

Development and Control of Epileptiform Bursting in Dissociated Cortical Cultures

Thesis by
Daniel A. Wagenaar

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2006

(Defended June 21, 2005)

© 2006

Daniel A. Wagenaar

All rights Reserved

Acknowledgments

The end of my years as a graduate student only weeks away, I am reflecting on how many different people have contributed to the completion of my thesis. Let me begin by thanking Jerry Pine and Steve Potter, the best advisers I could have wished for.

From the moment I first stepped into Jerry's lab, when I visited Caltech as a prospective, I knew it would be a good place to work. A great can-do atmosphere, dedicated people excited about their work, and remarkably little stress in the air. Jerry has been a wonderful adviser. He gave me tremendous freedom to choose my own directions, yet was always available when I needed help. He was always completely uninhibited in both his criticism and his praise, and while it hurt when he dismissed a draft for a paper as utterly obscure, he was usually right. Armed with his many suggestions, I would then rewrite the thing, until he gave it his sign of approval, "it reads like a novel."

Jerry always demanded absolute scientific integrity. If you tried to make a claim based on dubious data, he'd say, "that is nonsense, and you know it." And again, he was usually right. He held strong opinions, which he could state with enough force to blow over a tree, yet he was always willing to change his mind if presented with a solid argument. And without his unshakable confidence in my ability to make my project somehow work out, I do not know how I would have made it through the years in the middle, when my own confidence was sometimes rather low.

Without Steve's infectious enthusiasm for his *Animat* project, I might very well never have made the leap from physics to neuroscience. Accordingly, it was a great blow when he left Caltech to start his own lab at Georgia Tech, but fortunately he agreed to remain my mentor. It was through his generosity that I was able to attend many conferences in the US and abroad, unrivaled opportunities to expand my view of the larger field of neuroscience and to meet its practitioners. Over the years, our

student–mentor relationship has grown into a friendship, and I have greatly enjoyed several visits to Atlanta as his house guest. His broad interests in neuroscience and beyond, his perennial sense of wonder at the world, and his inexhaustible optimism have been continuing sources of inspiration.

Next, I'd like to thank my other committee members, Michael Cross, Gilles Laurent, and Pietro Perona. I hope that their insistence on quality during my qualifying exam, and their advice afterward, have made this a better thesis. Though we met only a few times, I have greatly benefitted from Scott Fraser's insights. I thank Ken Libbrecht, the executive officer for the physics program, for happily signing off on a thesis topic that has only the most tenuous connection with physics. I would also like to thank Henri Lester, for sharing space and equipment with us. Special thanks are due to Carole Worra for helping me with purchase orders and reimbursements, even when we at the Pine lab somehow never quite managed to do such things the way we were supposed to. I owe a great debt of gratitude to Sheri McKinney, who provided me with cortical tissue all these years. With her invariably high spirits, she was a joy to be around.

I thank the Burroughs-Wellcome Fund for funding during my second year. Other funding came from the National Institute for Neural Disease and Stroke (through grants NS38628 and NS44134), the National Institute for Biomedical Imaging and Bioengineering (through grant EB00786), the Whitaker Foundation and the NSF Center for Behavioral Neuroscience at Georgia Tech and Emory University. Much support came from Caltech's Physics Department in the form of Teaching Assistantships. Teaching freshman students lab physics was great, and I particularly enjoyed assisting Chris Adami with his *Artificial Life* courses.

For almost a year after Steve left, I was the only student in the lab, which really taught me to appreciate colleagues. Tom DeMarse and Axel Blau, and later Gary Chow and Jon Erickson were wonderful company. I'll particularly miss my lunch-hour trips with Jon to Ernie, the taco man—who, by the way, also deserves thanks.

I had colleagues not only here, but also at Georgia Tech. Though I met them in person only once or twice per year, we met on the phone in weekly video conferences. Radhika, Peter, Doug, Komal and Mark: thanks for accommodating me in my hard-of-hearing telepresence.

When you're in the middle of it, and buried in work, you don't always realize it, but Caltech is an exceptional community. My first experience with this was the evening I first arrived from the Netherlands. The airport shuttle had dropped me off on the far side of campus, and I was very jet-lagged and rather lost. A Caltech security officer drove me to my house, and even offered to carry my bags up the stairs. I don't remember his name, but I'm unlikely to ever forget how good it felt to be welcomed that way. And such a contrast to my first arrival in London, where the first thing I heard was "you can't park there."

This sense of being welcome was intensified over the next week, during International Student Orientation. Parandeh Kia and Tara Tram did an amazing job at creating a community for the hundreds of foreign students who arrive here. Later, Parandeh was a great source of support in difficult times, and I felt deeply guilty by association when the Caltech administration decided to let her go after 11 years of service. I have the greatest respect for Jim Endrizzi and Athena Trentin who picked up the pieces at the International Student Programs office afterward.

It was through ISP that I met most of the friends that made my first year at Caltech such a great time. Miro, Wilfred, Tammy and Andrej, Stefan, Kalo, Virginia, and of course Shao Yi. They joined me on many great hikes, that somehow always ended up being longer than I had advertised; yet they never complained. The party that Andrej organized for my birthday that year was a wonderful surprise.

For the past six years, I haven't seen nearly enough of my family and friends in the Netherlands. I know that my parents, Lodewijk and Erin, my siblings, Ellinor, Vera and Arthur, and my friend Steven have missed me. I missed them too, and though I often wished the distance between Amsterdam and Pasadena was less, their special welcome whenever I returned home for all too brief vacations was wonderful. I thank them for understanding my absence.

Finally, I would like to thank Jennifer. Ever since we first met, a year and two days ago today, she has been a source of light and happiness in my life. Without her, these last few months of finishing my thesis and finding a new job would probably have been a time of much tension. With her by my side, they were very enjoyable.

Pasadena, June 6th, 2005.

Abstract

Cortical cells in dissociated culture form densely interconnected networks. Within days after plating, neurons become electrically active, and soon after start to synchronize their activity into culture-wide bursts. By growing cultures on multi-electrode arrays (Petri dishes with a grid of substrate-embedded electrodes), their electrical activity can be recorded non-invasively.

I developed software, MEABench, for online visualization and analysis of multi-electrode data, and used it to follow the development of cultures obtained from (E18) embryonic rats. Globally synchronized bursting was observed in all but the most sparsely plated cultures. A remarkable range of bursting behaviors was observed, even in cultures with identical plating parameters. Activity patterns varied widely in terms of the frequency, intensity, duration, and degree of temporal clustering of bursts. During the 2nd week *in vitro*, bursts in many dense cultures clustered into well-defined trains, separated by long periods without bursts. The number of bursts within these ‘superbursts’ and their spatiotemporal structure were found to be stable for hours or days.

Cortical cultures on multi-electrode arrays are ideal for studying two-way communication between biological systems and computers. I designed and built hardware to deliver electrical stimuli in arbitrary patterns, developed software to remove stimulation artifacts from recordings, and studied the efficacy of several voltage-defined and current-defined stimulus waveforms. MEABench can control the stimulator in real-time. Thus, stimuli can be made dependent on a culture’s activity with only 15 ms lag-time.

We hypothesized that synchronized bursting can dominate activity patterns, because lack of external input puts cultures in a hypersensitive state. Indeed, by feeding cultures a steady stream of stimuli, distributed over many electrodes, bursting could be prevented completely. The number of electrodes required for successful burst control could be reduced by fine-tuning the stimuli with real-time feedback, to make each stimulus evoke the same number of spikes. Burst control could not be achieved with single electrode stimulation.

For the final chapter, I tested various protocols for inducing plasticity by tetanic stimulation. In contrast to earlier published reports, I found that none of them induced changes in burst patterns or responses to test pulses that exceeded spontaneously occurring changes.

Short Contents

Acknowledgments	iii
Abstract	vii
Table of Contents	xi
List of Figures	xvii
Lists of Tables and Boxes	xxi
1 Introduction	1
2 Effective Parameters for Stimulation	7
3 Real-Time Artifact Suppression	27
4 Development of Spontaneous Bursting Activity	41
5 Superbursts	63
6 Controlling Bursting with Multi-Electrode Stimulation	75
7 Plasticity, or the Lack Thereof	97
A Culturing Methods	129
B MEABench: A Toolset for Data Acquisition and Online Analysis	137
C An All-Channel Stimulator with Real-Time Control	149
D MEABench: User Guide	163
Bibliography	193

Table of Contents

Acknowledgments	iii
Abstract	vii
Table of Contents	xi
List of Figures	xvii
Lists of Tables and Boxes	xxi
1 Introduction	1
2 Effective Parameters for Stimulation	7
Introduction	7
Methods	9
Cell Culture	9
Pharmacology	9
Electrodes and Recording System	10
Stimulation System	10
Results	13
Current Pulses	14
Voltage Pulses	16
Number and Reliability of Responding Cells	20
Voltage vs. Current Control: Electrode and Leakage Currents	21
Discussion	22
Practical Consequences	23
Appendix: Identifying Response Components	24

3	Real-Time Artifact Suppression	27
	Introduction	28
	Methods	30
	Cell Culture and MEA Preparation	30
	Artifact Suppression	30
	Analysis	34
	Results	34
	Discussion	38
4	Development of Spontaneous Bursting Activity	41
	Introduction	42
	Methods	43
	Cell Culture	43
	Recording	44
	Data Analysis	44
	Results	45
	Burst Classification	49
	Development of Firing Rates	52
	Probe Responses	54
	Axon Outgrowth	54
	Development of Burst Duration, Propagation Speed, and Size	55
	Sensitivity to Movement	57
	Sources of Variability	59
	Discussion	61
	The Nature of Bursts in Culture	61
	Importance of Sampling from Multiple Platings	62
	Outlook	62
5	Superbursts	63
	Introduction	63
	Methods	64
	Spike Detection and Sorting	65
	Burst Identification	65

Similarity Indices	66
Results	67
Discussion	73
6 Controlling Bursting with Multi-Electrode Stimulation	75
Introduction	76
Methods	78
Cell Culture	78
Recording System	79
Stimulation System	80
Experimental Protocols	80
Quantifying the Level of Bursting	81
Tuning the Feedback	82
Immunostaining	83
Results	84
Spontaneous Bursting	84
Response to Stimulation	84
Burst Control by Distributed Stimulation	87
Burst Control by Closed-Loop Stimulation	91
Discussion	93
7 Plasticity, or the Lack Thereof	97
Introduction	97
General Methods	99
Overview of Results	101
Sanity Checks	104
Series I: Changes Induced in Burst Patterns	106
Protocols	106
Choice of Electrodes	106
Analysis	107
Results	109
Series II: Changes Induced in Stimulus–Response Maps	111
Protocols	111

Analysis	113
Results	114
Series III: Changes in Specific Responses	117
Introduction	117
Protocols	118
Selecting Stimulation Pairs	118
Results	120
Changes in the Probability of Evoking Bursts	125
Discussion	126
A Culturing Methods	129
Tissue Extraction	129
MEA Preparation	129
Cell Preparation	130
Plating	133
Maintenance	134
Recording	134
B MEABench: A Toolset for Data Acquisition and Online Analysis	137
Introduction	137
Methods	138
Software Engineering	138
Algorithms	140
Application	143
Results	144
Programs	144
Artifact Suppression	146
Spike Detection	146
Real-Time Operation and Stimulator Control	147
Discussion	147
C An All-Channel Stimulator with Real-Time Control	149
Introduction	150

Background	150
Design Philosophy	151
Methods	151
Cell Culture	151
Recording System	152
Stimulation System	152
Results	155
The Real-Time All-Channel Stimulator	155
Benchmarks	156
Discussion	160
D MEABench: User Guide	163
Introduction	163
Conventions	165
Compilation and Installation	165
Examples	166
Displaying Electrode Traces Online	166
Online Spike Detection	168
Triggered Recording	171
Using Commander	172
Basics of the Toolset	173
List of Components	174
Details of Each Component	175
Rawsrv	175
Spikedet	176
Filter60hz	178
Salpa	179
Record	181
Replay	182
Scope	182
Spikesound	183
Flexraster	183

Commander	184
Monitor	185
Neurosock and NSsrv	185
Spikedump	185
Doubletxt	185
Other Utilities	186
Matlab Functions	186
File Formats	188
Raw Files	188
Spike Files	189
Hints and Tips	189
None of the Programs Will Run	189
Client X Keeps Saying ‘Waiting for START from Y’ and Won’t Run . .	190
My Recordings Are Truncated or Empty	190
Shared Memory Problems	190
Abbreviating Commands	190
Passing Commands at Run Time	191
Interrupting Long Commands	191
Debugging Information	191
Contacting the Author	192
Reporting Bugs	192
Bibliography	193

List of Figures

1.1	Jerry Pine's <i>Neurochip</i>	2
1.2	Steve Potter's <i>Neuronally Controlled Animat</i>	3
2.1	Stimulation hardware	11
2.2	Array-wide responses to biphasic voltage pulses on a single electrode during first 30 ms after stimulation	15
2.3	Firing rate on one electrode as a function of latency after stimulus in control medium and with synapse blockers	16
2.4	Waveforms for current- and voltage-controlled stimulation	17
2.5	Responses to current stimulation	18
2.6	Responses to voltage stimulation	19
2.7	Reliability of direct responses	21
2.8	Comparison of the efficacy of voltage- and current-controlled stimulation waveforms	22
2.9	Classifying response components by FSMEM	24
3.1	Illustration of the fitting method used by SALPA	31
3.2	An example of an artifact in electrode recordings, with the output of SALPA and various other filters	35
3.3	Compromise between loss of recording time due to artifacts and reduction of spike peak-to-noise ratio due to artifact suppression	36
3.4	Comparison of various filter methods for artifact suppression	37
3.5	Examples of SALPA results on different artifact waveforms	37
3.6	Illustration of recordable responses with and without artifact suppression	39
4.1	Development of burst patterns in a dense culture	46

4.2	Phase contrast micrographs of central area of typical cultures of various densities at 1, 15, and 32 days <i>in vitro</i>	48
4.3	Classification of observed bursting behaviors	50
4.4	Development of firing and bursting activity in dense cultures	52
4.5	Comparison of the development of cultures of different sizes	53
4.6	Strength of responses to stimulation during the development of dense cultures	54
4.7	Measuring axonal outgrowth in cultures of different densities	55
4.8	Characterizing burst shapes in dense cultures	56
4.9	Comparison of burst sizes during culture development	57
4.10	Sensitivity of spiking activity to mild mechanical perturbation	58
4.11	Comparison of activity after moving a culture into the recording device with activity around 12 hours later	59
4.12	Quantifying sources of variability in activity levels	60
4.13	Distribution of inter-burst intervals and burst sizes	61
5.1	Development of bursting	67
5.2	Example of a superburst, and basic superburst parameters	68
5.3	Conservation of firing rates and activity propagation between superbursts	70
5.4	Return plots representing the temporal structure of burst propagation	72
5.5	Spatiotemporal structure of superbursts	73
6.1	A typical culture growing on an MEA at 2 and 34 div	77
6.2	Two-photon images of MAP2 and GABA immunocytochemically stained neurons	79
6.3	Performance of feedback protocol	82
6.4	Examples of different spontaneous bursting patterns	85
6.5	Array-wide responses to stimulation	86
6.6	Burstiness during single electrode stimulation, multi-electrode stimulation and spontaneous activity	88
6.7	Loss of responsivity to rapid stimulation on a single electrode	90

6.8	Rapid change of a culture's activity patterns after switching between stimulation protocols	90
6.9	Burstiness during closed-loop control of tonic firing rate	92
6.10	Assessment of the success rate of different burst suppression protocols	93
7.1	Reactions to pharmacological manipulations	103
7.2	Confirmation of the elasticity results of Eytan et al. (2003)	105
7.3	Responses to probing all electrodes at 100–900 mV	108
7.4	Examples of changes induced in spontaneous bursting	109
7.5	Results of significance tests for experiments on plasticity expressed in burst patterns	110
7.6	Examples of results of experiments on plasticity expressed in stimulus–response maps	114
7.7	Summary of results of experiments on plasticity expressed in stimulus–response maps	115
7.8	Significance tests for experiments on plasticity expressed in stimulus–response maps	115
7.9	Array-wide responses to individual stimuli in an experiment on plasticity induced in stimulus–response maps	116
7.10	Results of experiments on stimulus–response maps probed at different voltages	117
7.11	Response maps for several pairs of stimulation electrodes	120
7.12	Examples of changes induced in specific responses	121
7.13	Results of significance tests for plasticity in spike counts of specific responses	122
7.14	Results of significance tests for plasticity in latency-to-first-spike of specific responses	123
7.15	Small differences in the magnitudes of tetanus-coincident vs. spontaneous changes in experiments on plasticity in specific responses	124
7.16	Direct comparison of changes coincident with tetani and spontaneous changes in spike counts of specific responses	125
7.17	Changes in the probability that a test pulse evoked a burst	126

A.1	Peltier-cooled stage for recording in an incubator	135
B.1	Spike validation based on waveform shape	142
B.2	Screenshot of MEABench in operation	145
B.3	Tradeoff between spike detector efficiency and rate of false positive de- tection	146
C.1	Circuit diagrams of stimulator	153
C.2	Photographs of the main stimulator board and three modules connected to a MultiChannel Systems pre-amplifier with MEA in the center . . .	154
C.3	Pulse shapes	156
C.4	Neuronal activity recorded in response to stimulation on different elec- trodes	157
C.5	Stimulation artifacts for positive-then-negative biphasic stimuli	159
C.6	Histogram of loop times for delivering stimuli in response to recorded action potentials	160
C.7	A voltage-to-current convertor to adapt the RACS for current-controlled stimulation	161

Lists of Tables and Boxes

Tables

2.1	Summary of current-controlled pulses tested	12
2.2	Summary of voltage-controlled pulses tested	13
4.1	Plating parameters	49
5.1	Compendium of parameters for all five extended recordings	69
7.1	Details of experiments on plasticity expressed in burst patterns	107
7.2	Details of experiments on plasticity expressed in stimulus–response maps	112
7.3	Details of experiments on plasticity induced in specific responses . . .	119

Boxes

7.1	Overview of plasticity-inducing stimuli used by other researchers . . .	100
A.2	Vendor addresses	130
A.3	Poly-ethylene-imine solution (PEI)	130
A.4	Dulbecco’s modification of Eagle’s basal medium	131
A.5	Segal’s medium	132
A.6	Papain solution	132
A.7	Composition of culture media	133

1

Introduction

The human brain is arguably the most complex system known to science. From a structural point of view, it consists of 10^{11} cells, each with intricate internal structure, and some 10^4 connections to selected other cells. From a functional point of view, it has a capacity for a practically boundless range of highly complex behaviors—from directing fine hand movements while playing music, to recognizing the work of a favorite artist in a painting you have never seen before. Scientists have barely begun unravelling how the brain performs such behaviors, let alone how it supports the emergence of mind, which, from its own perspective, so conspicuously underlies many of the brain's most interesting functions.

The unparalleled complexity that makes the human brain such an attractive object of study also makes it extremely difficult to approach, so scientists have invented a variety of ways to reduce this complexity to manageable proportions. On one end of the spectrum, they have focused their attention on simpler brains, such as those of mollusks or insects, because they should be easier to understand. On the other end, they have used functional imaging techniques such as fMRI and PET to discover which parts of the human brain are involved in certain behaviors, if not how those behaviors are implemented at a neuronal level.

Our labs have taken a different approach. Stepping away from the workings of




Figure 1.1: Jerry Pine's *Neurochip*. **A** Scanning electro-micrograph of a single cage. Scale bar: 10 μm . **B** Nomarski micrograph of an array of cages, with several trapped neurons (arrows) in the process of forming a network. Scale: distance between cages is 110 μm . Photos courtesy Jon Erickson.

any particular brain, we have focused on the functional networks formed by mammalian cortical neurons in culture. Jerry Pine's *Neurochip* (Maher et al., 1999b; Erickson et al., 2005) traps individual neurons in an array of cages with embedded electrodes, to facilitate the study of the formation of connections in very small networks (Figure 1.1). For Steve Potter's *Neuronally Controlled Animat* project (DeMarse et al., 2001), larger cultures are grown on multi-electrode arrays (MEAs): Petri dishes with grids of flat electrodes embedded in the substrate. These cultures are made to serve as the brains of artificial animals ('animats') in computer simulated environments; more recently, robots have been used instead of computer simulations (Bakkum et al., 2004). A culture's electrical activity, recorded using the electrode array, is used to derive motor commands for the animat, while sensory information the animat obtains from its environment is fed back to the neuronal culture in the form of electrical stimulation (Figure 1.2). The intention is that, thus *re-embodied*, neuronal cultures become an easily accessible brain-like input-output system, in which biologically important behaviors such as memory formation can be studied at the network level.

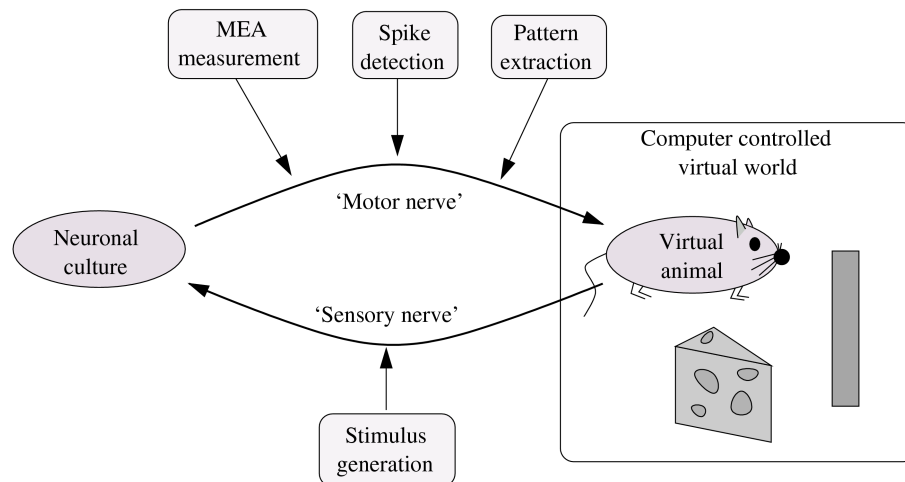


Figure 1.2: Steve Potter's *Neuronally Controlled Animat*. Drawing by the author.

For such a project to be feasible, several different technologies have to work together in a coordinated fashion: besides a means of keeping a culture alive and in contact with an array of electrodes, hardware is required to deliver stimuli to those electrodes, as well as to measure neuronal signals. Software is needed as well: drivers to control stimulation and measurement hardware, visualization tools for online inspection of neuronal activity, algorithms to detect action potentials, and code to interface all these different components. When I first joined the lab, only the culture system, hardware to record electrical activity from 60 electrodes, and a very basic stimulator were in place. My first task, then, was to create software that could make recorded neuronal activity accessible for online inspection, and for real-time processing for animat control. Software for recording electrode traces and threshold-detected spikes to disk already existed, but consisted of a large application running under Microsoft Windows. Adapting this software for real-time use as required for the *Animat* project was not practical, so I started afresh with a modular application framework under Linux. This eventually led to MEABench, a software suite for acquisition and online analysis of multi-electrode data with facilities for real-time data feed-through to external programs. MEABench is now in daily use by researchers in several labs.

Successful sensory feedback requires hardware to send electrical stimuli to a large number of electrodes independently, preferably without interfering with those electrodes' ability to record neuronal activity. We needed to learn how to effectively

use substrate electrodes to elicit neuronal responses, without damaging electrodes or neurons by electrochemistry. Extracellular stimulation involves currents several orders of magnitude larger than the electrode currents generated by action potentials, so artifacts induced in recorded signals as a result of stimulation were an important issue to be considered during hardware design, as well as for software signal processing.

Since the animat would be controlled by cultures' activity patterns, a second major task was to study those activity patterns. We discovered—in parallel with other research groups around the world—that the activity of dense cortical cultures growing in isolation is dominated by synchronized bursting events, in which all or most electrically active cells participate. During most of a culture's development, these bursts are largely unstructured, and hence not a good candidate signal to drive animat motor output. Worse, since they are synchronized across the entire culture and entail such a violent increase in firing activity, they likely disrupt more subtle, more information-rich, local processing; much like seizures do in an epileptic brain.

Preventing global bursts thus became a prerequisite for creating animats with interesting behavior. We hypothesized that the reason why cultures burst could be that they receive so little external input. Indeed, neurons are known to regulate their sensitivity to synaptic inputs to homeostatically control their firing rates. Thus, in the absence of external input, a neuronal network could end up in a state where the merest fluctuation can trigger an avalanche of activity—in other words, a culture-wide burst.

We reasoned that by continuously stimulating a culture electrically, thereby increasing its baseline firing rate above spontaneous levels, we could reduce its tendency to burst. Indeed, we found that while stimuli with long inter-stimulus intervals typically entrained bursts, high-frequency stimulation increased baseline activity levels and prevented bursting. Rapid stimulation through a single electrode ceased to be effective after about a minute, presumably because the axons in which the stimuli should induce action potentials suffered from ionic depletion. By distributing the stimuli across many electrodes—so that each individual electrode was stimulated at a relatively slow rate—we could maintain efficacy indefinitely.

The rest of this thesis is organized as follows.

Chapter 2, *Effective parameters for stimulation*, discusses a variety of voltage-defined and current-defined pulse shapes, and their efficacies at eliciting neuronal responses. One particular voltage-defined waveform shape is derived which is most effective; this waveform, scaled to different amplitudes, is used in the rest of the thesis. The chapter includes a discussion of how electrical stimuli delivered through extracellular electrodes evoke neuronal responses.

Chapter 3, *Real-time artifact suppression*, describes a software algorithm to remove stimulation artifacts from recorded electrode traces online. Such software is needed, because the difference in scale between recordable signals due to action potentials and pulses that are effective as stimuli makes it very difficult to prevent such artifacts entirely in hardware. Unfortunately, artifacts vary from array to array, and from electrode to electrode, and even depend on the recent stimulation history. The algorithm therefore dynamically models each artifact individually. After artifact removal, action potentials can be detected by voltage-thresholding.

Chapter 4, *Development of spontaneous bursting activity*, is concerned with the natural development of activity patterns, and especially bursting, in cortical cultures of various densities. Investigation of 963 recordings from 58 cultures reveals a remarkably rich spectrum of bursting behaviors. It appears that cultures have a substantial degree of individual ‘personality.’

Chapter 5, *Superbursts*, describes a particularly interesting developmental phase, in which burst patterns attain a high degree of temporal clustering. These burst clusters, or ‘superbursts,’ have a strongly conserved internal structure, both in terms of the number of constituent bursts, and in terms of the spatiotemporal propagation of those bursts.

In Chapter 6, *Controlling bursting with multi-electrode stimulation*, electrical stimulation is used to control bursting. The burst entrainment and short-lived control that can be achieved with single electrode stimulation is described first, then a multi-electrode stimulation protocol that successfully and indefinitely controls bursting. If twenty or more electrodes are used, bursting can be prevented completely. With only ten electrodes, burst suppression with fixed, pre-defined pulse amplitudes is not entirely satisfactory, but a substantial improvement can be attained by fine-tuning the amplitudes of individual stimuli using real-time feedback.

Finally, Chapter 7, *Plasticity, or the lack thereof*, discusses attempts to replicate a number of the plasticity results previously described in the literature. With few exceptions, I obtained negative results, despite 112 experiments on 18 cultures with 12 different protocols. Changes in responses to test stimuli do occur during tetanization (a common technique to induce plasticity), but are no larger than changes that occur spontaneously. The chapter closes with a discussion of the discrepancies between our results and previous reports.

After the main text, four appendices describe methodological and technological details. Appendix A, *Culturing methods*, details the procedures for preparing and maintaining cell cultures, and includes recipes for the various media we use. Appendix B, *MEABench: a toolset for data acquisition and online analysis*, provides an overview of MEABench, the software package I developed for the data acquisition, visualization, and online analysis required for this thesis work. Appendix C, *An all-channel stimulator with real-time control*, describes the design of our multi-channel stimulator. The complete user's guide to the most recent version of MEABench is included as Appendix D.

Portions of this thesis have been published before. Chapters that are based on published articles appear in this thesis unchanged except for typographic corrections and minor stylistic changes. Two other chapters (4 and 5) are based on manuscripts currently being finalized for submission. Details of previous publication are included with the relevant chapters.

2

Effective Parameters for Stimulation*

Electrical stimulation through multi-electrode arrays is used to evoke activity in dissociated cultures of cortical neurons. We study the efficacies of a variety of pulse shapes under voltage as well as current control, and determine useful parameter ranges that optimize efficacy while preventing damage through electrochemistry. For any pulse shape, stimulation is found to be mediated by negative currents. We find that positive-then-negative biphasic voltage-controlled pulses are more effective than any of the other pulse shapes tested, when compared at the same peak voltage. These results suggest that voltage control, with its inherent control over limiting electrochemistry, may be advantageous in a wide variety of stimulation scenarios, possibly extending to *in vivo* experiments.

Introduction

Multi-electrode arrays (MEAs) (Thomas et al., 1972; Gross, 1979; Pine, 1980; Potter, 2001; Heuschkel et al., 2002) have been used to record from a wide variety of neuronal preparations. Electrical stimulation through MEAs has been used to elicit spiking activity in dissociated cultures (Regehr et al., 1988; Jimbo and Kawana, 1992; Gross et al., 1993; Maher et al., 1999b; Jimbo et al., 1999), as well as brain

* Published as: Daniel A. Wagenaar, Jerome Pine, and Steve M. Potter, 2004: Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J. Neurosci. Methods* **138**(1–2), pp. 27–37. © 2004 Elsevier. Reprinted with permission.

slices (Heck, 1995; Egert et al., 1998; Novak and Wheeler, 1988; Echevarria and Albus, 2000; Harsch and Robinson, 2000; Tschertter et al., 2001) and isolated retina (Branner and Normann, 2000; Grumet et al., 2000). MEAs and related technology for multisite extracellular stimulation and recording, such as silicon probes (Wise and Angell, 1975; Bai and Wise, 2001) and multiwire probes (Nicolelis et al., 1998) have gained popularity because the relatively non-invasive nature of the technology allows for long-term interaction with healthy cells, and because they scale well to a large number of recording and stimulation channels.

When designing stimulation paradigms, researchers have to make many choices even after they have decided on electrode material: whether to use bipolar stimuli (between two electrodes) or monopolar stimuli (between one electrode and a large and usually distant ground electrode), whether to use voltage or current control, what pulse shape to use (monophasic, biphasic, perhaps even multiphasic or asymmetric). Compromises have to be found between efficacy of stimuli, harm to electrodes or cells, and stimulation artifacts that hamper recording of responses. For long-term experiments, it is crucial to prevent damaging electrodes and killing cells. Cell damage can result from high charge injection or high charge densities (McCreery et al., 1990), but our MEA electrodes cannot inject dangerous amounts of charge before exceeding electrolysis limits (Weiland et al., 2002). Electrolysis, which starts to play a role when electrode voltages exceed about one volt, directly damages electrodes, and is also harmful to cells. (This harm can be much reduced by employing charge-balanced stimuli, making such stimuli preferable for long-term or *in vivo* work when large voltages cannot be avoided (Lilly et al., 1955; Shepherd et al., 1991).) A secondary constraint can be the width of stimulus pulses: since recording is generally impossible for the duration of the stimulus pulse, short pulses are often desirable.

Here, we study electrical stimuli intended to evoke activity in dissociated cortical cultures on MEAs, with the aim of establishing robust two-way communication between such cultures and a computer system. Knowing a set of stimuli that are reliably and consistently effective is essential for long-term experiments on the development of functional networks, as well as for research on learning *in vitro* (DeMarse et al., 2001; Shahaf and Marom, 2001). While the electrical properties of MEA electrodes have been described in the literature (Kovacs, 1994; McAdams et al., 1995;

Buitenweg et al., 1998; McIntyre and Grill, 2001), the published knowledge on what kinds of stimuli are most effective at evoking responses is remarkably slim. A full quantitative understanding would require a detailed model of the electric fields that current pulses induce along axonal and somatic membranes, but in high-density cultures, the arrangement of neurons and glia is too complex to construct such a model. In this paper we hope to provide new practical information by identifying a range of stimuli that are effective, unharmed, and produce minimal artifacts. The results in this paper were obtained from dense neocortical cultures grown on MEAs with 30 μm titanium nitride electrodes. Qualitatively, the results should extend to other dissociated neuronal cultures, and to electrodes of different sizes and construction.

Methods

Cell Culture

Neocortex was dissected from rat embryos (E18) under sterile conditions. Cortices were cut into pieces of about 1 mm³, prior to dissociation using papain and trituration. Cells were plated at 5000 cells/mm² on multi-electrode arrays (MultiChannel Systems, Reutlingen, Germany) coated with poly-ethylene-imine (PEI) and laminin. Cultures were maintained for 2–3 weeks prior to recording, in a medium adapted from Jimbo et al. (1999): High glucose DMEM (Irvine Scientific cat. no. 9024) with 10% Horse Serum (HyClone), 0.5 mM GlutaMax (Gibco cat. no. 5050.061), 1 mM sodium pyruvate, and 2.5 $\mu\text{g}/\text{mL}$ insulin. To prevent evaporation and to stop infection, culture dishes were sealed using Teflon membranes (Potter and DeMarse, 2001). Cultures were kept in an incubator at 35 °C, 65% R.H., 5% CO₂, and 9% O₂. All experiments were performed inside this incubator, guaranteeing stability of environmental conditions. (At 65% relative humidity, our incubator is safe for electronics.)

Pharmacology

For blocking glutamatergic synaptic transmission, we replaced the culture medium by a medium pre-mixed with 800 μM D(-)-2-Amino-5-phosphonopentanoic acid (AP5, Sigma A-169), an NMDA-channel blocker, and 80 μM 6-cyano-7-nitroquinoxaline-2,3-

dione disodium salt (CNQX, Sigma C-239), an AMPA-channel blocker. For blocking all activity, we used a medium pre-mixed with 1 μ M tetrodotoxin (TTX, Sigma T-5651). Prior to either kind of blocker experiment, we obtained a baseline by replacing all of the culture medium with 1 mL of fresh medium and recording responses to voltage stimulation after allowing the culture to equilibrate for 30 minutes. We then prepared 1 mL of medium with either AP5 and CNQX or with TTX, and replaced the culture medium with this mixture, again recording after 30 minutes of equilibration. After that, we washed out the blockers by four full medium replacements spread out over five minutes, and recorded a final run through the stimulus set to confirm return to baseline conditions.

Electrodes and Recording System

We use glass MEAs with 30 μ m diameter titanium nitride electrodes and a silicon nitride insulation layer (MultiChannel Systems, Reutlingen, Germany). The electrodes have a rough surface which increases their capacitance for a given electrode diameter. Each MEA has 59 such electrodes laid out in a rectangular grid with 200 μ m inter-electrode spacing. Signals were amplified using an MEA60 pre-amplifier, and digitized using an MC_Card PCI board (both MultiChannel Systems). Data acquisition was controlled using MEABench software (D. A. Wagenaar, <http://www.its.caltech.edu/~wagenaar/meabench>). MEABench was also used for online data visualization and processing: stimulation artifacts were suppressed online using the SALPA algorithm (Wagenaar and Potter, 2002 [Chapter 3 in this thesis]), after which spikes were detected by thresholding at $5\times$ RMS noise. Candidate spikes were validated using a simple shape-based criterion (P. P. Mitra, personal communication).

Stimulation System

Current Pulses

We produced current-controlled stimuli by passing the output of a PCI-6216 digital-to-analog convertor (DAC) card (AdLink Technologies, Taiwan) through the voltage-to-current convertor depicted in Figure 2.1A (Horowitz and Hill, 1996). This circuit allows us to monitor the voltage generated by the current pulses. We tested both pos-

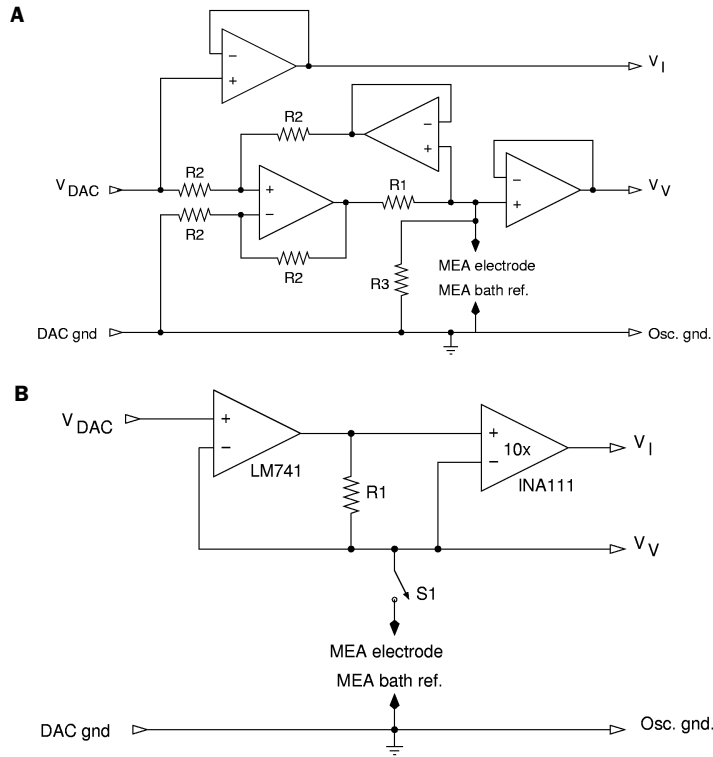


Figure 2.1: Stimulation hardware. **A** Circuit for current-controlled stimulation. V_{DAC} is converted to a current $I = V_{DAC}/R_1$ through the resistor $R_1 = 150 \text{ k}\Omega$. $R_3 = 6.8 \text{ M}\Omega$ acts as a shunt to prevent runaway electrode voltages. (Electrode impedances are 50–200 k Ω at 1 kHz). $R_2 = 150 \text{ k}\Omega$. Opamps are $\frac{1}{4}$ LM348. The terminals V_V and V_I are used to monitor the electrode voltage and the driving current (multiplied by R_1). **B** Circuit for voltage-controlled stimulation. The MEA electrode is driven at the voltage V_{DAC} . The voltage across $R_1 = 5.0 \text{ k}\Omega$, amplified and measured at V_I , directly reflects the current passing through the electrode. The electronic switch S_1 can be used to effectively disconnect the stimulator from the MEA, allowing neuronal signals to be recorded through the stimulated electrode shortly after stimulation. This circuit is part of a device for delivering stimuli to any electrode in an MEA with rapid switching between channels (Wagenaar and Potter, 2004).

itive (“anodic”) and negative (“cathodic”) monophasic currents of various strengths (1–20 μA) and duration (10–1000 μs) on 55 electrodes from five different MEAs. Additionally, we tested biphasic (positive-then-negative or negative-then-positive) current pulses on six electrodes from two MEAs. (The stimulus repertoire is summarized in Table 2.1.) Electrodes were randomly selected from the pool of electrodes that recorded activity during spontaneously occurring culture-wide barrages of action potentials. All stimuli were relative to a large electrode embedded in the MEA

Table 2.1: Summary of current-controlled pulses tested

Pulse shape	Amplitudes (μA)	Pulse widths (μs)
Monophasic, negative	1–10	20–1280
Monophasic, positive	1–10	20–1280
Biphasic, negative-first	5, 10	50, 200
Biphasic, positive-first	5, 10	50, 200

substrate which also served as a reference electrode for recording responses. Each pulse type was presented 50 times to each stimulation electrode, with one second between trials. Presentation order was randomized. Stimulation was fully automated using custom software controlling the hardware through real-time Linux. Responses to stimuli were measured using all non-stimulated electrodes, as described below.

Voltage Pulses

Voltage-controlled stimuli were generated using either one channel from the AdLink DAC used for current stimulation, or using a custom DAC based on a TLC-7628 (Texas Instruments). To monitor the current resulting from the stimulus, we used the circuit depicted in Figure 2.1B. We tested biphasic (positive-then-negative and negative-then-positive) pulses of various amplitudes (100–1000 mV) and duration (100–800 μs per phase) on 45 electrodes from five MEAs. Additionally, we tested monophasic (positive and negative) pulses on six electrodes from two MEAs. Finally, we tested “unterminated” pulses, that is, biphasic voltage pulses with the final phase not ended by driving the voltage back to zero but by opening the stimulation switch (‘S1’ in circuit diagram), thus cutting the driving current. (The stimulus repertoire is summarized in Table 2.2.) As before, a large substrate electrode doubled as return and reference.

Table 2.2: Summary of voltage-controlled pulses tested

Pulse shape	Amplitudes (mV)	Pulse widths (μ s)
Monophasic, negative	100–1000	400
Monophasic, positive	100–1000	400
Biphasic, negative-first	100–1000	100–900
Biphasic, positive-first	100–1000	100–900
‘Unterminated’ (see text)	500	100–400

Results

All electrodes tested could be used to evoke responses, given sufficiently strong stimuli (Figure 2.2). Responses to individual stimuli could be differentiated into three parts:

Direct responses Any response that does not depend on glutamatergic synapses. These occur in the first 10–20 ms post-stimulus, have less than 0.25 ms temporal jitter and can be close to 100% reliable (i.e., observed in close to 100% of trials). Direct responses most likely result from antidromic excitation through an axon near the stimulating electrode, though in principle they could result from a single axon passing by both the stimulation and the recording electrode, or from a cell with soma near the stimulation site and an axon passing the recording electrode. We did not study the potential role of gap junctions, but others have reported that it is minimal (Nakanishi and Kukita, 1998).

Early postsynaptic spikes Responses that do depend on glutamatergic synapses, occurring between 5 ms and 50 ms post-stimulus. Their temporal precision varies around 2 ms, and their reliability ranges up to 30%, though values around 10% are more common.

Culture-wide barrages In some fraction of trials, stimulation results in a barrage of activity spreading over the entire recording array. Depending on the electrode stimulated, such barrages can be initiated immediately after stimu-

lation, or at latencies exceeding 100 ms. In most cultures, barrages cannot be evoked more than once every 2–5 s.

For the present report, the first part is most important: the direct responses indicate how many cells are immediately stimulated by a given pulse. Using results from ten stimulation electrodes in two cultures, we found that blocking glutamatergic synapses inhibited all responses except for the very precisely timed ones (Figure 2.3). Thus, we found that direct responses are distinguishable from postsynaptic responses by their extremely reproducible timing: their temporal jitter is always less than 0.25 ms, compared to more than 1 ms for most postsynaptic responses. Based on this characteristic, we could distinguish direct responses even in medium without synapse blockers. We confirmed their biological origin by application of the sodium-channel blocker TTX, which reversibly abolished all spikes (data not shown).

Direct responses were observed even at electrodes distant from the stimulation site, but somewhat more frequently nearby the stimulus. Direct responses recorded at different electrodes were not correlated across trials of a given stimulus strength, indicating that those responses were due to different axons independently stimulated by the pulse, rather than due to a single axon passing by several recording electrodes.

Current Pulses

Many researchers prefer current pulses on theoretical grounds, because the electric field and potential near the electrode can be directly calculated (Maher et al., 1999a). Thus the effects of current pulse stimulation are thought to be better understood than the effects of voltage pulse stimulation, for which the (current-induced) electric field depends on the complex impedance of the stimulation electrode. In reality, a significant fraction of current applied to the electrode may be lost to leakage through the insulation of the electrode leads (see below, under *Voltage pulses*), reducing the reliability of such calculations.

For a typical new electrode with low leakage currents, 10 μ A current pulses of 50 μ s duration resulted in an electrode voltage of about 1 V (Figure 2.4A). Robust responses were observed after stimulating with negative currents (Figure 2.5). Positive pulses never resulted in any response. Increasing amplitude or duration of negative




Figure 2.2: Array-wide responses to biphasic voltage pulses on the marked electrode during first 30 ms after stimulation. Graphs are laid out following the physical geometry of the MEA—a rectangular array with 200 μm inter-electrode spacing. Each graph shows responses vs. time from one electrode. Within each graph, positive-then-negative stimuli are shown in the top half, with stronger stimuli towards the top, while negative-then-positive stimuli are shown in the bottom half, with stronger stimuli towards the bottom. Each dot represents a single spike. These data were accumulated over 50 trials at each of 36 voltage levels between -900 and $+900$ mV, with spikes from each trial in horizontal lines. Direct responses appear as very sharp vertical clusters, postsynaptic responses as larger clouds. This representation does not show the occasional occurrence of barrages.

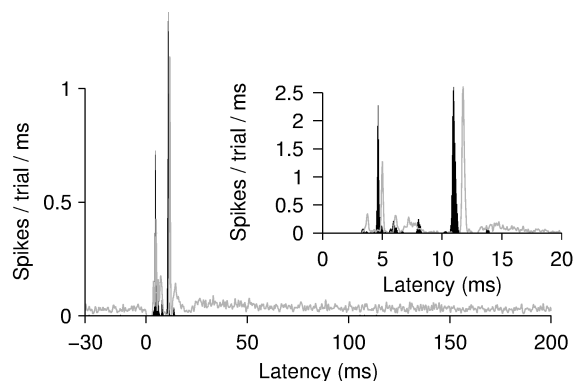


Figure 2.3: Firing rate on one electrode as a function of latency after stimulus. Baseline (*gray*) and in the presence of AP5 and CNQX (*black*). Inset shows detail at early latency. Data pooled from 475 trials. Synaptic blockers abolished all activity except for the precisely timed spikes ('direct responses'). Note the shift in latency of the direct responses. This is likely due to more sodium channels being ready to open when synapses are blocked—in baseline conditions, a certain fraction of sodium channels is always in the inactivated state due to ongoing spiking activity. The small peaks near 6 and 8 ms latency are due to events near the detection threshold.

currents improved the chance of eliciting responses, as expected. Very short (50 μ s) current pulses evoked as many spikes as long pulses with the same charge transfer, a surprising result, given that the time constant for the opening of sodium channels is about 0.8 ms (Sigworth and Neher, 1980). For long pulse widths (200 μ s), biphasic pulses were not significantly more or less effective than monophasic negative pulses (Figure 2.5A). For short pulse widths (50 μ s), biphasic pulses were less effective than monophasic negative pulses, presumably because the two phases partially cancel each other in the temporal integration performed by the target neuron's population of sodium channels (data not shown).

Voltage Pulses

Unfortunately, current-controlled stimulation has several practical drawbacks: good current sources are more complex than voltage sources, but more importantly, current pulses can cause high electrode voltages, which may harm cultures or damage electrodes through electrochemistry. For these reasons, most stimulation experiments on MEAs in our labs are actually performed under voltage control.

Positive-then-negative biphasic voltage stimuli were very effective at eliciting ac-

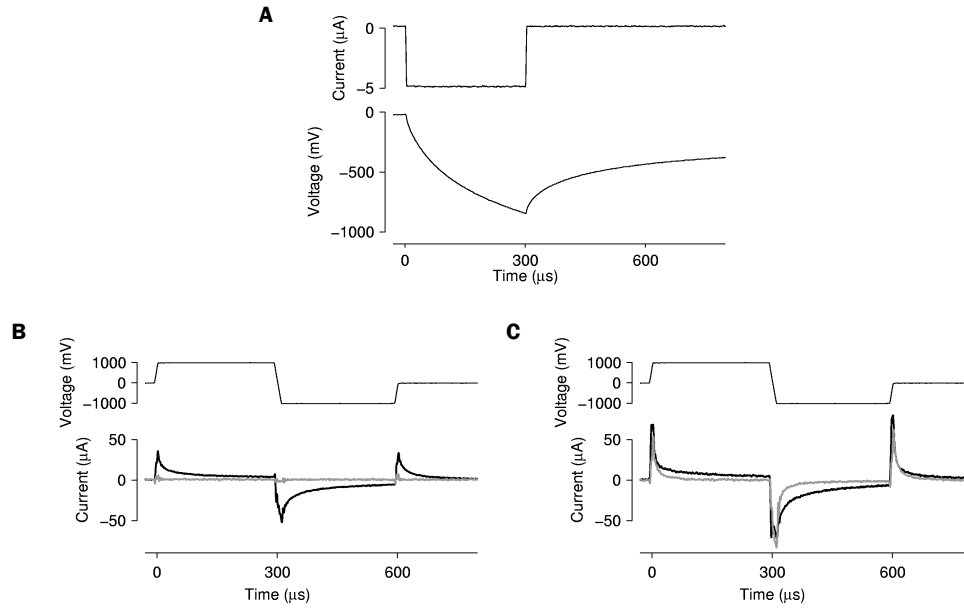


Figure 2.4: Stimulation waveforms. **A** Current-controlled stimulation: Driving current (*top*) and the resulting voltage (*bottom*). The discharge after cessation of driving current is through a shunt resistor (see Figure 2.1A). Waveforms were recorded using a Tektronix TDS2014 digital storage oscilloscope connected to the V_V and V_I terminals in Figure 2.1A. **B** Voltage-controlled stimulation: a typical driving voltage waveform (*top*) and the corresponding current response (*bottom*) for an electrode on an MEA that has been plated just once. We distinguish true electrode currents (*black*) from leakage currents (*gray*) (see text). Currents due to stray capacitance have been subtracted out. Waveforms were recorded from the V_V and V_I terminals in Figure 2.1B through 10x probes. **C** Voltage-controlled stimulation through an electrode on an older MEA, showing substantial leakage current, presumably through damaged insulation.

tion potentials, while sufficiently strong negative-then-positive stimuli also elicited action potentials (Figure 2.6). Responses were recorded across the entire array, although electrodes close to the stimulation electrode were slightly more likely to record responses. The voltage dependence of responses was the same for recording electrodes near to and far from the stimulation electrode (data not shown). Monophasic positive pulses did not elicit many responses, despite the fact that the downward transient at the end of the pulse corresponds to a strong negative current. Conversely, monophasic negative pulses *were* effective, but significantly less so than biphasic pulses.

