(2):1082–1095, Aug 1997.

[59] G Shahaf and S Marom. Learning in networks of cortical neurons. *J Neurosci*, 21(22):8782–8788, Nov 2001.

[60] W Singer. Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24:49–65, Sep 1999.

[61] N Spruston and D Johnston. Perforated patch-clamp analysis of the passive membrane properties of three classes of hippocampal neurons. *J Neurophysiol*, 67(3):508–529, Mar 1992. In Vitro.

[62] C A Thomas, P A Springer, G E Loeb, Y Berwald-Netter, and L M Okun. A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exp Cell Res*, 74:61–66, 1972.

[63] A Tooker. *Development of Biocompatible Parylene Neurocages for Action Potential Stimulation and Recording*. PhD thesis, Caltech, 2007.

[64] A Tooker, J Erickson, Y-C Tai, and J Pine. Robust and biocompatible neurocages. In *Proc. MicroTAS*, Malmo, Sweden, 2004.

[65] A Tooker, E Meng, J Erickson, Y-C Tai, and J Pine. Development of biocompatible parylene neurocages. In *Proc. IEEE-EMBS, San Francisco, CA, USA*, pages 2542–2545, September 2004.

[66] R W Turner, D E Meyers, T L Richardson, and J L Barker. The site for initiation of action potential discharge over the somatodendritic axis of rat hippocampal CA1 pyramidal neurons. *J. Neurosci.*, 11:2270–2280, Jul 1991.

[67] C Vicario-Abejon, C Collin, R D McKay, and M Segal. Neurotrophins induce formation of functional excitatory and inhibitory synapses between cultured hippocampal neurons. *J. Neurosci.*, 18:7256–7271, Sep 1998.

[68] D A Wagenaar, J Pine, and S M Potter. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J Neurosci Methods*, 138:27–37, Sep 2004.

[69] D A Wagenaar, J Pine, and S M Potter. An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci*, 7:11, 2006.

[70] D A Wagenaar and S M Potter. Real-time multi-channel stimulus artifact suppression by local curve fitting. *J Neurosci Methods*, 120(2):113–120, Oct 2002.

[71] J D Wiley and J G Webster. Analysis and control of the current distribution under circular dispersive electrodes. *IEEE Trans Biomed Eng*, 29(5):381–385, May 1982.

[72] C Wyart, C Ybert, L Bourdieu, C Herr, C Prinz, and D Chatenay. Constrained synaptic connectivity in functional mammalian neuronal networks grown on patterned surfaces. *J Neurosci Methods*, 117(2):123–131, Jun 2002.

[73] G Zeck and P Fromherz. Noninvasive neuroelectric interfacing with synaptically connected snail neurons immobilized on a semiconductor chip. *PNAS*, 98(18):10457–62, 2001.

[74] L Zhang and M M Poo. Electrical activity and development of neural circuits. *Nat. Neurosci. Suppl.*, 4:1207–1214, 2001.