At fixed duration, response reliability generally increased with increasing stimulus amplitude, though about 20% of precise responses decreased in reliability above

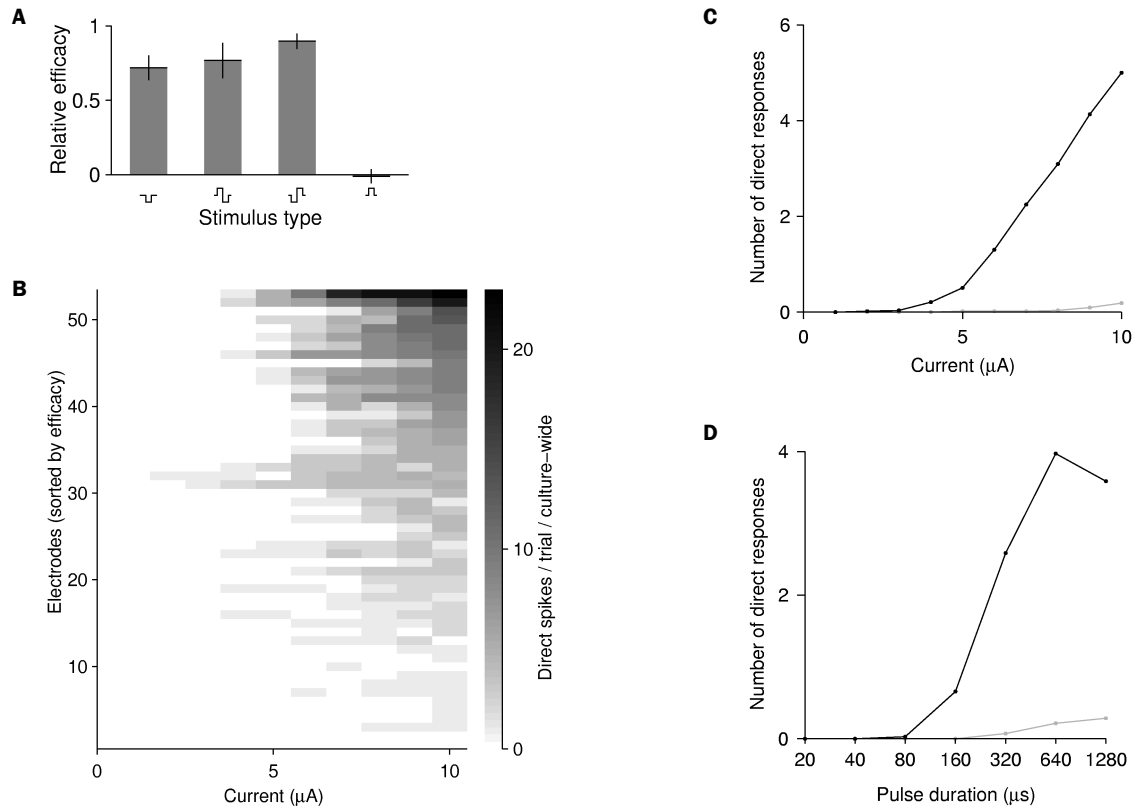


Figure 2.5: Responses to current stimulation. **A** Comparison of pulse shape efficacies. We measured the per-trial-average number of spikes recorded array-wide in the first 20 ms after stimulation, for various pulse shapes (left to right: monophasic negative; biphasic positive-then-negative; biphasic negative-then-positive; monophasic positive), with amplitude 5 μA and duration 200 μs per phase. Baseline firing rates were subtracted, and results were normalized by the value for the most effective pulse. Shown are averages and SEM ($N = 6$). **B** Number of direct responses (across the array) recorded after monophasic negative current pulses of 50 μs duration and varying amplitudes. Based on 55 electrodes in 5 MEAs. **C** Average across electrodes of data in **B** (black). Positive pulses (gray) did not excite spikes. **D** Average number of direct responses after negative (black) and positive (gray) current pulses of 1 μA amplitude and varying durations.

a certain point (Figure 2.7A). At fixed stimulus amplitude, the number of responses did not depend strongly on pulse duration, once the duration is above threshold. For short positive-then-negative pulses, the efficacy could be significantly improved by leaving the negative phase unterminated (cutting the stimulation current rather than explicitly driving the voltage back to baseline). This effect was much weaker for longer pulses.

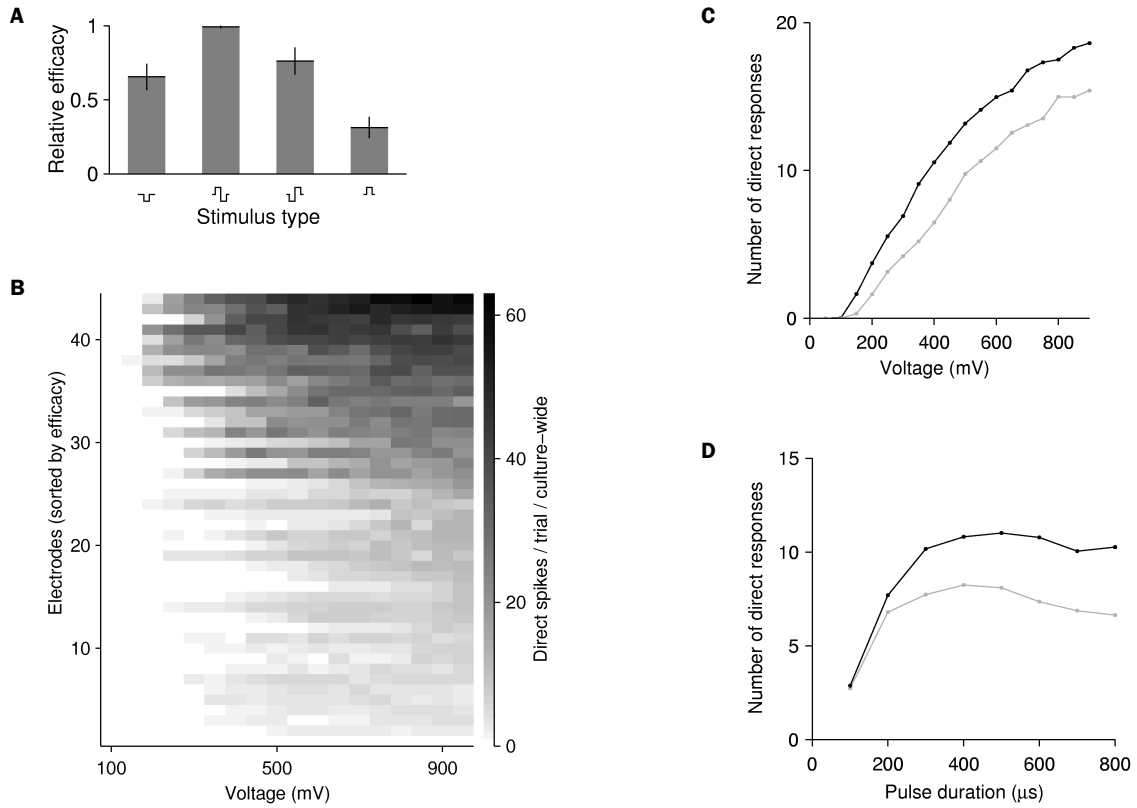


Figure 2.6: Responses to voltage stimulation. **A** Comparison of pulse shape efficacies. We measured the per-trial-average number of spikes recorded array-wide in the first 20 ms after stimulation, for various pulse shapes (left to right: monophasic negative; biphasic positive-then-negative; biphasic negative-then-positive; monophasic positive), with amplitude 500 mV and duration 400 μ s per phase. Baseline firing rates have been subtracted, and results were normalized by the value for the most effective pulse. Shown are averages and SEM ($N = 6$). **B** Number of direct responses (across the array) recorded after biphasic positive-then-negative voltage pulses of 400 μ s duration and varying amplitudes. Based on 45 stimulation electrodes in 5 MEAs. **C** Average across electrodes of data in B (black). Responses to negative-then-positive stimulation are slightly less strong (gray). **D** Average number of direct responses after positive-then-negative (black) and negative-then-positive (gray) voltage pulses of 500 mV amplitude and varying durations (same 45 electrodes). Note that very long pulses result in fewer recorded responses, because some early responses get obscured by stimulation artifacts.

To determine which part of a biphasic voltage pulse gives rise to the responses, we fitted the latency of well-isolated direct spikes as a linear function of per-phase-width of the stimuli. If the spikes are evoked by the initial transient, this function should have zero slope; if the spikes are evoked by the second transient, the slope should be one; and if the final transient is responsible, the slope should be two. (Slopes between zero and one indicate that the first constant phase, rather than the transient, is responsible, while slopes between one and two indicate that the second phase is responsible.) We found that most direct spikes are time-locked to the (first) downward transient in the voltage waveform: for positive-then-negative stimuli, the latency from stimulus onset typically increases directly with pulse width (median slope of latency vs. phase width is $1.01 \mu\text{s}/\mu\text{s}$), while for negative-then-positive stimuli, the latency does not change much (median slope is $0.23 \mu\text{s}/\mu\text{s}$). This confirms the observation that only negative current pulses are effective stimuli.

Number and Reliability of Responding Cells

In principle, when changing a parameter increases the number of direct spikes observed in the response, this can be due to either of two mechanisms: more neurons can become directly activated by the stimulus, or the activation of previously activated neurons can become more reliable. We isolated direct responses occurring in the first 30 ms post-stimulation using a clustering algorithm (see the Appendix at the end of this chapter) that identified sharp peaks in the post-stimulus time histogram (PSTH). From this, we calculated the ‘reliability’ of a direct response (at a given stimulus strength) as the fraction of trials in which it is observed. We defined the ‘relative reliability’ of a direct response as the absolute reliability of that response divided by the peak reliability of that response. Aligning the relative-reliability curves of many direct responses to their point of 50% reliability, we found that individual direct responses were initiated above sharply defined thresholds: a change of 10% in pulse amplitude could make the difference between no response and near-maximal response (Figure 2.7). This held for both voltage-controlled and current-controlled stimuli, suggesting that the increase in number of recorded spikes was due to an increased number of cells firing action potentials in direct response to the stimulus, and not due to an increase in the reliability of the firing of any individual cell.

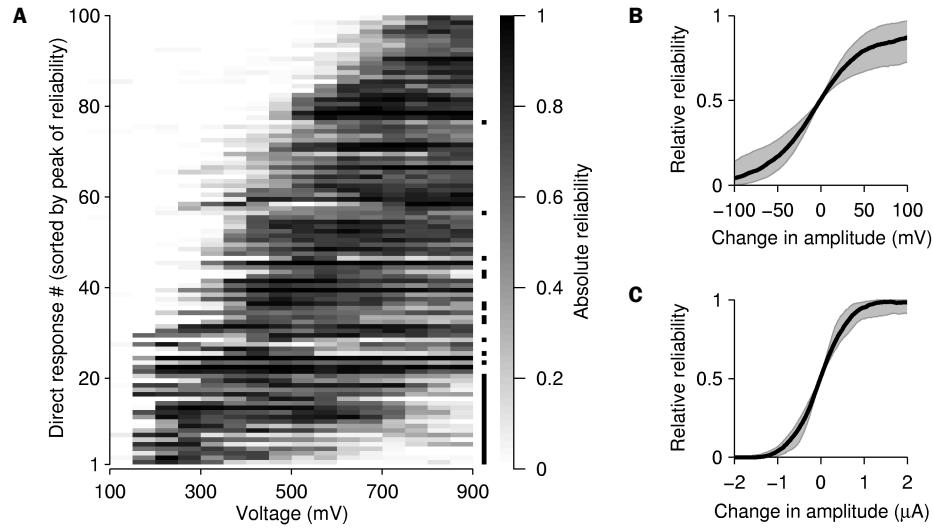


Figure 2.7: **A** Reliability of a random selection of 100 direct responses, sorted by the stimulation voltage at which their reliability peaks. While most direct responses increase in reliability with increasing voltage, a subset only occurs in a voltage range (indicated by black marks on right). **B** Relative reliability (see text) of direct responses evoked by voltage-controlled pulses near their threshold. Notice the sharpness of the onset. Curve is average of 757 direct responses, shaded area indicates quartiles. **C** Relative reliability of direct responses evoked by current-controlled pulses near their threshold. Curve is average of 362 direct responses, shaded area indicates quartiles.

Voltage vs. Current Control: Electrode and Leakage Currents

For comparison of current- and voltage-controlled stimuli, it is important to distinguish between currents flowing through the electrode, and currents lost through leakage or arising purely from stray capacitance in cabling. (Stray capacitance is especially significant for voltage-controlled pulses, because sharp voltage transients cause strong capacitive currents.) We measured the leakage currents explicitly by applying voltage pulses after fencing off the electrode area of three MEAs using the end of a wide-bore pipette tip held in place with vaseline. Thus we could flood the rest of the MEA dish and measure most of the lead-to-medium conductance as well as stray capacitance, in isolation from electrode conductance. Stray capacitance, at 450 pF, was responsible for most of the current peak during voltage transients. While we found leakage currents to be insignificant (less than 5% of total current) for brand




Figure 2.8: Efficacy (see text) of positive-then-negative voltage-controlled pulses of 400 μ s per phase (*solid*) and monophasic negative current pulses of 50 μ s duration for which the resulting voltage was measured (*dashed*). The voltage-controlled pulses are significantly more effective. (Average and SEM, $N = 6$.)

new MEAs, they could constitute up to 30% of the total current path in 6-month-old MEAs that had been plated several times (see Figure 2.6B). Upon physical inspection, such older dishes were often found to exhibit microscopic cracks in the insulation layer.

In all but one MEA tested, voltage pulses were much more effective than current pulses that reached the same peak voltage. We defined the ‘efficacy’ of a stimulus as the number of spikes that stimulus evoked on average, normalized by the number of spikes evoked by the most effective stimulus in the repertoire. To compare efficacies of voltage and current-controlled stimuli, we measured the peak electrode voltage reached during current-controlled stimulation, and plotted the experimental efficacies as a function of this peak voltage (Figure 2.8). One MEA in our data set exhibited anomalously strong responses to current stimuli. We found that this MEA had large electrode conductance, but also exceptionally large leakage (as shown in Figure 2.4C), presumably due to damaged insulation.

Discussion

From the great variety of possible stimulation pulse shapes that could be applied to MEA electrodes, we studied eight important families: monophasic and biphasic rectangular current pulses of either polarity, and monophasic and biphasic rectangular voltage pulses of either polarity. We found that the efficacy of stimuli in any of these

families can be attributed to the generation of negative electrode currents.

To explain why negative current pulses are effective stimuli while positive currents are not, others have modeled the intracellular currents resulting from extracellular fields. They found that stimulation efficacy follows the second spatial derivative of the potential along the length of an axon (Rattay, 1999; McIntyre and Grill, 1999). For stimulation of cell bodies or unmyelinated axons very near electrodes, we can consider the following simpler equivalent picture: negative current pulses lower the potential of the medium near the electrode. This can evoke an action potential if the pulse is long enough to significantly charge the local membrane capacitance, thus depolarizing that membrane. A simple model suggests that the relevant time constant for stimulation of axons that traverse an electrode is of the order of tens of microseconds, thus explaining how even the very short current pulses resulting from voltage-controlled stimulation can be effective stimuli. This picture also explains why positive pulses are essentially ineffective: for such pulses capacitive currents enter the axon more distally and over a much larger area, so that the resulting depolarization is slight.

The sheer number of direct responses strongly suggests that axons, not somata, are the main recipients of the stimuli, since geometry limits the number of somata that can be in close contact to a given electrode. Our results do not indicate whether the responses are *recorded* from axons or somata, but we believe they are from antidromically activated somata. The recorded pulse shapes resemble those from non-direct (synaptically mediated) responses.

Practical Consequences

We have determined a class of voltage-controlled stimulation waveforms that are effective at eliciting action potentials in nearby axons upon stimulation of a single substrate electrode. We were able to evoke activity, without ever requiring voltages that cause electrochemistry. By tuning the pulse amplitude, we could select how many cells to stimulate, since individual cells became responsive at sharply defined stimulus amplitudes. On about 20% of electrodes, tuning pulse amplitudes could even be used to select different subsets of cells to stimulate (Figure 2.7).

Surprisingly, we found that for almost all electrodes, positive-then-negative bipha-




Figure 2.9: Classifying response components by FSMEM. Raster plot of spikes recorded on one electrode in response to a distant electrode (*top*). FSMEM found five components (*bottom*) (alternate components are shown black and gray).

sic voltage-controlled stimuli were much more effective than either short or long negative current pulses that reached the same peak voltage. This is good news for long-term stimulation experiments, because voltage pulses can be explicitly designed to avoid electrochemistry, without having to monitor the results. We hope that this will encourage other researchers to explore the potential benefits of non-ionic stimulation using voltage-controlled pulses, not only in culture, but also in *in vivo* situations.

Appendix: Identifying Response Components

For each (stimulation electrode, recording electrode)-pair, we fitted a mixture of Gaussians to the first 30 ms of the peri-stimulus-time-histogram (PSTH), using a variation on the SMEM algorithm (Ueda et al., 2000) that automatically determines the appropriate number of Gaussian components in the mixture. Figure 2.9 shows an example of the input and output of the algorithm, which we dubbed FSMEM, for “Free Split/Merge Expectation-Maximalization.”

We construct a histogram of latencies of spikes on a given recording channel in response to a given stimulation channel, then fit it with a probability density function of the form:

$$p_{N,\vec{\mu},\vec{\sigma},\vec{\alpha}}(t) = \sum_{n=1}^N \frac{\alpha_n}{\sqrt{2\pi} \sigma_n} e^{-\frac{(t-\mu_n)^2}{2\sigma_n^2}}.$$

Here μ_n are the centroids (in time) of each of the components, σ_n are their widths,

and α_n are the mixing coefficients.

The fit is performed by maximizing an objective function which is the sum of two terms:

$$L = L_0 + L_{\text{MDL}}.$$

Here L_0 is the log-likelihood (Fisher, 1922),

$$L_0 = \sum_i \log p(t_i),$$

(where i sums over spikes, and t_i is the latency of the i -th spike in the dataset). The other term is a minimum-description-length correction (Rissanen, 1978):

$$L_{\text{MDL}} = -(3N - 1).$$

The inclusion of L_{MDL} allows the algorithm to autonomously select the number of components without bias.

After determining the distribution of responses, spikes were assigned to the most probable component in the histogram, so the reliability of each component could be determined in each stimulation condition.

The algorithm now consists of these steps:

1. Start with $N = 1$, and suitable default values for $\vec{\mu}$, $\vec{\sigma}$, and $\vec{\alpha}$.
2. Optimize $\vec{\mu}$, $\vec{\sigma}$, and $\vec{\alpha}$ using expectation-maximalization (EM) (Dempster et al., 1977).
3. For each component, n , consider whether splitting this component in two may improve the objective function $L = L_0 + L_{\text{MDL}}$:
 - (a) Create a new cluster $N + 1$ with $\mu_{N+1} = \mu_n + \alpha$ where α is a small random number. Set $\mu_n \leftarrow \mu_n - \alpha$. Set $\alpha_{N+1} \leftarrow \alpha_n/2$, and then $\alpha_n \leftarrow \alpha_n/2$.
 - (b) Run EM on the subset of spikes associated with the original cluster n , operating only on the parameters of the new clusters n and $N + 1$.
 - (c) Run EM on the entire dataset.
 - (d) If this improves L , accept the split, otherwise reject it.

4. If no split constitutes an improvement, consider whether merging any pair of clusters helps: For each pair:
 - (a) Replace the two clusters n and m by a single cluster n with $\mu_n \leftarrow (\mu_n + \mu_m)/2$, $\sigma_n \leftarrow (\sigma_n + \sigma_m)/2$, and $\alpha_n \leftarrow \alpha_n + \alpha_m$.
 - (b) Run EM on the subset of spikes associated with the original clusters n and m , operating only on the parameters of the new cluster n .
 - (c) Run EM on the entire dataset.
 - (d) If this improves L , accept the split, otherwise reject it.
5. If no merge constitutes an improvement, go back to considering merges.
6. If neither splits nor merges improve L , the algorithm stops.

A more detailed description of the algorithm, as well as an implementation in Matlab (MathWorks), is available upon request.

3

Real-Time Artifact Suppression*

We describe an algorithm for suppression of stimulation artifacts in extracellular multi-electrode array (MEA) recordings. A model of the artifact based on locally fitted cubic polynomials is subtracted from the recording, yielding a flat baseline amenable to spike detection by voltage thresholding. The algorithm, SALPA, reduces the period after stimulation during which action potentials cannot be detected by an order of magnitude, to less than 2 ms. Our implementation is fast enough to process 60-channel data sampled at 25 kHz in real-time on an inexpensive desktop PC. It performs well on a wide range of artifact shapes without re-tuning any parameters, because it accounts for amplifier saturation explicitly and uses a statistic to verify successful artifact suppression immediately after the amplifiers become operational. We demonstrate the algorithm's effectiveness on recordings from dense monolayer cultures of cortical neurons obtained from rat embryos. SALPA opens up a previously inaccessible window for studying transient neural oscillations and precisely timed dynamics in short-latency responses to electric stimulation.

* Published as: Daniel A. Wagenaar and Steve M. Potter, 2002: Real-time multi-channel stimulus artifact suppression by local curve fitting. *J. Neurosci. Methods* **120**(2), pp. 113–120. © 2002 Elsevier. Reprinted with permission.

Introduction

Multi-electrode arrays (MEAs) (Thomas et al., 1972; Gross, 1979; Pine, 1980) and related technologies such as tetrode probes (Gray et al., 1995), silicon probes (Bai and Wise, 2001) and multiwire probes (Nicoletis et al., 1998) offer great promise to record action potentials extracellularly from a large number of cells simultaneously (Meister et al., 1994; Potter, 2001), in cell culture, in slice or *in vivo* (Bragin et al., 2000). In addition, electrical stimulation through such arrays has been reported in a wide variety of preparations, such as murine spinal cord (Gross et al., 1993), rat cortex (Jimbo et al., 1999), cat sciatic nerve (Branner and Normann, 2000) and rabbit retina (Grumet et al., 2000). Simultaneously stimulating and recording through a single MEA is attractive for the study of input-output relationships (Novak and Wheeler, 1988; DeAngelis et al., 1998), but poses technical difficulties because the stimuli employed are often four or five orders of magnitude greater than extracellularly recorded action potentials ('spikes'). These may be as low as 10 μV (shown below), while stimuli are typically on the order of a volt (Pancrazio et al., 1998; Jimbo et al., 1999), causing substantial stimulation artifacts that corrupt the data or saturate the recording electronics.

Several factors contribute to these artifacts (Grumet, 1999). The stimulus induces pickup on other electrode channels by a combination of capacitive crosstalk between leads and conduction through the tissue or recording medium, saturating the amplification system. The nonlinear behavior of saturated amplifiers, together with the properties of the filters used for noise reduction, make this artifact last much longer than the stimulus that caused it, sometimes up to 100 ms (Maeda et al., 1995), even on channels not used for stimulation. In some cases this problem can be reduced by physically separating the recording site from the stimulation site (Grumet et al., 2000), or circumvented by using non-electronic means for either stimulation or recording, such as photo-uncaged glutamate (Wang and Augustine, 1995), optical recording (Obaid et al., 1996; Maher et al., 1999b) or muscle twitch response (Branner and Normann, 2000). In all other cases, careful design of the electronics is required to minimize pickup of stimulation artifacts.

One would like to stop large artifacts from entering the recording system in the

first place. To do so, Jimbo and Kawana (1992) recorded differentially between pairs of electrodes spaced at 10 μm , while stimulating between a similar, distant pair of electrodes. Sample-and-hold circuitry has also been used to prevent amplifier saturation (Novak and Wheeler, 1988; Jimbo et al., 1998; Grumet, 1999), but with mixed results. Jimbo et al. (1999) were able to record 5 ms after stimulation, even from the stimulated electrode, but the implementation details are not described.* In contrast, Grumet (1999) reports little or no reduction of artifacts with a sample-and-hold approach. Presently commercially available electrophysiology equipment for electrode arrays does not employ this strategy.

When artifacts cannot be entirely prevented in hardware, various forms of digital filtering can be used to reduce them. For example, if artifacts are the same across trials, template subtraction can be employed: Jimbo and Kawana (1992) estimated artifacts in stimulus responses by scaling the artifact recorded under sub-threshold conditions. Unfortunately, due to hysteresis at the electrode interface or in the electronics, artifact shapes often do vary between trials. Alternatively, Okajima et al. (1995) manually subtracted a linear baseline from recordings of muscle action potentials. This is too labor intensive for multi-channel recordings, and certainly cannot be applied in real-time as the data come in. As a last resort, blanking (digitally setting the signal to zero) can be used to eradicate any artifacts (O’Keeffe et al., 2001). Any action potentials occurring within the duration of the artifact are lost, so if one is interested in the early part of stimulus response, this is not an option. Our solution, which works in real-time, is SALPA, an algorithm for Subtraction of Artifacts by Local Polynomial Approximation. We show that its performance is superior to some possible alternatives: two simple high pass filters and one linear phase filter (see, e.g., Jackson, 1996).

* After publication of this article, they did make their hardware design public (Jimbo et al., 2003). (Note added, DAW 2005.)

Methods

Cell Culture and MEA Preparation

Dense cultures of dissociated rat cortical cells were grown in MEA dishes and maintained for several months. Culture methods have been detailed elsewhere (Potter and DeMarse, 2001). Very briefly, cortex from E18 Wistar rats was dissected under sterile conditions and dissociated using papain. MEA dishes from MultiChannel Systems (Reutlingen, Germany) with sixty 10 μm diameter electrodes arranged in a rectangular array with 200 μm inter-electrode spacing were coated with polyethylene imine (PEI) and laminin. A 15 μL droplet of Neurobasal medium containing 50,000 cells was applied to the electrode area, and the dishes were sealed with FEP Teflon lids (Potter and DeMarse, 2001). After 30 minutes, 1 mL of medium was added, and the dishes were transferred to an incubator (35 $^{\circ}\text{C}$, 5% CO_2 , 9% O_2 and 65% humidity). After one day, and thereafter every four days, the medium was replaced entirely by the medium adapted from Jimbo et al. (1999). Electrical activity of these cultures was recorded through the MEA, amplified and sampled with 12 bits resolution at 25 kHz (MultiChannel Systems). One electrode was used for stimulation, while all the others were used for recording. The data presented below were obtained by stimulating five-month-old cultures through one electrode with single biphasic voltage pulses of ± 600 mV, lasting 400 μs per phase, positive phase first.

Artifact Suppression

SALPA works by locally fitting (Hastie and Loader, 1993) a function to the recorded trace that has enough degrees of freedom to accurately model the artifact, but not enough to represent individual action potentials. By subtracting this fitted function from the recording, the artifact-free signal remains, and action potentials can be detected by setting a voltage threshold. We observed considerable variability in artifact shapes between electrode channels, and even on individual channels between trials of the same stimulus. Therefore we chose not to make assumptions about regularities in artifact shapes, and instead fit independent functions to each individual artifact.

For every time-point n_c in the recording from a given electrode, a third degree polynomial is fitted to a segment of $2N + 1$ samples centered around n_c . The seg-




Figure 3.1: Illustration of the fitting method used by SALPA. The first fit after depegging that has acceptable deviation (see text) is used to model $N + 1$ samples (bottom-most curve). Thereafter, each fitted polynomial is used to model one sample only (other dashed curves). The thin solid curve is the raw recording. For visual clarity, only one in ten fits is shown, and they have been vertically displaced. The thick solid curve is the resulting model of the artifact. Circles mark the centers of each fit. The dotted vertical lines indicate which parts of the fitted polynomials are used for modeling the artifact.

ment half-length N is experimentally tuned for optimal performance, as discussed below (see *Results*). The fitted value at the central point n_c is subtracted from the raw recording at that point to yield a cleaned signal. The rest of the fitted curve is discarded; to estimate the cleaned signal at $n_c + 1$, a new polynomial is fitted to the data in the segment of $2N + 1$ samples centered around $n_c + 1$.

The fitting process is different for the first N points of the raw data, starting when the stimulus-induced saturation of the electronics ends (‘depegging’; saturation of the electronics is determined by the digital signal having its minimum or maximum possible value). A polynomial is fitted to the data centered on the $(N + 1)$ -th point after depegging, and the artifact up to the center of that window is estimated using that single third-degree polynomial, as illustrated in Figure 3.1.

The raw electrode signal is represented as a sequence of sampled voltages, V_n , where n is the ordinal number of the sample. (We sample at 25 kHz, so the unit of n is 40 μ s in real time.) We decompose this raw signal into an estimated artifact A_n and

a cleaned signal v_n by assuming that in the vicinity of some central point n_c , we can approximate the artifact by a cubic polynomial expressed in terms of the distance $(n - n_c)$ from the central point*:

$$A_n^{(n_c)} = \alpha_0^{(n_c)} + \alpha_1^{(n_c)}(n - n_c) + \alpha_2^{(n_c)}(n - n_c)^2 + \alpha_3^{(n_c)}(n - n_c)^3.$$

The fit parameters α_k are found by minimizing the function

$$\chi_{(n_c)}^2 = \sum_{n=n_c-N}^{n_c+N} \left(V_n - A_n^{(n_c)} \right)^2$$

with respect to those parameters. We then estimate the corrected voltage at the central point n_c as:

$$v_{n_c} \equiv V_{n_c} - A_{n=n_c}^{(n_c)} = V_{n_c} - \alpha_0^{(n_c)}.$$

Next, we obtain a new fit centered around n_c+1 to estimate $v_{n_c+1} = V_{n_c+1} - \alpha_0^{(n_c+1)}$, and so on.

Fitting a new polynomial for every single data point might seem to be computationally very expensive, but it is not, because the fit parameters can be calculated recursively, as follows. Let us introduce the shorthands

$$T_k = \sum_{n=n_c-N}^{n_c+N} (n - n_c)^k, \quad \text{for } k = 0 \dots 6,$$

and

$$W_k^{(n_c)} = \sum_{n=n_c-N}^{n_c+N} (n - n_c)^k V_n, \quad \text{for } k = 0 \dots 3;$$

and the (4×4)-matrix S with entries $S_{kl} = T_{k+l}$, (for $k, l = 0 \dots 3$). The parameter values that minimize χ^2 can then be written as

$$\alpha_k^{(n_c)} = \sum_{l=0}^3 S_{kl}^{-1} W_l^{(n_c)}, \quad \text{for } k = 0 \dots 3,$$

which can be computed cheaply once the $W_l^{(n_c)}$ are known, since the entries of S^{-1}

* The notation $[\cdot]^{(n_c)}$ will be used throughout to represent the quantity $[\cdot]$ evaluated for the fit centered around n_c .

are constants depending only on N . The complexity is further reduced because T_k is identically zero for odd values of k .

A recursion relation for $W_k^{(n_c)}$ is obtained by straightforward algebraic manipulation of the expressions for $W_k^{(n_c+1)}$ and $W_k^{(n_c)}$:

$$W_k^{(n_c+1)} = \sum_{l=0}^k \frac{(-1)^{k-l} k!}{l!(k-l)!} W_l^{(n_c)} + N^k V_{n_c+N+1} - (-N-1)^k V_{n_c-N}.$$

It is the existence of this closed-form expression that makes the method viable for real-time operation.

As noted above, at the beginning of the recording, just after a channel depegs, we are forced to use a fit based on a window centered N samples ahead in time. Such a non-central window is likely to give a less accurate fit to the artifact, so it is important to assess the quality of the fit before trusting it. This assessment can be based on the *deviation*:

$$\mathcal{D}^{(n_c)} = \sum_{n=n_c-N}^{(n_c-N)+(\delta-1)} \left(V_n - A_n^{(n_c)} \right),$$

where the width of the estimator, δ , can be chosen to be some fixed fraction of N , e.g., $\delta = N/10$. For good fits, $\mathcal{D}^{(n_c)}$ is normally distributed with zero mean, and variance $\sigma_{\mathcal{D}}^2 = \beta^2 \delta \sigma_V^2$, where σ_V^2 is the variance of the recording, and β is a correction factor which is larger than unity if the noise in the recording is not white. For our equipment, $\beta^2 \approx 5$. We advance n_c until $\mathcal{D}^{(n_c)}$ attains an acceptably small (see below) absolute value relative to $\sigma_{\mathcal{D}}$, before declaring the artifact successfully suppressed.

The following results were obtained with $N = 75$ (corresponding to 3 ms at 25 kHz sampling rate), $\delta = 5$ and by rejecting fits after depegging until \mathcal{D}^2 no longer exceeds $3^2 \times 5 \times \sigma_V^2$. Spikes were identified by thresholding at five times RMS noise, and validated based on a test of their waveform shape: spikes were rejected if there were any peaks of either polarity larger than 90% of the main peak within ± 1 ms (P. P. Mitra, personal communication).

Our C++ implementation of SALPA is capable of processing 60 channels of MEA data at 25 kHz in real-time on an AMD Athlon 1.33 GHz processor, using just 75% of CPU time. This allows online spike detection entirely in software on an inexpensive PC system. The C++ source code is available upon request.

Analysis

To assess the quality of the algorithm, we compared its output on typical artifact-corrupted data with the output of three alternative filters: a three pole Butterworth high-pass filter with 400 Hz cutoff (BW-H), subtraction of the output of a three pole Butterworth low-pass filter with 600 Hz cutoff (BW-L), and a 39 pole linear phase high-pass filter with 500 Hz cutoff, designed using cosine expansion (LPC) (Jackson, 1996). The order of this filter was chosen such that we could compute it in real-time using a simple C++ program. The Butterworth filters were chosen because they are a computationally inexpensive simple alternative.

Two performance measures were used: *lost time*, the latency after depegging of the electronics at which the artifact is successfully suppressed, and *PNR loss*, the reduction of the ratio of action potential peak amplitude to noise (PNR) induced by the filter. Lost time was determined by computing 5 ms wide box-car averages of the signal, and rejecting data until the box-car average no longer exceeded the RMS noise. PNR loss was measured relative to raw data filtered through a single pole high-pass filter at 150 Hz, which, before developing SALPA, we used routinely to clean the data of DC drift and any low frequency local field potentials for the purpose of spike detection in recordings of spontaneous activity. In general, artifact suppression filters will reduce the ratio of spike amplitude to RMS noise, because there is substantial spectral overlap between artifacts and spike waveforms.

Results

Rat cortical cultures were stimulated with 600 mV biphasic pulses. Large dish-wide artifacts were observed in the resulting recordings. Figure 3.2 shows how SALPA and the other filters act on these stimulation artifacts and on action potential waveforms. Aside from reducing spike amplitudes, filters may distort spike waveforms in more subtle ways, exemplified by the positive ‘ghost’ phases induced by BW-L and BW-H. These may hamper subsequent spike sorting, and may even lead to spurious detection of non-existent upgoing action potentials. Fortunately, SALPA—being a linear phase filter (except in the initial N samples after depegging)—is free of such phase distortions.




Figure 3.2: An example of an artifact in electrode recordings, with the output of various filters (*left*) and the effect on spike waveforms (*right*). Notice the difference in scales between left and right: the artifact in the raw data is an order of magnitude larger than the spikes. Thin curves are raw data; thick curves are filter output. From top to bottom: SALPA, subtraction of low-pass Butterworth (BW-L), high-pass Butterworth (BW-H), and 39 pole linear phase filter (LPC). Notice SALPA's blanking of the output during saturation of the electronics. The Butterworth filters induce significant phase distortion while leaving much more lost time than SALPA. Even the linear phase filter leaves some echo of the artifact. The spike waveforms shown are from the same recording as the artifacts, but at longer latency, to allow direct comparison with non-corrupted raw data.




Figure 3.3: Lost time (left, solid; results are shown for 4 dishes separately) and PNR reduction (right, dashed) can be traded off by varying SALPA’s filter length, measured here in milliseconds. Note that the left-hand axis increases downwards, so ‘up’ means ‘better’ for both axes. Artifact sizes and duration vary by almost an order of magnitude between dishes, depending on electrode impedances (nominally 300 k Ω at 1 kHz). This is reflected in SALPA lost time, shown here for four different dishes. PNR drops dramatically when the filter half-width N approaches the duration of action potentials. The optimal choice of N must depend on the application, and on the PNR in the raw recordings. The results in the rest of this article were obtained with filter half-lengths of 3 ms. PNR change is measured relative to single pole high-pass filtering at 150 Hz (see text). The sample period, τ_{sample} , was 40 μs .

All filters used in this comparison have parameters that can be tuned to trade off lost time against PNR reduction. For SALPA, this is the segment half-width N ; for the other filters the cut-off frequency plays this role. Figure 3.3 presents the trade-off for SALPA. The optimal choice of N depends on the kind of experiment one is doing. In Figure 3.4, the performance of SALPA at $N = 75$ (3 ms) is compared with the other filters at the frequencies specified above.

One feature of SALPA that gives it an edge over the alternatives, is that it explicitly recognizes saturation of recording electronics, outputting zeros whenever the digital values of the recording are at the extreme ends of their range, and that it incorporates a statistic to test goodness of fit for the earliest time points, as detailed in *Methods*. As a result, SALPA performs well on a wide range of artifact sizes and shapes (Figure 3.5).




Figure 3.4: Comparison of various filter methods. PNR change is measured relative to a single pole high-pass filter at 150 Hz (see text). Notice that the reference 150 Hz filter also reduced spike amplitudes by a small fraction, so relative PNR *gain* resulted in some cases. ‘Restricted SALPA’ is SALPA without the third degree term. Lost time does not include the duration of amplifier and ADC saturation (1.04 ± 0.02 ms). Charted values are mean and standard deviation of data collected from 55 electrodes.




Figure 3.5: For widely different artifact waveforms (shown in insets), SALPA yields usable output as early as 2 ms post stimulus (less than 1 ms after depegging). Faint traces are raw data with two different vertical offsets added to show details. Bold is SALPA output. Notice the spikes riding on the slope of the artifacts which cannot be detected by thresholding the raw data.

The ultimate test for an artifact suppressor is whether it allows reliable detection of spikes at short latencies. Figure 3.6 shows detectable spikes on two electrodes in different MEA dishes containing mature cultures, using SALPA, using only a 150 Hz high-pass filter, and using each of the other filters mentioned above. SALPA reveals a structure of oscillations and very precisely timed spikes in the early response that would otherwise go undetected.

Discussion

We have presented an algorithm for stimulus artifact suppression that can be applied to sixty-channel electrode recordings in real-time on inexpensive PC hardware. SALPA does not cause phase distortion of spike waveforms unlike simple high-pass filters, and it is less computationally intensive than straightforward implementation of a generic linear phase filter of equivalent length. Perhaps more importantly, the algorithm covers the first few milliseconds of the artifact naturally, because it takes amplifier saturation into account explicitly. Other filters tend to suffer from ringing as a result of the sharp transient at the time of amplifier depegging.

In the bulk of the data, local regression of the form used in SALPA is just a special case of linear phase filtering. SALPA effectively functions as a high-pass filter with -3 dB cut-off frequency $f_0 = 0.6f_s/N$, i.e., $f_0 = 200$ Hz for $N = 75$ and sampling frequency $f_s = 25$ kHz. This is sufficient, since a spectrogram of the stimulation artifacts would reveal that high frequency power is mostly concentrated in the first few milliseconds—beyond that, the tail of the artifact is reasonably well spectrally separated from spike waveforms. Simple filters would have to find a difficult compromise between preventing ringing from the initial sharp transient of artifacts, and preserving signal shape. SALPA surmounts this problem by not using any samples from the sharp transient in its estimate of the shape of the tail, through the use of asymmetric fitting windows shortly after depegging (as illustrated in Figure 3.1).

One could attempt to improve the performance of the algorithm by increasing the order of the polynomials used. The computational expense of the algorithm would increase, and more dramatically than one might expect, because the higher powers of $(n - n_c)$ in the equations would make representation of these numbers as 32 bit




Figure 3.6: Evoked activity on two electrodes (from different dishes). Each set of raster plots shows action potentials detected using the reference 150 Hz high-pass filter ('raw') and each of the filters tested. Each raster contains results from 500 consecutive trials, with each dot representing one action potential. Stimulation was on electrodes 280 and 450 μm away from the recording site. Gray bars represent the time and duration (0.8 ms) of the stimulus. Spikes were detected by thresholding at five times RMS noise, and validated based on their waveform shape (see text). Notice that the lost time was estimated independently for each individual trial. This explains the limited reliability with which, e.g., the spike at 5 ± 0.5 ms in the bottom panel is detected by BW-H. Electrodes were chosen to illustrate a variety of short latency response types.

integers impossible. The main effect of increasing polynomial order is an increase of cut-off frequency, which can be achieved more economically by decreasing N . In any event, it is desirable to use odd order, because this gives the fit at the beginning of the trace one more degree of freedom than in the bulk, further improving the response to sharp transients without compromising signal shapes.

With current commercially available hardware, SALPA is less well suited for recordings from the stimulated electrode, because saturation on that channel lasts beyond the duration of the early phase of the response. Improvements in hardware, for example using sample-and-hold based artifact reduction, might bring the stimulated channel within SALPA's domain.

SALPA opens up a new window on very short latency multi-neuronal responses to electrical stimulation. The early post-stimulus neural dynamics comprise oscillations and action potentials timed with a precision not observed before. We are currently using SALPA to investigate the nature of these response components and their role in stimulus-induced plasticity. The precisely timed responses can be used to drive 'behaviors' in a neurally-controlled animat (DeMarse et al., 2001).

4

Development of Spontaneous Bursting Activity*

We have collected the most comprehensive set to date of multi-unit data on dissociated cortical cultures. We followed the first five weeks of the development of 58 cultures of different densities—3000 to 50,000 neurons on areas of 30 to 75 mm²—growing on multi-electrode arrays (MEAs). While the aggregate spike detection rate scaled linearly with density, as expected from the number of cells in proximity to electrodes, dense cultures started to exhibit bursting behavior earlier in development than sparser cultures. Analysis of responses to electrical stimulation suggests that axonal outgrowth likewise occurred faster in dense cultures. After two weeks, the networks' activity was dominated by population bursts in most cultures. In contrast to previous reports, development continued with changing burst patterns throughout the observation period. Burst patterns were extremely varied, with inter-burst intervals between 1 and 300 s, different amounts of temporal clustering of bursts, and different firing rate profiles during bursts.

These results are based on 963 half-hour-long recordings. To encourage further investigation of the rich range of behaviors exhibited by cortical cells *in vitro*, we are making the data available to other researchers, together with Matlab code to facilitate access.

* About to be submitted as: Daniel A. Wagenaar, Jerome Pine, and Steve M. Potter: Development of spontaneous bursting activity in cortical cultures.

Introduction

Dissociated cultures of cortical cells grown on multi-electrode arrays (MEAs) have been used in many studies of network physiology because of their superior accessibility compared to *in vivo* models, in terms of electrical recording and stimulation, pharmacological manipulation and imaging. These studies described fundamental properties of network activity patterns (Maeda et al., 1995; Gross and Kowalski, 1999; Segev et al., 2001; Beggs and Plenz, 2003; Van Pelt et al., 2004b), plasticity (Maeda et al., 1998; Jimbo et al., 1999), learning *in vitro* (Potter and DeMarse, 2001; DeMarse et al., 2001; Shahaf and Marom, 2001; Eytan et al., 2003; Ruaro et al., 2005), applications of cell cultures in pharmacological testing (Morefield et al., 2000), and models of epilepsy (Wagenaar et al., 2005b [Chapter 6 in this thesis]).

It is often stated that cortical cells in culture retain many of the properties found in their *in vivo* context, but this similarity is rarely quantified, and important differences assuredly exist (Segev et al., 2001; Potter and DeMarse, 2001). Therefore, the ‘natural’ development of dissociated cultures deserves to be documented, as a baseline against which the results of experimental manipulations can be compared. Previous investigations of this development have each focused on different aspects of the activity patterns exhibited by such cultures, resulting in a patchwork of descriptions that is difficult to integrate. Segev et al. (2001, 2002) observed population bursts, and studied the statistical properties of the distribution of the inter-burst intervals (IBIs). Mukai et al. (2003) also observed population bursts, but focused on development changes in IBI values and spatial extent of bursts, while Van Pelt et al. (2004a,b) focused on the temporal structure of the firing-rate envelope of bursts during development.

Most previous studies were based on observations of small numbers of cultures from unspecified numbers of plating batches. Therefore they may have underestimated the variety of activity patterns that different cultures can exhibit. Here, we present an in-depth study of the development of burst patterns over the course of the first five weeks *in vitro*, based on a dataset encompassing a total of 963 half-hour-long recordings from 58 cultures of five different sizes and densities. For the first time, we also systematically studied stimulation responses over the course of development in

36 cultures, by presenting probe pulses to each electrode in the array after recording spontaneous activity each day. The cultures exhibited a surprisingly wide spectrum of spontaneous activity patterns, characterized by population bursts of qualitatively different shapes, sizes and interval distributions.

Methods

Cell Culture

Cells (neurons as well as glia) were obtained from the cortices of (E18) rat embryos as described before (Wagenaar et al., 2005b). Briefly, timed-pregnant Wistar rats were sacrificed using CO₂, according to NIH approved protocols, at day 18 of gestation. Embryos were removed and decapitated, and the cortices were dissected out. Cortices were cut into 1 mm³ chunks, and dissociated using papain followed by trituration. Cells were spun down onto 5% bovine serum albumin (BSA) to remove debris, then resuspended in Neurobasal medium with 10% horse serum, and passed through a 40 µm strainer. Cell density was determined using a haemocytometer.

MEAs were pre-treated with poly-ethylene-imine (PEI) and laminin, as previously described (Potter and DeMarse, 2001). Laminin (0.02 mg/mL in Neurobasal medium) was applied directly to the center of the array, in drops of either 5 or 20 µL. To prevent premature evaporation, drops of 15 µL medium were spread around the inside edge of the culture dish, and the dishes were sealed with Teflon membranes impermeable to water (Potter and DeMarse, 2001). Laminin drops were removed by vacuum suction just prior to plating cells. The cell suspension was homogenized by pipetting with a wide-bore 1 mL tip (Hamilton, Reno, NV) before taking either 5 or 20 µL and plating it on the wet area left by the laminin. Several plating densities were used, as summarized in Table 4.1. Recordings were obtained from 58 cultures from eight dissections, performed over the course of nine months.

Cultures were maintained in Teflon-sealed dishes, in an incubator with 5% CO₂, 9% O₂, 35 °C and 65% relative humidity. After 24–36 h, the plating medium was replaced by a serum-containing DMEM-based medium adapted from Jimbo et al. (1999). Subsequently, half of the medium was replaced approximately every five days in most experiments. In some experiments ($N=3$), all medium was replaced every

seven days. This did not result in significantly different activity patterns compared with sister cultures. Feeding always took place after the day's recording session, to allow at least 12 hours for transient effects to disappear before the next recording.

Recording

Recording took place in the same incubator used for maintaining cultures, using a pre-amplifier from MultiChannel Systems (Reutlingen, Germany). Excess heat from the pre-amplifier was removed using a custom Peltier-cooled stage. Recording started immediately after transferring each culture into the recording device. MEABench (Wagenaar et al., 2005a [Appendix B in this thesis]) was used for data acquisition and online spike detection. After most recordings, cultures were probed using biphasic voltage pulses of ± 0.8 V, 400 μ s per phase (Wagenaar et al., 2004 [Chapter 2 in this thesis]), applied sequentially at 0.3 s intervals to all electrodes using our custom stimulator (Wagenaar and Potter, 2004 [Appendix C in this thesis]). Stimulation artifacts were removed in software using the SALPA algorithm (Wagenaar and Potter, 2002 [Chapter 3 in this thesis]). A total of 50 pulses were delivered to each electrode. Whether or not cultures were exposed to electrical stimuli did not result in significantly different activity patterns compared to sister cultures ($N=3$). In some cases ($N=36$), a recording was allowed to continue overnight, to collect a library of longer recordings, and to be able to test whether activity patterns observed shortly after moving a culture around were substantially different from activity patterns produced by the culture at other times.

To improve consistency between batches, all dissociations and handling of cultures was done by a single experimenter (DAW), while dissections were performed by one lab technician (Sheri McKinney).

Data Analysis

Burst detection and quantification Bursts were detected using the SIMMUX algorithm (Wagenaar et al., 2005a). Briefly, each electrode trace was searched for *burstlets*: sequences of at least four spikes with all inter-spike intervals less than a threshold (set to 1/4 of that electrode's inverse average spike detection rate). Any

group of burstlets across several electrodes that overlapped in time was considered a *burst*.

Burst detection is a prerequisite for quantifying burst shapes and burst patterns, but for merely describing the level of burstiness of a recording, it suffices to quantify the temporal clustering of spikes. This was done as described previously (Wagenaar et al., 2005b): We counted the array-wide number of spikes in non-overlapping 1 s windows, and determined what fraction of the total number of spikes was contained in the 15% most active windows. Since bursts always occupied fewer than 10% of 1 s windows, this number, f_{15} , is close to one if most spikes occur in bursts. Conversely, if spikes are evenly spaced in time, f_{15} is close to 0.15. We then defined our burstiness index as $BI = (f_{15} - 0.15)/0.85$. Thus, BI is normalized between 0 (no bursts) and 1 (all spikes in bursts).

Results

Dissociated neurons in culture began growing new neurites immediately after plating, and soon formed densely interconnected circuits. Starting from 3–4 days *in vitro* (div), we recorded half an hour of spontaneous activity on most days. In dense cultures (see Table 4.1), cells typically began firing action potentials around 4–5 div. Soon after, they began to synchronize their activity into culture-wide bursts, which mostly dominated the activity throughout the rest of development. During bursts, the array-wide spike detection rate (ASDR) could be elevated up to one hundred times the baseline rate. It appears that most or all cells participated in bursting.

The development of the activity of a typical dense culture is shown in Figure 4.1. It passed through various stages, characterized by different degrees of burstiness, different degrees of temporal clustering of bursts, different burst shapes, and different distributions of burst sizes. A particularly striking phenomenon was observed, in this case, at 7–11 div: culture-wide bursts clustered in short groups, or ‘superbursts’ (Chapter 5). Similar clustering was again seen at 34 div, although the internal structure of these latter superbursts was different: in early superbursts, successive bursts diminished in size, while in older cultures, burst sizes within a superburst often increased (compare panels B6 and B7 in Figure 4.3).





Figure 4.1: Development of burst patterns in a dense culture. Graphs show array-wide spike detection rates (ASDR) per second. The spike detection rate of individual electrodes is represented by gray-scale rasters for all 59 electrodes, stacked vertically below each graph. (Each horizontal line pertains to one electrode; gray values indicate firing rates. This representation is also used in subsequent figures.) Note the different vertical scales in the ASDR graphs; the gray-scale images all have the same density scale.

(Continued on next page)



(Figure 4.1, continued)




Figure 4.2: Phase contrast micrographs of central area of typical cultures of various densities at 1, 15, and 32 days *in vitro*. Note that even our ‘sparse’ cultures are considerably denser than those commonly used for investigating synaptic plasticity with intracellular electrodes (see Bi and Poo, 2001). Scale: electrode spacing is 200 μm .

Table 4.1: Plating parameters. The labels in this table will be used throughout the text when discussing cultures of a particular density.

	‘Dense’	‘Small’	‘Sparse’	‘Small & sparse’	‘Ultra sparse’
Plating volume (μL)	20	5	20	5	20
Density of suspension (cells/ μL)	2500	2500	625	625	156
Number of cells plated (nominal)	50,000	12,500	12,500	3,125	3,125
Culture diameter (mm) ^a	4.9 ± 0.4	3.1 ± 0.3	4.9 ± 0.4	3.1 ± 0.3	4.9 ± 0.4
Drop thickness (mm) ^{a,b}	1.69 ± 0.24	1.06 ± 0.23	1.69 ± 0.24	1.06 ± 0.23	1.69 ± 0.24
Density at 1 div ($\times 10^3$ cells/ mm^2) ^c	2.5 ± 1.5	1.6 ± 0.6	0.60 ± 0.24	0.30 ± 0.16	0.11 ± 0.06
Number of cultures followed	30	12	10	3	3
Number of batches	8	3	3	1	1

^a Mean \pm sample standard deviation, $N = 3$ measured drops each of 5 and 20 μL .

^b Measured by focusing an inverted microscope on surface of MEA and top of drop (made visible by sprinkling some dust on it), and correcting for the refractive index of the liquid.

^c Mean \pm sample standard deviation, based on $N = 8, 8, 7, 3, 3$ cultures. Cells in the central 0.36 mm^2 of each MEA were counted using digital images.

We followed cultures of five different densities (Table 4.1). The physical development of cultures of various densities is illustrated by photographs taken at 1, 15, and 32 div in cultures of three different densities (Figure 4.2).

Burst Classification

The quantitative details of the development of different cultures varied widely. Indeed, cultures from different platings could show qualitatively different patterns during development. For instance, superbursts were observed in only about half of all cultures. Therefore, we made analogues of Figure 4.1 for each of the cultures studied available online. However, a more concise overview of the activity patterns observed in different cultures during their development was deemed highly desirable. To summarize the wide range of activity patterns in a single figure (Figure 4.3A), we classified them according to the following criteria:

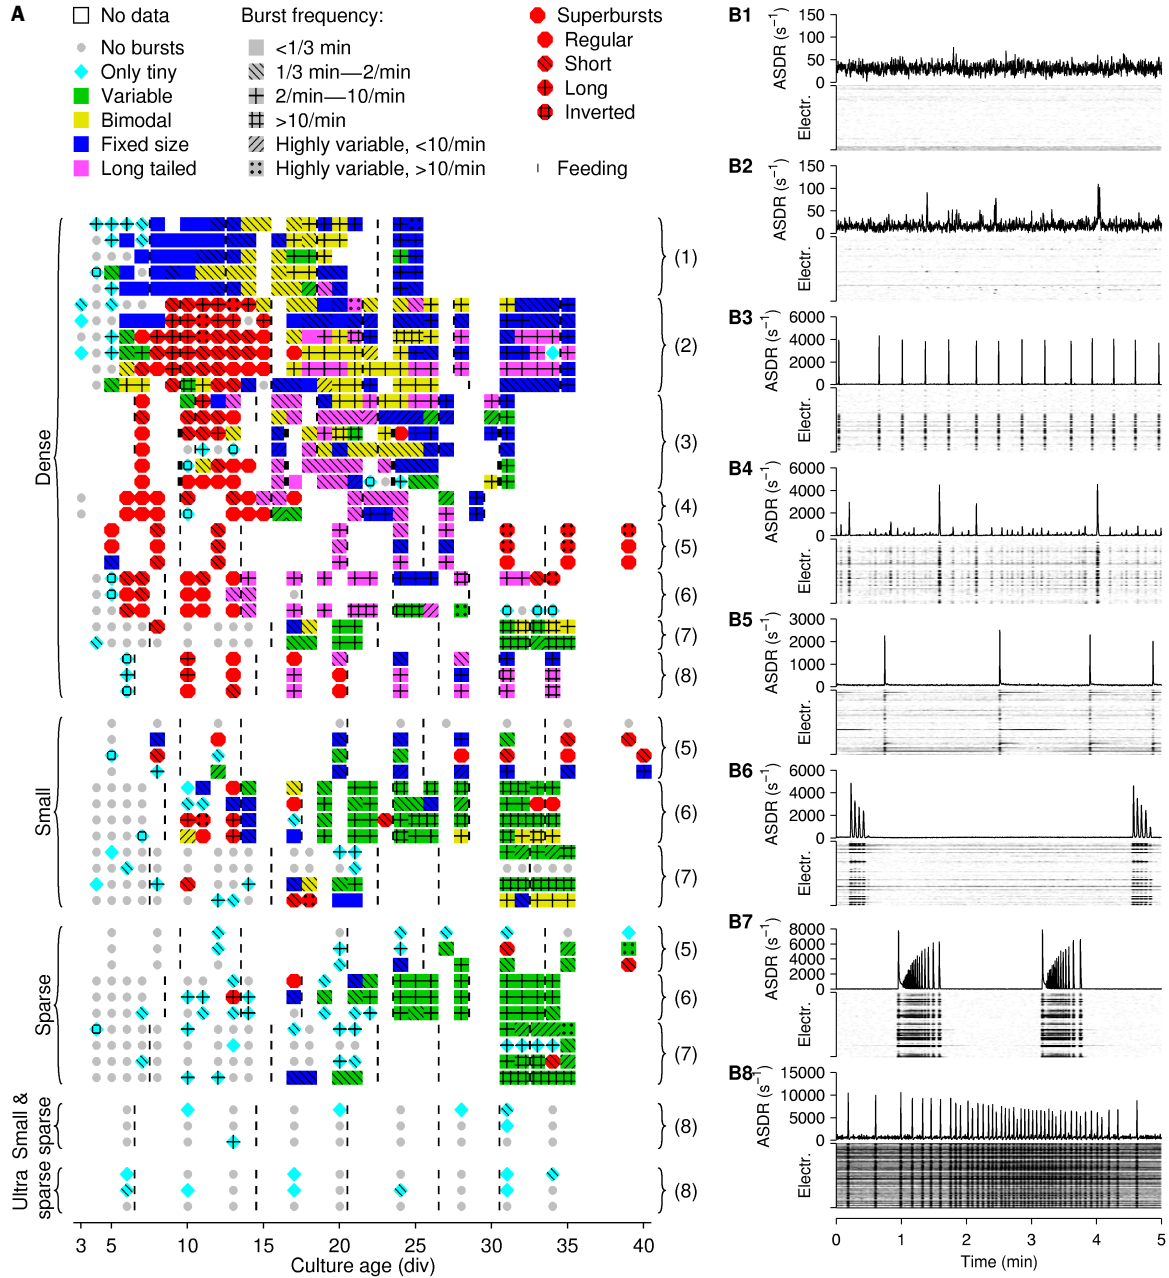


Figure 4.3: A Overview of the different classes of bursting behavior observed in our cultures. Numbers in parentheses indicate plating batch. Vertical bars indicate partial medium replacement times. Hash patterns indicate burst frequency for all types of burst patterns except superbursts. In batch 3, three cultures received full medium replacements (indicated by thicker bars). One culture in batch 6 got infected after 20 div, and had to be discarded. **B** Examples of burst pattern classes, with array-wide spike detection rates and gray-scale rasters for all electrodes, all taken from dense cultures. **B1** No bursting. **B2** Tiny bursts. **B3** Fixed size bursts. **B4** Variably sized bursts. **B5** Long-tailed bursts. **B6** Regular superbursts. **B7** Inverted superbursts. **B8** Dramatic burst rate variation.

Burstiness Any burst spanning fewer than 5 electrodes was termed *tiny*. If a recording contained only tiny bursts, it was considered *marginally bursty* (example in Figure 4.3B2). If it did not even contain any tiny bursts, it was considered *not bursty* (Figure 4.3B1).

Size distribution Let N^* be the number of spikes in the 3rd largest burst. Bursts with at least $\frac{3}{4}N^*$ spikes were termed *large*. Bursts with at least $\frac{1}{4}N^*$ spikes, but less than $\frac{3}{4}N^*$, were termed *medium*. Non-tiny bursts with fewer than $\frac{1}{4}N^*$ spikes were termed *small*. If there were more medium bursts than large bursts, the burst size was considered *variable* (Figure 4.3B4). Otherwise, if there were more small bursts than large bursts, the burst size distribution was considered *bimodal*. If there were more large bursts than medium or small bursts, the burst size was considered *fixed* (Figure 4.3B3).

Long-tailed bursts Non-tiny bursts with a ‘tail’ of at least 3 seconds during which the ASDR remained elevated by at least 50% above baseline levels were considered *long-tailed*. If at least half of all large and medium bursts in a recording were long-tailed, it was deemed dominated by long-tailed bursts (Figure 4.3B5).

Burst rates If the highest burst rate (determined from the shortest time span containing 10 inter-burst intervals) differed from the lowest burst rate (determined from the longest time span containing only 3 inter-burst intervals) by a factor 10 or more, the burst rate was considered *highly variable* (Figure 4.3B8).

Superbursts If at least 50% of all large and medium bursts occurred inside tight clusters, the recording was deemed to be dominated by *superbursts*. If the variance of the number of bursts per superburst was small (less than half of its average), the superbursts were considered *regular* (Figure 4.3B6). If not, they were considered *short* if the average number of bursts per superburst was less than 10, or *long* otherwise. If the number of spikes decayed in successive bursts inside superbursts, the superburst shape was considered normal. If it grew, the shape was considered *inverted* (Figure 4.3B7).

Array-wide synchronized bursting usually began after 5–7 div in dense cultures, and later in sparser cultures. Local bursts involving 1–5 electrodes were often ob-

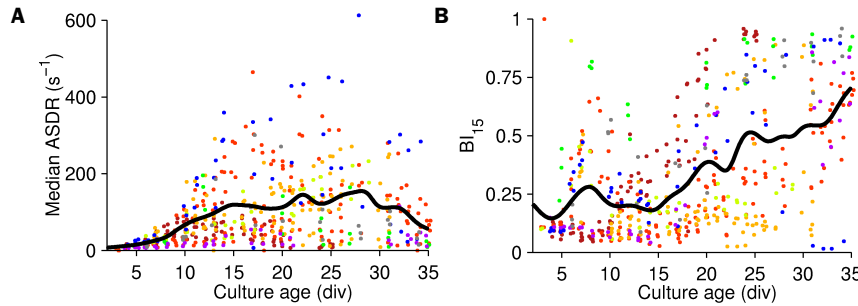


Figure 4.4: Development of activity in dense cultures. **A** Median ASDR across 1800 one-second-wide time bins. (Since bursts occupy a small fraction of time bins, a culture’s median ASDR over time is a good indication of the baseline ASDR outside of bursts.) **B** Burstiness index. Dots are measurements from individual cultures, colored by plating batch. Black lines are interpolated averages across all cultures, using a Gaussian window with a half-width of 1 day. Dots were horizontally jittered by ± 0.25 days for visual clarity.

served several days before global synchronization. Burst patterns changed with culture age, and these changes were still on-going after 30 div. Thus it does not appear that cultures were truly mature at this age, in contrast to a previous report (Kamioka et al., 1996). Figure 4.3A reveals that cultures from the same plating batch developed along strikingly parallel lines, while some large differences existed between batches both in terms of development speed and in terms of the type of burst patterns exhibited. It is further worth noting that the details of the medium replacement schedule appear to be unimportant, since the ‘feeding’ times (indicated in the figure by vertical thin black bars) did not coincide with marked developmental changes. Finally, the sparsest cultures we studied had very low burstiness, consistent with observations in low-density cultures used by other researchers for patch-clamp experiments in studies of synaptic plasticity.

Development of Firing Rates

The activity level (median ASDR) in dense cultures steadily increased during the first three weeks *in vitro*, then leveled off (Figure 4.4A). By contrast, the degree to which culture-wide bursts dominated the activity kept increasing (Figure 4.4B). If individual cells fire at a fixed rate, the ASDR of a sparse culture should be smaller than that of a dense culture, because in a sparse culture, there are fewer cells in the

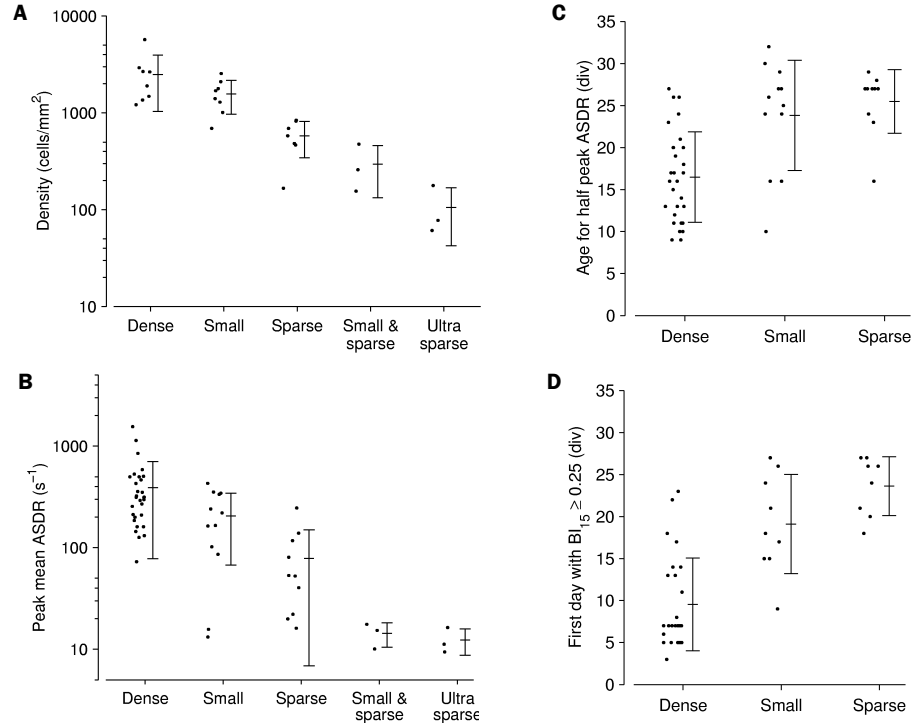


Figure 4.5: Comparison of the development of cultures of different sizes. **A** Actual density of cultures at 1 div. **B** Maximum ASDR observed in first 35 div. **C** First day on which ASDR reached half of its maximum. **D** First day with burstiness index greater than 0.25. Error bars indicate the mean and the sample standard deviation. Horizontal jittering of dots is for visual clarity only. Note that the vertical scale is logarithmic in A and B, which explains why the error bars appear asymmetric. In small-and-sparse and ultra-sparse cultures, the ASDR remained so low that the age at which half of the maximum was reached could not be measured accurately, and the BI never reached 0.25. Therefore, no data are shown.

vicinity of electrodes. The ASDRs in sparser cultures were indeed smaller than in denser cultures, commensurately with their smaller cell densities (Figure 4.5A–B). Beyond this expected result, we found that the development of sparse cultures was delayed compared to dense cultures, both in terms of how fast their ASDR increased (Figure 4.5C), and even more so in terms of their burstiness (Figure 4.5D). This observation cannot be explained by the fact that fewer cells are in contact with electrodes in sparser cultures, and instead indicates that cell density regulates the development of firing rates of individual cells (by a cellular or network-level mechanism).

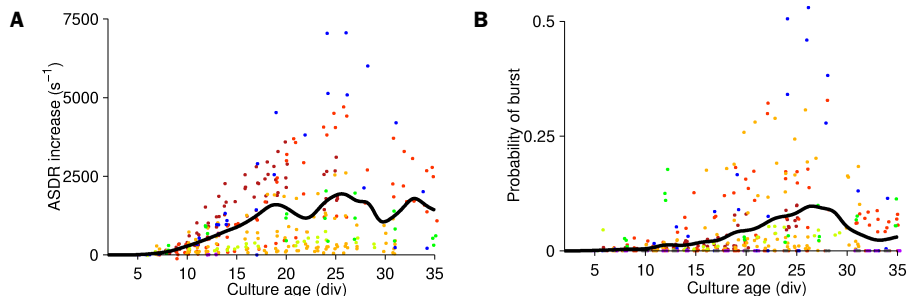


Figure 4.6: Strength of responses to stimulation during the development of dense cultures. **A** Increase in ASDR during the first 50 ms post-stimulus, averaged over all 50 presentations and 59 electrodes. Each dot represents a set of stimuli delivered to one culture on a given day, colored by plating batch and horizontally jittered for visual clarity. The line is the interpolated average across all cultures, using a Gaussian window with a half-width of 1 day. **B** Fraction of stimuli that evoked a burst.

Probe Responses

For plating batches 1–7, we recorded responses to probe pulses after each recording of spontaneous activity. Probe pulses were applied to each electrode in the array sequentially, with 0.3 s between pulses. The sequence was repeated 50 times. We found that responses to stimuli appeared later in development than spontaneous activity (compare Figure 4.6 to Figure 4.4). In cultures older than two weeks, stimulation often elicited bursts, but most cultures could not be driven to burst more frequently than once every few seconds. Interestingly, the probability that stimuli elicited bursts decreased after about four weeks *in vitro*, even as the burstiness in spontaneous activity increased, and while the immediate response to stimulus pulses remained constant.

Axon Outgrowth

We previously reported that a monopolar biphasic stimulus pulse on one electrode typically evokes very precisely timed responses on a number of other electrodes that are insensitive to synapse blockers (Wagenaar et al., 2004). We concluded that stimulation most likely evokes action potentials in axons, which then cause recordable action potentials elsewhere along the axon, or in the cell body by antidromic transmission. By detecting these ‘direct responses’ and measuring their distances from




Figure 4.7: Measuring axonal outgrowth in cultures of different densities. Median distance of non-synaptic responses to stimulated electrode (*solid, dots*) and 90th percentile of distance distribution (*dashed, circles*) in individual recordings from dense culture, small cultures, and sparse cultures (*left to right*). The diameter of the MEA (maximum electrode distance) is 1.72 mm. Axons likely continued to grow beyond this length, especially in the dense cultures, but our method is incapable of following that development. (Interpolation was performed with a window half-width of 1 day for panel A, and 2 days for panels B and C, to obtain a smooth curve for the smaller data sets.)

the stimulation site, one can investigate the distances reached by axonal projections in the culture in which the responses were recorded, up to the point where axonal projections reach beyond the extent of the electrode array.

We determined the distribution of distances, and normalized it by the number of (stimulation site, recording site)-pairs existing in the array for each distance. The median distance of direct responses to stimuli applied to all electrodes of a given culture provides a measure of the typical length of axons in that culture. The 90th percentile of the distribution is a robust lower bound for the maximal length of axons in a culture. We found that axons grow rapidly during the first week *in vitro* in dense cultures, reaching across the entire array within 15 days (Figure 4.7). Outgrowth was slower in small and sparse cultures, and the typical length axons ultimately attained was shorter.

Development of Burst Duration, Propagation Speed, and Size

We measured the time it took for bursts to spread across a dense culture (a burst's onset phase), as well as the time it takes for bursts to be extinguished (its offset phase) and the total duration of bursts over the course of development (Figure 4.8). In contrast to Habets et al. (1987), we found substantial developmental changes: The aver-

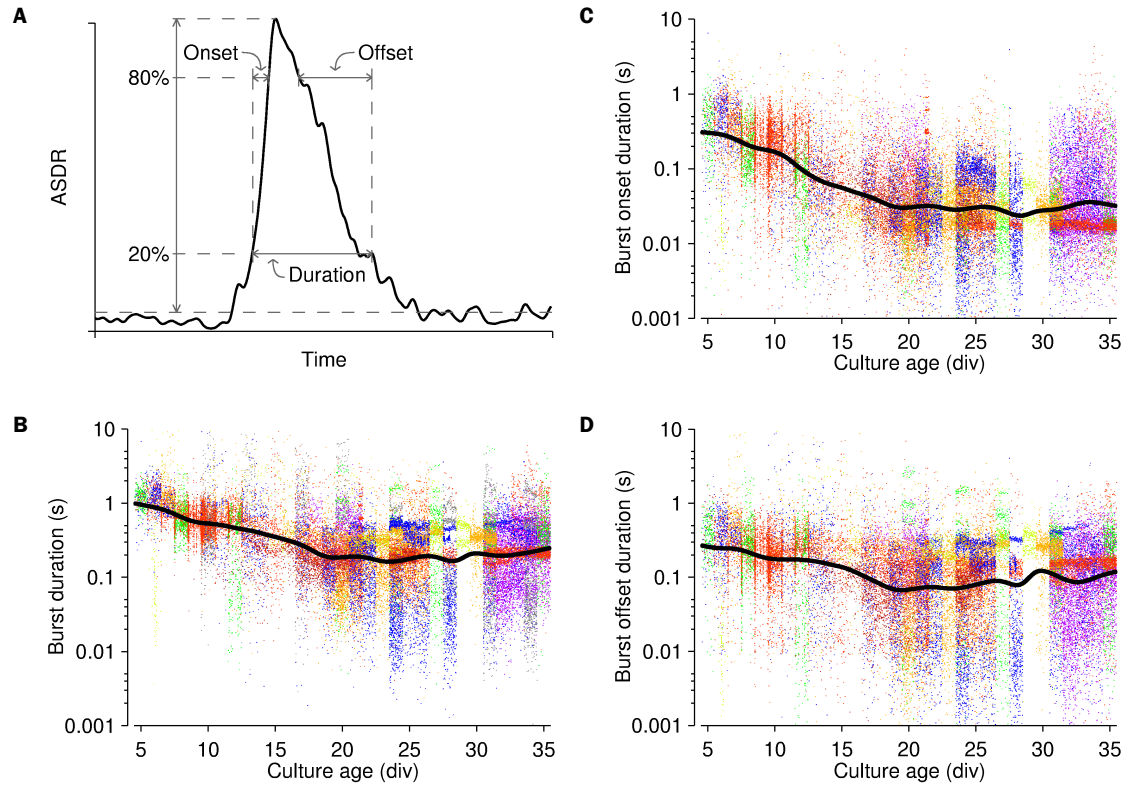


Figure 4.8: Characterizing burst shapes in dense cultures. **A** Parameters that define burst shape. We measured the ASDR as a function of time during the burst, smoothed with a 10 ms Gaussian filter. The 20% and 80% points between baseline and peak ASDR were determined and used to define the various phases. **B** Total duration. (Note log scale on y-axis.) **C** Duration of onset phase. **D** Duration of offset phase. Dots represent individual bursts, horizontally jittered for clarity, and colored by plating batch. Lines are interpolated averages, computed in log-space, using a Gaussian window with a half-width of 1 day.

age total burst duration decreased from 1 s when bursts first appeared, to less than 200 ms after 20 div (Figure 4.8B). Simultaneously, the burst onset phase decreased from 300 ms to only 30 ms (Figure 4.8C), while the burst offset phase changed much less on average (Figure 4.8D). Note that substantial differences existed in both the values and variances of burst parameters of different cultures.

Based on the observation that much axonal outgrowth occurs between 5 and 10 div in dense cultures, one might expect that the spatial extent of network bursts also develops in that time frame. Indeed, at 5 div, bursts were mostly small, spanning at most 20 electrodes, while three days later many bursts spanned the entire array (Figure 4.9). Remarkably, in mature cultures, the number of spikes in a burst was

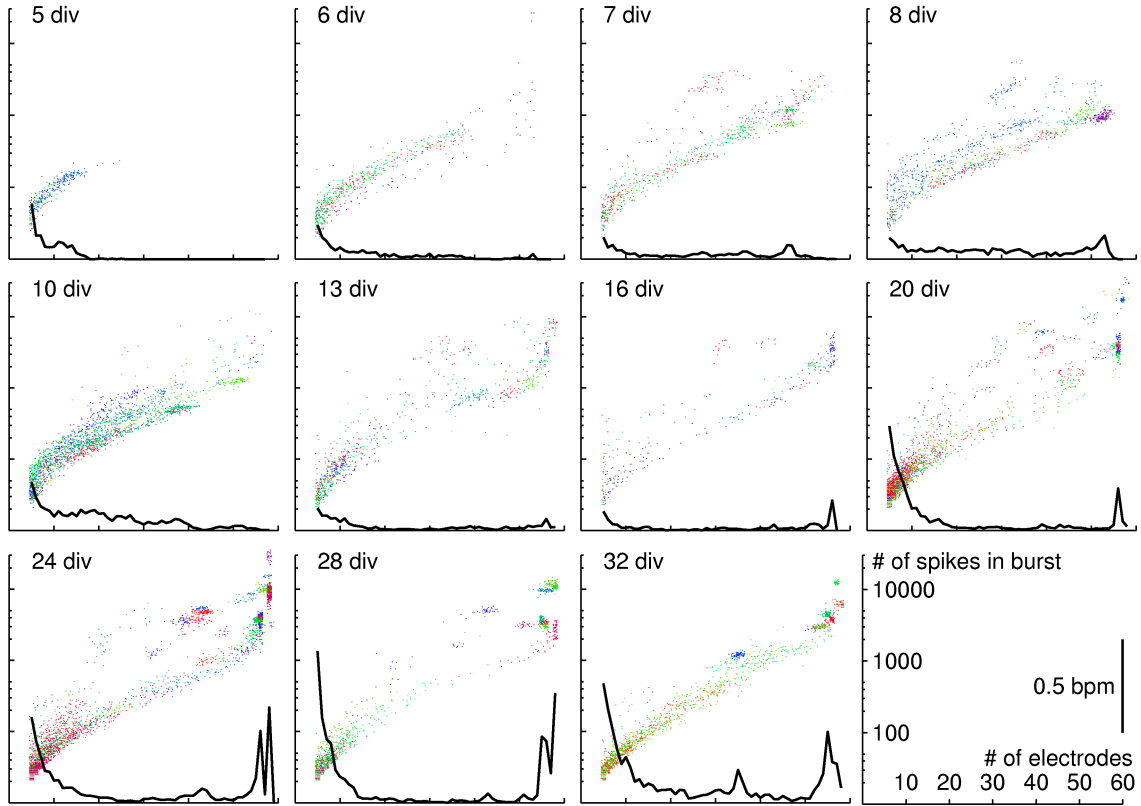


Figure 4.9: Comparison of burst sizes during culture development. Scatter plot of total number of spikes in burst and number of participating electrodes. Colors represent bursts from different (dense) cultures. Black traces are histograms of electrode counts, averaged across all cultures represented. Note log scale on y-axis.

well approximated by an exponential in the number of electrodes involved over much of the range. Least-squares fits for each of the ages shown in Figure 4.9 resulted in

$$\#spikes = (33 \pm 13)e^{\#electrodes/(12 \pm 2)},$$

where uncertainties represent differences between days.

Sensitivity to Movement

Of necessity, when recording from multiple cultures on a day, MEAs needed to be moved in and out of the recording device. We tested whether this mild mechanical perturbation had an effect on a culture's activity, and found that indeed it did. Young cultures (<20 div) often responded by firing a volley of bursts that lasted 1–2 min-

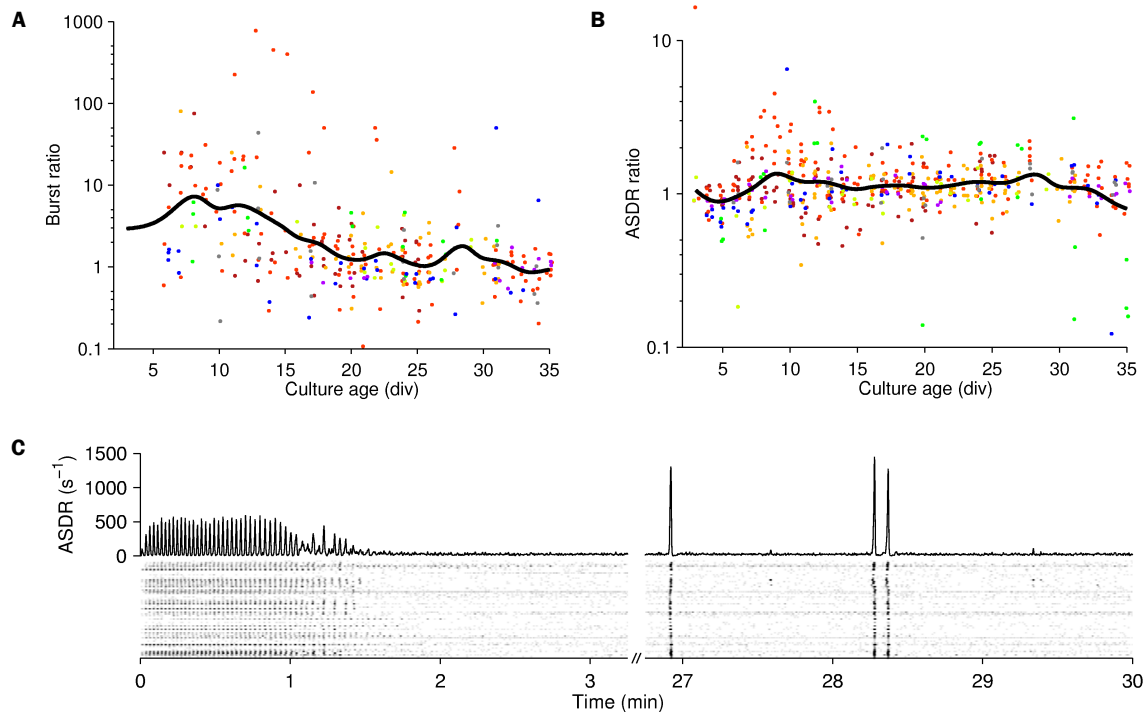


Figure 4.10: Sensitivity of spiking activity to mild mechanical perturbation (movement of the culture dish). **A** Number of bursts in first minute after culture was moved into the recording device, normalized to burst rate 10–30 minutes later. **B** Mean ASDR in first minute after culture was moved into the recording device, normalized to ASDR 10–30 minutes later. **C** Example of mechanically-induced bursting, recorded at 8 div from a dense culture. Lines are interpolations of the data, using a Gaussian window with a half-width of 1 day.

utes (Figure 4.10). Interestingly, the total rate of spikes fired in the first few minutes after moving an MEA was only slightly elevated. We confirmed that the effect was mechanically induced—rather than by subtle differences between the recording environment and the storage shelf, 30 cm lower in the same incubator—by lifting the recording device after 15 minutes of recording and putting it back down. This resulted in another volley of bursts (data not shown). Thus, the effect was not due to environmental influences such as a possible change in temperature. We also ruled out that exposure to light might cause the increased bursting, by moving cultures in total darkness and finding no difference between light and dark conditions (data not shown).

The observation that mechanical perturbation affects activity patterns led to the concern that perhaps recording only immediately after such a perturbation gives

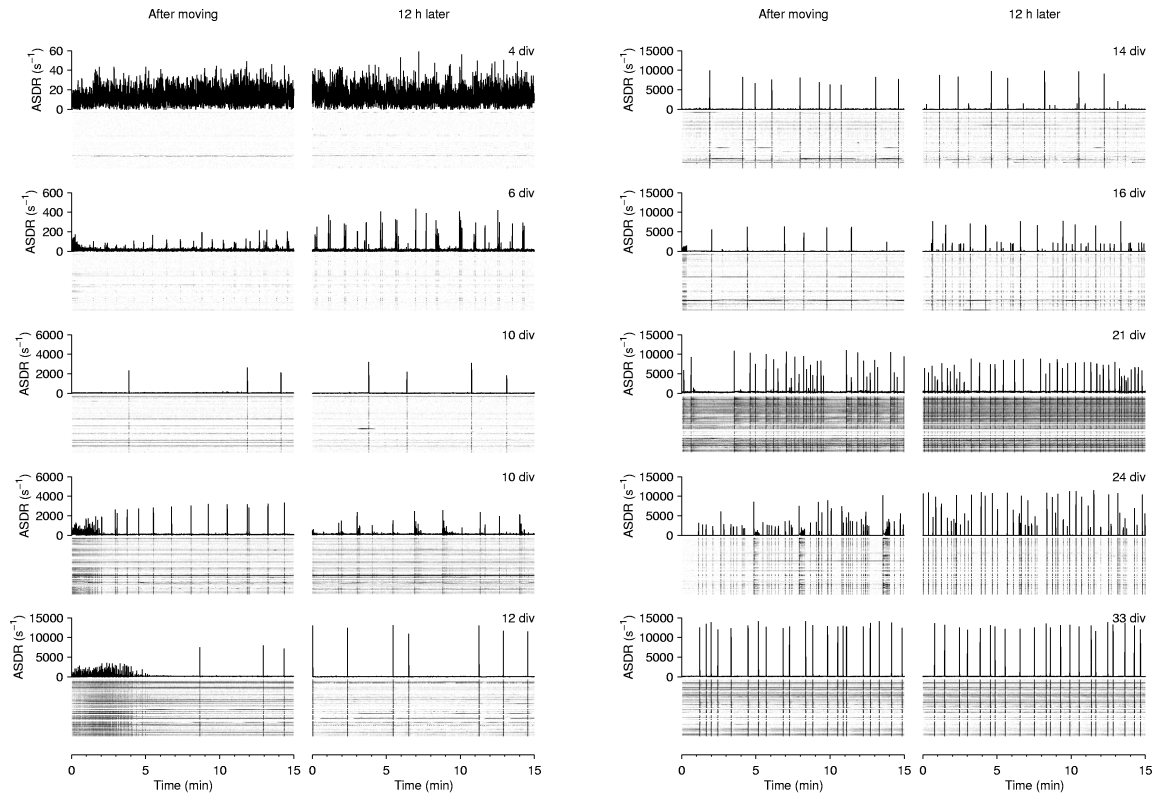


Figure 4.11: Comparison of activity in first 15 minutes after moving a culture into the recording device (*left-hand graphs*) with activity in the same culture around 12 hours later (*right-hand graphs*), in cultures of various ages *in vitro*. Examples are from different cultures.

a distorted view of the ‘typical’ activity of cultures at a given age. We compared the activity patterns immediately after moving a culture with the activity about 12 hours later in our overnight recordings, and found that substantial differences can indeed be observed in some young cultures (Figure 4.11). However, the effect was mostly limited to the first 5 minutes of recordings.

Sources of Variability

Figure 4.3 suggests that different cultures had some degree of maintained personality, that is, that differences between cultures exceeded the day-to-day differences in the behavior of individual cultures. Moreover, even batches seem to have distinguishing features. This observation was quantified using a difference index (DI) defined as

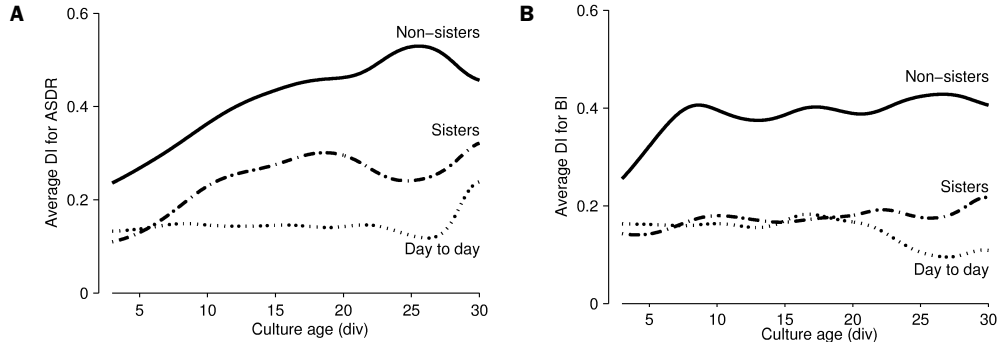


Figure 4.12: Quantifying sources of variability in activity levels. **A** Variability of median ASDR. **B** Variability of burstiness index. We compared the day-to-day variability for individual dense cultures with variability between sister cultures, and with variability between different platings. Interpolated using a Gaussian window with a half-width of 2 days.

follows:

$$DI_{ASDR}(f_1, f_2) = \frac{|f_1 - f_2|}{f_1 + f_2}.$$

Here, f_1 and f_2 are the mean ASDRs in the two recordings. This DI is normalized to lie between 0 (if the ASDRs are the same) and 1 (if one is much larger than the other). Analogously, we computed a DI of the burstiness index from each pair of recordings. We used DIs for a number of comparisons:

Day to day DIs computed between all possible pairs of recordings made from the same culture on consecutive days. The mean DI at a given age is a quantitative measure of day-to-day variability at that age.

Sister cultures DIs computed between all possible pairs of recordings made on the same day from cultures from the same plating batch.

Non-sister cultures DIs computed between all possible pairs of recordings made at the same developmental age from cultures from different plating batches.

This revealed that same-day differences between sister cultures were not much larger than the day-to-day differences between recordings from the same culture, and that the differences between cultures from different platings were substantially larger (Figure 4.12).

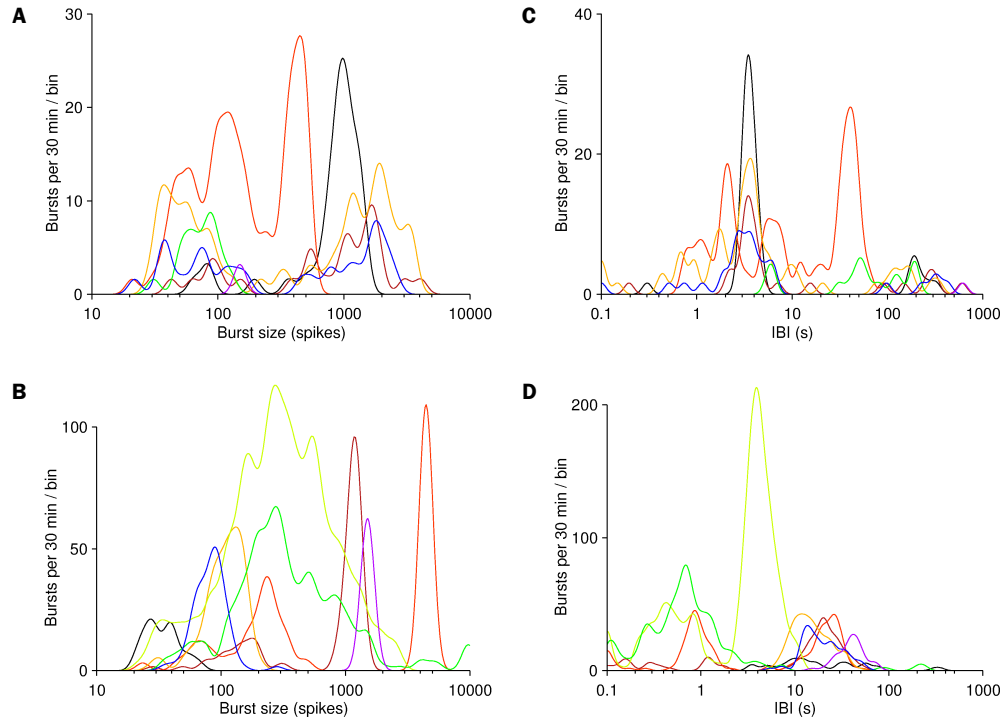


Figure 4.13: **A–B** Distribution of burst sizes. **C–D** Distribution of interburst intervals. Data are shown from 8 recordings from arbitrarily selected cultures (*different colors*), at around 7 div (**A** and **C**), and around 35 div (**B** and **D**). Bimodality in the IBI distribution results from temporal clustering of bursts. Data were binned using a Gaussian window in logarithmic space, bin size was 5% of a decade.

Discussion

The Nature of Bursts in Culture

It has previously been reported that cortical networks in organotypic culture exhibit burst patterns with sizes (number of constituent spikes) governed by a Lévy distribution (Beggs and Plenz, 2003). The dissociated cultures we followed showed very different burst patterns: in a given recording, there would be a clear distinction between tiny bursts and global bursts, the latter often having a relatively constant size (Figure 4.3B3 and Figure 4.13A–B). Furthermore, IBIs were quite narrowly distributed in most recordings (Figure 4.13C–D), in contrast to previous reports that they follow a scale-free distribution (Segev et al., 2002). Moreover, burst patterns often had rich temporal structure. Examples include superbusts as well as dramatic minute-scale changes in burst rate (Figure 4.3B6–8).

Importance of Sampling from Multiple Platings

Despite our best efforts to keep conditions stable, we found substantial differences between the development of cultures from different platings. These differences may have been due to characteristics preserved from the animals from which the cultures derived, or they may have originated later during development. The fact that cross-plating variability increased with age *in vitro* (Figure 4.12) supports the second possibility. Whatever the root cause, the observation that cross-plating variability was larger than variability between sister cultures implies that it is crucial to use cultures from several different platings to obtain unbiased results. We suggest it would be appropriate to report not only the number of cultures used, but also the number of platings from which these cultures stem, whenever dissociated cultures are used in network physiology experiments.

Outlook

By following the development of a larger number of cultures than any previous report based on MEA recordings, we have found that the range of behaviors exhibited by networks of cortical cells *in vitro* is much more complex than previously publications suggest. While the activity of all dense cultures became dominated by array-wide bursts as the cultures matured, the sizes, shapes, and temporal patterns of these bursts varied widely. Indeed, the range of behaviors of these cultures is so rich that it is tempting to spend countless hours analyzing individual cases. Constraints on both space and time prevent us from exploring such avenues in this report, but we intend to continue the analysis of this data set in the future. Moreover, we invite others to join us in the study of activity patterns of networks of cortical cells *in vitro*. To this end, we have made the entire data set used for this paper publicly available on the web at <http://neuro.gatech.edu/~potter/public-data>. Researchers may download our recordings of spike waveforms (a total of 45 GB, compressed), or reduced files containing only time stamps and electrode IDs (a total of 4 GB). Example Matlab code to efficiently access the files, and documentation are available as well.

5

Superbursts: Persistent Dynamic Attractors in Activity Patterns*

The potential to generate complex spatio-temporal patterns of neuronal activity is believed to underlie diverse functions in the nervous system, from stereotypical motor behavior to sensory persistence. Theory suggests that complex spatio-temporal patterns can result from dynamic attractors emerging spontaneously from a sufficiently strongly interconnected network. We report that population-wide bursts in cultured cortical networks spontaneously cluster into groups of up to twelve members, with highly stereotyped spatio-temporal structure. During these ‘superbursts,’ the firing sequence of the culture periodically converges to a dynamic attractor orbit. Lasting up to 30 seconds, superbursts are among the longest stereotyped patterns reported in any neural system to date.

Introduction

In models of neural networks, attractor dynamics displaying complex reverberations emerge naturally if there are sufficient feedback connections (Wilson and Cowan, 1973; Amit, 1989; Wang, 2001). Donald Hebb proposed that such reverberations may be used to encode and maintain information in the nervous system (Hebb, 1949).

* To be submitted as: Daniel A. Wagenaar, Zoltan Nadasdy, and Steve M. Potter: Superbursts: Persistent dynamic attractors in activity patterns of cultured neuronal networks.

Recurring short spatio-temporal patterns of action potentials recovered from simultaneous recordings of multiple neurons *in vivo* and *in vitro*, variously called ‘sequences’ (Prut et al., 1998; Nádasdy et al., 1999; Ikegaya et al., 2004), or ‘synfire chains’ (Abeles, 1989), may be subsamples of such dynamics. Evidence for recurring spatiotemporally complex activity patterns has recently begun to emerge in sensory systems (Stopfer et al., 2003) as well as motor systems (Leonardo and Fee, 2005). We studied the emergence of dynamic attractors in dissociated cultures of cortical neurons, by recording their spontaneous activity using microelectrodes arrays for several weeks *in vitro*. The electric activity of such cultures is dominated by culture-wide bursts of high frequency action potential firing, separated by periods of low firing rates (Kamioka et al., 1996; Segev et al., 2002; Van Pelt et al., 2004a). Bursting in culture is reminiscent of bursting observed in the developing cortex (Chiu and Weliky, 2001) and LGN (Weliky, 1999) *in vivo*, as well as of sleep spindles in the thalamic reticular nuclei (Contreras et al., 1997) and subthalamic nucleus during slow wave activity (Magill et al., 2004). We find that burst patterns can have a precisely defined spatiotemporal structure that recurs with great fidelity over an interval of many hours. This shows that dissociated cortical networks in culture are capable of generating complex stereotypical behaviors that were previously believed to require specific network architecture.

Methods

Dense cultures of rat cortex were prepared on multi-electrode arrays (MEAs) as described before (Wagenaar et al., 2005b [Chapter 6 in this thesis]). Briefly, cortices from E18 rat embryos were dissected and dissociated using papain and trituration. Cells—neurons and glia—were plated at a density of 2500/mm², on MEAs coated with poly-ethylene-imine (PEI) and laminin. Cultures were maintained in a serum-containing DMEM-based medium. We recorded daily from 27 cultures from day 3 to day 35 *in vitro*. Five cultures were followed for 2–3 days continuously.

Spike Detection and Sorting

Electrical signals from 59 electrodes were sampled at 25 kHz. Putative spikes were detected by thresholding the electrode traces at $4.5\times$ estimated RMS noise. Spikes were then validated based on their shapes (P. P. Mitra, personal communication): the detected peak must be the highest peak of either polarity within a ± 0.5 ms window, and no secondary peaks of the same polarity and more than 50% of the amplitude of the detected peak must exist within the same window. This prevents double detections of multi-phasic spikes, as well as many false positives due to noise.

Most subsequent analysis was performed using multi-unit activity (MUA), obtained from 59 electrodes in a square grid with 200 μm spacing. In two cultures, we analyzed single-unit activity (SUA). SUA was obtained by using an unsupervised spike sorting method (Quiroga et al., 2004) with cross- and autocorrelation verification. Only the largest four spike clusters per electrode (i.e., four neurons with highest firing rates) were included in the data analysis. To ensure stability of clusters over time, the sorting was done in data segments of 400 s with 40 s overlaps. The redundant clustering on the overlaps allowed us to match spike clusters consistently across segments. Spike clusters were tested for refractoriness. In both cultures the sorting resulted in 236 putative neurons (59 electrodes \times 4 clusters).

Cross-correlation analysis revealed that inter-electrode spacing was such that cells did not evoke potentials on more than one electrode. This also implied that using multi-unit data does not compromise the spatial resolution of the analysis.

Burst Identification

Bursts were detected using the SIMMUX algorithm (Wagenaar et al., 2005a). Briefly, each electrode trace was searched for *burstlets*: sequences of at least four spikes with all inter-spike intervals less than a threshold (set to $1/4$ of that electrode's inverse average spike detection rate). Any group of burstlets across several electrodes that overlapped in time was considered a *burst*.

Similarity Indices

A global similarity index (GSI) was computed based on the (multi-unit) firing rate summed over all electrodes. For each superburst n , we computed this firing rate, $f_n(t)$, in 50 ms (Gaussian) sliding windows (sampled at 500 Hz). $f_n(t)$ was set to zero for $t < 0$ or $t > (\text{the duration of superburst } n)$. The similarity index $GSI(n, m)$ between two superbursts n and m was then defined as the correlation coefficient between the functions f_n and f_m , optimally time-shifted:

$$GSI(n, m) = \max_{\tau} \left(\frac{\int (f_n(t) - \bar{f}_n) (f_m(t + \tau) - \bar{f}_m) dt}{\sqrt{\int (f_n(t) - \bar{f}_n)^2 dt} \sqrt{\int (f_m(t) - \bar{f}_m)^2 dt}} \right),$$

where \bar{f}_n is the average of $f_n(t)$ over the duration of the superburst.

A subburst similarity index (SSI) was based on the (MUA) burst onset times. To determine the onset time of a neuron or an electrode, we measured both its baseline firing rate, and its peak increase above this baseline during a given burst. The onset time $t_{\text{on}}^c(n, k)$ of electrode or cell c in the k^{th} subburst of the n^{th} superburst was defined as the moment when the firing rate first increased to 25% of the peak increase. (This use of relative thresholds ensured that differences in firing rates between electrodes (or cells) did not cause a systematic bias in onset time estimation. We tested the independence of onset time estimates and firing rates by calculating the Pearson correlation coefficient, and found it was negligible ($r = -0.05$, $p = 0.13$; $N = 845$.) The similarity index $SSI(n_1, k_1; n_2, k_2)$ between two subbursts (n_1, k_1) and (n_2, k_2) was then defined as the correlation coefficient between onset times across electrodes (or cells):

$$SSI(n_1, k_1; n_2, k_2) = \frac{\sum_c \left(t_{\text{on}}^c(n_1, k_1) - \bar{t}_{\text{on}}(n_1, k_1) \right) \left(t_{\text{on}}^c(n_2, k_2) - \bar{t}_{\text{on}}(n_2, k_2) \right)}{\sqrt{\sum_c \left(t_{\text{on}}^c(n_1, k_1) - \bar{t}_{\text{on}}(n_1, k_1) \right)^2} \sqrt{\sum_c \left(t_{\text{on}}^c(n_2, k_2) - \bar{t}_{\text{on}}(n_2, k_2) \right)^2}},$$

where $\bar{t}_{\text{on}}(n_i, k_i)$ is the mean onset time of the k_i^{th} subburst of the n_i^{th} superburst across electrodes. Only electrodes with peak firing rates of at least 75 spikes per second were used in this calculation (typically: 40 electrodes).

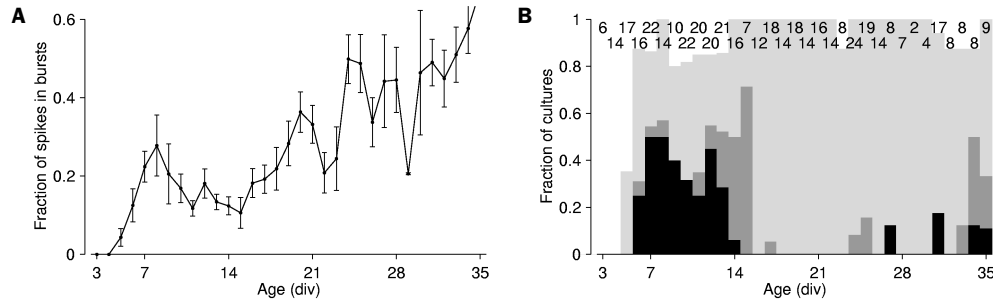


Figure 5.1: Development of bursting. **A** The fraction of spikes that occur in large bursts (rather than during tonic dispersed firing) grows with culture age. Here, ‘large’ means at least 5 participating sites with a total of at least 50 spikes. **B** Fraction of cultures that fire superbursts exclusively (*black*) or superbursts mixed with other bursts (*dark gray*). Light gray indicates fraction of cultures that exhibits any kind of bursts. Numbers on top indicate number of cultures studied at each age.

Results

Bursting in dissociated cultures commenced after 5–8 days *in vitro* (Figure 5.1), and persisted throughout a culture’s lifetime (over one year (Potter and DeMarse, 2001)). During most of a culture’s life, burst patterns were relatively unstructured. Burst frequencies ranged from 1 to 30 per minute, and appeared to be generated by a Poisson-like process modulated by a refractory period of 1 to 5 seconds. However, a small majority of cultures passed through a developmental period lasting 3–5 days during the 2nd week *in vitro*, during which burst patterns acquired a large degree of structure.

During such epochs, bursts clustered into sequences of 5 to 12, with inter-burst intervals of 1 to 2 seconds (Figure 5.2A).^{*} These sequences, which we call ‘superbursts,’ were separated by 1 to 10 minutes without bursts of any kind. The intervals between superbursts were consistent with a Poisson process modulated by refractoriness (Figure 5.2B). In contrast, the intervals between the constituent bursts (‘subbursts’) within superbursts were highly stereotyped (Figure 5.2C). The number of bursts per superburst was likewise strongly conserved over long periods of time, though it varied considerably from culture to culture. Typically, the first burst in a su-

^{*} A movie showing array-wide activity during a superburst is available online at <http://www.its.caltech.edu/~pinelab/wagenaar/superbursts>.

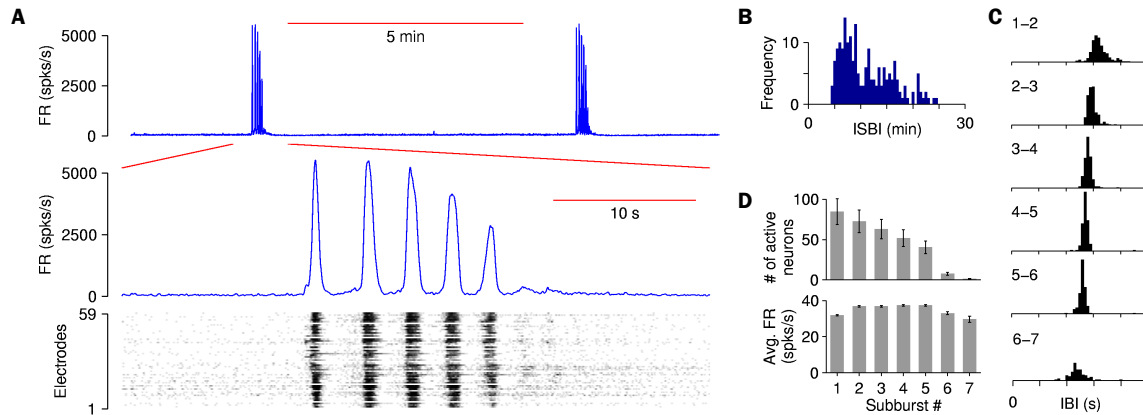


Figure 5.2: **A** An example of a 10 minute data segment illustrates the typical two-level temporal organization of population activity: ‘superbursts’. Firing rates (FR) are culture-wide aggregates. Simultaneous raster plots from 59 electrodes reveal that nearly all electrodes record from neurons participating in this structure. **B** The distribution of intervals between 195 superbursts recorded over a 35 hour period. **C** The distributions of the intervals between subbursts within superbursts. Histograms show all subburst intervals at a fixed ordinal position (indicated on top-left) in their superbursts. **D** Number of active neurons (*top*) and average firing rate per active neuron (*bottom*), per subburst. Spike sorting was performed using super-paramagnetic clustering (Quiroga et al., 2004).

perburst contained the largest number of spikes, followed by a gradual decline. This was due to a reduction in the number of participating neurons, with single-neuron firing rates remaining nearly constant during most of the superburst (Figure 5.2D). In Figure 5.2 as well as in the following, we concentrate on results obtained from the longest recorded superburst epoch (63 hours). Results from all extended recordings are summarized in Table 5.1.

To quantify the conservation of firing rate profiles across superbursts, we measured the array-wide aggregate of the firing rate in 50 ms sliding windows. This yielded a firing rate profile for each superburst (Figure 5.3A). We defined a ‘global similarity index’ (GSI) between a pair of superbursts as the correlation coefficient between their profiles (aligned to maximize GSI, but not time-warped; see *Methods*). The GSI between consecutive superbursts was very high (>90% on average), and remained high (>80%) between pairs of superbursts separated by dozens of other superbursts (Figure 5.3C). The GSI matrix of our longest recording (Figure 5.3B) is characterized by a block-diagonal structure, indicating that changes in superburst

Table 5.1: Compendium of parameters for all five extended recordings.

	Culture #1	#2	#3	#4	#5
Age (div)	10	19	9	8	12
Duration of superbursting (hr)	63	11 ^a	41	49	3 ^a
Number of superbursts	292	49	94	154	24
Average number of subbursts	7	12	3	7	5
Global similarity index (see text)	0.89	0.81	0.73	0.76	0.81

^a These cultures were still superbursting when the recording was terminated.

shape occurred in discrete steps of varying size.

Like the global activity profile, the way the activity spread across the culture was also preserved within and across superbursts. We quantified this by taking the relative times at which individual electrodes started to record each subburst, and combining those into a (59-dimensional) vector, which constitutes an ‘onset-time profile’ for the subburst. We defined a ‘subburst similarity index’ (SSI) as the correlation coefficient between pairs of such vectors. This revealed considerable similarity between subbursts within a superburst (Figure 5.3D), particularly between the 2nd to 5th subbursts. Moreover, homologous (like-numbered) subbursts had very similar onset profiles between consecutive superbursts (Figure 5.3E). Between the 2nd to 5th subbursts, this ‘between-superburst’ SSI exceeded the ‘within-superburst’ SSI. Comparing pairs of superbursts with longer intervals between them, the SSI between the 1st subbursts was much reduced, indicating a gradual change of the state of the network. In striking contrast, the SSI between the 2nd to 5th subbursts remained high, suggesting that after the first bursts of each superburst, the network dynamics locked into an attractor orbit.

After spike sorting (Quiroga et al., 2004), we further analyzed conservation of burst onset profiles at the cellular level using return plots. Return plots elucidate higher order temporal relationship between successive events, by recursively plotting the latency of the n^{th} event against the latency of the $n+1^{\text{st}}$. The appearance of clusters in return plots indicates a conserved temporal pattern in successive events, and the spread of clusters reflects the precision of conservation.

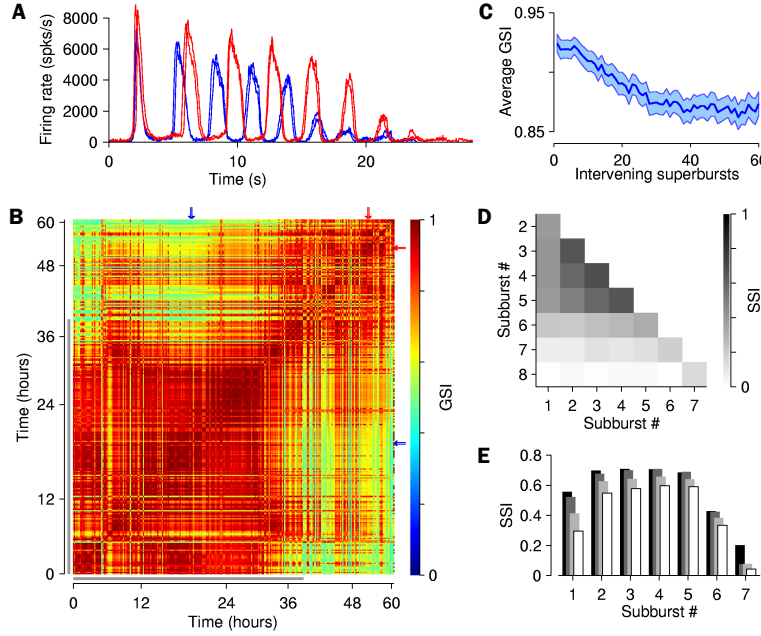


Figure 5.3: Conservation of firing rates and activity propagation between superbursts. **A** Aggregate firing rates of two pairs (one pair *red*, one pair *blue*) of consecutive superbursts. In the 30 hours between the two pairs, 150 other superbursts occurred (not shown here). Consecutive superbursts are seen to be almost indistinguishable in shape. **B** The GSI matrix between superbursts recorded over a 63 hour period shows a fractal-like structure. Two main blocks can be distinguished. Blue and red arrows mark examples shown in A. Gray bars mark portion of data used in C–E. **C** Even with dozens of intervening superbursts, the GSI between superbursts separated by many hours remains very high. (Mean \pm SEM for 160 superbursts.) **D** SSI between bursts within a superburst, averaged over 170 superbursts. A conserved structure is observed between subbursts 2–5. **E** SSI between homologous subbursts across superbursts. Between consecutive superbursts (*black*), the 2nd through 5th subbursts are more conserved than the 1st subburst. Between superbursts 30–60 minutes apart (*dark gray*), 1–2 hours apart (*light gray*) or 6–24 hours apart (*white*), this effect is even more pronounced.

We performed return plot analysis on the onset latencies of individual neurons in successive bursts, defined as $\lambda_{n,k}^c \equiv t_{\text{on}}^c(n, k) - \overline{t_{\text{on}}}(n, k)$. After spike sorting, we repeated this analysis at the level of single neurons. We constructed return plots both between consecutive subbursts within a superburst, and between homologous subbursts across superbursts. If individual neurons play conserved roles in different bursts, their relative burst onset latencies should be conserved from burst to burst, causing the latencies to line up along the diagonal of the return plot. Moreover, the latencies of an individual cell should cluster in a confined region along the diago-

nal. Both effects are indeed evident in Figure 5.4. The relative latencies of different electrodes were consistent across successive component bursts ($r = 0.55$; $p < 0.01$; Figure 5.4A). Individual neurons engage in the successive bursts with even more precise latency relative to other neurons ($r = 0.58$; $p < 0.01$; Figure 5.4B). When comparing correlation coefficients in return plots at electrode level with those at single cell level, we balanced the sample size by random sub-sampling the population spikes. Thus we obtained an unbiased estimate of reproduction fidelity.

A neuron that started bursting earlier than the population would always be earlier than a neuron that started bursting later (Figure 5.4C). The gross conservation of latencies was complemented by a systematic drift in the onset latencies for a given neuron across successive component bursts (Figure 5.4D). Relative onset latencies of different neurons were also strongly preserved between homologous components of consecutive superbursts ($r = 0.57$; $p < 0.01$; Figure 5.4E). The observation that the latency profile is consistent across subbursts and superbursts implies that the transition from tonic to burst-firing propagates across the culture following a similar path each time. Since the burst onset order at cellular level was slightly more consistent than at electrode level ($r = 0.58$ vs. $r = 0.52$), we concluded that this path must be dependent on the transmission between individual neurons. Given that this difference was small, we mostly used electrode-level dynamics for the subsequent analysis, since that level allowed for higher precision (due to larger spike counts).

After burst onset, the subsequent firing rate dynamics were also conserved. We visualized the temporal aspect of superburst dynamics with a phase plot of the aggregate firing rate (in 100 ms sliding windows) during superbursts (Figure 5.5A). The tightness of the bundles of the orbits of the 2nd through 5th subbursts demonstrates the conservation of the firing rate structure of these subbursts. The orbits of the 1st subbursts show considerably more variation, re-affirming that the initial stages of the superbursts were affected by gradual drift in the state of the culture, but that the superburst self-organized into an activity trajectory—or dynamic attractor—which was stable for hours or days.

To gain more insight into the spatial structure of superburst dynamics, we defined the horizontal differential firing rate of a culture as the aggregate firing rate in the right half of the array minus the aggregate firing rate in the left half of the array

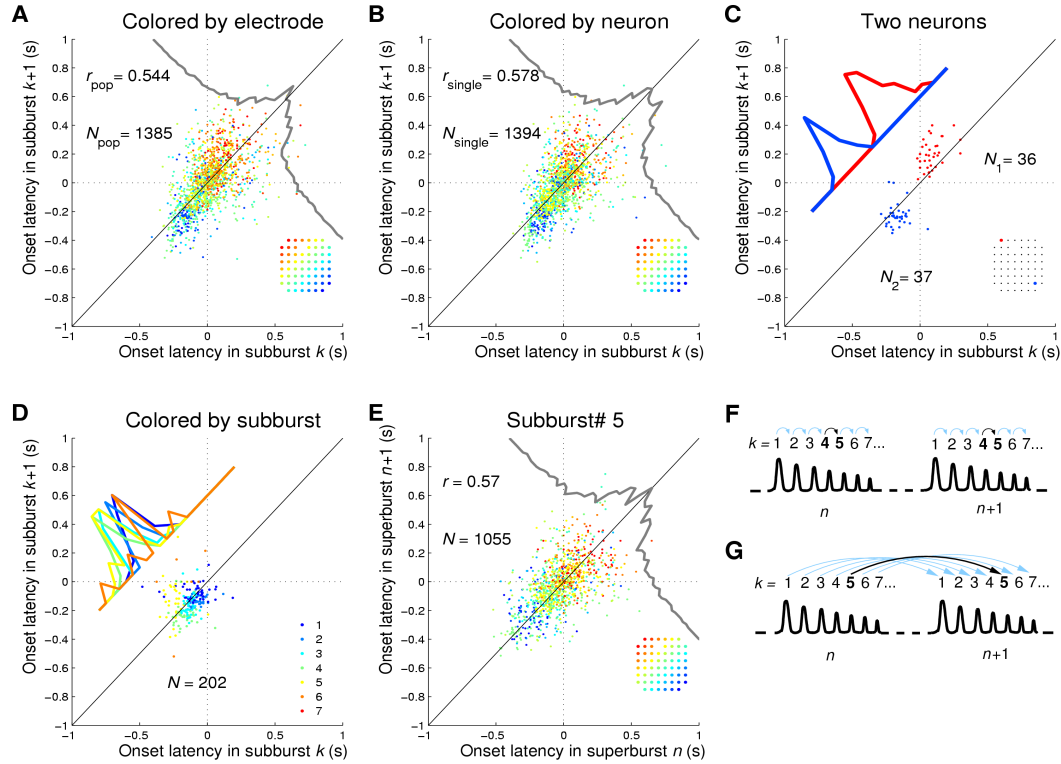


Figure 5.4: Return plots representing the temporal structure of burst propagation by recursively plotting the latencies at which a cell or electrode starts to participate in one burst against its latency in the next (or next homologous) burst. **A** Electrode level return plot of burst onset latency from the 4th to the 5th subburst in successive superburs. The diagonal represents exact latency preservation. Electrodes are color-coded according to the inset. **B** The same return plot as in **A**, but instead of combining all spikes from a given electrode, we isolated the most active single unit from each electrode. Color code as in **A**. Note that the correlation is slightly better than that of the electrode level. **C** The burst onset latency return plot for two neurons, extracted from **B**. One neuron (*blue*) was selected that tended to burst early, and one (*red*) that tended to burst late. Inset shows the locations of the two neurons. **D** Single-neuron return plots of burst latency across different subburs, for the blue neuron in **C**. Interval number is color-coded. **E** Single neuron level return between the 5th component-bursts across successive superburs. **F–G** Explanation of return plots: **F** Return plots for successive component bursts within a superburs (for **A–D**) and **G** Return plots for homologous subburs across successive superburs (for **E**).

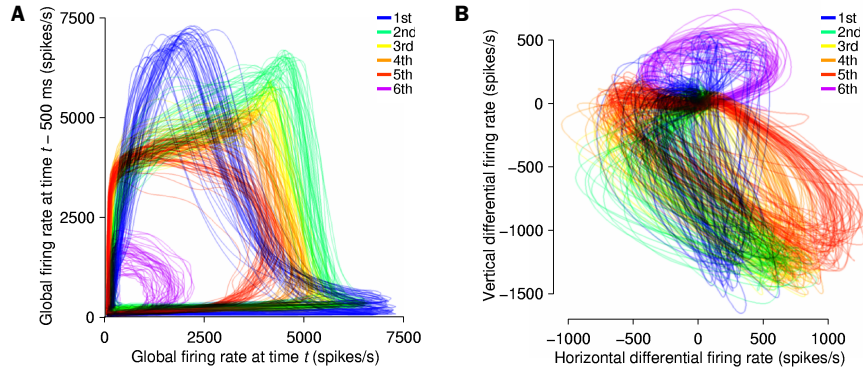


Figure 5.5: Spatiotemporal structure of superbumps. **A** Phase plots of the aggregate firing rate for 50 consecutive superbumps. Obtained using a 100 ms sliding window. Differently numbered subbumps are drawn in different colors. The delay, 500 ms, is short enough to not mix subbumps. **B** X-Y plot of differential firing rate (see text) for the same 50 superbumps.

(in 200 ms sliding windows). A vertical differential firing rate was similarly defined. The orbits of superbumps in the state space of differential firing rates show that the preservation of burst shape increases from the 1st to 5th subburst (Figure 5.5B). The 6th subbursts, which mark the end of the superbump structure, have orbits of distinctly different shapes than the earlier subbursts.

Discussion

In different cultures, the precise locations of neurons and their connectivity varied considerably. As a result, the number of subbursts per superbump and the spatiotemporal propagation varied between cultures. Nevertheless, self-organization of activity patterns into a two-level structure of subbursts and superbumps was consistently observed (Table 5.1). While superbumps appeared at irregular intervals, their internal structure was highly regular and strongly conserved for hours or days: once a superbump had been initiated, it generated a constant number of subbursts that each had its own well-preserved geometry of propagation and temporal dynamics. This preservation was found to be precise at the single-neuron level. The spontaneous occurrence of superbumps shows that neurons and glia retain an ability to self-organize into multicellular ensembles with non-trivial functional structure, even

when taken out of their physiological context.

In vivo, short conserved patterns of action potentials have been described in several preparations (Prut et al., 1998; Nádasdy et al., 1999; Ikegaya et al., 2004), and fixed-point attractors have been observed in the form of up/down state transitions (Cossart et al., 2003). Superbursts constitute much more extensive patterns, and are among the longest conserved activity patterns observed in any neural system to date. In *in vivo* experiments, such patterns may have remained hidden because only a small fraction of the neurons from a large ensemble were monitored, or because recordings were too short. By contrast, the use of dissociated cultures permitted us to monitor and evenly sample an entire intact network for weeks.

The coordination of cellular dynamics in dissociated culture at the superburst level indicates that information is maintained by a dynamic renewal process which persists orders of magnitude longer than the time constants of synaptic processing. This allows the information to be protected from the interference of local information processing: individual neurons can engage in multiple functions without disrupting the recurring motif reverberating in the larger scale circuitry of the culture. Such globally organized and tightly orchestrated activity is of critical importance for neuronal tissue that generates any kind of highly stereotyped sequential behaviors, from locomotion to language. The same mechanism may also support a sensory persistence and memory that does not require synaptic plasticity. *In vitro* systems are ideal for studying in detail the conditions that allow such activity patterns to emerge.

6

Controlling Bursting with Multi-Electrode Stimulation*

One of the major modes of activity of high-density cultures of dissociated neurons is globally synchronized bursting. Unlike *in vivo*, neuronal ensembles in culture maintain activity patterns dominated by global bursts for the lifetime of the culture (up to two years). We hypothesize that persistence of bursting is due to a lack of input from other brain areas. To study this hypothesis, we grew small but dense monolayer cultures of cortical neurons and glia from rat embryos on multi-electrode arrays (MEAs), and used electrical stimulation to substitute for afferents. We quantified the burstiness of the cultures' firing in spontaneous activity and during several stimulation protocols. While slow stimulation through individual electrodes increased burstiness due to burst entrainment, rapid stimulation reduced burstiness. Distributing stimuli across several electrodes, as well as continuously fine-tuning stimulus strength with closed-loop feedback, greatly enhanced burst control. We conclude that externally applied electrical stimulation can substitute for natural inputs to cortical neuronal ensembles in transforming burst-dominated activity to dispersed spiking, more reminiscent of the awake cortex *in vivo*. This non-pharmacological method of controlling bursts will be a critical tool for exploring the information processing capacities of neuronal ensembles *in vitro*, and has potential applications for the treatment of epilepsy.

* Published as: Daniel A. Wagenaar, Radhika Madhavan, Jerome Pine, and Steve M. Potter, 2005: Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *J. Neurosci.* **25**(3), pp. 680–688. © Society for Neuroscience, 2005. Reproduced with permission.

Introduction

The mammalian cortex has been studied *in vitro* in the form of dissociated monolayer cultures for several decades. Such cultures retain many morphological, pharmacological and electrical properties of cortical networks *in vivo* (Dichter, 1978) and allow much more detailed observation and manipulation than intact brains, at the molecular, cellular, and network levels (Droge et al., 1986; Emery et al., 1991; Curtis et al., 1992; Wilkinson, 1993; Bove et al., 1994; Rhoades et al., 1996; Bove et al., 1997; Canepari et al., 1997; Gross et al., 1997; Liu et al., 1997; Harsch et al., 1998; Honma et al., 1998; Jimbo et al., 1998, 1999; Turrigiano, 1999; Harsch and Robinson, 2000; Zhu et al., 2000; Keefer et al., 2001; Streit et al., 2001; Shahaf and Marom, 2001; Corner et al., 2002).

The most prominent feature of the electrical activity of high-density dissociated cortical cultures is their propensity for synchronized bursting (Murphy et al., 1992; Gross et al., 1993; Wong et al., 1993; Kamioka et al., 1996; Canepari et al., 1997; Voigt et al., 1997; Gross and Kowalski, 1999). The cells in these cultures (Figure 6.1) begin to start firing after about 4 days *in vitro*, and soon after synchronize their activity globally across the culture. This synchronization takes the form of intense bursts of activity, 0.5–2 s in duration, that recur several times per minute. During global bursts, a large fraction of cells in the culture rapidly increase their firing rates by a factor of 10 or more. Bursting persists for the lifetime of the culture, although the fully synchronized bursts of young cultures are gradually replaced by more spatially localized bursts in maturity (Maeda et al., 1995; Corner et al., 2002). Globally synchronized bursting is an extremely robust phenomenon. Suppressing it using pharmacological agents like glutamate receptor blockers (Furshpan and Potter, 1989; Gross et al., 1993; Kamioka et al., 1996) also abolishes most or all other spontaneous electrical activity.

In vivo, bursting occurs during development and plays a role in establishing appropriate connections (Meister et al., 1991; Ben-Ari, 2001; Zhang and Poo, 2001; Leinekugel et al., 2002). However, this phase only lasts for days or at most weeks. The persistence into maturity of bursting in culture may then be interpreted as a sign that cultures are arrested in their development (Corner et al., 2002). Bursting




Figure 6.1: A typical culture growing on an MEA. **A** At 2 div. **B** At 34 div. At this age, glia have formed a carpet covering the culture. Both photographs show the same central part of the MEA. The electrodes are clearly visible. [This figure appeared as supplemental (online) material with the published article.]

in cultures has also been likened to spindles observed in the EEG of sleeping brains (Krahe and Gabbiani, 2004), as well as to epileptic activity (Furshpan and Potter, 1989; Litt and Echauz, 2002). Techniques that reduce bursting in culture are therefore of potential importance for the treatment of epileptic patients.

We hypothesize that the persistence of global bursts in dissociated cortical cultures is a result of deafferentation. Deafferentation has two effects. Firstly, the lack of (thalamocortical) input might lead to increased strength of connections within the network. Indeed, Turrigiano (1999) showed that blocking the inputs to cortical neurons using TTX during development significantly increased the strength of excitatory connections. Secondly, the lack of structured input and presence of strong excitatory connections puts the network in a highly unstable state in which positive feedback between excitatory cells can easily lead to synchronized bursts of activity (Corner and Ramakers, 1992). Latham et al. (2000) found that bursting results when too few cells in the network are tonically active. This auto-regulation may be due to slow after-hyperpolarization or regulation of intracellular Ca^{2+} (Darbon et al., 2002). We propose that substituting multi-electrode stimulation for sensory input (Heck, 1995) has the same effect as an elevated tonic firing rate, and should therefore reduce the predominance of global bursts, favoring more locally differentiated neuronal activity.

Methods

Cell Culture

Neocortical cells were dissociated from the brains of E18 rats and plated on multi-electrode arrays (MEAs). Timed-pregnant Wistar rats were euthanized using CO₂, according to NIH-approved protocols. Embryos were removed and euthanized by chilling and decapitation. The entire neocortex, excluding the hippocampus, was dissected under sterile conditions. Cortices were cut into 1-mm³ cubes in Segal's medium (Segal et al., 1998). (In mM: MgCl₂: 5.8; CaCl₂: 0.25; HEPES: 1.6; Na₂SO₄: 90; K₂SO₄: 30; Kynurenic acid: 1; DL-2-Amino-5-phosphonovaleric acid (APV): 0.05. pH-ed to 7.3 using NaOH and 0.001% Phenol Red.) After enzymatic digestion for 30 minutes by 2.5 U/mL Papain (Roche 108014) in Segal's medium, cells were separated by 6 or 9 trituration passes using a 1 mL pipette tip, in Neurobasal medium with B27 (Gibco; Brewer et al., 1993), 0.5 mM GlutaMax (Gibco) and 10% equine serum (HyClone). After every 3 passes, the cells already in suspension were transferred to a separate tube to reduce stress on them. Cells were centrifuged at 160×*g*, onto 5% bovine serum albumin (BSA) in phosphate buffered saline (PBS), resuspended by very gentle trituration and passed through a 40 μm cell strainer (Falcon) to remove large debris. 50,000 cells were plated in a 20 μL drop of Neurobasal, on MEAs pre-coated with poly-ethylene-imine (PEI) and laminin as previously described (Potter and DeMarse, 2001). This led to a plating density of 2500 cells/mm² in a monolayer. After 1 h of incubation, 1 mL of Neurobasal was added to each culture dish. After 24 h, the plating medium was replaced by a medium adapted from Jimbo et al. (1999): Dulbecco's modified Eagle's medium (DMEM, Irvine Scientific 9024) with 0.5 mM GlutaMax and 10% equine serum, but no antibiotics or antimycotics. Cultures were maintained in an incubator with 5% CO₂ and 9% O₂ (Brewer and Cotman, 1989). We replaced half the medium every 5–7 days. Glial growth was not suppressed, since glia are essential to long-term culture health. As a result, glia gradually formed a carpet over the neurons. Our use of Teflon-sealed dishes (Potter and DeMarse, 2001) allowed us to maintain the incubator at 65% relative humidity, making it an electronics-friendly environment. Thus we could perform all experiments inside the incubator, ensuring long-term stability of recording conditions. Experiments took place at 25–45 days *in*




Figure 6.2: Two-photon images of immunocytochemically stained neurons. **A** MAP2. **B** GABA. The two images show the same field of view. Arrow indicates a GABA-positive cell. Negative controls showed no visible signals. Images were taken with a Carl Zeiss LSM510 multiphoton microscope. Scale bar: 20 μm .

vitro (div). At this age, over 90% of electrodes recorded spikes. Only cultures that fired at least three bursts in 10 minutes of pre-experimental screening were used.

To determine the fraction of inhibitory cells, we stained two cultures at 16 div for microtubule-associated protein-2 (MAP-2) and γ -amino butyric acid (GABA), as described under *Immunostaining* below. In the two cultures, 10 randomly selected images showed anti-GABA immunoreactivity in 29 out of 499 neurons (5.8%) and 16 out of 440 neurons (3.6%), respectively (Figure 6.2).

Recording System

Electrical activity was recorded with a square array of 59 titanium-nitride electrodes, 30 μm in diameter, embedded in the substrate at 200 μm spacing (MultiChannel Systems, Reutlingen, Germany; www.multichannelsystems.com). After 1200 \times amplification, signals were sampled at 25 kHz using a MultiChannel Systems data acquisition card, controlled through our MEABench software.* MEABench's digital filtering system for reducing stimulus artifacts (Wagenaar and Potter, 2002 [Chapter 3 in this

* Software available for free public download: <http://www.its.caltech.edu/~wagenaar/meabench.html>.

thesis]) allowed us to detect action potentials as early as 2 ms after stimulation.* Spikes were detected online by thresholding at $5\times$ RMS noise, and later validated based on the shapes of their waveforms (P. P. Mitra, personal communication).

Stimulation System

Stimuli were generated using our custom-made 60-channel stimulator (Wagenaar and Potter, 2004 [Appendix C in this thesis]). We used biphasic rectangular voltage pulses, positive phase first, since these were found to be the most effective stimulus at any given voltage (Wagenaar et al., 2004 [Chapter 2 in this thesis]). We used stimulus pulse widths of 400 μ s per phase and voltages between 100 and 900 mV. Higher voltages were not used, to prevent possible electrochemical damage to electrodes and nearby cells. The stimulator was switched to high impedance output 100 μ s after each pulse using the built-in switches of our stimulator.

Experimental Protocols

Before experimenting on any MEA, we probed each electrode in the array with voltage pulses between 100 and 900 mV, in random order. For each electrode, we determined the voltage V^* at which the response was five times the spontaneous firing rate. Typically, 40–50 electrodes per dish were in sufficiently close contact with the culture to attain that level of response by voltages in the range tested. For each experimental series, we selected either individual electrodes or groups of 2 to 25 electrodes randomly from this pool.

We used three stimulation protocols:

‘S’ — Single electrode stimulation: One electrode was stimulated repeatedly at its voltage V^* . We used this protocol at ten different frequencies between 0.05 and 50 stimuli per second (stim/s).

‘M’ — Multi-electrode stimulation: A group of 2 to 20 electrodes was stimulated cyclically at 2 to 20 stim/s, such that each electrode received stimuli once per second, or 25 electrodes were stimulated cycli-

* Except on the electrode used for stimulation, which remained saturated by stimulation artifacts for 50–150 ms.

cally at 50 stim/s, so each received two stimuli per second. Each electrode was stimulated at its own previously determined V^* . Five different group sizes with corresponding stimulation rates were tested with this protocol.

‘FB’ — Closed-loop feedback stimulation: Ten electrodes were stimulated cyclically at 10 stim/s (so that again each electrode was stimulated once per second), with voltages continuously tuned to maintain a constant tonic firing rate, as described under *Tuning the feedback*, below. With this protocol we could stably maintain firing rates between spontaneous levels and 500 spikes per second array-wide (spsa) using voltages not exceeding 900 mV (Figure 6.3).

Experimental runs lasted 5 minutes each and were randomly interleaved with each other and with control runs in which we recorded spontaneous (unstimulated) activity. Protocols ‘M’ and ‘S’ were performed on $N=11$ cultures from 3 platings; protocol ‘FB’ was performed on $N=10$ cultures from 2 platings. In all cases, each condition was tested 10 times on each culture, with a new random selection of electrodes each time.

Quantifying the Level of Bursting

Bursts come in different forms, so simply tallying up the number of bursts is not sufficient to describe the burstiness of a culture: it is essential to account for the size of bursts, measured in terms of the number of participating neurons, the aggregate number of spikes, or burst duration. Fortunately, we found that it is not necessary to identify individual bursts in order to quantify the level of burstiness of a recording. Instead, we used the following method: divide a 5-minute recording into 300 one-second time bins, and count the number of spikes (total across all electrodes) in each bin. Compute the fraction of the total number of spikes accounted for by the 15% of bins with largest counts. If the firing rate is tonic, this number, f_{15} , will be close to 0.15. Conversely, if a recording is so bursty that the majority of spikes are contained in bursts, f_{15} will be close to one, since even at the highest burst rates observed during these experiments, bursts did not occupy more than 45 one-second bins (15%)

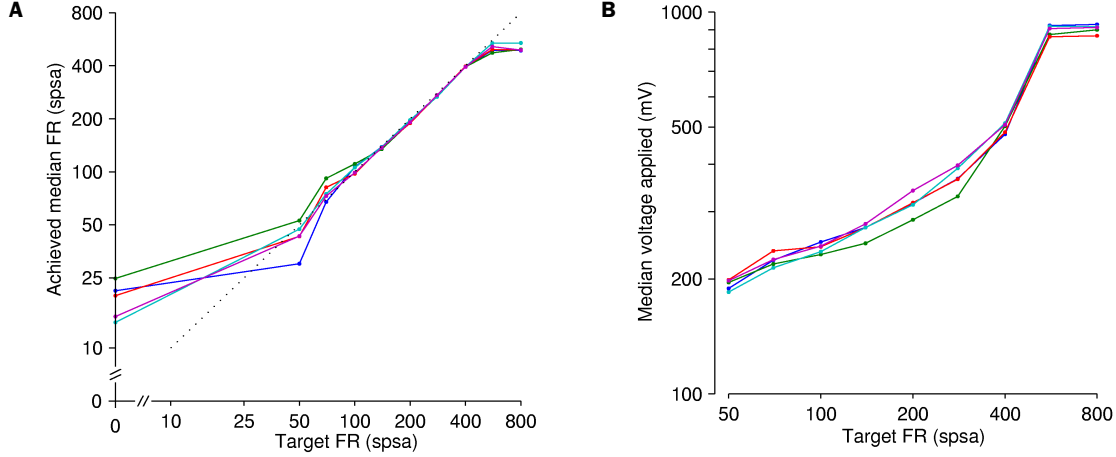


Figure 6.3: Performance of feedback protocol. **A** Median firing rate (dishwide) achieved vs. target. Any firing rate between the spontaneous rate and 500 spsa could be stably maintained. Higher rates were not achievable in this culture without exceeding the safe voltage limit of 900 mV. Dotted line marks equality of achieved and target rates. Data are from 5 series on different sets of electrodes; recorded at 45 div. **B** Stimulus voltage used to control firing rate at different levels.

in a 5-minute recording. We then defined a *burstiness index*, normalized between 0 (no bursts) and 1 (burst dominated) as $BI = (f_{15} - 0.15)/0.85$. (Statistical fluctuations make BI deviate slightly from zero even in complete absence of bursts.)

Tuning the Feedback

For the closed-loop stimulation protocol, we continuously monitored a culture's actual firing rate, and adjusted the stimulation voltages for each electrode to maintain the target firing rate, f_0 , as follows. Initially, we used a base voltage, $\bar{V} = 200$ mV, applied to all electrodes. We then measured the (culture-wide) firing rate, \bar{f} , in 2-s sliding windows, and used this to update \bar{V} every 100 ms according to:

$$\bar{V} \leftarrow \bar{V} \left(1 - \varepsilon \frac{\bar{f} - f_0}{f_0} \right),$$

where ε is a gain factor which determines how fast \bar{V} reacts to changes in \bar{f} . We set $\varepsilon = 0.02$, corresponding to a time constant of 5 s. This ensured rapid feedback, while preventing oscillations due to overcompensation.

To account for variations in stimulation efficacy between electrodes, we measured

the firing rates in the first 100 ms after each stimulus individually. For each electrode k , we used these measurements to maintain a running average, f_k , of the firing rates after the most recent 20 stimuli to that electrode. Every 100 ms, we recalculated fine-tuning factors, α_k :

$$\alpha_k \leftarrow \mathcal{N} f_k^{-1},$$

where \mathcal{N} is a normalization factor to make the average of all α_k -s be 1. We then set the voltage for the next stimulus on electrode k to

$$V_k = \alpha_k \bar{V}.$$

Thanks to MEABench and our custom stimulator (Wagenaar and Potter, 2004), these adjustments could be made in real-time without interrupting the stimulation process.

Since we wanted to control the tonic firing rate, updating \bar{V} and α_k was suspended during putative bursts, detected using a simple heuristic: any 100 ms windows that had a spike count higher than 5 \times the target were considered potential bursts, and thus excluded for the estimate of the tonic firing rate.

Immunostaining

The fraction of GABAergic neurons was determined as follows. Cultures were fixed with 4% paraformaldehyde at room temperature for 30 minutes. After treatment with 0.1% Triton-X-100 in PBS for 20 minutes, they were incubated in 2% goat serum for 1.5 hours and then in the primary antibodies anti-MAP2 (mouse, 1:200; MAB378 from Chemicon, CA) and anti-GABA (rabbit, 1:100; AB131 from Chemicon, CA) overnight at 4 °C. After washes, cells were incubated with secondary antibodies (Alexa Flour 488 goat anti-mouse, 1:200; Alexa flour 594 goat anti-rabbit, 1:1000; and Hoescht, 1:1000; all from Molecular Probes, CA) for 1 hour at room temperature. Fluorescence images were obtained from a Sony digital camera on a Nikon TE300 fluorescence scope and a Carl Zeiss LSM510 multiphoton microscope.

Results

Spontaneous Bursting

Before developing a method to control bursting, we needed to characterize the different kinds of bursting encountered in the spontaneous activity of cultured cortical networks. The frequencies of bursts as well as their sizes were highly variable between cultures from different platings, and even between cultures within platings. Additionally individual cultures showed large variations from day to day. Cultures spontaneously exhibited a wide range of bursting behaviors, from short single cell bursts, to small local bursts involving 2–5 electrodes, to long global bursts (Figure 6.4). Some cultures exhibited ‘superbursts’: stereotyped sequences of global bursts, separated by several minutes devoid of bursts (Wagenaar, Z. Nadasdy, and Potter, *in preparation*). Global bursts were typically first observed at around 7 div, after which burstiness steadily increased until they dominated the activity at around 20–25 div. After that, burstiness fluctuated somewhat, but remained high for as long as we looked (up to 45 div in these experiments).

Response to Stimulation

The immediate response to stimulation at any electrode consisted of three phases (Wagenaar et al., 2004): (1) Direct, non-synaptically-propagated responses, with very precise timing (typical jitter: 100 μ s), and latencies of 3 to 10 ms; (2) Postsynaptic responses, mostly with latencies between 5 and 50 ms; (3) Bursts, often evoked by strong or low-frequency stimuli. Such bursts were time-locked to the stimulus pulse with latencies characteristic of the local network around the electrode stimulated—usually in the range of 50–200 ms—but were otherwise similar to spontaneous bursts. Examples of early responses are shown in Figure 6.5.

During slow single electrode stimulation (0.05 stim/s), most or all stimuli entrained bursts as previously reported by Gross et al. (1993) and Maeda et al. (1995). At slightly higher frequencies (0.1–0.5 stim/s), bursts were elicited less consistently, depending on stimulation electrode. At still higher frequencies (1–5 stim/s), most stimuli did not elicit bursts, and in fact the burstiness began to drop below spontaneous levels. Increasing the stimulation rate further (10–50 stim/s) did not reduce

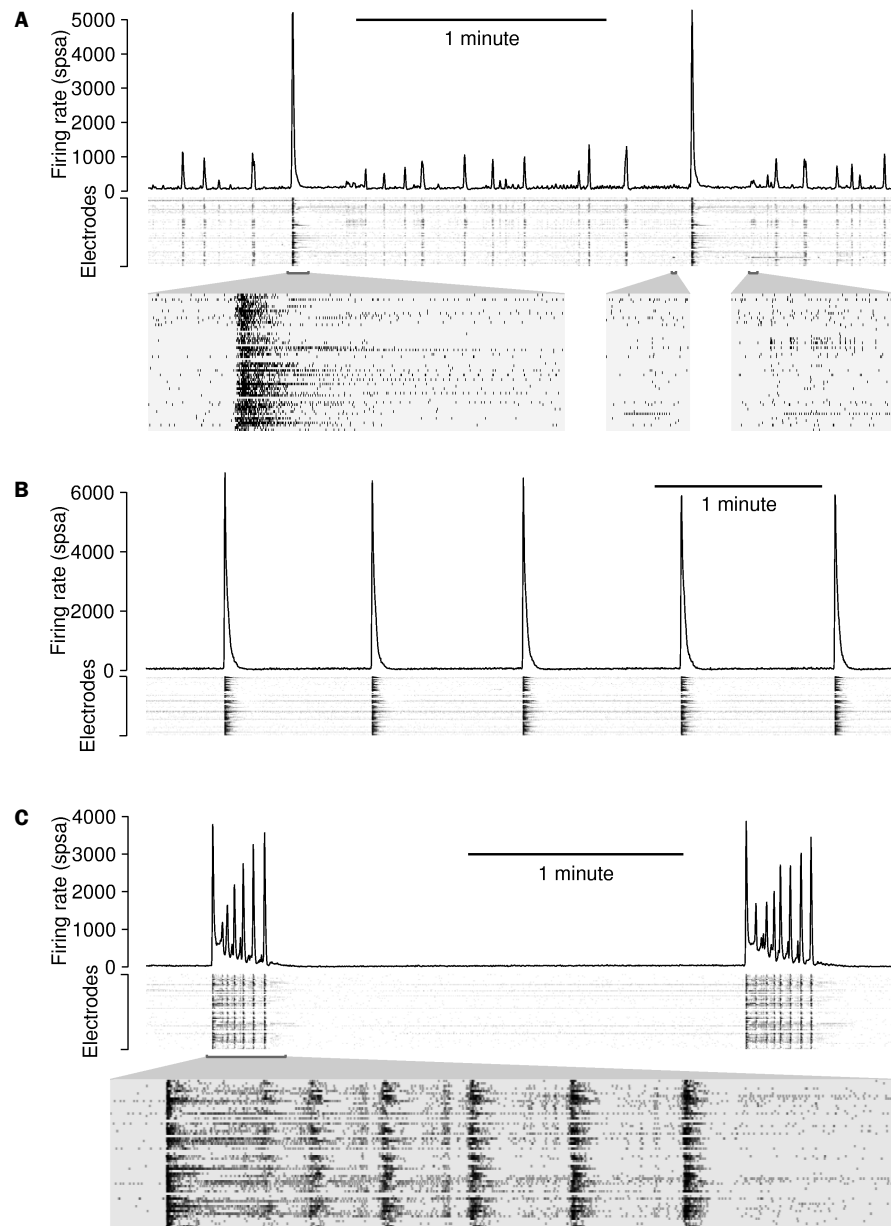


Figure 6.4: Examples of different spontaneous bursting patterns, with array-wide firing rate (*line graphs*) as well as per-electrode firing rates (*gray-scale plots*). **A** Chaotic bursting. Insets below show spike raster plots for a large global burst, a single channel burst and a small local burst, at 20x magnification. Recorded at 25 div. **B** Spontaneously regular bursting. Recorded at 39 div. **C** Superbursts. Inset shows spikes at 10x magnification. Recorded at 34 div.



Figure 6.5: Array-wide responses to stimulation. Each graph shows the responses on one electrode, represented according to the geometry of the array. The stimuli were delivered to the marked electrode. Vertical line indicates time of stimulation. Spikes were detected after artifact suppression (Wagenaar and Potter, 2002); TTX control confirmed the biological origin of all detected spikes.

burstiness more (Figure 6.6A). The best burst control on average was achieved at 10 stim/s: $BI = 0.19 \pm 0.02$ (mean \pm standard error of the mean (SEM), $N=105$ runs using different electrodes in 11 cultures; range of per-culture means: 0.04–0.55). This level of burstiness was significantly below the average spontaneous level $BI = 0.48 \pm 0.02$ ($N=199$ runs, same 11 cultures; range: 0.19–0.86).

When stimuli were applied through a single electrode at high rates, the immediate response to stimulation (spikes recorded 2–20 ms post-stimulus) dramatically decreased with increasing stimulation frequencies (Figure 6.7). This was likely responsible for the lack of improvement of burst control at those high frequencies. However, the responses to infrequent stimulation through one electrode were not affected by rapid stimulation through another electrode, so this reduction of efficacy was due to a mechanism local to the stimulated electrode, and not to a network-level fatiguing effect.

Burst Control by Distributed Stimulation

Based on the observation that rapidly stimulating single electrodes reduced the efficacy of those stimuli but not of stimuli to other electrodes, we proceeded to test whether better burst control could be achieved by distributing the stimulus load across several electrodes, using protocol ‘M’ (see *Methods*). At intermediate frequencies (2–10 stim/s), this protocol resulted in somewhat higher burstiness than single-electrode stimulation, but at frequencies above 10 stim/s, multi-electrode stimulation resulted in greatly improved burst reduction (Figure 6.6B–C). At the highest stimulation rate tested, 50 stim/s distributed across 25 electrodes, bursts were completely suppressed in all cultures tested. In all cases, a change of stimulation protocol rapidly affected burstiness, and bursting resumed as soon as stimulation was stopped (Figure 6.8).


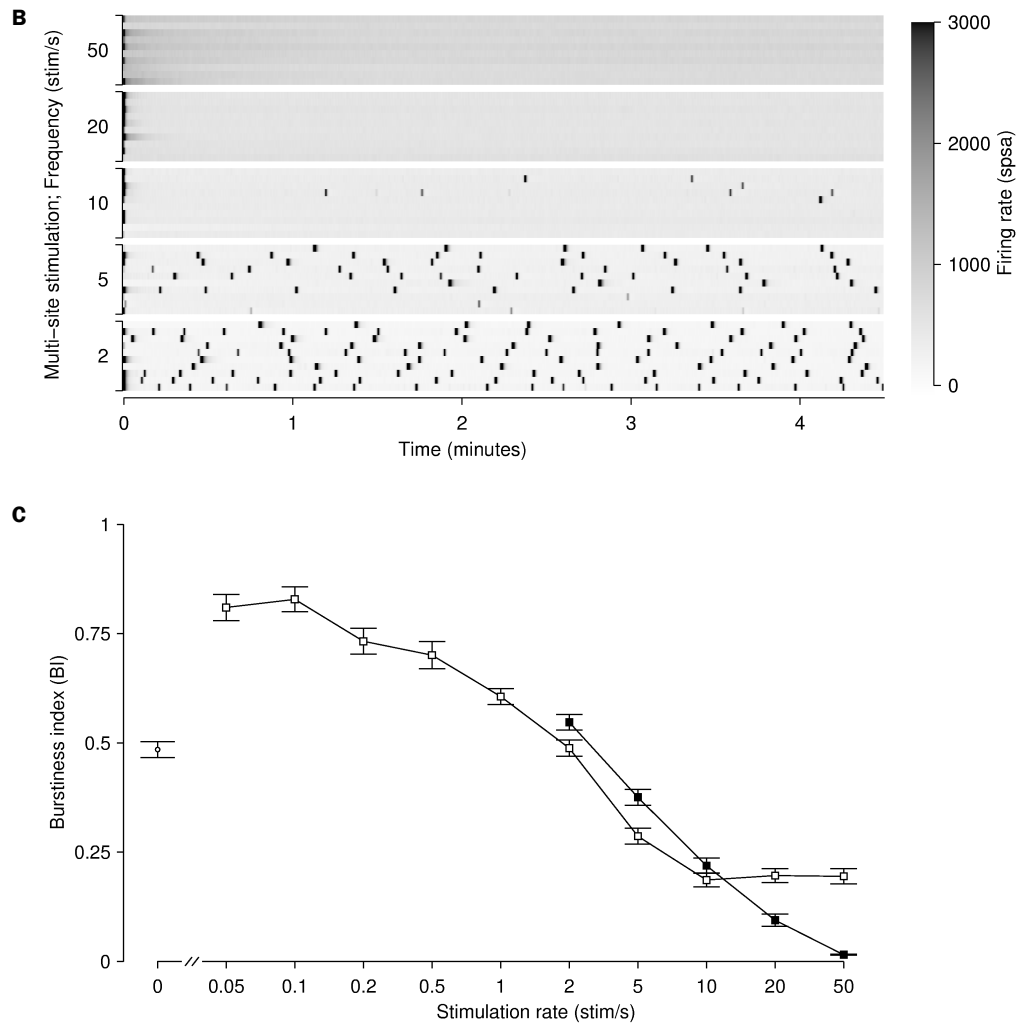


Figure 6.6: A Burstiness during single electrode stimulation (protocol 'S') and spontaneous activity (no stimulation). Each row shows the array-wide firing rate (coded by the gray scale at right) as a function of time during one 5-minute experimental run. In the 10 examples of spontaneous activity shown (*bottom*), bursts occurred irregularly about once per minute. In the 10 examples of stimulation at 0.05 stim/s, bursts were perfectly aligned with stimuli, except in a few cases where a spontaneous burst just preceded the stimulus. (The stimulating electrode was different in each of the 10 rows.) At 0.1–0.2 stim/s, bursts underwent period doubling. Bursts during stimulation at 1–5 stim/s were less frequent, but still mostly stimulus-locked. In the 10–50 stim/s runs, burst control was perfect for the first 45 s, after which a spontaneous-like pattern returned. Data from a culture at 39 div. Note that experimental runs were executed in random order.

(Continued on next page)



(Figure 6.6, continued)

B Bursting during multi-electrode stimulation (protocol 'M'), same culture as A. Perfect and sustained burst control is attained at the higher stimulation frequencies. Note the increase in tonic firing rate (reflected in the background shading) as the stimulation frequency is increased. **C** Burstiness index as a function of stimulation frequency, for single-electrode stimulation (*open squares*) and multi-electrode stimulation (*filled squares*). Slow single-electrode stimulation elevates the burstiness over spontaneous (unstimulated) levels (*open circle*), while rapid stimulation reduces it. Values are mean \pm SEM from $N = 100$ runs on 10 cultures. The most effective protocol tested, 50 stim/s distributed across 25 electrodes, suppressed bursts completely ($N = 60$ runs, 6 cultures).




Figure 6.7: Stimuli presented to a single electrode ('primary stimulation electrode', *filled circles*) yielded much reduced responses in the first 20 ms post-stimulus when the stimulation rate was increased. (We focused on short-latency responses, because the majority of response spikes occurred at short latencies, and because responses cannot be unambiguously defined beyond one inter-pulse-interval, i.e., 25 ms for the highest stimulation frequency.) In fact, at a stimulation frequency of 40 stim/s, the response was not much higher than the spontaneous firing rate (*arrow at left*). Each stimulation series lasted five minutes, and we discarded the responses recorded during the first 30 seconds so as to measure the sustained response rate. Results are mean \pm SEM from 53 electrode pairs in 4 cultures. During these experiments, we presented stimulus pulses to a second electrode every five seconds. The responses in the first 20 ms after these latter stimuli (*open circles*) were not affected by the rate at which the first electrode was stimulated. Response strength in all cases was normalized to the results obtained from single-electrode stimulation at 0.2 stim/s. The response strengths are plotted as a function of the frequency at which the primary electrode was stimulated. Inset and associated arrows: Explanation of stimulation protocol. Irrespective of the frequency of the primary stimulation electrode, the secondary electrode was stimulated once every 5 seconds.




Figure 6.8: When switching between stimulation protocols, a culture's activity pattern rapidly changed to match the new stimulation context. Here we show switches from rapid single electrode stimulation to slow single electrode stimulation, to rapid multi-electrode stimulation, to no stimulation.

Burst Control by Closed-Loop Stimulation

Perfect burst control was achieved using protocol ‘M’, but only at very high stimulation rates and using a large number of electrodes. If good burst control could be attained using fewer electrodes or lower stimulation rates, this would have practical advantages. We noted that the bursts that occurred in protocol ‘M’ at intermediate stimulation rates were mostly entrained by only one of the electrodes used in a given run, indicating that the calibration of stimulus efficacy performed before the experiment (see *Methods*) was not a very good predictor of efficacy in the context of much more intense multi-electrode stimulation (data not shown). Thus we hypothesized that the level of burst control attained by pre-defined voltage pulses could be further improved by tuning the stimulation voltages in real-time to obtain a constant level of response. We used feedback control (protocol ‘FB’; see *Methods*) to regulate the median firing rate at 9 fixed levels between 50 and 800 spikes per second array-wide (spsa). Increasing the median firing rate above spontaneous levels reduced burstiness monotonically (Figure 6.9). At the highest target rate of 800 spsa, this protocol was significantly more effective than either single-electrode or multi-electrode stimulation compared at the same stimulation rate (10 stim/s). This held despite the fact that the spontaneous *BI* was 50% higher on average for those cultures on which we tested feedback stimulation compared to those tested with single or multi-electrode stimulation. (This difference in spontaneous behavior was due to variability between cultures, not to our intervention.)

A final comparison of the various protocols was made by counting in what percentage of cultures each protocol suppressed bursts completely during a 5-minute run (Figure 6.10). Any run with $BI < 0.05$ was considered burst-free for this assessment. The most intense protocol ‘M’ stimulation (50 stim/s distributed across 25 electrodes) suppressed bursts in all cultures, independent of the selection of stimulation electrodes. Although a set of electrodes could be found to suppress bursts at 10 stim/s with fixed voltages in over 50% of cultures (white bars), a random selection of electrodes suppressed bursting in only 1 in 5 cultures (gray bars). Closed-loop feedback did much better: a random selection of electrodes suppressed bursting in over 50% of cultures (gray bar), and in 30% of cultures, all 10 random selections of electrodes

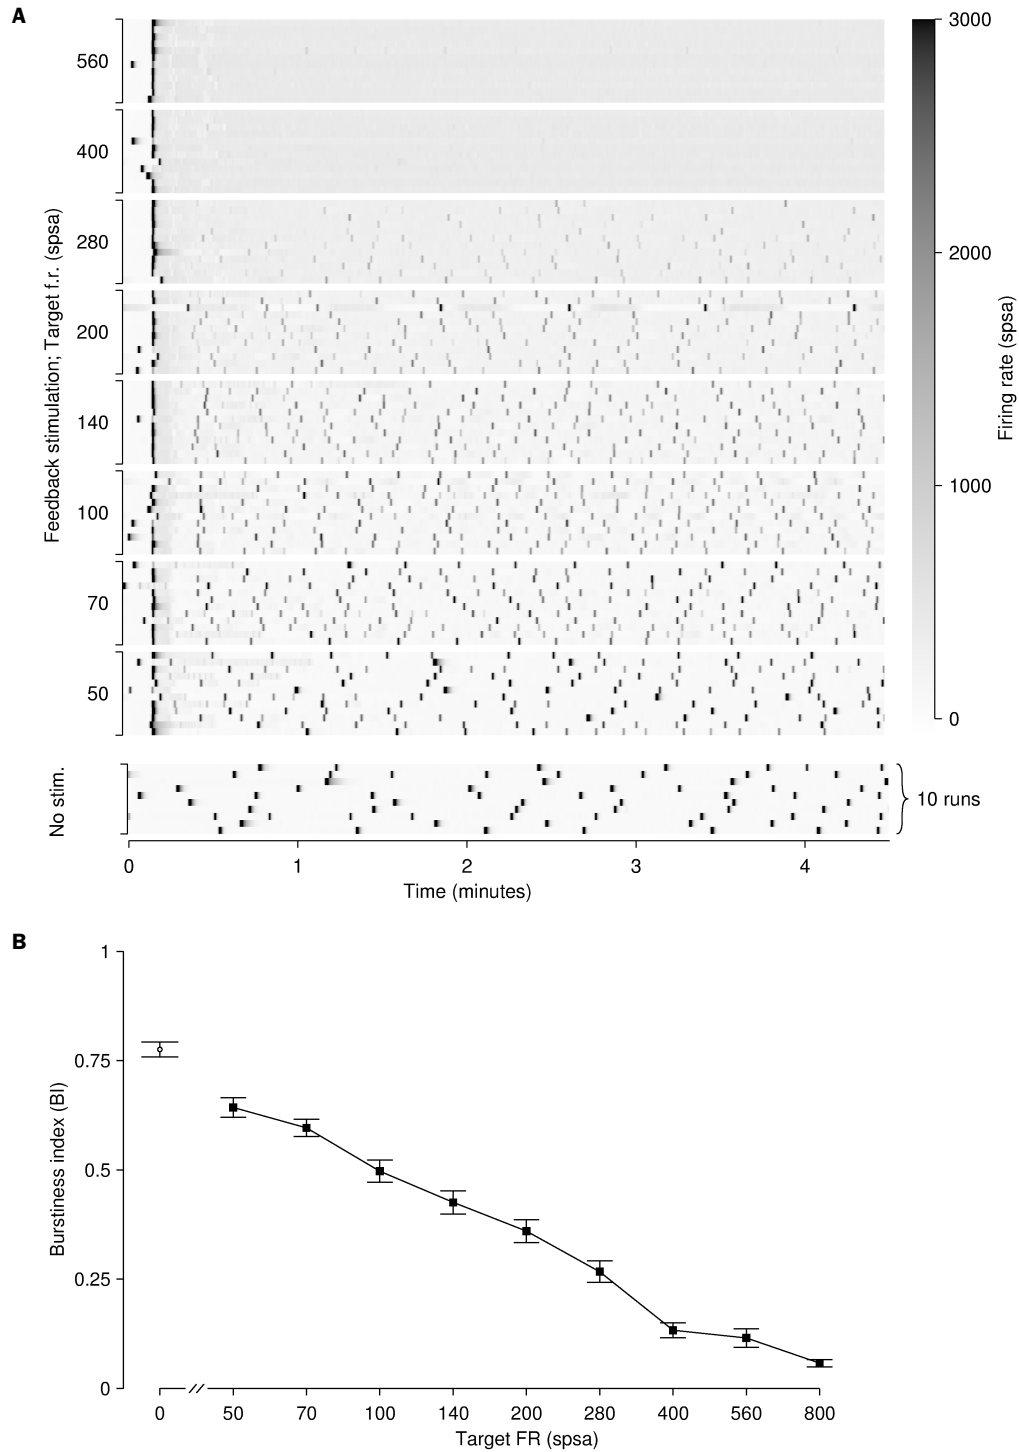


Figure 6.9: Burstiness during closed-loop control of tonic firing rate. **A** After an initial period of about 15 s during which the feedback algorithm settles, burst control was perfect at the higher target firing rates. From a culture at 43 div. (This culture was not tested at 800 sp/s.) **B** Burstiness index decreased monotonically with the target rate, and was always below the spontaneous level (*open circle*). Values are mean \pm SEM from $N = 85$ runs using different sets of electrodes, on 10 cultures.




Figure 6.10: Assessment of the success rate of different burst suppression protocols. Bars show the percentage of cultures in which each protocol successfully suppressed bursting, by random electrode selection (gray), by at least one of 10 selections tested (white), or by all of 10 selections tested (black). None of the cultures used in these experiments were burst-free in spontaneous activity. Protocols compared are: protocol ‘S’ at its optimal stimulation rate (10 stim/s, ‘S10’); protocol ‘M’ at the same rate (‘M10’); protocol ‘FB’ at its optimum (target 800 spsa, ‘FB’); and protocol ‘M’ at its optimum (50 stim/s distributed across 25 electrodes, ‘M50’).

tested suppressed bursts (black bar).

Discussion

Several years ago, Latham et al. (2000) showed that networks with a large fraction of intrinsically active neurons have a reduced tendency to burst. We extend this finding by demonstrating that increasing the tonic activity above spontaneous levels by high-frequency multisite electrical stimulation also reduces or suppresses bursting. Strikingly, complete suppression of bursts was achieved by a combination of stimuli that entrained bursts when applied singly. Rapid stimulation through single electrodes yielded fewer bursts than slow stimulation, not just per stimulus, but per unit time: stimulation at 5 stim/s or more reduced burstiness to below spontaneous levels. Distributing the stimuli across 20 or more electrodes proved highly effective to reduce it even further, and with 50 stim/s distributed across 25 electrodes, bursting was suppressed completely in all cultures tested, independent of the selection of electrodes. However, such a high stimulation rate may be undesirable in some

applications, or that many electrodes may not be available. When the number of electrodes used for burst control was limited to 10, stimulating with closed-loop feedback was found to be the optimal solution: this protocol completely suppressed bursts in over 50% of cultures using 10 stim/s distributed across randomly selected groups of 10 electrodes. With careful selection of electrodes, fixed-voltage stimulation through single or multiple electrodes suppressed bursting in a similar fraction of cultures as feedback stimulation. However, feedback stimulation was far more robust: in 30% of cultures it worked regardless of electrode selection. Electrode independence was never seen for fixed-voltage stimulation at 10 stim/s. To extend burst control beyond 5-minute runs, such robustness is highly desirable.

Synchronized bursting is fundamentally a network phenomenon, emerging from the synaptic interactions between a large number of cells. Whether these cells would endogenously burst in the absence of synaptic input is probably not essential for this phenomenon. The cellular and network mechanisms of bursting and burst suppression are not yet understood in detail. There is some controversy about the origin of the refractory periods for spontaneous bursts: Opitz et al. (2002) reported synaptic depression immediately after population bursts, while Darbon et al. (2002) found no evidence of synaptic depression: no depletion of vesicles, and no desensitization of postsynaptic receptors.

It has been suggested that the persistence of global bursting in mature cultures is evidence that such cultures are in a state of arrested development as a result of lack of sensory input. Our experiments support this view, since we found that substituting for thalamic inputs with distributed electrical stimulation reduced bursting dramatically. Given that the developmental fine structuring of several primary cortical sensory areas *in vivo* is known to be determined by the pattern of inputs, it is tantalizing to ask whether persistently applied stimulation *in vitro* might similarly influence network topology.

In contrast to burst suppression by (partially) blocking excitatory synaptic transmission, e.g., using AP5, CNQX (Jimbo et al., 2000), magnesium or kainic acid (Furshpan and Potter, 1989), distributed stimulation does not reduce the ability of the culture to respond to additional stimuli. Continuously applying distributed stimulation to suppress bursts is thus compatible with studies of use-dependent modifi-

cation of activity in cultured networks. Additional stimuli can be superimposed on a background of burst-quieting stimuli, to tetanize particular pathways or to probe network activity. Moreover, distributed stimulation mimics more natural modes of activation in which sensory signals are continuously coming in to the network. Bursts are known to have an effect on tetanus-induced synaptic plasticity (Maeda et al., 1998). Therefore, we suggest that burst suppression may lead to more stable connections, and more predictable results of tetanization (Chao et al., 2005). We expect that burst control will make these networks more useful for the study of distributed information processing, robotic control, and network plasticity related to learning and memory (Potter, 2001; DeMarse et al., 2001; Potter et al., 2006).

If distributed stimulation so effectively reduces bursting *in vitro*, it might also work *in vivo*. Epileptic seizures in human cortex, while probably due to very different causes, have a strikingly similar phenomenology: ensemble bursts extending over large areas of neural tissue (Lopes da Silva et al., 2003). Electrical stimulation has been used in several experimental therapies for epilepsy; stimulation of the vagus nerve is the most well-known example (Penfield and Jasper, 1954; Hammond et al., 1990; Ben-Menachem et al., 1994; Fisher et al., 1997; Handforth et al., 1998; Koo, 2001). Alternatively, animal and modeling studies suggest that focal stimulation at the site of the seizure can terminate seizures after they have started (Lesser et al., 1999; Franaszczuk et al., 2003; Slutzky et al., 2003). In humans, focal stimulation in the cortex or hippocampus has indeed been found effective in a number of studies (Cooper et al., 1973, 1976, 1977; Lüders et al., 1988; Shulz et al., 1997; Velasco et al., 2000, 2001; Motamedi et al., 2002; Vonck et al., 2002). Stimulation through a single electrode offered protection against seizures, but only if the stimulus was strong enough that the entire seizure-prone area was reached (Motamedi et al., 2002; Kellinghaus et al., 2003), which was difficult in practice. Distributing stimulation across multiple electrodes might be attractive for several reasons. Firstly, the amplitude of pulses delivered to each electrode could be much lower, reducing the risk of side effects (Wheless, 2001; Schachter, 2002), tissue damage (Shepherd et al., 1991; Tehovnik, 1996), or electrode damage commensurately. Secondly, the system would be more fault tolerant (Davis, 2000), as losing one or two electrodes from a large ensemble would hardly compromise efficacy. Thirdly, unlucky placement of a single

electrode can result in poor burst control, while with multi-electrode stimulation, the result is much less dependent on exact placement. Finally, if the electrodes were connected to a recording system equipped with seizure prediction software, stimulation parameters could be tailored to the predicted locus of impending seizures (Iasemidis, 2003). Our real-time controlled stimulator (Wagenaar and Potter, 2004) could be a starting point for developing such a system.

7

Plasticity or the Lack Thereof

This final chapter describes the results of a number of experiments in which I attempted to induce synaptic plasticity by tetanization of one or several substrate electrodes. I looked for plasticity expressed in changes in spontaneous burst patterns, in array-wide response patterns to test pulses, and in the strength of specific stimulus–response pathways that were deemed, based on pre-experimental probing, to be likely loci of plasticity. Experiments were performed under baseline conditions, as well as with spontaneous bursting suppressed by either distributed electrical stimulation or by elevated extracellular magnesium concentrations. Except in a few isolated cases, changes coinciding with tetanization were no larger in magnitude than changes that occurred spontaneously. The chapter concludes with a discussion of the discrepancy between these results and previously published studies that did report plasticity induced in dense cultures of cortical cells on MEAs.

Introduction

Changes to the connections between cortical neurons are thought to play an essential role in learning. Such changes can take the form of the extension or retraction of neurites, accompanied by the formation or elimination of synapses, or they can take the form of strengthening or weakening of existing synapses (e.g., Buonomano and Merzenich, 1998; Poirazi and Mel, 2001; Malinow and Malenka, 2002). In culture,

plasticity in individual synapses can be induced by patching on to pre- and postsynaptic cells, and repeatedly forcing the postsynaptic cell to fire either just before or just after the synapse has been activated (Bi and Poo, 2001). By cleverly manipulating visual inputs, Fu et al. (2002) have shown that such ‘timing dependent plasticity’ can also be made to occur in the cortex *in vivo*, though whether it plays a dominant role under natural circumstances remains a conjecture.

Cultures on MEAs are a potentially very attractive model for studying both structural and synaptic plasticity, since they make it possible to record from the same set of cells for several months—as opposed to mere hours for patch clamp experiments. Furthermore, it is much easier to image the same region of a culture over time than it is to do the same thing in an intact brain. Accordingly, it is no surprise that several research groups have been interested in inducing and studying plasticity in the connectivity of neuronal networks using MEA electrodes.

In the 1990s, the research group of Akio Kawana at NTT in Japan published a series of papers in which they reported that tetanic stimulation through one or several electrodes resulted in plasticity: Maeda et al. (1998) observed a change in the probability of evoking bursts by test pulses, as well as a change in the rate of spontaneous bursting, as a result of repeatedly evoking bursts using strong tetani. The next year, Jimbo et al. (1999) reported that tetanizing a single electrode resulted in changes in the responses to test pulses to other electrodes. Culture-wide responses to a particular stimulation electrode were either all up-regulated or all down-regulated, a phenomenon they called ‘pathway-dependent plasticity.’ Individual pathways were up-regulated or down-regulated depending on the correlation between (pre-tetanus) responses to stimuli applied to the test electrode and to the tetanization electrode.

In the first paper to claim actual learning in cultured neuronal networks, Shahaf and Marom (2001) reported that networks could be made to learn to respond in specific ways to test pulses, by repeatedly stimulating until the desired response was obtained. Later, Ruaro et al. (2005) reported that networks could learn to ‘recognize’ complex stimuli that had been presented repeatedly. However, to date, none of these experiments have been independently verified by other research groups. The only example of long-term plasticity in cortical networks on MEAs that has been confirmed by more than one group, is a form of adaptation first observed in Marom’s group (Ey-

tan et al., 2003): if two electrodes are repeatedly stimulated, one at a very slow rate (0.02 Hz) and one at a faster rate (0.2 Hz), the responses to the ‘slow’ electrode are enhanced while the responses to the ‘fast’ electrode are weakened, effects which are fully reversible. This result was verified by Thomas DeMarse when he used a cortical culture to control a flight simulator (DeMarse, personal communication; University of Florida Press release, 2004, <http://www.napa.ufl.edu/2004news/braindish.htm>).

Meanwhile, others have tried and failed to reproduce the other reported forms of plasticity, leading Van Staveren et al. (2005) to question whether cortical cultures can, in fact, learn. In our labs, we have performed many different experiments on plasticity, each time finding that apparently positive results did not stand after careful analysis of experiments and controls, or that positive results could not be replicated in other cultures, suggesting that initial optimism was due to overfitting or wishful thinking, or perhaps that some cultures are somehow special in as yet unexplained ways. After a plethora of $N=1$ experiments with unconvincing results, the concern that published plasticity results cannot be replicated became quite real. Thus I decided to perform a series of carefully controlled experiments, each on a larger number of cultures. These experiments should have had enough power to discover plastic changes if any of the effects previously reported occurred in our cultures. In the following, I shall describe (mostly negative) results obtained with three paradigms for detecting plasticity. After that, I will discuss possible causes of the discrepancy between our findings and those reported by other groups. The protocols they used are summarized in Box 7.1.

General Methods

Cultures were prepared and maintained as before (see Appendix A for details). Cortical cells—neurons and glia in natural proportions—were obtained from rat embryos, and 50,000 cells in a 20 μL drop were plated over the center of MEAs. This resulted in monolayer cultures of 5 mm diameter—three times larger than the diameter of the electrode array—with a density of about 2,500 cells/ mm^2 after one day *in vitro* (div). Experiments were performed after 10–32 div, in the same incubator in which cultures were maintained.

Box 7.1: Overview of plasticity-inducing stimuli used by other researchers

The following is a very brief synopsis of the methods and main results of five previous studies that reported plasticity in dense cortical cultures on MEAs. I cannot possibly do justice to any of these papers in this space; the reader is encouraged to refer to the original papers for more information.

Ref.	Induction stimuli	Test stimuli	Results
Maeda et al., 1998	Trains of 20 pulses at 20 Hz simultaneously to each of 5 electrodes, repeated 5–10x at 10–15 s intervals.	Trains of 20–30 pulses at 1 kHz or stronger single pulses, to 1 or 5 electrodes, repeated every 15–30 s.	Increased probability of evoking bursts by test stimuli after tetanization.
Jimbo et al., 1999	Trains of 10 pulses at 20 Hz to one electrode, repeated 20x at 5 s intervals.	Individual pulses to each of 64 electrodes, repeated 10x at 3 s intervals.	‘Pathway-dependent’ plasticity.
Shahaf and Marom, 2001	Bipolar stimulation between a pair of electrodes, at 1–3 s intervals, repeated until the desired response was seen, or for 10 min max.	<i>Induction stimuli served as test stimuli.</i>	Desired responses (increased spike rate 50–60 ms post-stimulus) obtained after fewer trials on successive test series.
Eytan et al., 2003	One electrode stimulated every 50 s (‘slow’) and another every 5 s (‘fast’).	<i>Induction stimuli served as test stimuli.</i>	Responses to ‘slow’ electrode enhanced, to ‘fast’ electrode reduced; both reversibly.
Ruaro et al., 2005	Trains of 100 pulses at 250 Hz simultaneously to each of 15 electrodes in an L-shape, repeated 40x at 2 s intervals.	Stimuli, simultaneously to several electrodes, in an L- or Γ -shape.	Responses to L-shape enhanced relative to Γ -shape.

Stimuli were delivered using our custom stimulator (Appendix C). All pulses were biphasic, 400 μ s per phase, positive first, in accordance with the results of Chapter 2. Care was taken to limit voltages to less than 1 V, to avoid electrochemically damaging cells or electrodes. All 59 electrodes in the array could be used for stimulation, but due to a broken wire in one pre-amplifier channel, only 58 could be used for recording.

To deliver chemicals to a culture, we used a stock solution of at least 20x concen-

tration in DMEM and added drops of at most 50 μL directly to the 1 mL of culture medium present in the dish, to avoid any transients due to medium exchange. To ensure equal distribution of the applied substance through the entire medium, 0.5 mL of medium was taken out near the spot where the drop was added, and gently returned on the opposite side of the dish. This technique was verified by applying a small drop of strong acid to a culture dish with pH-indicator medium. After mixing in the manner described, the color of the medium became a uniform orange.

Overview of Results

I looked for plasticity induced by electrical stimulation in three series of investigations: *Changes induced in burst patterns*, *Changes in stimulus–response maps*, and *Changes in specific responses*. Within each series, I performed experiments with several different protocols. I will devote separate sections to the results obtained in each series below, but first provide an overview:

Series I: Changes induced in burst patterns If tetanization has an effect on many synapses, it should have an effect on a culture’s overall activity, and in particular on its spontaneous burst patterns. Strong tetani, delivered through several electrodes in parallel, should have the best chance of inducing such global plasticity. To test this hypothesis, I recorded spontaneous activity before and after attempting to induce plasticity using strong tetani, and measured burst frequencies, sizes, and the total number of spikes in bursts per unit time. Out of 28 experiments, only one showed clearly tetanus-induced changes, and those changes had the opposite sign of those reported by Maeda et al. (1998). In separate experiments, I tested whether a change could be induced in the probability that test stimuli evoked bursts, but found no significant effect.

Series II: Changes in stimulus–response maps According to Jimbo et al. (1999), tetanization through a single electrode can induce changes that are stimulation-site specific, that is, array-wide responses to test stimuli on a given electrode (not necessarily the tetanized electrode) are either all up-regulated or all down-regulated. To test this hypothesis, I recorded responses to test pulses

delivered sequentially to each electrode in the array before and after tetanization. Then I asked two questions: (1) Is there any change in how strongly individual recording sites respond to particular stimuli? (2) Are such changes stimulation-site specific (as reported by Jimbo et al. (1999)), recording-site specific, or more complexly distributed? Stimulation-site specific changes did occur, but they occurred spontaneously, and were not magnified by tetanization.

Series III: Changes in specific responses From patch clamp experiments, it is known that tetanizing a pair of cells can strengthen or weaken synapses between those cells depending on the timing of the tetanizing stimuli. MEA electrodes do not provide direct access to pairs of cells with known synaptic connectivity, but if one electrode records responses both after stimulation to electrode *A* and to electrode *B*, it is likely that shared synaptic pathways exist. Therefore, tetanizing the pair *A* and *B* can be expected to affect the responses on the shared target. To test this hypothesis, I selected pairs of stimulation electrodes that both evoked responses at a third site, recorded those responses, and compared them before and after paired-pulse tetanization. No significant effects were found, though in one protocol tetani appeared to increase the rate of change of spike counts in responses slightly compared to the spontaneous rate of change.

We previously hypothesized that ongoing spontaneous bursting activity may interfere with inducing plasticity and maintaining changes (Chapter 6; see also Maeda et al. (1998)). Therefore, in addition to experiments under baseline conditions, I used two different methods to reduce bursting. One was to add 1 or 2 mM magnesium chloride to the medium, which transiently reduced or abolished spontaneous bursting (Figure 7.1). The other was distributed electrical stimulation (Chapter 6), which completely suppressed bursting for as long as it was applied.

In all experiments, spontaneous or test-pulse-evoked activity was recorded for (at least) two hours before and two hours after the induction sequence. The activity in the first hour after induction was then compared to the activity in the last hour before, to determine the changes *coincident* with the induction sequence. Importantly, the activity in the hour before induction was also compared to the activity

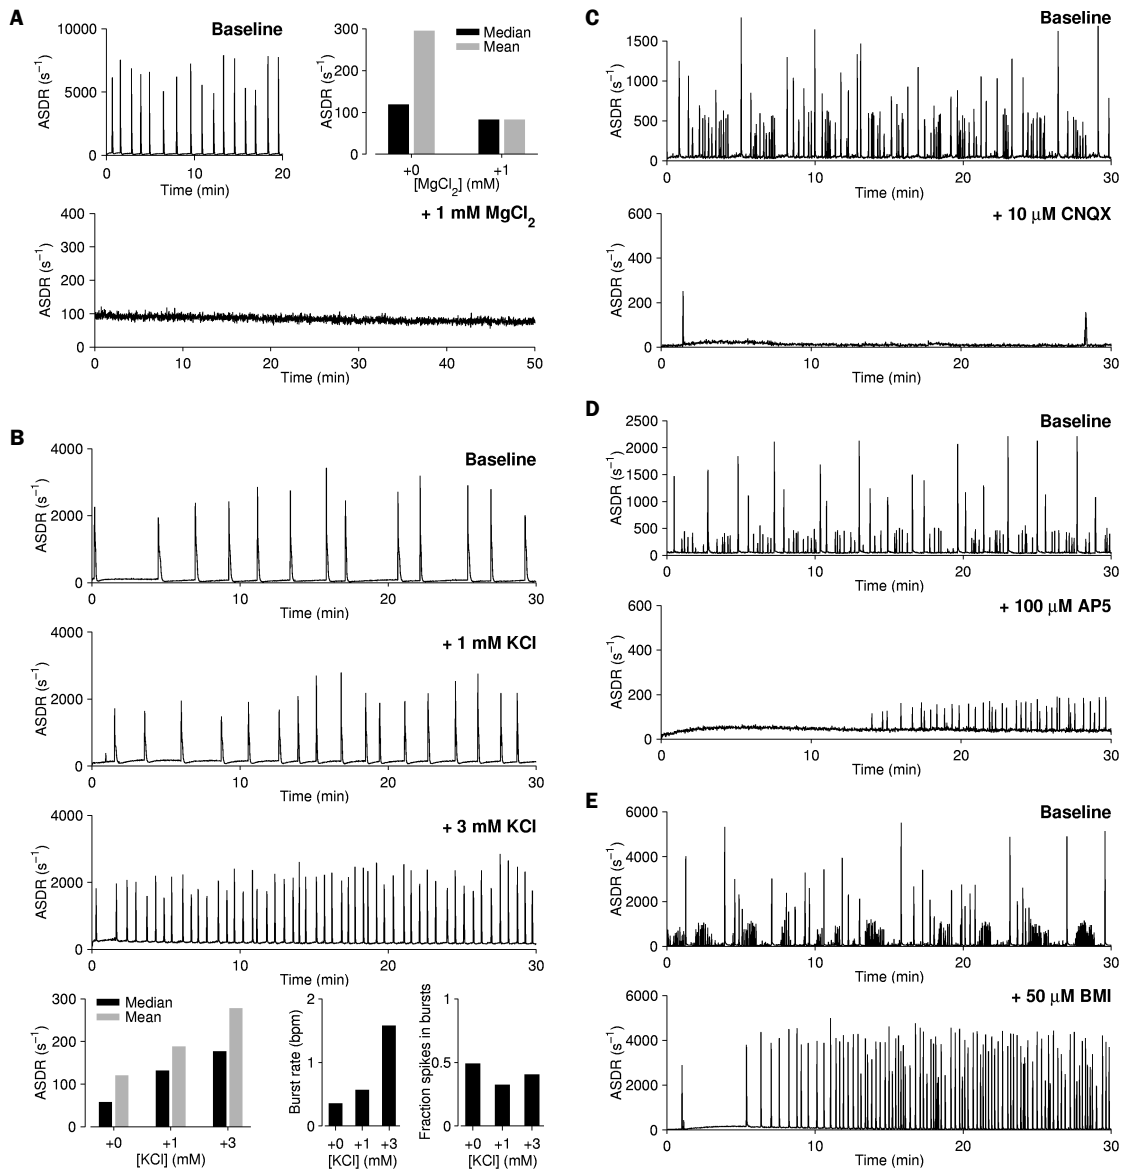


Figure 7.1: Reactions to pharmacological manipulations. **A** Adding 1 mM Mg²⁺ (to the baseline of 0.8 mM) stopped spontaneous bursts, and reduced the array-wide spike detection rate (ASDR) outside of bursts slightly. (ASDR outside of bursts is measured by median ASDR.) **B** Adding 1 or 3 mM K⁺ (to the baseline of 5.4 mM) increased burst rates and inter-burst firing rates. The fraction of spikes that occurred inside bursts (as opposed to between bursts) remained similar. **C** CNQX, an AMPA channel blocker, inhibited bursting and reduced baseline ASDR. **D** AP5, an NMDA channel blocker, inhibited bursting for a limited period of time. **E** Bicuculline methiodide (BMI), a GABA channel blocker, increased burstiness. (Data for A–E were obtained from different cultures, $N = 1$ for each substance. Baselines were recorded immediately prior to adding drugs. Since the results were fully consistent with expectations, a more in-depth investigation was deemed unnecessary.)

one hour before that, to estimate the magnitude of *spontaneous* changes. Statistical tests were applied to determine whether coincident changes were larger than spontaneous changes.

Sanity Checks

Since this chapter describes mostly mostly negative results, it was critical to make sure that positive results *could* have been obtained. That is, the stimulation and recording systems must be working, the preparations healthy, and their spontaneous activity and responses to test pulses comparable to those observed in cultures in which induced plasticity has been reported by others. Similarity in reaction to common pharmacological agents should also be confirmed.

Our cultures passed each of these checks:

Spontaneous activity The spontaneous activity of our cultures consisted of interspersed firing of several cells at low rates, interrupted by culture-wide bursts at varying intervals (Chapter 4). This is compatible with the behavior of the cultures used by the NTT group (Maeda et al., 1995).

Responses to test pulses As detailed in Chapter 2, we observed individual spikes and short trains of spikes on many electrodes in response to electrical stimulation on a single electrode, just as the NTT group did (Jimbo et al., 1999). In addition, culture-wide bursts were observed in response to some stimuli, compatible with the findings of Maeda et al. (1998).

Reactions to pharmacological manipulations An increased magnesium concentration in the medium reduced or abolished burstiness, presumably by blocking the calcium binding sites on NMDA receptors (Figure 7.1A). An increase in burst frequencies and inter-burst spike rates was obtained by adding potassium (Figure 7.1B), presumably through shifting the resting membrane potential: adding 3 mM K^+ (to the baseline of 5.8 mM; see Box A.7) should result in a depolarization by about 11 mV. With NMDA receptors blocked by AP5 (100 μ M), bursting ceased (Figure 7.1C). Blocking AMPA receptors with CNQX (10 μ M) also prevented bursting, and reduced inter-burst spike rates (Figure 7.1D). Con-




Figure 7.2: Confirmation of the elasticity results of Eytan et al. (2003). One electrode was initially stimulated at 1 Hz for one hour (*solid symbols*), while another was stimulated at 1/60 Hz (*open symbols*). Then, the roles were reversed. The graph shows the number of spikes recorded array-wide, 15–30 ms after a stimulus, normalized to the value at the beginning of the experiment. ‘Start’ refers to the first stimulus to the ‘slow’ electrode, or the average of the first 5 stimuli to the ‘fast’ electrode; ‘Early’ refers to the average of the first 5 stimuli to the ‘slow’ electrode, or the average of the 5×4 surrounding stimuli to the ‘fast’ electrode; ‘Late’ refers to average of the last 20 stimuli to the ‘slow electrode, or the average of the last 1200 stimuli to the ‘fast’ electrode. (This slightly unusual way of organizing the data was used to balance the need to collect sufficient statistics with the desire to measure as close as possible to the beginning of the experiment.) Data are mean \pm SEM (in log-space) from 16 experiments on 4 cultures. The sequence of open and closed symbols near the top of the graph are a cartoon of the stimulation sequence; the actual number of stimuli was much greater.

versely, bicuculline, a blocker of GABA receptors, increased burst rates at a concentration of 50 μ M (Figure 7.1E).

As a final sanity check, I tested whether my cultures exhibited elastic changes in response strength when simultaneously presented with fast and slow stimuli on two electrodes, as Eytan et al. (2003) reported. One electrode, *A*, was stimulated at 1 Hz for one hour, while another, *B*, was stimulated at 1/60 Hz. Indeed, responses to electrode *A* decreased significantly ($p < 0.001$; $N = 16$), while responses to electrode *B* increased slightly ($p \sim 0.06$; Figure 7.2). Then, the roles were reversed for one hour—*B* was stimulated at 1 Hz, and *A* at 1/60 Hz—and soon responses to *A* increased ($p < 0.001$) to marginally above baseline levels ($p \sim 0.2$), while responses to *B* decreased significantly ($p < 0.05$), in agreement with Eytan et al.

Series I: Changes Induced in Burst Patterns

Protocols

I tested whether strong tetanization could induce changes in spontaneous bursting behavior in $N = 10$ cultures. I measured the number of bursts spanning at least 10 electrodes in one-hour windows before and after tetanization, as well as the number of spikes in those bursts. Very strong tetani, consisting of volleys of stimuli to 5–10 electrodes, were used in these experiments. Within a volley, each electrode received one stimulus, and successive electrodes were stimulated at 2–5 ms intervals (inter-electrode interval; IEI). Such volleys had a high probability of evoking bursts, which, according to Maeda et al. (1998) is essential for affecting later spontaneous bursting. Volleys were either delivered singly, or in sets of 4 or 20 with an inter-volley interval (IVI) of 50–500 ms. A pause of 5–10 s was interposed between sets, so that each set had a good chance of evoking bursts. (In general, evoking bursts was subject to a refractory period.) The full tetanization sequence lasted 8–17 min. In most cases, several experiments were performed consecutively on one culture, with several hours between experiments. Details for all protocols are listed in Table 7.1.

Choice of Electrodes

Effective tetanization requires good contact between the tetanized electrode and the culture: if a tetanization pulse cannot elicit a response, it is unlikely it will induce plasticity. Therefore, all experiments began with probing each electrode in the array with pulses of several voltages: each electrode was stimulated 27 times at 100–900 mV; the sequence was fully randomized; inter-stimulus intervals were 0.3 s. Electrodes with strong contact with the culture should elicit a graded response: with increasing voltage, more and more cells should be recruited. For tetanization, electrodes were chosen that clearly showed this effect, and which could be used to increase the array-wide spike detection rate (ASDR) to at least twice the baseline level, using pulses no larger than 900 mV (Figure 7.3). In most cultures, plenty of electrodes fulfilled these requirements. When multiple experiments were performed on a single culture, I selected electrodes from different regions of the array for each new experiment, to maximize independence between experiments.

Table 7.1: Details of experiments on plasticity expressed in burst patterns

Proto-col	Tetanus	Conditions	No. and ages of cultures	Tet. per culture	Inter-vals
I.1	Sets of 4 volleys (IVI: 500 ms) to 10 geometrically close electrodes (IEI: 5 ms), repeated every 5 s for 15 min.	Baseline medium, spontaneous bursting.	2x2 ^a ; 10–19 div	1	4–8 h
I.2	Single volleys to 5 electrodes (IEI: 2 ms), repeated every 10 s for 17 min.	Baseline medium, spontaneous bursting.	4; 13–16 div	4	4 h
I.3a	Single volleys to 8 electrodes in a vertical column (IEI: 2 ms), repeated every 10 s for 15 min.	Elevated magnesium (1–2 mM) to reduce spontaneous bursting.	3; 18–20 div	2	2 h
I.3b	Sets of 20 volleys (IVI: 50 ms) to 8 electrodes in a vertical column (IEI: 2 ms), repeated every 5 s for 8 min.	Elevated magnesium (1–2 mM) to reduce spontaneous bursting.	1; 17 div	2	2 h

^a Two cultures were each used twice, 6 days apart, resulting—for practical purposes—in 4 independent experiments.

Theoretically, this mapping procedure could cause plasticity that might make a culture less amenable to other plasticity, but the choice of electrodes was made up to 12 h before some experiments, so it seems unlikely that negative effects persisted. Also, in earlier experiments I had found no systematic differences between cultures that had never before been stimulated and those that had (data not shown).

Analysis

To test whether tetani had an effect on spontaneous bursting, I counted the number of bursts in the hour immediately before the tetanus (N_0), as well as in the hour after the tetanus (N_+). In order to be able to test whether the change coincident with the tetanus was larger than changes that occurred spontaneously, I also counted bursts in the second hour before the tetanus (N_-). I then computed the absolute value of




Figure 7.3: Responses to probing all electrodes at 27 voltages between 100 and 900 mV, normalized to average ASDR (horizontal line). Marks (+), and numbers in top left, indicate the stimulation voltage at which the spike rate was first doubled according to a linear fit. Numbers are gray if the baseline ASDR was not doubled for any usable stimulus voltage.

the change coincident with the tetanus, $\Delta N_{\text{coinc}} \equiv |N_+ - N_0|$, as well as the spontaneous change, i.e., the change attributable to drift, $\Delta N_{\text{spont}} \equiv |N_0 - N_-|$. To test whether ΔN_{coinc} was significantly larger than ΔN_{spont} , I made the assumption that burst counts can be approximated by a Poisson process, so that $\sigma_{N_x} = \sqrt{N_x}$. (Here x is $-$, 0 or $+$.) This made it possible to test for statistical significance by computing

$$t_N = \frac{\Delta N_{\text{coinc}} - \Delta N_{\text{spont}}}{\sqrt{\sigma_{\Delta N_{\text{coinc}}}^2 + \sigma_{\Delta N_{\text{spont}}}^2}} = \frac{\Delta N_{\text{coinc}} - \Delta N_{\text{spont}}}{\sqrt{N_+ + 2N_0 + N_-}}.$$

If $t_N \geq 1.63$, one can state with 95% confidence that the coincident change exceeds the spontaneous change. Conversely, if coincident changes are no greater than spontaneous changes, the values of t_N for different experiments should be symmetrically distributed around zero (with a variance of 1).

Similar tests were applied to the aggregate number of spikes in bursts per unit

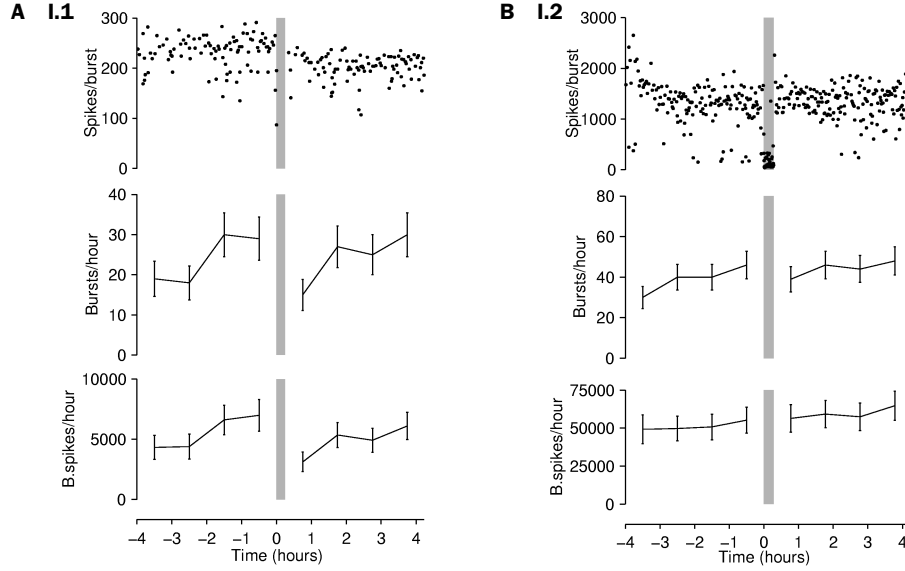


Figure 7.4: Examples of changes induced in spontaneous bursting by tetanization using several electrodes. **A** An exceptional case from protocol **I.1**, where the tetanus resulted in reduced burst rates and sizes. **B** A typical example from protocol **I.2**. Tetani are marked by gray bars. Top to bottom: number of spikes in individual bursts; number of bursts in successive one-hour time windows (with error bars based on assumed Poisson statistics); total number of spikes in bursts in successive hours.

time before and after tetanization, S_x . To estimate uncertainties, the average number of spikes per burst, μ_x , was calculated within each one-hour window, as well as its standard error, σ_{μ_x} . Then, since $S_x \equiv N_x \mu_x$, the uncertainty in S_x is

$$\sigma_{S_x} = \sqrt{\mu_x^2 \sigma_{N_x}^2 + N_x^2 \sigma_{\mu_x}^2}.$$

This made it possible to look for significant changes by comparing $\Delta S_{\text{coinc}} \equiv |S_+ - S_0|$ with $\Delta S_{\text{spont}} \equiv |S_0 - S_-|$ using

$$t_S = \frac{\Delta S_{\text{coinc}} - \Delta S_{\text{spont}}}{\sqrt{\sigma_{\Delta S_{\text{coinc}}}^2 + \sigma_{\Delta S_{\text{spont}}}^2}}.$$

Results

Only one experiment out of 28 showed significantly larger changes coincident with the tetanus than in spontaneous activity; this is the example shown in Figure 7.4A. Contrary to the observations by Maeda et al. (1998), these changes consisted of a de-




Figure 7.5: Results of significance tests for all 28 experiments on plasticity expressed in burst patterns. **A** Significance of changes in burst counts. **B** Significance of changes in counts of total number of spikes in all bursts occurring within an hour. In both panels, each point corresponds to one tetanus sequence, and represents the calculated t_N (in A) or t_S (in B) for the activity before and after that tetanus. For each protocol, points are spread out horizontally for visual clarity. Dotted lines represent (one-tailed) 95% confidence intervals, so on average out of every 20 data points, one would be expected to fall above the upper limit, and one below the lower limit. (Negative t -values indicate that the spontaneous change happened to be larger than the coincident change.) NB: The surprisingly close clustering around zero of the data for protocol **I.2** is real, and indicates that burst rates were more constant than Poisson statistics would predict.

crease in burst rates. Significance tests (Figure 7.5) revealed that, overall, changes coincident with tetani were no larger than spontaneous changes. In fact, in several cases, the magnitude of spontaneous change exceeded the coincident change (although in the case of protocol **I.3**, this might be due to transient effects of elevated magnesium). Thus, these data fail to confirm the plasticity results reported by Maeda et al. (1998).

It is not clear why one culture did show plasticity; apart from its reaction to tetanization, nothing set it obviously apart from its sister cultures. Certainly, the top panel of Figure 7.4A looks quite convincing, so it is attractive to hypothesize that something special happened. However, it could also have been a statistical fluke—testing at the $p < .05$ -level, one positive result out of 28 is not unexpected.

Series II: Changes Induced in Stimulus–Response Maps

Protocols

For two hours before and after tetanization, each of the 59 electrodes in the array were probed with test stimuli. Probes were delivered sequentially to each of 59 electrodes, with 3 s between pulses. The firing rates of each of 58 functional recording electrodes were observed, 10–50 ms after a test pulse to one of the 59 stimulation electrodes. By averaging the responses recorded within one hour (separately for each stimulation electrode–recording electrode pair), a response map was constructed. Changes in this map that occurred coincident with a tetanus were then compared to spontaneous changes.

Details of all experiments are summarized in Table 7.2. In most experiments, test probe amplitudes were fixed at 0.8 V. In some (protocol **II.2**), they were reduced in an attempt to define test pulses that would not evoke bursts, as follows. Before each experiment, each electrode in the array was probed at 19 voltages between 40 and 580 mV. Stimuli were presented in random order at 3 s intervals. For each electrode, the lowest voltage that ever evoked a burst was determined, and the amplitude of test pulses used in the main experiment was set to two thirds of this voltage. Unfortunately, it transpired that in the absence of stronger stimuli many of these relatively weak stimuli could still evoke bursts. (Note that this does not imply that plasticity was induced. Instead, it can likely be understood as follows. If stimuli are presented at a rate faster than the spontaneous burst rate of the culture, only some of the stimuli will trigger bursts, due to burst refractoriness. If the stimulation is a mix of strong and weak pulses, most bursts will be entrained to the strong pulses, because the refractory period for evoking bursts with weak pulses is likely to be longer than for evoking bursts with strong pulses. Thus, weak pulses never have a chance to evoke bursts. When there are no strong pulses, the network is never made refractory to the weaker pulses, which may thus get their chance to evoke bursts. This mechanism makes the entrainment of bursts context-dependent, but does not imply plasticity in the sense used by most researchers of long-term synaptic plasticity.)

In most experiments, tetanization consisted of several trains of stimuli delivered to a single electrode. Each train consisted of 20 pulses, at 50 ms intervals. A complete

Table 7.2: Details of experiments on plasticity expressed in stimulus–response maps

Proto-col	Tetanus target	Probe amplitude	Conditions	No. and ages of cultures	Tet. per culture	Inter-vals
II.1	Single electrode.	Fixed, 0.8 V.	Baseline medium, spontaneous bursting.	4; 17–22 div	2	2 h
II.2	Single electrode.	Fixed, 0.8 V.	Bursts completely suppressed by 50 Hz background stimulation distributed over 20–40 electrodes, except during tetanization.	3 ^a ; 17–22 div	2	2 h
II.3	Single electrode.	Fixed, 0.8 V.	Spontaneous bursts suppressed by 1 mM magnesium.	3; 26–28 div	2	2 h
II.4	Single electrode.	Reduced (see text).	Spontaneous bursts suppressed by 2 mM magnesium.	4; 29–32 div	4	2 h
II.5a	8 electrodes, as in I.3a .	Range of voltages, 100–900 mV.	Spontaneous bursts suppressed by 1–2 mM magnesium.	3; 18–20 div	4	2 h
II.5b	8 electrodes, as in I.3b .	Range of voltages, 100–900 mV.	Spontaneous burst suppressed by 2 mM magnesium.	1; 17 div	4	2 h

^a In a 4th experiment, burst suppression did not work sufficiently well. Those data were excluded from further analysis.

tetanization sequence consisted of 20 trains, with 2 s between trains. Before experiments, the relation between stimulation voltage and array-wide response strength was determined first, for each electrode (see *Choice of electrodes* under Series I). For tetanization, I then chose electrodes that evoked strong culture-wide responses.

In a final set of experiments (protocol **II.5a** and **b**), tetanic stimulation was applied to clusters of electrodes, as in **I.3a** and **b**. In an attempt to get to a finer

level of detail, I probed for test responses using many different pulse amplitudes. I recorded for two hours before and after each tetanus. In each hour, 9 test sequences of 27 pulses to each of the 59 electrodes were presented. (The same sequences were used as for choosing tetanization electrodes.)

As in Series I, several experiments were usually performed on each culture, with several hours between experiments.

Analysis

For comparing responses to test stimuli before and after tetanization, a similar approach was used as for comparing burst rates in Series I. Within each one-hour window, I computed the mean number n^{SR} of spikes detected on electrode R (for ‘Recording’), 10–50 ms after a test stimulus on electrode S (for ‘Stimulation’), as well as its SEM. Significance tests were then based on $\Delta n_{\text{coinc}}^{SR} \equiv |n_+^{SR} - n_0^{SR}|$ and $\Delta n_{\text{spont}}^{SR} \equiv |n_0^{SR} - n_-^{SR}|$. The mean latency τ^{SR} of the first spike on response electrode R after stimulation on S was also computed, with its SEM. For this calculation, trials without responses were ignored.

I wanted to know not only whether significant changes in individual responses occurred (as measured through n^{SR}), but also whether such changes were specific to stimulation sites, as reported by Jimbo et al. (1999). In that case, responses on all or most recording sites to one given stimulation site should be up- or down-regulated together, i.e., the tetanus-coincident changes $\Delta n_{\text{coinc}}^{SR}$ and $\Delta n_{\text{coinc}}^{S'R'}$ should have the same sign when $S = S'$. I also considered the converse hypothesis: changes might be specific to recording sites, in other words, responses on a given recording site could be up- or down-regulated independently of which stimulation site was used to evoke the response. In that case, $\Delta n_{\text{coinc}}^{SR}$ and $\Delta n_{\text{coinc}}^{S'R'}$ would have the same sign when $R = R'$. To test these hypotheses, I calculated

$$\Delta n_{\text{coinc}}^{\text{stim}} \equiv \sum_S \left| \sum_R \left(n_+^{SR} - n_0^{SR} \right) \right|,$$

which should deviate significantly from zero if Jimbo’s hypothesis holds, as well as

$$\Delta n_{\text{coinc}}^{\text{rec}} \equiv \sum_R \left| \sum_S \left(n_+^{SR} - n_0^{SR} \right) \right|,$$

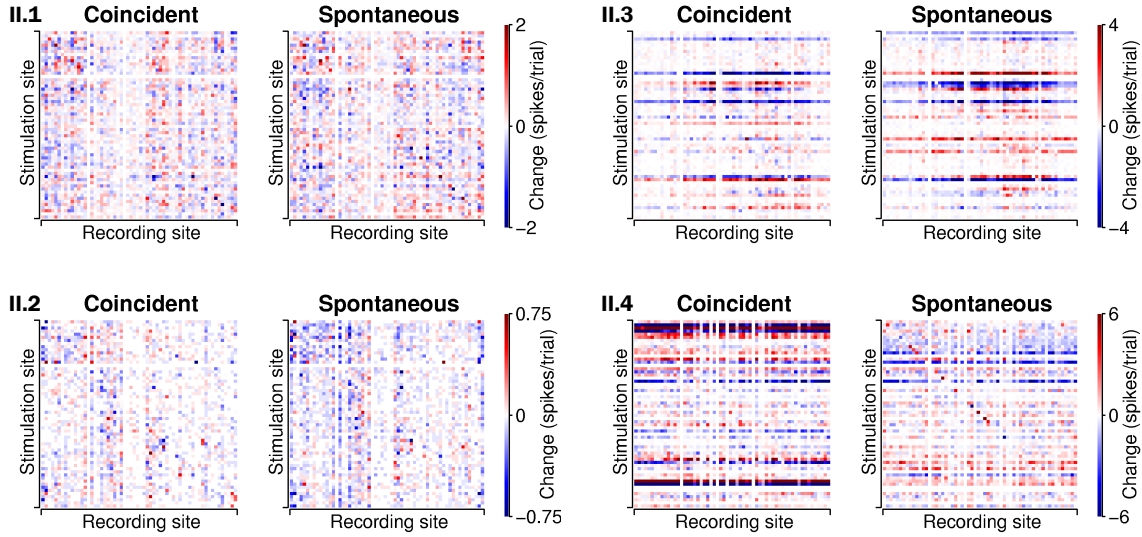


Figure 7.6: Examples of results of experiments on plasticity expressed in stimulus–response maps. One example is shown for each protocol (indicated in top left). Colored pixels represent changes in the average number of spikes on a given recording electrode 10–50 ms after a test pulse to a given stimulation electrode. Examples **II.3** and **II.4** show changes that are stimulation-site specific—note the horizontal stripes of similar coloration—but in all cases spontaneous changes (*right sub-panels*) are comparable in magnitude to changes coincident with tetani (*left sub-panels*). Examples were picked blindly.

which should deviate significantly from zero if changes were recording-site specific. (If changes are randomly distributed, both inner sums will have a roughly equal number of positive and negative terms, and hence will not be very large.) The uncertainties in these numbers were also calculated, as were $\Delta n_{\text{spont}}^{\text{stim}}$ and $\Delta n_{\text{spont}}^{\text{rec}}$, so that a significance test could be used to assess whether changes coincident with tetani indeed exceeded spontaneous changes.

Results

In many cases, especially in protocols **II.3** and **II.4**, stimulation-site specific changes exceeded recording-site specific changes, in agreement with Jimbo et al. (1999); see Figure 7.6 for examples. However, spontaneous stimulation-site specific changes were also observed, and no significant difference was seen between changes coincident with tetani and spontaneous changes (Figures 7.7 and 7.8).

Stimulation-site specific changes were mostly confined to experiments with extra

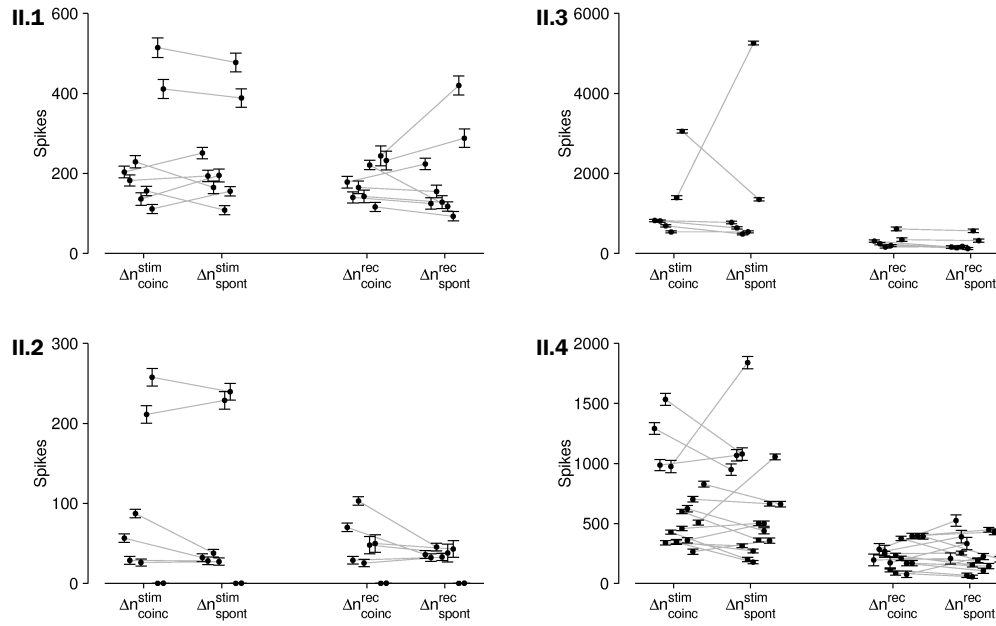


Figure 7.7: Summary of results of all 38 experiments on plasticity expressed in stimulus-response maps. Stimulation-site specific changes coincident with tetani ($\Delta n^{\text{stim}}_{\text{coinc}}$) are compared with spontaneous changes ($\Delta n^{\text{stim}}_{\text{spont}}$). Recording-site specific changes ($\Delta n^{\text{rec}}_{\text{coinc}}$ and $\Delta n^{\text{rec}}_{\text{spont}}$) are also compared. Stimulation-site specific changes were more pronounced than recording-site specific changes. However, changes coincident with tetani did not systematically exceed spontaneous changes. Each panel contains data from one protocol (indicated in top left). Data points were slightly jittered horizontally for visual clarity.

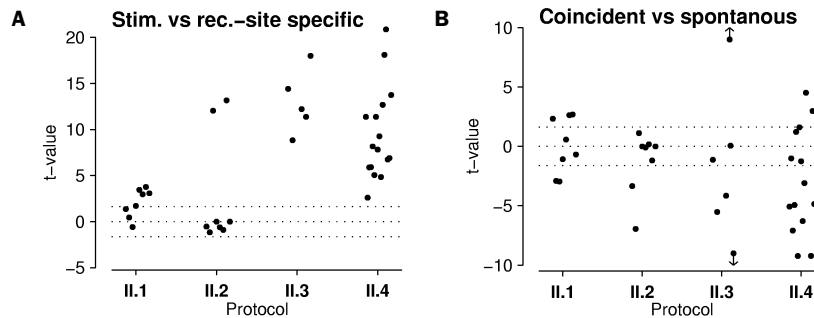


Figure 7.8: Significance tests for experiments on plasticity expressed in stimulus-response maps. **A** Stimulation-site specific changes ($\Delta n^{\text{stim}}_{\text{coinc}}$) were significantly larger than recording-site specific changes ($\Delta n^{\text{rec}}_{\text{coinc}}$) in most protocols. (This corresponds to the observation that horizontal lines are more prominent than vertical lines in Figure 7.6.) **B** Stimulation-site specific changes coincident with tetani ($\Delta n^{\text{stim}}_{\text{coinc}}$) did not systematically exceed spontaneous changes ($\Delta n^{\text{stim}}_{\text{spont}}$). In fact, spontaneous changes were larger than coincident changes in at least as many cases. (This corresponds to the observation that horizontal lines in Figure 7.6 are no more prominent in the 'Coincident' panels than in the 'Spontaneous' panels.) Dots with arrows represent two cases with off-scale t-values (+58 and -29 respectively).




Figure 7.9: Array-wide responses to individual stimuli in an experiment from protocol **II.3**. Spikes recorded on any electrode, 10–50 ms after a stimulus, are counted together and represented on a gray scale. Each pixel is one stimulus response; stimulation order is top to bottom, then left to right. Separating white bars indicate division into 1-hr windows; the dashed line represents the tetanus. Stimuli that evoked bursts show up as black pixels. The arrow points to a burst that, on its own, was responsible for the top-most horizontal blue stripe in Figure 7.6, example **II.3**, right sub-panel.

magnesium in the medium, and did not occur often in baseline conditions or with bursts quieted by electrical background stimulation. I hypothesized that this could be due to the differences in burst patterns between these three conditions: Inspection of several examples had revealed that evoked bursts in the presence of magnesium were rare events that involved a very large number of spikes compared to ‘typical’ responses. By contrast, evoked bursts could have a continuum of sizes in baseline medium, and were absent where electrical background stimulation was used. In the presence of magnesium, individual evoked bursts could entail more than 30 spikes on each of 20–50 electrodes in the time window used for measuring responses, 10–50 ms post-stimulus. Thus, if by random circumstance stimuli to a given electrode evoked one extra burst in an hour, that could amount to a significant change in spike counts, all linked to that electrode. For instance, the response indicated by the arrow in Figure 7.9 contained 1113 spikes in the time-window 10–50 ms after the stimulus, which amounts to 0.96 spikes per recording electrode, even after averaging with the other 19 trials in the same hour that did not evoke bursts. This is sufficient to yield a very distinct horizontal stripe in the representation of Figure 7.6. In baseline conditions

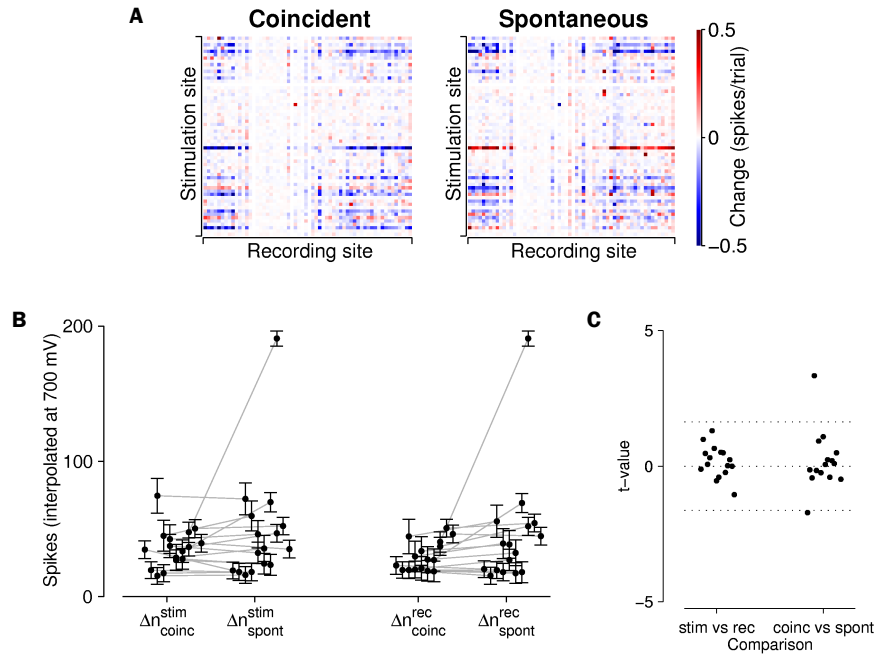


Figure 7.10: Results of experiments on stimulus–response maps probed at different voltages (protocol II.5). **A** An example (cf. Figure 7.6). **B** Comparison of changes (cf. Figure 7.7). **C** Significance tests for stimulation-site specific vs. recording-site specific changes (‘stim vs. rec’), and for coincident vs. spontaneous changes (‘coinc vs. spont’).

this effect was invisible, because bursts were more frequent and not as large.

In protocols II.5a and b, where stimuli of many different voltages were used on each electrode, I considered each of the ~ 3400 stimulus–response pairs in turn, and fitted a straight line to the response 10–50 ms post-stimulus vs. voltage, independently for each hour. The fit value at 700 mV was then compared before and after tetanization, just as n^{SR} was in other protocols. This protocol did not yield any positive results (Figure 7.10).

Series III: Changes in Specific Responses

Introduction

Paired-pulse stimulation of a presynaptic and a postsynaptic cell with sharp intracellular or patch electrodes is a well established protocol for inducing plasticity. Depending on the timing between the pulses, both long-term potentiation (LTP) and

long-term depression (LTD) can readily be obtained. I tested whether a similar protocol could be used in cultures growing on MEAs by stimulating pairs of electrodes that shared some response sites. If test pulses to electrode S_1 and test pulses to electrode S_2 both evoked responses at electrode R , those responses would likely be affected if paired-pulse stimulation of S_1 and S_2 modified any synapses. Moreover, if there are several electrodes that are responsive to both S_1 and S_2 , it should be more likely that S_1 and S_2 share synaptic pathways which may be modifiable. This directed search for plasticity is more sensitive than a global assay of changes in responses anywhere evoked by stimulation through any electrode, because for a global assay so many tests have to be performed that a very tight probability bound must be used to avoid a deluge of false positives. Pietro Perona (personal communication) has argued that that might easily hide small but significant effects on specific electrodes.

Protocols

I looked for plasticity in specific responses using four protocols (Table 7.3). Test pulses with fixed amplitude (0.6 or 0.8 V) were delivered to the two electrodes used for tetanization and to four other electrodes used as controls. Test pulses were presented in cyclic order, with 1 or 5 s between pulses. Where electrical burst quieting was used, quieting was suspended for 50 ms before and 200 ms after a test pulse, so that the responses to test pulses could be measured without interference. Tetanization consisted of trains of pulse pairs: one pulse to the ‘leader’ electrode, one pulse to the ‘follower’ electrode, 5 or 10 ms later (inter-electrode interval; IEI). Each train consisted of 20 pairs, with 50 or 100 ms between pairs (inter-pair interval, IPI). A complete tetanization sequence contained 20 or 150 trains, at 2 or 6 s intervals (inter-train interval; ITI). During tetanization, quieting with electrical stimulation was suspended, except in protocol II.4.

Selecting Stimulation Pairs

To find pairs of electrodes that shared synaptic targets, I started with the set of candidate stimulation electrodes identified in initial probing (see *Choice of electrodes* under Series I), and delivered 50 pulses of fixed amplitude to each of them at 3 s in-

Table 7.3: Details of experiments on plasticity induced in specific responses

Proto-col	Tetanus	Probing	Conditions	No. and ages of cultures	Tet. per culture	Inter-vals
III.1	20 trains (ITI: 2 s) of 20 pulse pairs (IPI: 50 ms; IEI: 5 ms).	Single pulses of 0.8 V, at 5 s intervals.	Baseline medium, with spontaneous bursts.	4; 13–16 div	2	2 h
III.2	20 trains (ITI: 2 s) of 20 pulse pairs (IPI: 50 ms; IEI: 5 ms).	Baseline medium, with spontaneous bursts.	Electrical burst quieting as in II.2 .	3 ^a ; 13–16 div	2	2 h
III.3	20 trains (ITI: 2 s) of 20 pulse pairs (IPI: 50 ms; IEI: 5 ms).	Baseline medium, with spontaneous bursts.	Spontaneous bursts suppressed by 1 mM magnesium.	4; 25–28 div	2	2 h
III.4	150 trains (ITI: 6 s) of 20 pulse pairs (IPI: 100 ms; IEI: 10 ms).	Single pulses of 0.6 V, at 1 s intervals.	Electrical burst quieting as in II.2 .	4; 20–23 div	2	2 h

^a In a 4th experiment, burst suppression did not work sufficiently well. Those data were excluded from further analysis.

tervals. For each stimulation electrode, I determined the set of recording electrodes that responded with a spike rate elevated above baseline by at least 5σ , 10–50 ms post-stimulus. I then selected from about 2000 available stimulation pairs those that had the largest overlap in their response sets, but that also had a number of electrodes responding to one member of the pair but not to the other (Figure 7.11). For each experiment, I adjusted the selection criteria (size of overlap, and number of electrodes that responded to only one member of the pair) until the number of candidate pairs was no larger than 20, and at least 10. This procedure typically yielded overlap sets of 5–20 electrodes, and 5–10 electrodes that responded to one but not the other member of the pair. Up to three pairs were then selected for experimental use that did not share members, and that had distinct overlap sets as far as possible. The




Figure 7.11: Response maps for several pairs of stimulation electrodes that evoke activity on overlapping sets of recording electrodes. Dots represent electrodes, laid out according to array geometry. Black and gray circles mark stimulation sites. Black and gray dots are electrodes that respond to only one of the stimulus electrodes, and not the other; electrodes marked by black '+' respond to both.

distance between stimulation electrodes within a pair was not a consideration. In most experiments, tetanic sequences were delivered to two different pairs, two hours apart, while a third pair served as an internal control. Only those six electrodes were probed during the main experiment.

Results

Examples of results show substantial spontaneous changes (Figure 7.12). Only in protocol **III.3** did changes coincident with tetani exceed spontaneous changes judging from the graphs, but even most of those changes were not significant (Figure 7.13). Changes in latency to first spike were not significant either (Figure 7.14). In short, no significant plasticity was induced by paired-pulse tetanization.

Still, even if the changes coincident with tetani are not significantly larger than spontaneous changes in individual cases, it could be that coincident changes are

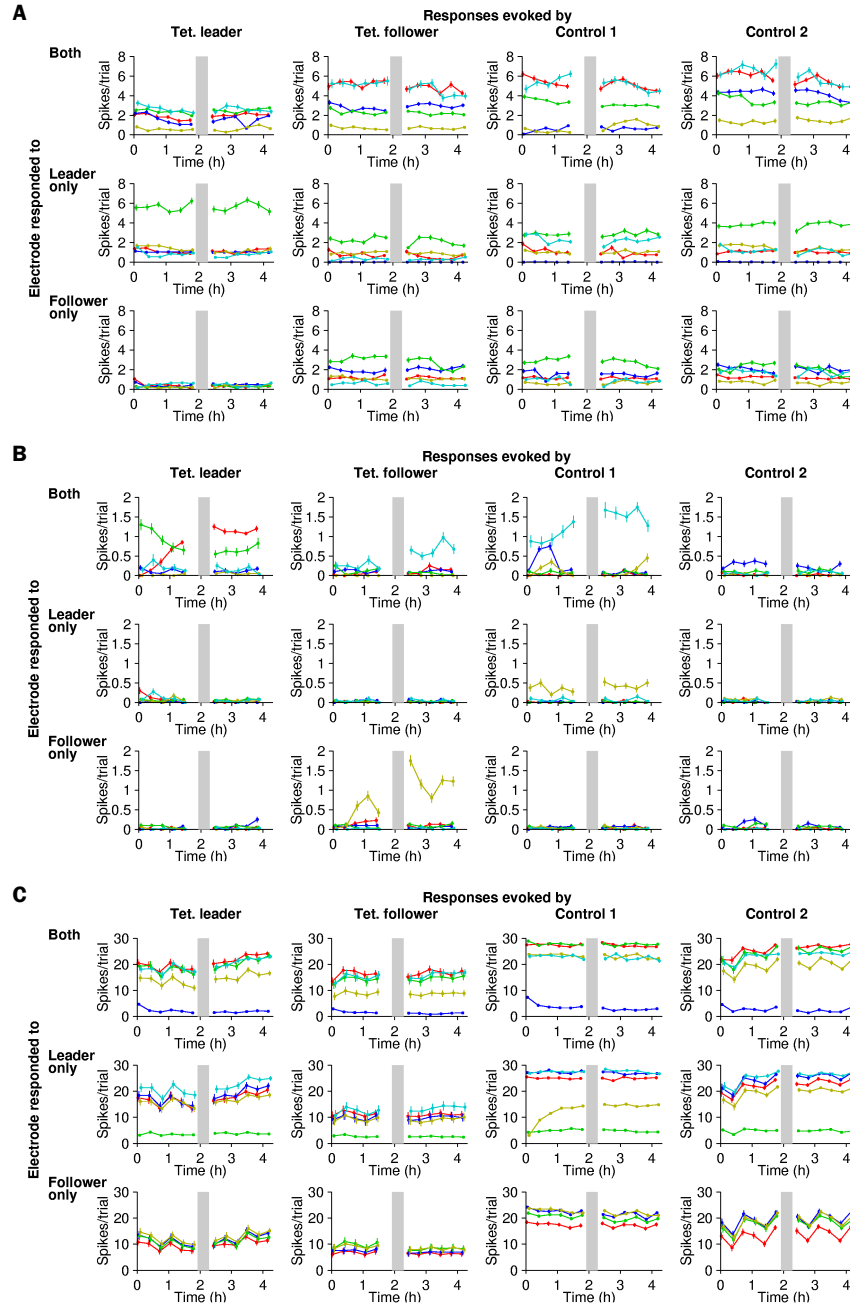


Figure 7.12: Examples of changes induced in specific responses by paired-pulse tetanization. Each curve represents responses recorded on a different electrode. Data points are mean and SEM for 40 consecutive stimuli. Tetani are marked by gray bars. Graphs are organized left to right by the electrode that evoked the responses: ‘Tet. leader’ is the first-stimulated electrode in the paired-pulse tetanus; ‘Tet. follower’ is the second; ‘Control 1’ and ‘Control 2’ are non-tetanized electrodes. Top to bottom, graphs count spikes evoked on recording sites that responded to both leader and follower; only to the leader; and only to the follower. **A** A typical example from protocol III.1. **B** A typical example from protocol III.2. **C** An example from protocol III.3 with positive results.

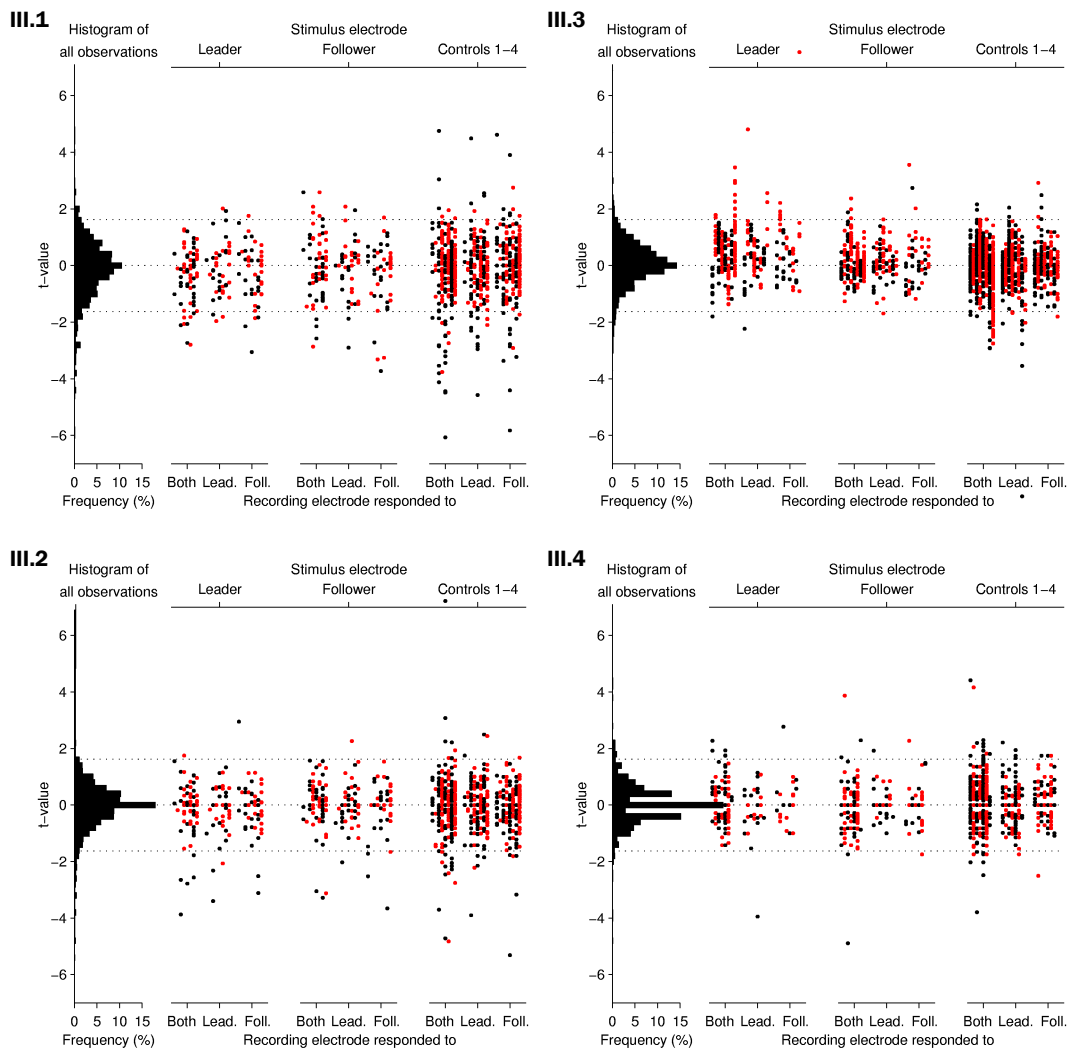


Figure 7.13: Results of significance tests for plasticity in spike counts of specific responses. Each dot represents a t-value comparing tetanus-coincident changes in average spike counts on one electrode R , 10–50 ms after a test pulse to another electrode S ($\Delta n_{\text{coinc}}^{SR}$) with spontaneous changes in the same quantity ($\Delta n_{\text{spont}}^{SR}$). Averages are taken over 120 trials in one hour. Data presentation is organized by protocol (*bold numbers*), by the role of the electrode that evoked the response ('Leader': the electrode that received the first of each of the pulse pairs that made up the tetanus sequence; 'Follower': the electrode that received the second pulses; and 'Controls': electrodes that were not used for tetanization), and by how the recording electrode responded to test pulses in pre-experimental probing (see *Methods*). Within each cluster, data points are organized by culture. Black dots relate to the first tetanus experiment on a culture, red to the second (2 hours later). Histograms are aggregates for all experiments of a given protocol. Only recording electrodes that responded significantly to either or both of the tetanus electrodes in pre-experimental probing are represented in this figure. Dotted lines mark the 90% confidence interval. Given the total number of data points, the number of points outside this interval is as expected.

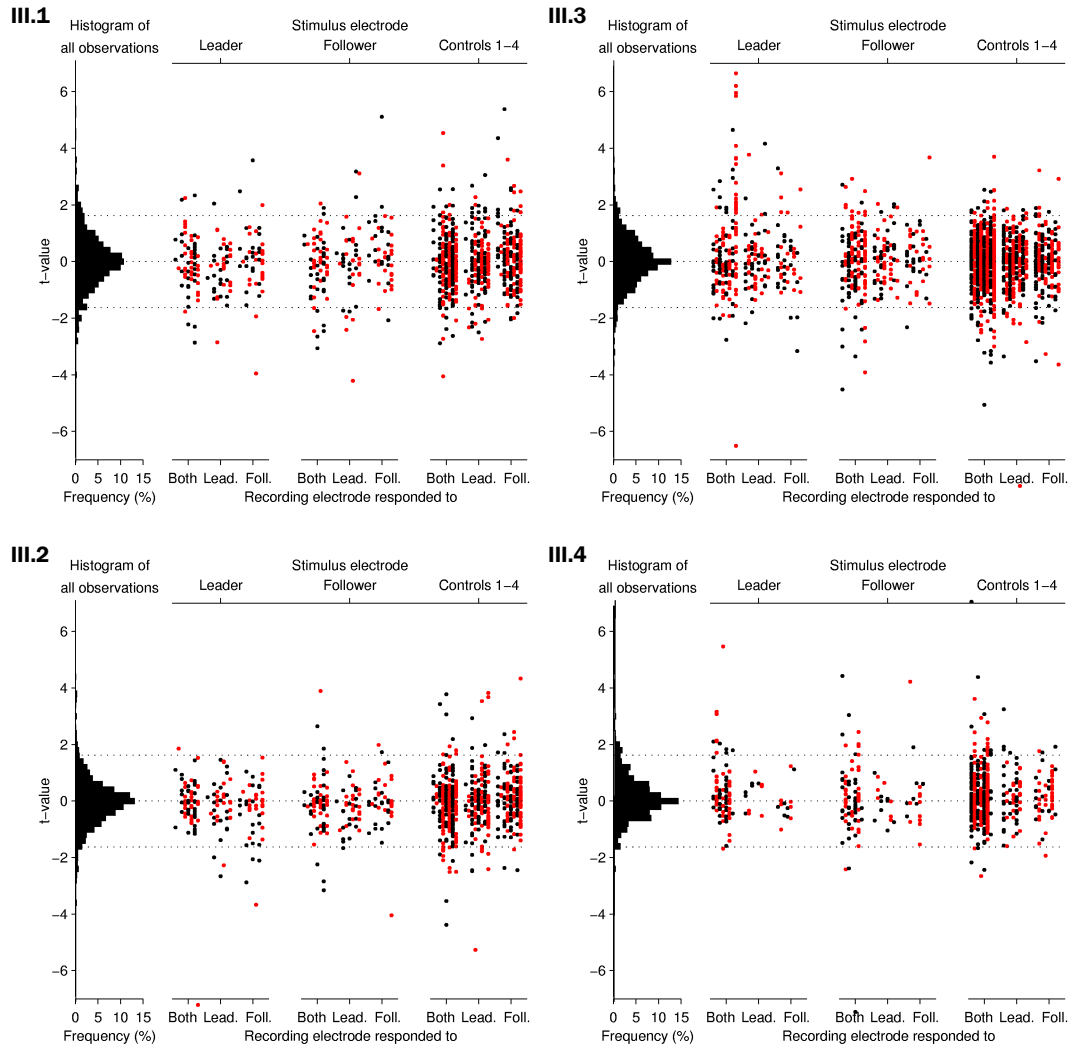


Figure 7.14: Results of significance tests for plasticity in latency-to-first-spike of specific responses. Each dot represents the significance of changes in the average latency to the first spike on an electrode R after a test pulse to another electrode S . Averages are taken over 120 trials in one hour; trials in which no spikes were elicited on electrode R were discarded. Data presentation is organized as in Figure 7.13.

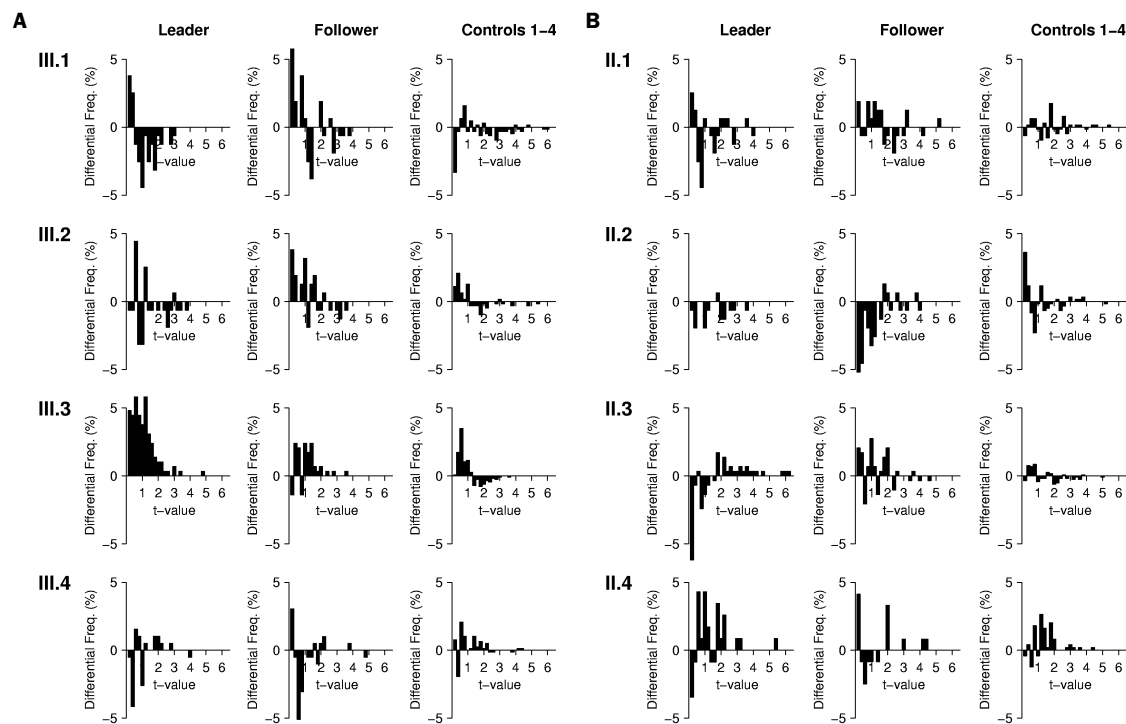


Figure 7.15: Small differences in the magnitudes of tetanus-coincident vs. spontaneous changes in experiments on plasticity in specific responses. Difference in number of stimulus–response pairs for which change coincident with tetani exceeded spontaneous change at a given t -value and number of pairs for which the spontaneous change exceeded the coincident change at the same (but opposite) t -value. Data presentation is organized by protocol (numbers in corners) and by the role of the electrode that evoked the response, as in Figures 7.13 and 7.14. **A** Changes in average spike counts (cf. Figure 7.13). **B** Changes in average latency to first spike (cf. Figure 7.14).

slightly larger on average than spontaneous changes. If that were the case, positive t -values should be more common than negative t -values. To test this idea, I counted the number of stimulus–response pairs that had a positive t -value t , and subtracted the number of pairs that had an equal but opposite t -value $-t$. Positive t -values were indeed over-represented in several protocols. (Figure 7.15). The effect was most notable for changes in spike count in protocol **III.3**, so I looked at the coincident and spontaneous changes in all experiments of that protocol in detail (Figure 7.16). The differences were not large, but in most experiments the coincident changes were indeed slightly larger in magnitude than the spontaneous changes.

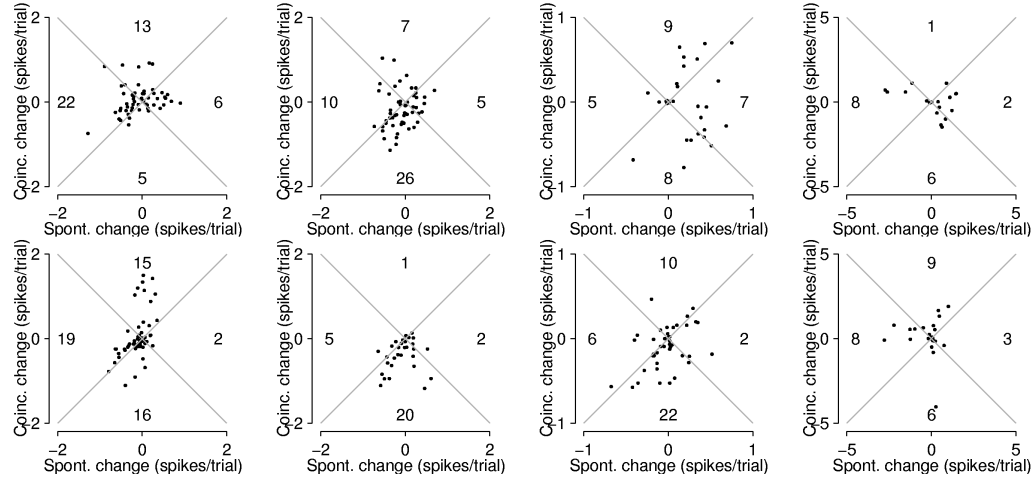


Figure 7.16: Direct comparison of changes coincident with tetani and spontaneous changes in spike counts of specific responses in protocol III.3. Numbers in top and bottom quadrants count cases in which the change coincident with a tetanus ($\Delta n_{\text{coinc}}^{SR}$) exceeded the spontaneous change ($\Delta n_{\text{spont}}^{SR}$). Numbers in left and right quadrants count cases where $\Delta n_{\text{spont}}^{SR}$ exceeded $\Delta n_{\text{coinc}}^{SR}$. Points on the gray lines have coincident and spontaneous changes of the same magnitude. Each column of graphs shows data relating to two tetanus sequences applied to one of four cultures.

The observation that changes occurred in responses to the ‘leader’ electrode but not to the ‘follower’ electrode was somewhat surprising, given that the short inter-electrode delay in the tetanus could easily be compensated by differences in axonal and synaptic propagation delays. Still, inspection of several individual cases revealed that postsynaptic responses with a timing precision of 5–10 ms were fairly common, making it possible for a 5 ms IEI in tetanization to make a substantial difference.

Changes in the Probability of Evoking Bursts

In addition to testing for changes induced in stimulus–response maps, I investigated whether tetanization had an effect on the ability of test pulses to evoke bursts. I counted spikes across the array 100–500 ms after each stimulus, and found a clearly bimodal distribution in each experiment, making it very easy to distinguish trials that evoked bursts from those that did not. For each stimulation electrode, I determined the fraction of stimuli that evoked bursts in one-hour windows. I calculated spontaneous and tetanus-coincident changes in this fraction, and found that they

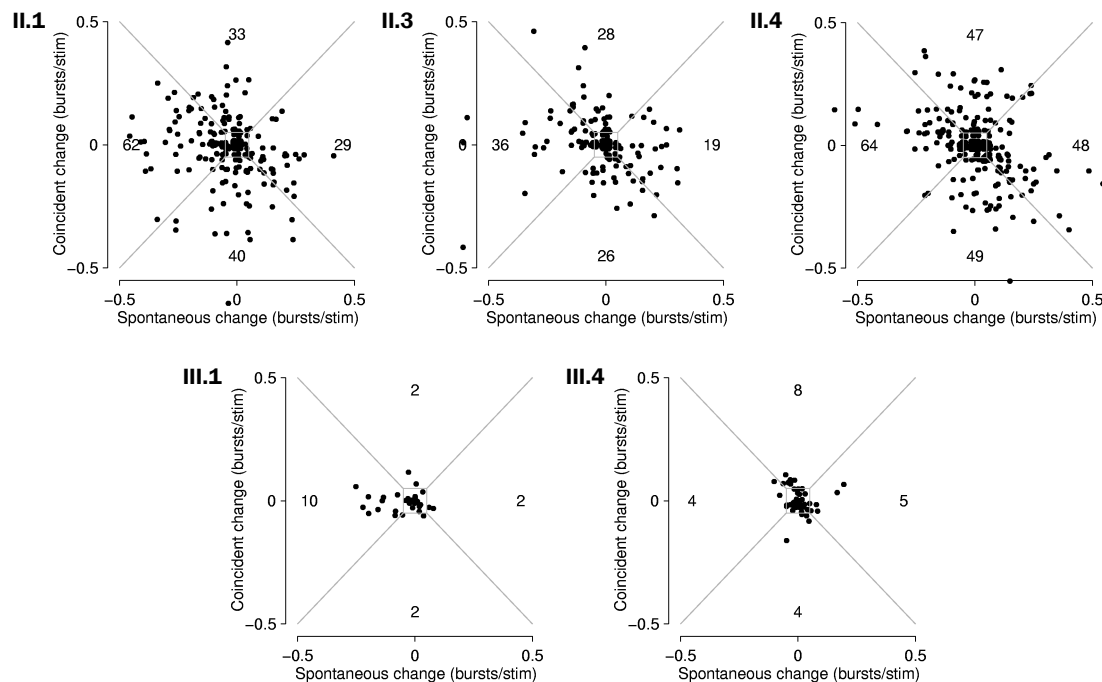


Figure 7.17: Changes in the probability that a test pulse evoked a burst. For each protocol in which test pulses could evoke bursts, we compare changes coincident with a tetanus to spontaneous changes. Each dot represents results for one stimulation electrode in one experiment. Thus, protocols **II.1–4**, which involved probing all electrodes, yielded far more data points than protocols **III.1–4**, in which only 6 electrodes were probed. Numbers in the figures count the observations in each quadrant (changes less than $\pm 5\%$ were not counted). The upper and lower quadrants contain cases in which the changes coincident with a tetanus exceed the spontaneous change in absolute value; in the left and right quadrants spontaneous changes were larger. Thus, if tetani induce significant changes, data clouds should resemble vertically stretched ellipses. In protocols **II.1–4**, measurements are based on only 20 probes each, so changes are implicitly quantized in steps of 5%. For visual clarity in these cases, dots were randomly displaced by $\pm 2\%$ horizontally and vertically.

were equally large (Figure 7.17). In conclusion, tetanization did not affect the probability of evoking bursts by test pulses.

Discussion

Why did I not see so many of the plastic changes reported by others? One possibility is that subtle differences in culturing conditions made our cultures less amenable to inducing plasticity. Certainly, substantial differences existed even between different plating batches created within our lab (see Chapter 4). However, the negative results

reported here are based on 112 experiments on 18 cultures from 4 dissections, so it is unlikely that I simply had an unlucky pick. It remains possible that Jimbo et al. (1999) had an exceptionally lucky pick, since they recorded from only 8 cultures (from an unspecified number of platings, possibly only one), and concentrated on only one. (They explicitly state that they chose the one ‘which demonstrated [their] conclusions most clearly’.) Such a ‘low-N’ explanation is less plausible for the discrepancies between the present results and those of Maeda et al. (1998), since they tested 26 cultures.

With elevated extracellular magnesium concentration, the results on stimulus–response map plasticity obtained here resemble the results of Jimbo et al. (1999) superficially (compare our Figure 7.12 with their Figure 2), but without underlying true plasticity. Instead, the present results could be attributed to the stochasticity involved in evoking bursts. Accordingly, the changes observed coincident with tetani were no larger than spontaneous changes; Jimbo et al. do not report on spontaneous changes. It is not clear whether this same mechanism contributed to their results. It is unlikely that it was solely responsible, since their observed changes depended on the correlation between the responses to test stimuli and the responses to stimulation of the tetanized electrode. I saw no such dependence. In any case, the present results fail to support their report of tetanus-induced plasticity.

Indirectly, my results are also in conflict with the more recent ‘learning’ results of Shahaf and Marom (2001) and Ruaro et al. (2005): if strong tetani that are known to induce plasticity when applied intracellularly fail to have any effect when applied extracellularly through MEA electrodes, even though they do evoke bursts, it is difficult to conceive how less strong stimulation sequences applied through MEA electrodes could be more effective.

These experiments also failed to confirm our own hypothesis that spontaneous bursts interfere with plasticity (Wagenaar et al., 2005b [Chapter 6 in this thesis]), as burst quieting with distributed electrical stimulation did not result in positive plasticity results. It remains possible that the thousands of bursts cultures experienced before these experiments took place had already done irreparable damage to synapses’ readiness to change in reaction to tetanization, but at present we have no data to put that idea to the test.

It should be noted that I am not concluding that cortical networks in dissociated culture are not plastic. First of all, I was able to reproduce the only reported form of plasticity that was previously seen by more than one group (Eytan et al., 2003). Several other figures in this chapter also show evidence that responses to test pulses changed considerably over time. Additionally, the burst suppressing effect of elevated magnesium wore off within several hours, suggesting a general up-regulation of excitability, since it seems unlikely that the cells somehow removed the magnesium from the medium. However, these plastic changes occurred spontaneously, and were not caused by our tetanic stimulation sequences.

Finally, it remains possible that tetani did induce plasticity, but that my probes are not sensitive enough to detect it. For instance, even though burst counts and firing rates did not change significantly in most experiments with protocols **II.1–3**, it is conceivable that the spatiotemporal extent of bursts changed in more subtle ways. Modeling studies (Chao et al., 2005) suggest that this may be a realistic scenario. Additionally, it may be that so many cells contributed to the spike count on each individual electrode, that one cell would be drowned out. If that were the case, spike sorting might reveal small plastic changes that are invisible currently.

Certainly, I would not be so confident about these negative results if I were the only one unable to replicate plasticity studies. However, other students in Steve Potter's group have tried and failed to reproduce the results of Marom (Douglas Bakkum, personal communication) and Jimbo (Radhika Madhavan, personal communication). Moreover, Van Staveren et al. (2005) obtained similar negative results independently. A thorough collaborative investigation into the sources of these discrepancies could be very helpful to move the field forward.

Appendix A

Culturing Methods

The experiments described in this thesis were performed on cortical cell cultures from embryonic Wistar rats. This appendix describes the details.

Tissue Extraction

Timed-pregnant Wistar rats (Simonsen Laboratories, see list of vendors, Box A.2) were euthanized using CO₂, according to NIH-approved protocols, by our lab technician, at day 18 of gestation (E18). The embryos were removed and euthanized by chilling and decapitation. Cortices were extracted, while hippocampi were set aside for use by other researchers. Cortices were stored at 4 °C in calcium and magnesium-free Hanks' balanced salt solution (CMF Hanks; Gibco, 14170-120) for up to 3 hours while preparing other materials.

MEA Preparation

Multi-electrode arrays were cleaned in 3% "BM solution" (ALA Scientific) for 15 minutes, then rinsed and left in water overnight. MEAs, lids (see below) and dissection tools were then autoclaved at 122 °C for 10 minutes followed by 20 minutes of drying. After MEAs had cooled to room temperature, they were treated with poly-ethyleneimine solution (PEI; see Box A.3): 750 mL PEI solution was applied to the MEA and

Box A.2: Vendor addresses

Vendor	Address	On the web
ALA Scientific	Westbury, NY	www.alascience.com
Falcon / BD Biosciences	San Jose, CA	www.bdbiosciences.com
Fisher Scientific	Hampton, NH	www.fishersci.com
Gibco / Invitrogen	Carlsbad, CA	www.invitrogen.com
Hamilton	Reno, NV	www.hamiltoncompany.com
HyClone	Logan, UT	www.hyclone.com
Irvine Scientific	Santa Ana, CA	www.irvinesci.com
MultiChannel Systems	Reutlingen, Germany	www.multichannelsystems.com
Sigma-Aldrich	St. Louis, MO	www.sigmaaldrich.com
Simonsen Laboratories	Gilroy, CA	www.simlab.com

Box A.3: Poly-ethylene-imine solution (PEI)

PEI solution consists of 0.05% by weight poly-ethylene-imine (from 50% w/v aq. sol., Sigma P3143), in borate buffer solution (BBS): 3.10 g boric acid (Fisher, A73-500), and 4.75 g borax (Sigma, B0127) in 1 L ddH₂O. To make the BBS, adjust pH to 8.4 prior to and after stirring. Dissolving the ingredients takes a lot of effort and is best done on a stirrer/hot plate. PEI solution is stored at 4 °C.

left for 45–60 minutes at room temperature. MEAs were then quadruply cleaned by removing the liquid with vacuum suction and replacing it with 1–2 mL ddH₂O. After removing the fourth wash, MEAs were left to air-dry in the culture hood for 30 minutes. A drop of 20 µL laminin solution (0.020 mg/mL laminin, Sigma, L2020; in DMEM, Box A.4; stored in aliquots at –20 °C) was applied to the center of the MEA. The laminin drop size always matched the drop size of the cells added later. After laminin application, the MEA was left in the hood, inside a glass Petri dish, for 30 minutes. Laminin was removed with vacuum suction immediately prior to plating.

Cell Preparation

The following procedure was modified from a protocol by Segal et al. (1998). Cell preparation was started 15 minutes after application of PEI solution to the MEAs, to

Box A.4: Dulbecco's modification of Eagle's basal medium (DMEM; Irvine Scientific, 9024)

Component	mg/L	Component	mg/L
Sodium Chloride	6400	L-Arginine HCl	84
Potassium Chloride	400	L-Cystine 2 HCl	63
Glucose	4500	L-Glutamine	— ^a
Folic Acid	4	Glycine	30
Inositol	7	L-Histidine HCl · H ₂ O	42
Nicotinic Acid Amide	4	L-Isoleucine	105
Riboflavin	0.4	L-Leucine	105
Thiamine HCl	4	L-Lysine HCl	146
Ferric Nitrate	0.1	L-Methionine	30
Phenol Red, Na salt	15	L-Phenylalanine	66
Sodium Phosphate, monobasic (NaH ₂ PO ₄ · H ₂ O)	125	L-Serine	42
Pantothenic Acid, Ca salt	4	L-Threonine	95
Pyridoxine HCl	4	L-Tryptophan	16
Calcium Chloride, anhyd.	200	L-Tyrosine 2 Na · 2 H ₂ O	104
Magnesium Sulfate, anhyd. (MgSO ₄)	98	L-Valine	94
Choline Chloride	4		
Sodium Bicarbonate	3700		

^a To be added in the form of GlutaMax (Gibco, 35050-061).

be ready 30 minutes after application of laminin solution. Cortices were transferred to 3 mL Segal's medium (Box A.5) using a wide-bore pipette tip (Hamilton 11003-57), and cut into $\sim 1 \text{ mm}^3$ pieces using small scissors under a dissection scope. DNase (500 μL of 1 mg/mL, Sigma D7691; in CMF Hanks; stored in aliquots at -20°C) was added prior to cutting to prevent bits of tissue sticking to each other. Tissue bits were then transferred to papain solution (Box A.6) and incubated for 30 minutes, gently stirring every 5 minutes. At this point, the medium looks slightly cloudy. Papain solution was then replaced by 1 mL plating medium (Box A.7). Next, the cortex bits were dissociated by trituration: with a (narrow) pipette tip 1 mL of plating medium

Box A.5: Segal's medium

Segal's medium (Segal et al., 1998) contains:

Component	mM	FW (g/mol)	mg for 500 mL
MgCl ₂ · 6 H ₂ O	5.8	203.31	590
CaCl ₂ · 6 H ₂ O	0.25	147.02	18.4
HEPES ^a	1.6	238.3	191
Phenol Red	(0.001%)		5
Na ₂ SO ₄ · 10 H ₂ O	90	322.21	14500
K ₂ SO ₄	30	174.26	2610
Kynurenic acid	1	189.2	95.6
APV ^b	0.05	197.1	4.92

Use about 1 mL NaOH at 0.1N to pH to 7.3, before adding APV and Kynurenic acid, then again after. Bring up to 500 mL after final pH. NB: Kynurenic acid takes a lot of stirring to dissolve, plus a fair bit more NaOH added while dissolving to keep the pH reasonable. Sterile filter, aliquot, freeze in LN₂, and store at −20 °C.

^a Sigma, H4034. ^b DL-2-Amino-5-phosphonopentanoic acid (Sigma A5282).

Box A.6: Papain solution

200 µL of papain suspension (Roche, 108014) is added to 2 mL of Segal's medium (Box A.5) and pH-ed to 7.3 (judged by color). The result is left to fully mix for 15–30 minutes at room temperature, then sterile filtered. Papain solution is made fresh immediately before use, because papain is not immune to its own protease action.

was added, then the cortex bits were sucked in and squirted out three times. This dissociates a good fraction of the cells. After the remaining clumps settled, the suspended cells were transferred to another tube. This procedure was repeated twice more. At this point, most if not all clumps had disappeared. The cell suspension was then spun down onto bovine serum albumin (BSA, Sigma, B3156; 5% in plating medium) at 160×g for 5 minutes. The supernatant—containing cell debris—was discarded, and the cells were resuspended in plating medium. The suspension was then passed through a 40 µM cell strainer (Falcon, 35-2340) to remove any left-over clumps. Cortices from two embryos usually yielded ~6 million cells in ~1.7 mL of

Box A.7: Composition of culture media

The compositions of culture media are as follows:

<i>Plating medium</i>		<i>Maintenance medium</i>	
Neurobasal medium ^a	90 mL	High-glucose DMEM ^e	90 mL
Horse serum ^b	10 mL	Horse serum ^b	10 mL
GlutaMax ^c	0.5 mM	GlutaMax ^c	0.5 mM
B27 ^d	2 mL	Sodium pyruvate ^f	1 mM
		Insulin ^g	6 IU

Both media are sterile filtered, quick-frozen in LN₂, then stored at −20 °C. No antibiotics or antimycotics are used in any of our media.

Assuming that horse serum has similar ion concentrations as DMEM, the ion concentrations in the maintenance medium are^h:

<i>Cations</i>	mM	<i>Anions</i>	mM
Sodium	154.6	Chloride	118.5
Potassium	5.4	Phosphate	1.0
Calcium	1.8	Sulfate	0.8
Magnesium	0.8	Bicarbonate	44.0

^a Gibco, 21103-049. ^b HyClone, SH30074.03.

^c Gibco, 35050-061; stored at −20 °C as 200 mM in DMEM. ^d Gibco, 17504-044.

^e Irvine Scientific, 9024; see Box A.4. ^f Sigma, P4562 stored at −20 °C as 100 mM in DMEM.

^g Sigma, I5500, has 24.5 IU/mg; stored at −20 °C as 2.5 mg/L in DMEM.

medium (as determined by counting 0.1 µL with a haemocytometer). The suspension was diluted to 2500 cells/µL with plating medium.

Plating

Drops of 20 µL of cell suspension were plated on MEAs treated with PEI and laminin as described before. Cell density in the suspension was 2500 cells/µL. This gave a plated area density of 2500 cells/mm², for a total of 50,000 cells/culture. (In some experiments in Chapter 4, lower densities or smaller drops were used. In those cases,

the laminin drop size was adjusted to match to plating drop size.) MEAs were immediately sealed with FEP Teflon membranes (Potter and DeMarse, 2001). When small drops were used, additional larger drops of plating medium (without cells) were placed along the inside edge of the MEA dish, to mitigate evaporation. Cultures were transferred to an incubator (Napco 7101) maintained at 5% CO₂ and 9% O₂. (A below-ambient O₂ level was chosen to more closely match conditions inside the brain, though we do not in fact have data to support that this aids cell survival.) The relative humidity inside the incubator was kept at 65%. This low humidity level prevents incubator infections, and makes the incubator a safe environment for electronics. The cultures do not dry out thanks to the Teflon covers, which allow gas exchange, but not water. Cultures were left for 2 hours to allow cell attachment; then, 1 mL of plating medium was added.

Maintenance

One day after plating, the plating medium was replaced by maintenance medium. Thereafter, every 4–6 days, half of the medium was replaced. Glial cell proliferation was not restricted, since glia are essential to long-term culture health. Whenever an MEA was opened, the Teflon cover was replaced with a freshly autoclaved one, to prevent infection. With this method, only one out of several hundred cultures got infected over the course of the last five years. It allowed us to maintain cultures healthy for as long as we wanted (longest: two years).

Recording

All recordings were made inside the incubator in which the cultures were kept. Placement of recording equipment inside the incubator was found to be critical—in many parts of the incubator recording noise or 60 Hz pickup was strongly elevated, presumably due to the incubator's pumps and electronics. In one sweet spot near the top of the incubator on the front-right, these effects were minimal. Indeed, with careful placement, noise levels were lower than out on the bench top.

Even if the temperature of the MEA was only 1 °C higher than ambient, medium

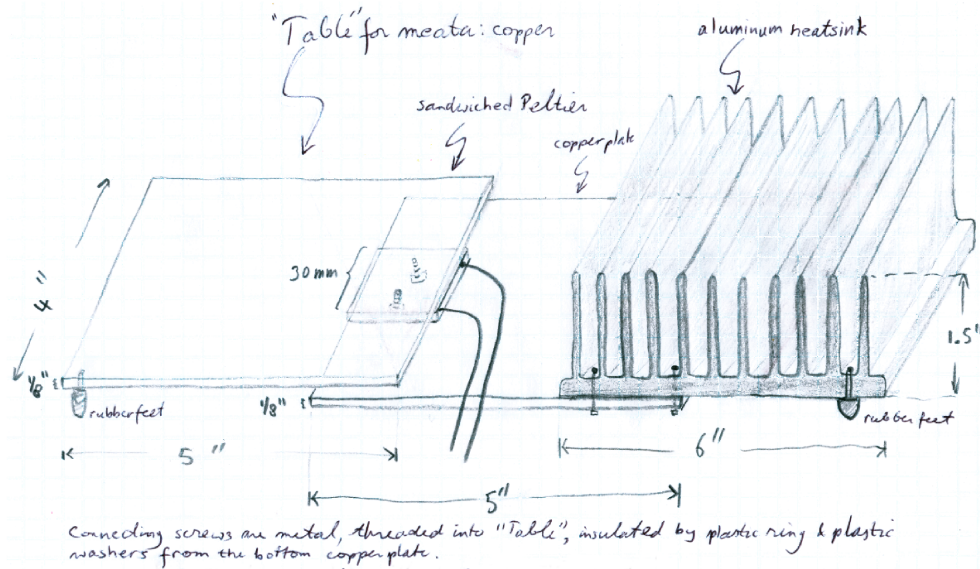


Figure A.1: Peltier-cooled stage for recording in an incubator.

vapor would condense on the inside of the Teflon membrane. After 24 h of continuous recording, this could lead to a 50% increase in the osmolarity of the remaining medium. Therefore, a custom-made Peltier-cooled stage was used to remove excess heat from the pre-amplifier (Figure A.1). Since conditions inside an incubator are quite constant, we found that closed-loop control of the cooler was not required. A constant current to the Peltier element of 0.8 A (at 1.5 V) was enough to reliably keep the temperature of an MEA in the recording device at 0–0.5 °C below ambient. (This was confirmed using a thermistor glued to an old MEA, and held even if placement of the pre-amplifier on the stage varied by an inch or two.) With the power supply for the cooler (HP 6214A) outside the incubator, no additional 60 Hz pickup or other noise was produced.

Appendix B

MEABench: A Toolset for MEA Data Acquisition and Online Analysis*

We present a software suite, MEABench, for data acquisition and online analysis of multi-electrode recordings, especially from multi-electrode arrays. Besides controlling data acquisition hardware, MEABench includes algorithms for real-time stimulation artifact suppression and spike detection, as well as programs for online display of voltage traces from 60 electrodes and continuously updated spike raster plots. MEABench features real-time output streaming, allowing easy integration with stimulator systems. We have been able to generate stimulation sequences in response to live neuronal activity with less than 20 ms lag time. MEABench is open-source software, and is available for free public download at <http://www.its.caltech.edu/~pinelab/wagenaar/meabench.html>.

Introduction

Recording from large numbers of electrodes has become increasingly common in neuroscience over the last 30 years. Multi-electrode arrays (MEAs) (Thomas et al., 1972;

* Published as: Daniel A. Wagenaar, Thomas B. DeMarse, and Steve M. Potter, 2005: MEABench: A toolset for multi-electrode data acquisition and online analysis. *Proc. 2nd Intl. IEEE EMBS Conf. on Neural Eng.*, pp. 518–521. © 2005 IEEE. Reprinted with permission.

Gross, 1979; Pine, 1980) with 60 or more electrodes have been used to study many *in vitro* preparations including cortical cultures (e.g. Jimbo et al., 1999; Potter, 2001; Shahaf and Marom, 2001; Wagenaar et al., 2005b [Chapter 6 in this thesis]), spinal cord cultures (e.g., Gross and Kowalski, 1999) as well as intact retina (e.g., Meister et al., 1994), while silicon probes (Wise and Angell, 1975) and multiwire probes (e.g., Nicolelis et al., 1997) have been used extensively *in vivo*. Most labs have used commercial recording systems, which ship with dedicated software. Such software is typically not user-extendible, and not well suited for real-time applications involving bidirectional communication between computers and biological tissue. For our research on learning in embodied cultures (Potter, 2001; Bakkum et al., 2004; Potter et al., 2006), both are critical requirements. Here we present a software suite that fills these needs. MEABench is a free, open-source, set of programs for multi-electrode data acquisition (DAQ) and online analysis. MEABench is different from previously described multi-electrode data analysis software such as MEA-Tools (Egert et al., 2002), in that it directly communicates with DAQ hardware and provides real-time visualization. This makes it especially suitable for online operation. Thanks to its ability to communicate in real-time with stimulator hardware, it can be used in closed-loop stimulation experiments.

Methods

Software Engineering

MEABench was programmed as a set of semi-independent programs sharing a common library. These programs communicate through standard Linux* inter-process communication (IPC) facilities such as pipes and shared memory. This loose modular approach was adopted to help make the software easily maintainable. It also makes it easy for third parties to add their own MEABench programs, e.g., to add new visualization or data export methods. The core of MEABench was coded in C++. Some utility programs were written in perl, and several additional data analysis tools were

* MEABench was designed primarily for Linux. It should work on other Unix-like operating systems with minimal changes. In particular, a port for Mac OS X (Apple Computer, Cupertino, CA) is available (from TBD, tdemarse@bme.ufl.edu).

written for Matlab (The Mathworks, Natick, MA).

Data Acquisition

MEABench has a modular interface for communication with data acquisition (DAQ) cards. Currently, a module to acquire data from the 'MC_Card' hardware (Multi-Channel Systems, Reutlingen, Germany) is well-supported. This module uses DMA transfer to minimize CPU load.

Inter-Process Communication

MEABench programs use a client-server model to transfer data. A server creates a shared memory block with a header describing its contents. Clients can then independently read from this memory. To obviate the need for clients to continuously poll to check whether data are available, servers notify clients through pipes when new data become available, and when a run starts or ends. This model is used to transfer both electrode voltage data and spike information.

Multi-Computer Setup

More advanced and more demanding algorithms for data analysis become available every year. Therefore computer hardware needs to be upgraded regularly to keep up. Yet, it is not attractive to replace a working DAQ setup, since getting DAQ hardware to work in new computers can be highly non-trivial. Therefore, MEABench provides the option of using one dedicated computer just for data acquisition, and a second computer for all analysis tasks. This second computer can then be upgraded whenever faster processing is required, without upsetting the DAQ computer. The two computers communicate through a (local) ethernet network.

Visualization

Online data visualization is a main feature of MEABench. To show acquired voltage traces and spike rasters online, a set of widgets was written for use with the (open source) Qt library (Trolltech, Oslo, Norway). Qt-Designer was used to create graphical user interfaces (GUIs) for visualization programs.

Data Streaming

MEABench provides an interface for external programs to read data directly from MEABench shared memory. Additionally, MEABench comes with a program to output data as a posix stream, which can then be ‘piped’ to user’s programs using standard posix shell facilities.

Algorithms

MEABench includes programs for the suppression of line noise and stimulation artifacts, as well as for spike detection. Matlab code for (multi-electrode) burst detection is included as well. The algorithms implemented by these programs are described in the following.

Line Suppression

Pickup at 60 Hz (or, outside the USA, at 50 Hz) is a common problem for electrophysiology. Even with careful hardware design, some pickup is usually present in multi-electrode recordings. We used an adaptive template filter to suppress this pickup digitally. This approach works best if a synchronization signal is provided from the mains on a dedicated DAQ channel. The line suppressor divides the period of the pickup into 128 time bins. A separate template of 128 values is maintained for each electrode channel. Templates are updated continuously, with a user-controlled decay time constant (typical value: 1.5 s). Subtraction of the template from the recording yields a signal cleaned of the 60 Hz line pickup and all its harmonics.

Stimulation Artifact Suppression

Since stimulation pulses for multi-electrode arrays typically have amplitudes around 0.5 V (Wagenaar et al., 2004 [Chapter 2 in this thesis]), and recorded spike waveforms typically reach at most 100 μ V, stimulation artifacts often occur even with carefully designed recording hardware. MEABench incorporates an algorithm for stimulation artifact suppression that locally fits low-order polynomials to the recording, and subtracts the fit results (Wagenaar and Potter, 2002 [Chapter 3 in this thesis]) This removes most artifacts that do not saturate the input stage of the recording system.

Spike Detection

After removal of artifacts, detecting spikes (action potentials) is often the next step in the analysis of multi-electrode recordings. To facilitate online operation, we opted for a simple amplitude-threshold detector. Thus a detection threshold that optimizes detection efficiency while controlling the rate of false positives must be established. In many cases it is not possible to temporarily switch off the biological signal, so noise levels must be estimated from the recorded superposition of signal and noise. An extra confounding factor is that noise levels often drift on a time scale of hours, so for long-term recordings, estimates must be adapted continuously.

MEABench implements the following algorithm for detecting spikes. First, the recorded voltage trace is band-pass filtered between 100 Hz and 3 kHz. Then, the data stream is split into non-overlapping 10-ms windows, and the 2nd and 30th percentiles of the distribution of voltages in each such window are determined. Call these $V_{.02}$ and $V_{.30}$. (Note that both are usually negative because of the filtering, which sets $V_{.50} \sim \langle V \rangle \sim 0$.) Then, two tests are performed:

- Is the ratio of $V_{.02}$ over $V_{.30}$ less than 5?
- Is the absolute value of $V_{.30}$ (significantly) non-zero?

The first test makes sure that there was no actual spike in the window; the second test makes sure that the data in the window was not blanked out (e.g., by artifact removal filters). If both tests are passed, the window is considered ‘clean’, and $V_{.02}$ is used to update the current noise level estimate. Specifically, the noise level estimate is the output of passing the absolute values of $V_{.02}$ from all ‘clean’ windows through a low-pass filter with a time constant of 100 windows (i.e., 1 s if all are clean). Spikes are detected whenever the absolute value of the voltage exceeds the current noise estimate by a user-settable factor.

This algorithm adapts rapidly to changing noise situations, while not desensitizing during bursts of spikes.

Spike Validation

Spike waveforms are often multi-phasic, so some additional processing must be performed to prevent double detections of unitary events. We accept a spike only if its




Figure B.1: Spike validation based on waveform shape. A putative spike is rejected if its waveform crosses any of the validation threshold lines, which are based on the spike's peak amplitude.

detected peak is the highest peak of either polarity within a ± 1 ms window, and no secondary peaks of the same polarity and more than 50% of the amplitude of the detected peak exists within the same window (P. P. Mitra, personal communication; see Figure B.1). This validation step also strongly reduces the rate of false positive detections.

Burst Detection

Culture-wide bursts or waves are observed in multi-electrode recordings from many neuronal preparations. To detect such ‘global’ bursts, as well as local bursts which may be of interest to researchers, MEABench includes an algorithm (implemented in Matlab) which first detects bursts on individual electrodes, and then groups together bursts that overlap in time. To aid in the following discussion, the term *burstlet* will be used to refer to a burst on an individual electrode. A burstlet is not necessarily a true single-electrode event: it may well co-occur with burstlets on other electrodes. A group of temporally overlapping burstlets will be called a *burst*. Bursts may be array-wide bursts, or more localized events, even including single cell bursts.

Let f_c be the average firing rate on electrode c , that is, the total number of spikes recorded on that electrode divided by the duration of the recording. We then define threshold inter-spike intervals (ISI), τ_c , for each electrode c . This τ_c is set to $\frac{1}{4f_c}$ or to 100 ms, whichever is smaller. The factor four ensures that only spikes that succeed each other faster than four times the average firing rate can be considered burstlets.

Initially, the algorithm considers each electrode independently. For a given electrode, it searches for sequences of four or more spikes with all internal ISIs less than τ_c . After these ‘core’ burstlets have been found, they are extended into the past and the future to also contain spikes that have ISIs less than 200 ms, (or less than $\frac{1}{3f_c}$, whichever is smaller). Thus, a burstlet consists of a core of at least four very closely spaced spikes, with an ‘entourage’ of any number of slightly less closely spaced spikes, all on one electrode. Once all burstlets on all electrodes have been found, they are sorted in temporal order. A burst is then simply a sequence of one or more burstlets that have non-zero temporal overlap.

In many cases, a small number of electrodes record strongly elevated firing rates for extended periods after a global burst, sometimes until the next one. If that happens, several global bursts would all be grouped together according to the algorithm as described so far. This problem is corrected in a post-processing stage. Each detected burst is considered in turn, and a graph of the number of simultaneous burstlets vs. time is constructed. If a putative burst corresponds to several global bursts, this graph will have more than one hump. The algorithm finds these humps, and splits the bursts accordingly.

Application

We have used MEABench to acquire and analyze multi-electrode array (MEA) recordings from dissociated cultures of cortical neurons. Dense cultures of cortical neurons from embryonic (E18) rats were prepared on MEAs as described before (Potter and DeMarse, 2001; Wagenaar et al., 2005b). Very briefly, we used MEAs with sixty 30 μm diameter electrodes spaced at 200 μm (MultiChannel Systems) in an 8x8 grid with missing corners. MEAs were pre-coated with poly-ethylene-imine and laminin. Cortices were dissociated using papain and trituration. Cells (both neurons and glia) were plated at a density of 2,500 cells/mm³. Cultures were maintained in a DMEM-based medium (Jimbo et al., 1998) in Teflon-sealed dishes (Potter and DeMarse, 2001) in an incubator with 65% relative humidity. This low humidity made the incubator safe for electronics, allowing us to perform all recordings inside the incubator. The Teflon seals prevented evaporation. Partly replacing media every 5–7 days, we could maintain cultures indefinitely.

Electrode signals were amplified and digitized using MultiChannel Systems hardware, controlled by MEABench. For stimulation, we used a custom device (Wagenaar and Potter, 2004 [Appendix C in this thesis]) which connects to the MultiChannel Systems pre-amplifier, and which can generate arbitrary spatio-temporal stimulation sequences under real-time control.

Use of TTX for Noise Estimation

To obtain an unbiased estimate of the recording noise, we bath-applied 1 μ M tetrodotoxin (TTX) to a culture and recorded 300 s of voltage traces after equilibration. By feeding these traces to the spike detector, we determined the rate of false positive spikes* (per second, per recording channel) as a function of the detection threshold. By subtracting this number of spikes from the firing rate in the active culture (recorded before TTX application), we could estimate the true firing rate, and hence the detector efficiency as a function of the detection threshold.

Results

Programs

MEABench consists of the following main programs:

- Rawsrv The grandmother server. It reads voltage traces from the hardware and makes them available to other MEABench programs.
- Spikedet A threshold-based spike detector. It reads from a voltage-trace stream, and publishes a spike information stream. Includes algorithms to adapt to fluctuating noise levels.
- 60hz Template filter to remove 60 Hz pickup.
- Salpa Stimulation artifact filter (Wagenaar and Potter, 2002).
- Record Records voltage or spike data to disk.

* While bathed in TTX, neurons cannot generate sodium action potentials, so any detected spikes must necessarily be false positives.

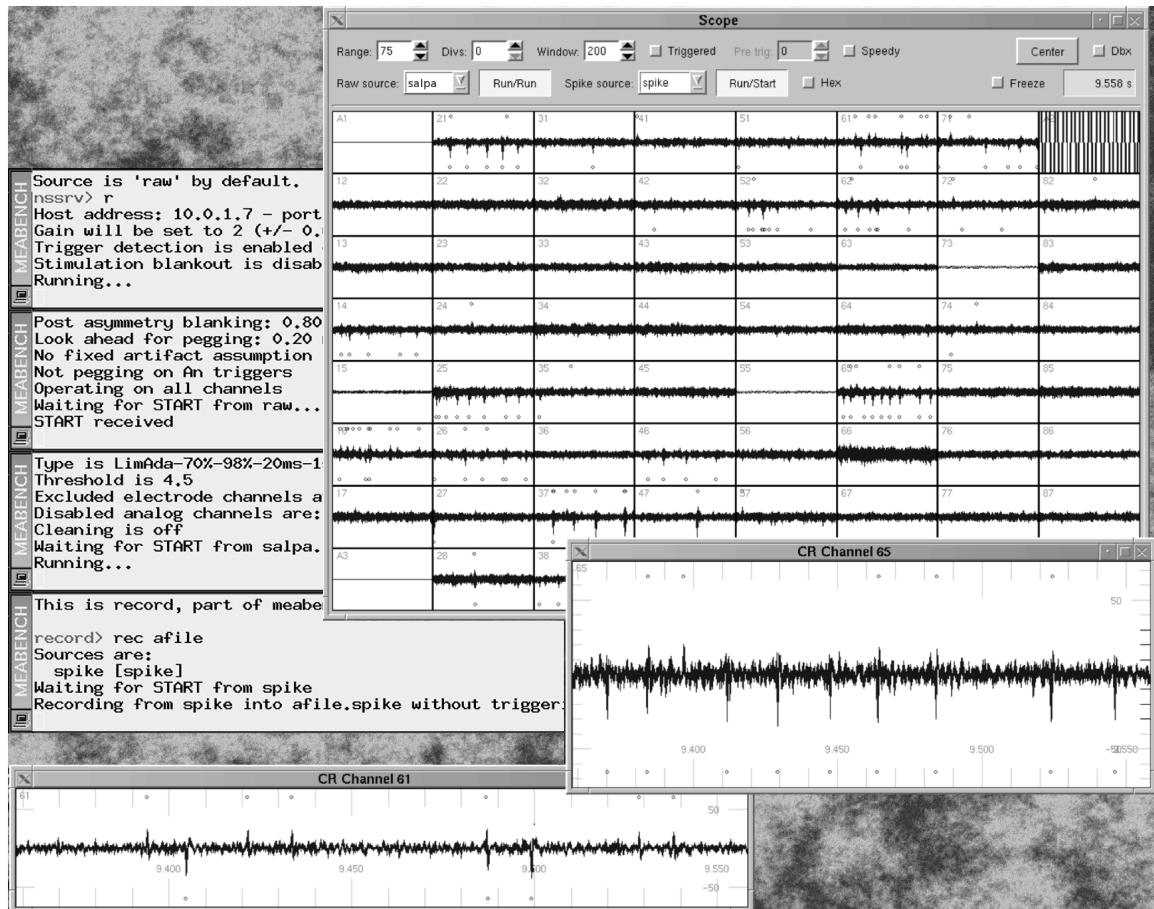


Figure B.2: Screenshot of MEABench in operation. The large window displays voltage traces from all electrodes, in MEA geometry. Small windows on the side show the command-line interfaces of (top to bottom) Rawsrv, Salpa, Spikedet, and Record.

- **Replay** Replays files created by Record.
- **Scope** GUI program for online display of voltage and spike data. Scope includes a ‘freeze’ feature for instant replay of the last 5 s of data.
- **Spikesound** GUI program for online sonification of spike data.
- **Flexraster** GUI program for online generation of spike raster plots. Flexraster allows zooming and scrolling through an entire recording.
- **Neurosock and Nssrv** An alternative to Rawsrv that allows one to dedicate one computer to data acquisition, and another to online analysis (see *Methods*).

In addition, a number of scripts are provided for off-line processing and for automating data acquisition tasks. Full details may be found in the MEABench User




Figure B.3: Tradeoff between spike detector efficiency and rate of false positive detection.

Guide (included with the software). An example MEABench session is depicted in Figure B.2.

Artifact Suppression

Stimulation artifacts were suppressed within 1 ms after the recording system returned from saturation, allowing spike detection within 2 ms after stimulation on most electrodes (Wagenaar and Potter, 2002). (Typically, the stimulated electrode remained saturated for 50–100 ms.)

Spike Detection

What constitutes an acceptable rate of false positive detections typically depends on the experimental situation. To be able to make an educated decision about appropriate spike detector threshold settings, we determined the detection rate of false positives in a culture quieted by TTX (see *Methods*). If one false positive per second per channel is acceptable, the detection threshold could be set at 4.25x estimated RMS noise (Figure B.3). Increasing the detection threshold reduced the rate of false positive detections, but also reduced the detector efficiency.

Real-Time Operation and Stimulator Control

A typical MEABench online data processing chain consisting of stimulation artifact suppression, spike detection, recording, and visualization could be run on a Linux system with an AMD Athlon XP 2800+ CPU. Artifact suppression took about 20% of CPU time, spike detection 10%, and recording about 3%. Thus, sufficient CPU power was available to display continuous voltage traces at a 5 Hz frame rate, as well as continuously updated raster plots.

Thanks to the modular structure of MEABench, it required only a few lines of perl code to make it control our custom stimulator (Wagenaar and Potter, 2004). In this way, we could generate stimulation sequences as a function of observed neuronal activity in real-time, with less than 20 ms latency between activity and generated stimuli*. Thus, MEABench allowed real-time bidirectional communication with neuronal cultures. Recently, we used this to control and suppress culture-wide bursting by providing stimulation controlled by a real-time feedback loop to maintain a stable tonic firing rate (Wagenaar et al., 2005b).

Discussion

MEABench was first conceived to facilitate the study of learning in embodied neuronal networks (Potter, 2001). We needed a data acquisition system that could communicate in real-time with arbitrary code to transform outputs from a neuronal culture to motor commands for a simulated animal (animat) or robot. MEABench has since developed into an extensive and very stable framework for multi-electrode data acquisition and analysis, and has become a core component of our experimental setups. In our labs, MEABench is used in conjunction with MEAs, but its modular structure makes it straightforward to adapt it to different hardware, including *in vivo* multi-electrode probes. MEABench is available for free public download at <http://www.its.caltech.edu/~pinelab/wagenaar/meabench.html>.

* Using a special low-latency version of the Linux kernel. With standard Linux, the latency was about 50 ms.

Appendix C

An All-Channel Stimulator with Real-Time Control*

Over the last few decades, technology to record through ever increasing numbers of electrodes has become available to electrophysiologists. For the study of distributed neural processing, however, the ability to stimulate through equal numbers of electrodes, and thus to attain bidirectional communication, is of paramount importance. Here, we present a stimulation system for multi-electrode arrays that interfaces with existing commercial recording hardware, and that allows stimulation through any electrode in the array, with rapid switching between channels. The system is controlled through real-time Linux, making it extremely flexible. Stimulation sequences can be constructed on-the-fly, and arbitrary stimulus waveforms can be used if desired. A key feature of this design is that it can readily and inexpensively be reproduced in other labs, since it interfaces to standard PC parallel ports and uses only off-the-shelf components. Moreover, adaptation for use with *in vivo* multi-electrode probes would be straightforward. In combination with our freely available data acquisition software, MEABench, this system can provide feedback stimulation in response to recorded action potentials within 15 ms.

* Published as: Daniel A. Wagenaar and Steve M. Potter, 2004: A versatile all-channel stimulator for electrode arrays, with real-time control. *J. Neural Eng.* 1, pp. 39–44. © 2004 IOP. <http://jne.iop.org>.

Introduction

Background

Two-way communication between brains and computer systems has been an important goal in neural engineering for several decades, since it can significantly broaden the horizons in many research areas, such as cortical population coding (e.g., Salzman et al., 1990) or long-term plasticity in vivo (e.g., Staubli and Lynch, 1987). Over the last thirty years, technology to record from a large number of cells has been developed and applied to a wide range of model systems: multi-electrode arrays (MEAs) (Thomas et al., 1972; Gross, 1979; Pine, 1980) have been used to record from many preparations from dissociated cortex (Potter, 2001) to intact retina (Meister et al., 1991), while their *in vivo* counterparts, silicon probes (Wise and Angell, 1975) and multiwire probes (Nicolelis et al., 1997) are also gaining popularity. Stimulation technology has not kept equal pace. Commercially available systems presently are limited to a relatively small number of channels (typically 10 or less), or require programming ahead of time, making true two-way real-time communication impossible. Accordingly, many researchers have built their own custom devices: one of the first such devices used a manual switchboard to select 8 out of 61 channels for stimulation (Regehr et al., 1989). Jimbo and Kawana (1992) describe a complete system with 18 stimulation channels. Pancrazio et al. (1998) describe a 16-channel stimulation system for cardiac myocytes implemented in VLSI, while Zeck and Fromherz (2001) use FET technology to construct a similar system for invertebrate neurons. Another system that combines recording and stimulation capabilities for 64 electrodes was recently described (Jimbo et al., 2003).

Our research focuses on re-embodied neural cultures (Potter, 2001; DeMarse et al., 2001), and the developmental impact of persistent stimulation on network formation. For both these projects, we needed a stimulator with the following properties:

- Ability to stimulate through any of the electrodes of our MEAs;
- Rapid switching between stimulation and recording through the same electrode;

- On-the-fly specification of stimulation sequences as a function of live neuronal activity;
- Compatibility with existing recording technology.

The ‘Real-time All-Channel Stimulator’ (RACS) we present here has all these properties, and has the additional advantage that its assembly from off-the-shelf components is straightforward.

Design Philosophy

Conceptually, stimulators are relatively simple devices: they consist of a voltage or current source, some logic to route the signal to the appropriate electrodes, and a set of isolator switches to allow recording from electrodes whenever they are not being used for stimulation. Most complexity comes in the form of the logic that controls the timing of the system. Commercial systems currently use micro-controllers for this task, but we have chosen an alternative design route: our stimulator is controlled externally by a computer running a real-time operating system, RTLinux (Barabanov, 1997). This eliminates the need for a dedicated microprocessor on the stimulator board, and makes for much easier programming. Perhaps surprisingly, it is also extremely cost-effective, since even an old 100 MHz CPU is fast enough to provide the required real-time control. The total cost of our stimulator is about USD 250: \$150 for components, and \$100 for printed circuit board fabrication. Assembling the system takes about one day and requires only basic electronics skills.

Methods

Cell Culture

Details have been described before (Potter and DeMarse, 2001). Briefly, cortices from E18 rat embryos were dissected and dissociated using papain and trituration. Cells—neurons and glia—were plated at a density of 5,000 cells/mm², on MEAs coated with poly-ethylene-imine (PEI) and laminin. Cultures were maintained in a serum-containing DMEM-based medium, in Teflon-sealed dishes. Stimulation experiments

were performed from *in vitro* day 8 onward.

Recording System

Signals were recorded through 30 μm titanium nitride electrodes spaced at 200 μm . Each MEA (MultiChannel Systems, Reutlingen, Germany) has 59 such electrodes, and one large ground reference electrode that also served as return electrode for stimulation. Signals were pre-amplified 1200x, then sampled at 25 kHz using MultiChannel Systems hardware. Data acquisition and online analysis, including visualization, artifact suppression, spike detection, and storage were controlled through MEABench (D. A. Wagenaar, <http://www.its.caltech.edu/~wagenaar/meabench>). Artifact suppression was performed by subtracting third order polynomials locally fitted to the recorded voltage trace, using the SALPA (“Suppression of Artifacts by Local Polynomial Approximation”) algorithm (Wagenaar and Potter, 2002 [Chapter 3 in this thesis]). A C++ implementation of this algorithm is available upon request from the authors. Further analysis was performed using Matlab (MathWorks).

Stimulation System

In order to keep the fabrication as simple as possible, we used only standard 0.1" DIP technology, avoiding surface-mount technology. Furthermore, we opted for a modular design, which made it easier to test a variety of isolation switches as well as several types of digital-to-analog convertors (DACs). This also facilitates adaptation to other recording hardware. Circuit layout was performed using software from ExpressPCB (www.expresspcb.com), who also machined the printed circuit boards.

The RACS consists of one main board containing a DAC and some interface logic, as well as four modules that each contain isolation switches to gate signals to one of 16 electrodes (Figures C.1 and C.2). It is controlled from a PC’s parallel port, which provides four “Control” pins and eight “Data” pins. The Control pins are used to route signals from the Data pins to either of two DAC channels—one for the stimulus proper, one for an auxiliary analog output—or to either of two latches—one to control the stimulation switches, one for 8 auxiliary digital outputs. Potential uses for the auxiliary outputs include tagging stimulus identities and triggering external

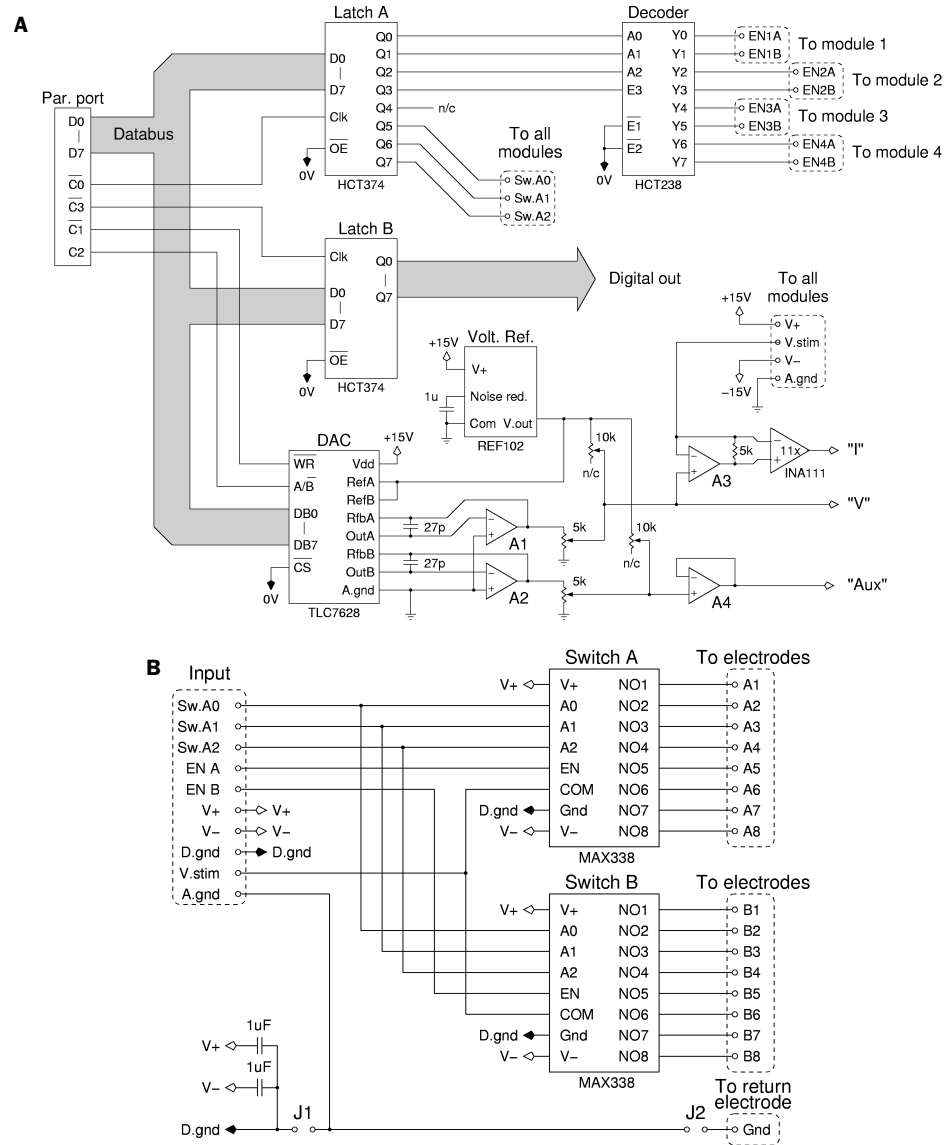


Figure C.1: Circuit diagrams. **A** Main board. The system is controlled through a standard PC's parallel port, and takes power from an external $\pm 15\text{V}$ supply. Terminals "V" and "I" monitor the voltage and current to the currently selected electrode. "Aux" provides a buffered auxiliary analog output. \downarrow represents digital ground; \downarrow_{a} is analog ground. Op-amps A1–A4 are $\frac{1}{4}$ LM348 (Fairchild Semiconductor); REF102 is a precision 10V voltage reference (Texas Instruments); INA111 is a high-speed instrumentation amplifier (Burr-Brown). Other ICs are described in the text. For clarity, the +5V digital supply and decoupling capacitors have been omitted from this diagram. **B** One of four identical modules which deliver stimuli to the MEA. In our setup, they plug into a set of terminals on the pre-amplifier (MultiChannel Systems) which in turn directly connect to the MEA. Jumpers J₁ and J₂ can be used to chose different grounding schemes. We leave both open on all four modules. Power lines are connected as indicated by the arrows (\rightarrow and \rightarrow_{a}).




Figure C.2: Photographs showing the main board (*left*) and three modules connected to a MultiChannel Systems pre-amplifier with MEA in the center (*right*). (The fourth module has been unplugged to allow a better view.) To facilitate stimulation and recording experiments lasting several months, the MEA is sealed with a Teflon membrane which prevents evaporation and infection (Potter and DeMarse, 2001).

equipment. The electrodes are grouped into eight banks of eight, each bank being serviced by one 1-to-8 electronic switch with very low leakage and charge injection (MAX338, Maxim). A 3-to-8 decoder (74HCT238, Philips Semiconductor) is used to select which bank (if any) is activated. A copy of the stimulation voltage is available at the “V” monitor terminal. The stimulation current passes through a 5 k Ω resistor, and the voltage across this resistor, amplified 11x, is available at the “I” monitor terminal. Multi-turn potentiometers are included to tune the range of stimulation voltages as well as to center the range to zero. The maximal selectable range is ± 10 V; we trim the range to ± 1.0 V to prevent electrolysis and damage to electrodes. At that setting, the 8-bit DAC (TLC7628, Texas Instruments) allows specification of stimulation voltages to 8 mV precision. On-board latches (74HCT374, Texas Instruments) and switches are all fast enough (<100 ns propagation delays) that operation speed is effectively limited only by the controlling computer. The circuitry around the DAC, by contrast, was purposefully designed to have a relatively slow slew rate of 110 mV/ μ s. This ensures that the signal generated by the DAC can be accurately reproduced by the output op-amps, without distortion due to large capacitive currents at sharp voltage transitions.

A kernel module for RTLinux ('Open' version 3.1, FSMLabs, www.fsmlabs.com) was written to allow control over output voltages and switching with microsecond-level precision. Driver architecture was based on code examples in the RTLinux documentation, as well as in Rubini and Corbet (2001). To enhance maintainability, only minimal functionality was implemented in the driver, with most of the higher-level control left to user-space programs. Software allowing experimenters to specify pulse-shapes in more conventional terms was written in perl (www.perl.org).

PCB designs, part lists, and driver and application software are available on request by email to the authors (wagenaar@caltech.edu).

Results

The Real-Time All-Channel Stimulator

The RACS provides:

- Stimulus outputs for direct connection to 64 electrodes, all driven from a single DAC, with high-quality isolation switches to select stimulation channels with microsecond timing;
- One independently controlled auxiliary analog output channel that may be used, e.g., to provide stimulus markers;
- Eight digital output lines that may be used, e.g., to trigger external equipment;
- Op-amp buffered outputs that allow monitoring the actual voltage and current being delivered to the currently selected electrode.

The major anticipated use of the device is to output (multiphasic) rectangular stimulation pulses, but it is possible to construct waveforms of arbitrary shape (Figure C.3). While the single-DAC design does not allow for truly simultaneous stimulation through more than one electrode, different electrodes can be stimulated with less than 10 μ s between stimuli. We wrote software that provides several levels of abstraction of the stimulation hardware. At the lowest level, one controls the switches and DACs directly, saying, e.g., "at time $t = 500$ ms, switch to channel 37; 50 μ s later,




Figure C.3: Pulse shapes. **A** Rectangular voltage pulse (*top*) with current response (*bottom*). **B** Sawtooth voltage pulse (*top*) with current response (*bottom*). Stray currents due to cabling capacitance have been subtracted off from these graphs.

set the DAC to 700 mV; 400 μs later, set the DAC to -700 mV; 400 μs later, set the DAC to 0 mV; 50 μs later, release all switches.” At a higher level, one could specify the same command as “at time $t = 500$ ms, send a biphasic pulse with amplitude 700 mV and width 400 μs to the electrode at position (6,3).” Since we found it convenient to use a special recording channel to mark the time and identity of stimuli, the highest level software automatically provides such markers through the auxiliary analog output.

Stimuli like those shown in Figure C.3 could be used to evoke neuronal responses through almost any of the electrodes on a densely plated MEA. Stimuli delivered to different electrodes each elicited distinct array-wide spatio-temporal response patterns (Figure C.4). More details on the efficacy of pulses of various shapes for eliciting action potentials may be found in a separate report (Wagenaar et al., 2004 [Chapter 2 in this thesis]).

Benchmarks

We tested the performance of the RACS in several key areas, including noise injection, stimulation artifacts, timing precision and integration with recording hardware. This section describes the results of those tests.

Switching time, real-time control The timing of individual switching events could be controlled with very high precision: timing accuracy was 0.5 μs RMS, with

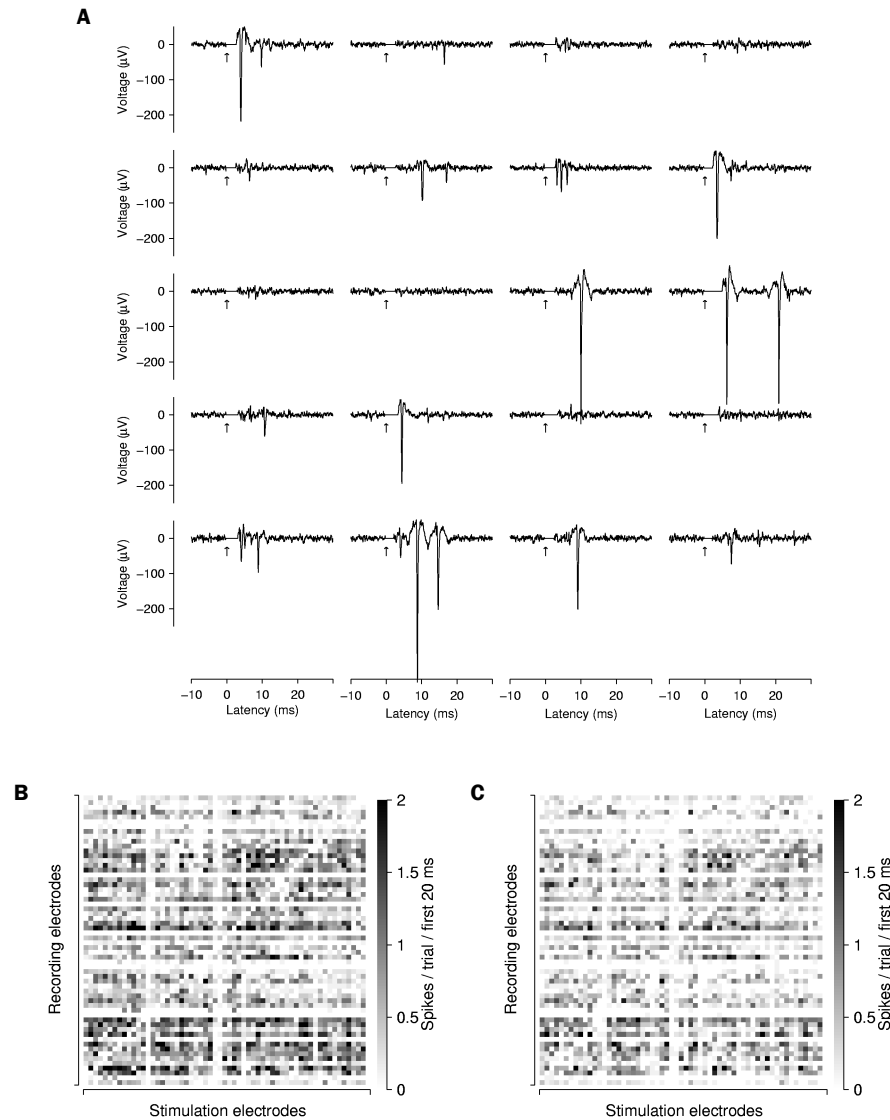


Figure C.4: **A** Neuronal activity recorded on 5 selected electrodes (*top to bottom*) in response to stimulation on 4 different electrodes (*left to right*). Arrows indicate time of stimulation. SALPA (Wagenaar and Potter, 2002) was used to suppress stimulation artifacts. Examples of recordings with and without artifact suppression may be found in that paper. **B** Number of spikes recorded in the first 20 ms after stimulation with biphasic rectangular pulses (as shown in Figure C.3A) through 59 different electrodes ($N = 10$ trials/electrode). (The diagonal is white, reflecting the fact that an electrode cannot be used for recording so quickly after stimulation.) **C** Same, for sawtooth pulses (as shown in Figure C.3B). The pattern is similar, but rectangular pulses elicit 77% more spikes on average.

a worst-case deviation of 2.0 μs ($N = 5000$). Smooth stimulation waveforms could be approximated by controlling the DAC output voltage at a maximum rate of 130 kHz.

Noise injection The RACS did not add a significant amount of noise to the recordings, in contrast to stimulation systems without high-quality isolation switches, which often increase the noise to well above useful levels. With the RACS connected to the MEA, we measured $2.32 \pm 0.28 \mu\text{V}$ RMS noise (mean \pm sample std. dev., $N = 46$ electrodes) in the frequency range 150–2500 Hz used for spike detection. Without the RACS, the baseline noise was $2.26 \pm 0.27 \mu\text{V}$ RMS on the same electrodes, not significantly different (t-test).

Switching and stimulation artifacts Because stimulation involves voltages between 100 mV and 1 V while recorded neuronal signals are typically around 50 μV , even the most careful design cannot completely prevent stimulation artifacts. Such artifacts can be attributed to two sources: the stimulation hardware itself, and the electrode, which undergoes surface charging and electrochemistry. Both kinds of artifacts may affect the signal recorded from the stimulated electrode, as well as signals from other electrodes in the array.

We measured the switching artifact directly generated by the RACS by closing and opening a stimulation switch while outputting a 0 V signal through the DAC. This caused minor artifacts on the other channels: signals remained within the amplifier's dynamic range throughout the stimulus in >99% of trials, and the absolute value of the artifact 1 ms after the end of the stimulus was $10.6 \pm 15.6 \mu\text{V}$ (mean \pm sample std. dev.). These artifacts could be entirely suppressed in software using SALPA (Wagenaar and Potter, 2002). The stimulated channel itself did record significant artifacts: in 55% of trials the signal was driven outside of the amplifiers' dynamic range ($\pm 683 \mu\text{V}$) for 10 ms or more. This artifact is the combined effect of charge injection by the isolation switch and the fact that the DC electrochemical equilibrium potential of the electrode is not necessarily precisely 0 V, so that a 0 V stimulus may still involve non-zero currents.

The combined artifact of stimulator and electrode caused by actual stimuli is of more direct relevance to research. To measure it, we presented biphasic pulses of




Figure C.5: Stimulation artifacts for positive-then-negative biphasic stimuli of 0.5 V amplitude, 400 μ s duration per phase. **A** Amount of time the stimulated electrode cannot be used for recording because the signal is driven outside the dynamic range of the amplifier. The histogram shows a bimodal distribution, because the recorded signal sometimes swings to the other rail after recovering from the first phase of the artifact. **B** Histogram of artifact sizes on other electrodes, measured 1 ms after the end of stimulation. These artifacts were well inside the dynamic range of the pre-amplifier, and could be suppressed in software.

500 mV and 400 μ s per phase, as commonly used during experiments. We found that the signal on the stimulated electrode transiently exceeded the amplifier’s dynamic range in all cases, for 61 ms on average (Figure C.5A). On other channels, the artifact was outside the dynamic range only during the stimulus itself, and had absolute values of $16.8 \pm 17.3 \mu$ V at 1 ms after the stimulus (Figure C.5B). Spikes could be detected within 1–2 ms after the stimulus by using the SALPA artifact suppressor.

Feedback loop time Together with MEA, pre-amplifier, data acquisition board and MEABench software, the RACS forms a feedback loop allowing us to generate stimuli as a function of the observed activity in the culture, as required for our experiments with neurally controlled animats (Potter, 2001). Thanks to the open, modular and extensible structure of MEABench, it took only about 10 lines of new code to generate stimuli in response to action potentials recorded through a particular electrode of the MEA, thus closing the feedback loop. We tested the speed of this loop and observed loop times of 12.2–17.7 ms (98% confidence interval), and a worst case of 24.8 ms ($N=5873$ trials) (Figure C.6). Even the worst case is easily fast enough to provide feedback on biologically relevant time scales, since it corresponds to only a few typical neuron-to-neuron propagation delays.




Figure C.6: Histogram of loop times for delivering stimuli in response to recorded action potentials. The loop consists of data acquisition, spike detection and identification, and stimulus generation. This quick feedback time makes it possible to construct stimulus sequences as a function of observed neuronal activity with loop delay times equivalent to only a few typical cortical synaptic delays.

Discussion

We have described a stimulation system for multi-electrode arrays that interfaces to existing recording systems, and which can be reproduced readily in other labs. While we designed this stimulator for use with MultiChannel Systems pre-amplifiers, it can be adapted to any other recording system that allows direct access to the MEA electrodes. The system could trivially be expanded to handle more electrodes by adding Decoder chips (top right in Figure C.1A). Furthermore, adaptation for use with *in vivo* multi-electrode probes would involve nothing more major than miniaturizing the modules, probably by replacing the MAX338 switches by their surface-mount versions. For applications where current-controlled stimulation is preferable, the voltage output stage formed by opamp A3 and instrumentation amplifier INA111 (bottom right in Figure C.1A) can be replaced by a voltage-to-current convertor (Figure C.7).

The RACS allows stimulation of all electrodes in the array, with arbitrary patterns and rapid switching between stimulation and recording through individual electrodes. Combining the RACS with our artifact suppression algorithm, SALPA, we could detect spikes as early as 1–2 ms post-stimulation, using any electrode except the stimulated electrode itself. On stimulated electrodes, spikes could be detected after 40–160 ms—as soon as artifacts no longer saturated the pre-amplifier. To further




Figure C.7: A voltage-to-current convertor to adapt the RACS for current-controlled stimulation. R_1 (10–100 k Ω) converts the voltage to current. R_3 (6.8 M Ω) acts as a shunt to prevent runaway voltages when all switches are off. R_2 are 150 k Ω . Opamps are $\frac{1}{4}$ LM348 as before. Based on a design by Horowitz and Hill (1996).

reduce that time, sample-and-hold technology can be employed (Novak and Wheeler, 1988). A simpler approach would be to isolate the amplifier from the electrode during stimulation using an additional switch. Unfortunately, we found this was insufficient to prevent artifacts, most likely because artifacts are primarily due to the electrode slowly returning to its electrochemical equilibrium potential after stimulation. We chose not to use sample-and-hold circuitry here, because that would require an integrated design of stimulation and recording systems (Jimbo et al., 2003), and one of our goals was to maintain independence between stimulator and recording system designs.

By allowing arbitrary stimulation patterns which can be modified continuously during the course of an experiment, the RACS opens the way for a new experimental paradigm, in which cultures are probed continuously with naturalistic patterns of distributed stimulation. We formed a feedback loop that allowed stimulation in response to recorded action potentials within 15 ms on average, by combining this stimulator with MEABench data acquisition software. When cultures are used as the brain of artificial animals (or ‘animats’) as in DeMarse et al. (2001), this fast feedback will allow for much more direct interaction between the animat and its environment, which should result in greatly enhanced performance.

Appendix D

MEABench: User Guide*

MEABench is a set of command line and GUI utilities to process data from the Multi Channel Systems MEA60 amplifier. Adaptation for different hardware should be straightforward. MEABench was designed with extendibility in mind, and is fully modular. That means that filters can be inserted anywhere in the data processing stream. This document describes the basics of the toolset and the usage of the individual programs.

MEABench is free software. You can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version. However, I encourage you to contact me if you wish to do so.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

You may redistribute and/or modify this documentation under the terms of the GNU Public Documentation License as published by the Free Software Foundation. However, as for the software, I encourage you to contact me if you wish to do so.

The latest version of this document and of the software described in it, is available for public download from <http://www.its.caltech.edu/~wagenaar/meabench.html>.

Introduction

MEABench is an open-source suite of programs for acquisition and analysis of multi-electrode array (MEA) recordings. MEABench was developed by Daniel Wagenaar at Caltech, drawing

* This is version 1.0.16, dd August 2004, of the Meabench user guide.

on the excellent example set by MultiChannel Systems' MC_Rack suite*.

The software runs under Linux and other Unix variants, and is freely distributable under the terms of the GNU Public License (see <http://www.gnu.org/copyleft/gpl.html>). It offers the following functionality:

- Acquisition of raw electrode data from MultiChannel Systems' MCard;
- Complete removal of mains (60 Hz) interference using template filtering;
- Removal of stimulation artifacts using the SALPA algorithm (Wagenaar and Potter, 2002);
- Online and offline detection of spikes;
- Online visualization of electrode data and spikes;
- Online sonification of spikes;
- Continuous or windowed saving of raw data and spikes;
- Saving of spike waveforms, for later spike sorting and analysis;
- Replaying of raw and spike files, at any speed;
- Instant-replay buffer for easy analysis of recent events;
- Online generation of raster plots;
- Continuous monitoring of varying noise levels;
- A variety of utilities for analysis and data format conversion, including:
 - Averaging of electrode recordings over trials;
 - Conversion of binary spike files to ASCII representation;
 - Filtering of spike files based on any mathematical expression involving shape or timing parameters;
 - Extraction of single channels from 64 channel streams;
 - Splitting of long data files into trials;
 - Splitting of long data files into channels;
 - Computing spike rates;
 - Detecting culture-wide bursts.
- Matlab functions to import MEABench data[†];
- A program to allow easy scripting of MEABench modules for offline processing.

MEABench is fully modular, and any user with some Unix programming experience can extend it to fit her or his needs. Since MEABench can stream live data to your extension modules, it is well suited, for example, to drive real-time feedback systems. In fact, the ability to communicate with other software or hardware in real-time was one of the primary motives for the conception of MEABench. It allowed a reliable, sub-100 ms feedback loop time in our Neurally Controlled Animat (DeMarse et al., 2001).

MEABench was written primarily for use with the MultiChannel Systems MEA hardware, and a driver is included for their MCard data acquisition card, written by Thomas B. DeMarse

* See <http://www.multichannelsystems.com>.

[†] Users may also be interested in Uli Egert's comprehensive set of matlab code for MEA data analysis; freely available at: <http://www.brainworks.uni-freiburg.de/projects/mea/meatools/overview.htm>.

with advice from MultiChannelSystems. If you use different data acquisition hardware, you may still find MEABench useful, because, due to its modular nature, it is possible to write plug-in modules to read data from your hardware. An experimental driver for one such board (manufactured by United Electronic Industries, but not endorsed by us at its state of development as of Nov 2002) is included as well.

MEABench has been in constant use in the Pine lab at Caltech for over four years, and at Steve Potter's group at Georgia Tech* since its beginning. MEABench remains a work in progress; we welcome suggestions for improvement (and bug reports). Please join in the development by submitting your code (patches and improvements) for inclusion in future releases.

Conventions

Throughout this document, the names of MEABench programs are typeset **Like this**. When running MEABench programs, the capitalization should be omitted. Commands defined within MEABench programs are set **like this**. The names of MEABench streams look *like this*, and their types LIKE THIS. **Ctrl-C** means holding the Ctrl key while pressing 'C'.

In examples, text you type is set like this. Text the computer produces is set like this, with prompts highlighted **like this**. The unix prompt is represented as \$.

In definitions of MEABench commands, parameter names are written *like this*, optional parameters are enclosed in brackets [*like this*], and alternatives are separated by a vertical pipe: *this|that*.

I would like to acknowledge valuable input and support from Steve Potter, Tom DeMarse and Jerry Pine. We are all grateful for financial support from the NIH-NINDS and the Burroughs-Wellcome Fund, and cooperation, technical support, and equipment from MultiChannel Systems.

Compilation and Installation

If you intend to use MEABench with MultiChannel Systems hardware, installation is very straightforward. The following is a step by step guide.

- Make sure you have gcc 2.95 or later.
- Make sure you have Qt 3.0 or later.
- Make sure your kernel is 2.4.10 or later.
- Download the latest version of MEABench, and unpack it:

```
$ tar xzf meabench-1.0.16.tar.gz
```

- Enter the directory:

```
$ cd meabench-1.0.16
```

- Choose the directory were you want to install MEABench, e.g., /opt/meabench, and configure:

* Public website: <http://www.neuro.gatech.edu/groups/potter/index.html>.

```
$ ./configure --prefix=/opt/meabench --with-hardware=mcs
```

(The `--with-hardware` chooses the particular DAQ hardware you are using. Currently defined values are `mcs` for MultiChannel Systems' MCard, and `uei` for United Electronic Industries' PD2-MF64-14H cards. This second driver is currently “experimental”.)

- Compile:

```
$ make
```

This will take a while. You may see various compiler warnings, such as:

```
SD_BandFit.C:97: warning: assignment to 'short int' from 'float'
```

These can safely be ignored. Actual compiler *errors* are another story, of course: If you see something like:

```
SD_BandFit.C:52: 'sqr' undeclared (first use this function)
make: *** [SD_BandFit.o] Error 1
```

you have discovered a bug, which I would like to hear about (see *Reporting bugs* below).

- If compilation went well, you may now complete the installation by typing:

```
$ make install
```

Depending on the installation location you have chosen using `--prefix=...`, you may have to become super user (root) before executing `make install`.

- If the installation location you chose is in your *path* already, you can run MEABench programs simply by typing their name. If not, you can either modify your path variable, or copy the program **Mea** from the MEABench 'bin' directory to a location in your path. That will allow you to run MEABench programs by typing, e.g., `mea rawsrv`.

You may have noticed that the installation procedure closely matches the standard GNU style. Generic information about GNU style installation is provided in the file `INSTALL` in the top directory of the MEABench source tree.

Examples

In this section you will learn how to do most common tasks with MEABench; subsequent sections will provide a reference guide to individual components. I will assume that you have already compiled MEABench and installed it in a place where your unix shell can find it.

Displaying Electrode Traces Online

This first example shows how to display incoming data from an MEA in real-time.

Start by checking that all hardware is connected properly, with an MEA locked in the pre-amplifier. Then open two terminal windows.

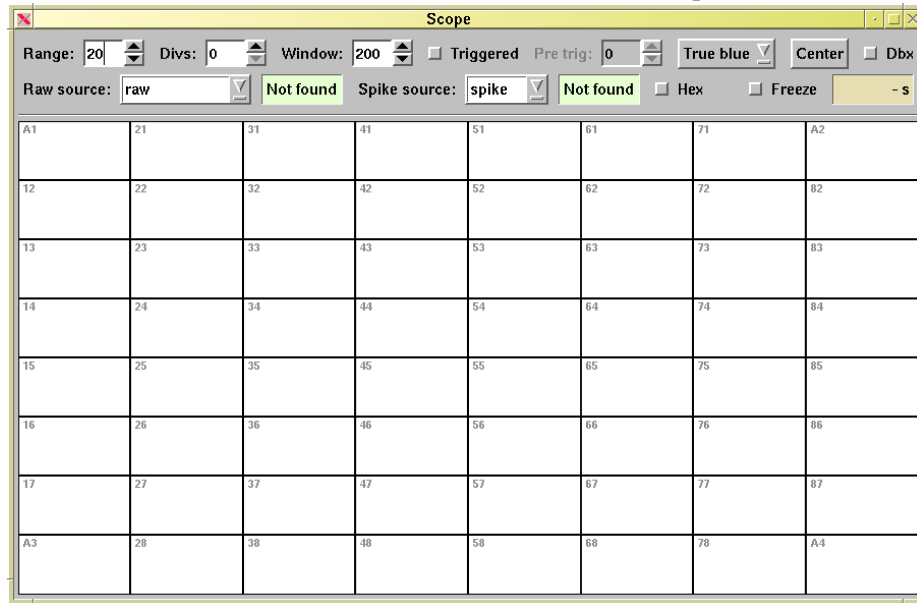
- Run the visualization program, **Scope** from the first terminal:

\$ scope

Error from Sockclient: Constructor failed at connect [Connection refused]

Error from Sockclient: Constructor failed at connect [No such file or directory]

The warnings are normal, and indicate that the scope is not yet ‘connected’ to any input. We will fix that in a moment. Let us first take a look at the scope window:



The top area contains a lot of controls, which we will explore shortly. The main part of the window displays an 8 by 8 grid of rectangular panels, blank for now, which will show electrode traces in a moment. The panels are laid out in the same shape as the physical electrodes on an MCS MEA, with the corner positions occupying the auxiliary channels A1 through A3. The bottom right corner is not connected to anything on our data acquisition board. If you are using a ‘hex’ MEA, you can click the ‘Hex’ button in the second row of controls to change the layout appropriately.

- Start the core data acquisition program, **Rawsrv** in the second terminal:

\$ rawsrv

This is rawsrv, compiled for use with the following hardware:

Pre-amp: MEA1060 by MultiChannel Systems

A/D: MCCard by MultiChannel Systems

Driver: MCCard.o by Thomas DeMarse

Plugin: MCS by Daniel Wagenaar

rawsrv>

The information printed by **Rawsrv** may be different if you configured MEABench for different hardware.

- Put the cursor in the ‘Raw source’ box of **Scope**, and hit return, or re-select *raw* from the pull-down menu. The label ‘Not found’ should be replaced by ‘New/Ready’, indicating that

the scope and **Rawsrv** are now connected. You will also notice that the graphs now are adorned with a zero-line.

- Return input focus to the terminal with **Rawsrv**, and set the data acquisition running:

```
rawsrv> run
Gain setting is 2 (+/- 0.683 mV full range)
Trigger detection is disabled.
Stimulation blankout is disabled.
Stimulated channels: none
Running...
```

Notice that the time displayed at the far right of the second row of controls in the scope runs.

- Click the ‘Center’ button of the scope to remove DC offsets from the traces. Now may be the time to explore the other controls in the top row of the scope. Help balloons pop up when you hover the mouse cursor over any item.
- To stop data acquisition, bring input focus to **Rawsrv**, and press **Ctrl-C**.

```
(interrupt)
rawsrv>
```

- The scope may be terminated using the ‘Close’ button provided by your window manager. **Rawsrv** should be terminated by pressing **Ctrl-D**, or by typing ‘quit’:

```
rawsrv> quit
$
```

Online Spike Detection

The next example will extend the previous one, by adding a spike detector, and a recorder to save the data to hard disk.

- Open four terminal windows, and start **Scope** and **Rawsrv** as before. In the following, I will let you figure out in which window to type from the prompts shown. For example, I will write

```
rawsrv> run
```

and leave it implicit that you need to bring keyboard focus to the terminal window in which you started **Rawsrv**, and then type ‘run’.

- Start the spike detector, and configure it:

```
$ spikedet
spikedet> source raw
Source is raw
spikedet> type 3
Type is BandFit-25
spikedet> thresh 5
Threshold is 5
```

spikedet>

The threshold is specified in terms of an estimate of RMS noise. The spike detector can only base this estimate on observing the noisy signal and guessing which part is noise. This is not an exact science, and **Spikedet** implements more than one algorithm. For details, see under *Spikedet* in the *Details of each component* section below. If you can think of a better scheme, please contribute it!

- Set **Rawsrv** running:

```
rawsrv> run
```

```
Gain setting is 2 (+/- 0.683 mV full range)
```

```
Trigger detection is disabled.
```

```
Stimulation blankout is disabled.
```

```
Stimulated channels: none
```

```
Running...
```

(Is the scope aware of the existence of **Rawsrv** and **Spikedet**? If not, use the ‘Raw source’ and ‘Spike source’ controls to remedy the situation.)

- The spike detector needs to be ‘trained’ on the amount of noise in the source:

```
spikedet> train
```

```
Type is BandFit-25
```

```
Threshold is 5
```

```
Training...
```

- A few seconds will pass, and then the spike detector will return:

```
Training complete
```

```
spikedet>
```

- Interrupt **Rawsrv** as before.
- Start and configure the recorder:

```
$ record
```

```
record> cd
```

```
Working directory: /home/wagenaar/tmp
```

```
record> source raw spike
```

```
Sources are:
```

```
raw [raw]
```

```
spike [spike]
```

(The recorder tells you that it is about to record from **Rawsrv**, which provides a stream called *raw* of type RAW, and from **Spikesrv**, which provides a stream called *spike* of type SPIKE.)

- Prepare the recording:

```
record> record firstdata
```

```
Sources are:
```

```
raw [raw]
```

```
spike [spike]
```

Waiting for START from raw
Waiting for START from spike

The recorder will output files called 'firstdata.raw', and 'firstdata.spike', but for now it is waiting for the data to arrive.

- Start the spike detector:

```
spikedet> run
Source is raw
Type is BandFit-25
Threshold is 5
Quick recording disabled
Excluded electrode channels are: none
Disabled analog channels are: None
Waiting for START from raw...
```

As you can see, the spike detector is quite chatty, and tells you lots of things you already know. It helps to prevent surprises later...

- Let us record precisely 100 seconds:

```
rawsrv> run 100
Gain setting is 2 (+/- 0.683 mV full range)
Trigger detection is disabled.
Stimulation blankout is disabled.
Stimulated channels: none
Running...
```

- The spike detector will confirm that the run started:

Running...

As will the recorder:

Recording from raw into firstdata.raw without triggering
Recording from spike into firstdata.spike without triggering

It is important to set the different programs in motion in the right order: downstream first. The programs will automatically wait for a START signal from their source, and they will wait forever if the source is already running when they open the communication channel.

- After 100 seconds, **Rawsrv** will stop, and the spike detector will issue an end-of-run report:

```
STOP received - 1511 spikes detected
Buffer use percentages: 1
spikedet> (Client 'record' lost)
spikedet>
```

- The recorder will also report some statistics:

```

Recording from raw ended
Buffer usage: 9
Recording from spike ended
Buffer usage: 1 4
record>

```

Those ‘buffer usage’ numbers are shown for each of the communication channels between pairs of MEABench programs. If any number is above 80 (percent), the recorder will warn of the risk of buffer overruns.

- Take a look at the ‘description’ files that **Record** generated. They are called ‘firstdata.raw.desc’ and ‘firstdata.spike.desc’. They contain more statistics about the run, which may be helpful for you or your computer to interpret the data later on.

Triggered Recording

This final fully worked example explains how to perform event-triggered recording. Let us assume that we have a stimulator connected to the MEA, which delivers current pulses on one of the channels once every few seconds. We want to record the responses.

- The first step is to connect a TTL trigger from the stimulator to channel A1 on the data acquisition card.
- Then, set up **Rawsrv** for trigger detection:

```

$ rawsrv
rawsrv> trigchannel 1
Trigger detection is disabled
rawsrv> trigthresh 3000
Trigger detection is disabled
rawsrv> usetrig 1
Trigger detection is enabled on channel A1 – threshold is 3000

```

If you are using MCS hardware, zero volts is represented as digital value 2048, and a TTL trigger will max out the amplifier at digital value 4095, so a threshold halfway is appropriate. If you are using different hardware, use the following trick: set **Rawsrv** running, and open a **Scope** on it. Control-double-click on the panel displaying A1. In the controlling terminal, you will see a lot of numbers scrolling past. Those are the digital values read from channel A1. 50 ms of data is shown, one millisecond per line (if the terminal window is wide enough). If you catch a trigger pulse using the ‘Freeze’ button, and display it in a 50 ms window, you can figure out both the baseline value and the peak value, and set a threshold based on those. Here is a part of the output I obtained doing this experiment:

```

36.746 2383 2384 2383 2383 2383 2383 3534 4095 4095 4095 4095 4095 4095 4095 4095 4095
36.747 4095 4095 4095 4095 4095 4095 4095 4095 4095 4095 4095 4095 4095 4095 2793 2446 2392
36.748 2378 2371 2370 2373 2376 2379 2380 2381 2381 2381 2382 2382 2382 2382 2381 2383 2382

```

The trigger pulse is clearly visible as a sequence of maxed-out values (4095).

- Set up the recorder for triggered recording:

```
$ record
record> source raw
Sources are:
raw [raw]
record> trecord secondfile 50 450
Sources are:
raw [raw]
Waiting for START from raw
```

- Start the data acquisition:

```
rawsrv> run
Running...
```

- Now set the stimulator going, and watch the progress using **Scope**.
- After the stimulation program has ended, stop the recording by pressing **Ctrl-C** in **Rawsrv**'s terminal window. **Do NOT terminate the recorder, instead, terminate the source.** That way you are guaranteed that the final trigger window is saved to disk properly.

Using Commander

Here is an example script for commander that reads a RAW file from 'noisy.raw', filters 60 Hz noise out of it using a reference signal on channel A2, and detects spikes using BandFilt at five times RMS noise. The results are saved as 'denoised.spike'.

```
# Start programs
new replay so:noisy.raw
new filter60hz/60hz so:reraw lock:a2 nper:100
new spikedet so:60hz ty:3
new record so:spike

# Check whether they came up OK
expect replay 1 replay>
expect 60hz 1 60hz>
expect spikedet 1 spikedet>
expect replay 1 replay>

# Let's train the spike detector
tell filter60hz cont
tell replay run
sleep 5
tell spikedet train
flush spikedet
expect spikedet 20 spikedet>
dieif spikedet STOP received before training complete
```

```

flush replay
intr replay
expect replay 5 replay>

# Good, let's go.
tell spikedet run
tell record record denoised
expect record 1 Waiting
expect spikedet 1 Waiting
tell replay run
flush record
expect record 1000 record>

# All done.
quit

```

Basics of the Toolset

MEABench consists of a number of independent linux command line programs. Some of these programs are *servers*, that make the result of their computations available to others. Others are simply *clients*, that read data from one (or several) of the servers, but do not make their results available. Many programs are both client and server. Such programs can be thought of as generic *filters*.

This section describes some of the internals of MEABench, essential for potential developers, and hopefully helpful for users who want to understand how things work. When first reading of this document, you may wish to skip ahead to the next section, *List of components*.

The core of MEABench is a library of C++ classes that can be used to easily construct new components. This library is stored in the meabench/base subdirectory. Presently, the only library documentation is contained in the source (header) files.

Components of MEABench communicate with each other in a standardized way. Servers publish a *shared memory stream*, from which clients can read asynchronously. Currently, two data types are supported: RAW, which contains raw digital data as read from the driver, and SPIKE, which contains information about spikes. Internally, RAW data are represented by C++ datatype *Sample*, while SPIKE data are represented as *Spikeinfo*. These datatypes are defined formally in meabench/common/Types.H.

Associated with each stream is some header information, from which clients can find out how much data are ready to be read from the stream.

Servers never check whether clients are keeping up—it is up to the individual clients to detect overruns. This philosophy was adopted because some clients may not care too much about overruns (e.g., display programs) whereas others do (filters and recorders). All current core clients detect buffer overruns and report buffer usage at the end of a data taking run.

In addition to shared memory streams, servers publish a *wakeup socket*, from which clients can receive wakeup calls whenever a given amount of data is available in the stream. The wakeup

socket also notifies clients when a run starts or ends, and when triggers are detected.

List of Components

These are the programs that currently make up MEABench. Following sections detail each program.

Rawsrv The grandmother server. It reads data from the hardware using Tom DeMarse's driver and makes it available as a RAW stream.

Spikedet Basic spike detection. It reads from a RAW stream, and publishes a SPIKE stream.

Filter60hz Template filter to reduce 60 Hz pickup.

Salpa Stimulus artifact filter.

Record Records RAW or SPIKE data to disk.

Replay Replays files created by **Record**.

Scope GUI program to display RAW and SPIKE data online.

Spikesound GUI program for online sonification of SPIKE data.

Flexraster GUI program to display raster plots of SPIKE data online.

Monitor A debugging aid, it shows the status of all servers.

Neurosock and **Nssrv** An alternative to **Rawsrv** that allows one to dedicate one computer to data acquisition, and another for online analysis.

The following is a set of utility programs that can be used with MEABench or on their own right.

Runmeab Opens a set of xterms from which MEABench programs can be launched.

Spikedump Converts a SPIKE file into human readable form.

Doubletxt Takes a SPIKE stream and a RAW stream, and tags additional context on to the SPIKE information from the RAW channel at which the spike occurs.

Noisehisto Takes a RAW stream and outputs a histogram of voltages observed in each channel.

Noiseshape Takes a RAW stream and outputs the first few central moments of the voltage distribution for each channel.

Uniquespike Output spikes found in one but not in another file.

Trigvar Computes the variance in a triggered RAW stream as a function of time post stimulus.

A growing set of additional utilities remains to be documented.

Note that command names are usually capitalized in this manual, but must always be spelled all lower case on the command line.

Details of Each Component

This section explains the core MEABench components in more detail. Most of these components have a command line interface. Thus, the following entries focus mostly on the available commands. In addition to the commands listed in the individual descriptions below, the following commands are common across components:

- ?** Provide help in the form of a list of commands with brief descriptions.
- quit** Terminates the program gracefully. The same can be effected by pressing **Ctrl-D** at the prompt.

Servers that are capable of loading and saving data to disk support:

- cd** [*directory*] Change or report current working directory.
- ls** [*arguments*] Directory listing as per `/bin/ls`.
- mkdir** [*directory*] Create a new directory.
- !** *command* [*args*] Execute an arbitrary shell command.

Client programs support:

- source** [*stream-name*] Specify from which other MEABench program the data are to be taken by specifying its stream-name.

Rawsrv

Rawsrv reads RAW data from the hardware and publishes it as a shared memory stream called *raw*. It is an extremely straightforward piece of software. Other than providing a nice large buffer to prevent overruns, it is able to detect trigger signals on any of the analog channels (A1, A2 or A3) and to blank out the electrode channels for some time during and after a trigger. These commands are supported:

- run** [*time-in-s*] Starts a run. Optional argument limits the duration of a run to the given time. Otherwise, press **Ctrl-C** to stop a run.
 - usetrig** [0/1] Enables (1) or disables (0) reporting trigger events on the wakeup socket and in the auxilliary data of the RAW stream.
 - trigchannel** [1/2/3] Selects which of the three analog channels to monitor for triggers.
 - trigthreshold** [*digivalue*] Specifies the (digital) value of the threshold above which a trigger is detected.
 - autothresh** [*multiplier*] Sets the threshold for trigger detection at *multiplier* standard deviations above the channel mean value.
 - gain** [*gain-step*] Sets the gain of the MCS amplifier. Type `gain ?` to list possible values.
- For our MultiChannel Systems card, the values are as follows:

Value	Full range (mV)	Step (μ V)
0	3.410	1.665
1	1.205	0.588
2	0.683	0.333
3	0.341	0.167

blankout [*period-in-ms* | 0] Enables or disables blanking of electrode channels during a trigger. Blanking is performed by replacing the data by the average of four samples obtained just prior to the stimulus. Blanking only works if trigger detection is enabled. Signal blanking is largely outdated by **Salpa**.

Spikedet

Spikedet detects spikes on a RAW data stream and publishes them as a SPIKE stream called *spike*. Several different types of spike detection may be supported by **Spikedet**. Currently, these are the fully supported spike detectors:

BandFilt Straightforward threshold detector;
AdaFilt Threshold detector with adaptive threshold;
LimAda Better threshold detector with adaptive threshold;
SNEO Detector based on instantaneous energy in signal.

BandFilt is thoroughly tested and quite stable. AdaFilt (7/12/01) seems to work well except that its threshold varies during bursts; LimAda solves this problem. SNEO has never worked as well as I would hope. See below for details on each detector.

These commands are supported:

source See general description.
type [*detector-name*] Changes or reports the detector being used. Use **type ?** to query supported detectors.
run Starts a single run. If the detector hasn't been trained yet, the first few seconds of the input data are used for training.
cont Same as **run**, but restarts automatically after receiving a STOP command from the server.
train Trains the detector on the current source. Unlike other commands, **train** does not wait for a START message from the source to begin operation. It is recommended to let the source run for a few seconds before commencing training, to ensure that any transients have died out.
info Reports the result of training.
threshold [*value*] Sets or reports the threshold for spike detection, in units particular to each detector. Thresholds are automatically scaled for each channel.
disableanalog [*channel ...* | -] Disables spike detection on the given set of analog channels. Typical use is to exclude the 60 Hz reference signal from being recorded in spike files.
excludechannels [*RC ...* | -] Marks a set of electrode channels as 'dead'. Useful to prevent spurious spike detection on clamped down or flaky channels.
outputfilt [0/1] Enables or disables the creation of a shared memory stream called *spraw*, on which the filtered raw data are reported.
savenoise *filenamebase* Saves the current training results as '*filenamebase.noise*'. This data consists of estimated RMS noise values, and can be used by **Salpa** as well.

loadnoise *filenamebase* Loads a previous set of training results from '*filenamebase.noise*'.

Salpa generated noise files may be loaded as well.

alias A fake command. Reminds the user of the existence of the '-alias' flag (see below).

In some circumstances it may be useful to run more than one spike detector simultaneously. To make that work, **Spikedet** can be made to publish its results on a differently named stream, by starting it as '`spikedet -alias streamname [cmds]`'.

BandFIt

BandFIt passes the RAW data through a band pass filter. Currently a first order filter with cutoffs at 150 Hz and 2.5 kHz are used, but this may be changed in '`spikedet/Filters.H`'. It detects spikes if the filtered stream exceeds a given multiple of the estimated RMS noise in each individual channel. Useful threshold values are 4 to 6. It should be noted that noise in RAW data is far from Gaussian, so future versions may be changed to employ more relevant noise measures. In the current implementation, the noise is estimated by a slightly unusual method, which is intended to minimize the effect of spikes and stimulus artifacts on the estimate. Three hundred 10-ms windows of electrode data are read. For each of these windows the RMS value is calculated. The results are sorted, and the final estimate of RMS noise is taken to be the 25th percentile of the measurements. While I recognize that this method finds an underestimate of the RMS noise, this algorithm is much more useful than straightforward RMS measurement, for the stability reasons mentioned above.

AdaFIt

AdaFIt uses the same initial band pass filter and also collects 128 windows of length 10 ms from the beginning of the recording. From then on, it proceeds differently: it measures the minimum and maximum values in each of those windows, and finds the 40th percentile of both collections of extrema. The initial thresholds for upward and downward spikes are based on the result. While running, it keeps collecting minima and maxima in 10 ms windows, although it uses only one in ten windows*. Whenever 128 windows have been collected, the thresholds for that channel are updated.

For every detected spike, the ruling threshold at the time of detection is written into the SPIKE stream.

Useful threshold values are 1.3 to 2 (multiples of the extrema).

LimAda

After band pass filtering as for BandFIt, LimAda splits the data stream into 10 ms windows, and determines the 2nd and 30th percentiles of the distribution of voltages found in each such window. Call these voltages $V_{.02}$ and $V_{.30}$. (Note that both are usually negative because of the filtering, which sets $V_{.50} \sim \langle V \rangle \sim 0$.) It then performs two tests:

* More precisely: for every window it collects extrema for only six out of 60 electrode channels, to spread the computational load.

- Is the ratio of $V_{.02}$ over $V_{.30}$ less than 5?
- Is the absolute value of $V_{.30}$ (significantly) non-zero?

The first test makes sure that there was no actual spike in the window; the second test makes sure that the data in the window was not blanked out (e.g. by **Rawsrv** or **Salpa**). If both tests are passed, the window is considered ‘clean’, and $V_{.02}$ is used to update the current noise threshold estimate. Spikes are detected whenever the absolute value of the voltage exceeds the current threshold, which is the output of passing the absolute values of $V_{.02}$ from all ‘clean’ windows through a low-pass filter with a time constant of 100 windows (1 second if all are clean). This algorithm adapts rapidly to changing noise situations, while not desensitizing during bursts.

The threshold settings are normalized to estimated RMS noise, so values of 4 to 6 are reasonable. As of August 26, 2004, this is my favorite spike detector.

SNEO

SNEO also passes the RAW data through a band pass filter, but then computes the instantaneous energy in each electrode stream:

$$E_c(t) = V'_c(t)^2 - V_c(t)V''_c(t).$$

This energy is smoothed over 5 samples and spikes are detected if it exceeds a given multiple of the RMS value of the energy. Although Kim and Kim (2000) report that SNEO works very well at signal-to-noise ratio as low as unity, I am less convinced.

Useful threshold values seem to be 5 to 20.

Filter60hz

Filter60hz provides a template filter to reduce 60 Hz line pickup in raw data and publishes the results as *60hz*. It works best if an external lock in signal is provided on one of the analog lines. If such a signal is not available, fast adaptation should be chosen to reduce the effects of gradual desynchronization. Templates are collected for each electrode channel independently. These commands are supported:

source See general description.

run Starts a single run. At the start of the run, a small amount of data is used to train the filter. During this period no output is generated.

cont Same as **run**, but restarts automatically after receiving a STOP command from the server.

nperiods [*periods*] Sets the adaptation time of the filter. Old contents are decayed by a factor $1/e$ after the given number of periods. (A period is 16.67 ms.)*

* Note for non-US users or programmers: the period is specified in units of the sample period by the variable `REALPERIOD` in `60hz/Defs.H`. When changing this variable, please be aware that MCCard’s sampling frequency, although very constant, is not exactly 25.000 kHz.

- templsize** [*size-of-template*] Number of data points to use for each template. This value is rounded internally to a power of two. Low values reduce the efficacy of the filter, but high values require longer training times and may make the system less stable. The default value, 128, should normally be adequate.
- wait** [0/1] Enables (1) or disables (0) waiting for a START command from the server. Operation with waiting disabled is poorly tested and may result in desynchronized and useless recordings. Not recommended for normal use.
- lockin** [- | *An*] Enables or disables the use of an external synchronization pulse on a given analog channel. Rising edge on the specified channel will be used to synchronize the filters to the physical 60 Hz signal.
- limit** [*adaptation-period-in-seconds* or 0 for unlimited] It may be desirable to stop adaptation all together after a certain amount of time. For example, if very strong signals are expected occasionally on the electrode channels. Such signals might be picked up by the template and cause echoes at 16.67 ms intervals. In practice, **blockonmark** provides a better solution for most cases.
- blockonmark** [- | *An* [*block-ms* [*thresh-digi* [*lookahead-ms*]]]] If your data contains (stimulation) artifacts, the adaptive filter tends to create echoes of those artifacts. This command can be used to temporarily suspend the adaptation (but not the filtering) during artifacts. When enabled, **Filter60hz** will detect upward threshold crossings on the specified analog channel (*An*), and disable adaptation on all electrode channels for the given period (*block-ms*). The threshold is specified in digital units (*thresh-digi*). The final argument (*lookahead-ms*) can be used if the marker may occur (a fraction of) a millisecond after the start of an artifact*.

For off-line usage, the command **Posthoc60hz** will read from a file and output to a unix pipe. It has the additional benefit of skipping artifacts when training. (This requires some command line switches — try ‘posthoc60hz --help’.)

Salpa

Salpa is the artifact suppression algorithm described in Wagenaar and Potter (2002). Please refer to that paper for functional details.

Salpa supports the following commands:

- run** Starts a single run. At the start of the run, a small amount of data is used to train the filter. During this period no output is generated.
- cont** Same as **run**, but restarts automatically after receiving a STOP command from the server.
- digithresh** [*digital-threshold*] Sets the threshold for acceptable asymmetry in digital units, or reports the current value. In the paper, this ‘asymmetry’ is referred to as the *deviation D*.

* Even if the marker is timed to exactly coincide with the start of the artifact, *lookahead-ms* can be used to provide a safety margin of a few samples.

noisethresh [*threshold-in-units-of-RMS-noise*] Sets the threshold for acceptable asymmetry in units of the estimated RMS noise, or reports the current value.

halfwidth [*halfwidth-in-ms*] Sets the SALPA filter halfwidth. Except for a factor τ_{sample} , this is the number N in the paper.

asymduration [*asymmetry-window-width-in-ms*] The window over which the asymmetry is measured. Except for a factor τ_{sample} , this is the number δ in the paper.

blankduration [*blanking-duration-in-ms*] Determines how much signal is blanked even after the asymmetry (deviation) test has been passed successfully. This is a bit of a hack, which was not used in the paper.

lookaheadwindow [*look-ahead-window-in-ms*] The last few samples before the signal pegs are probably not entirely artifact-free. This command allows you to blank a little bit of data just before the stimulus artifact onset.

digirails [*digi-rail1* [*digi-rail2*]] Specifies which digital values constitute the rails of the ADC. Defaults 0 and 4095 are for MultiChannel Systems hardware with **Rawsrv**.

fixedperiod [*period-ms* *delay-ms* [*blank-ms*]] In some cases, artifacts may occur that do not quite peg the channel. If this happens in a triggered recording, **Salpa** can still work: just specify the length of the trigger window (*period-ms*), the amount of time before the onset of the artifact in each window (*delay-ms*), and the duration of the fast part of the artifact (*blank-ms*). Typical values for *blank-ms* would be one or two milliseconds. **Salpa** will kick in after that and determine the end of the irreparable part of the artifact using the asymmetry (deviation) test as usual.

pegontrigger [- | An [*blank-ms* [*thresh-digi*]]] A more flexible solution to the problem explained above. **Salpa** can treat a positive threshold crossing on one given channel (usually A1) as a signal to consider all channels pegged. *Blank-ms* has the same meaning as for **fixedperiod**, and *thresh-digi* specifies the threshold (in digital units) for the detection of stimulation markers.

channels [- | + | CR ...] Normally, **Salpa** operates on all electrode channels. Using this command you can restrict operation to any subset. This is useful when the artifacts on most channels reliably last less than 1 or 2 ms, so the SALPA algorithm doesn't improve things. Say 'channels CR1 CR2 ...' to limit operation to channels *CR1*, *CR2*, ...; or 'channels -' to operate on all channels. Alternatively, say 'channels + CR ...' to add channels to an existing list, or 'channels - CR ...' to remove channels from an existing list. 'channels +' restores a list previously removed by 'channels -'. Channels excluded from operation are still subject to blanking by **fixedperiod** or **pegontrigger**. This is usually desirable. If not, you can set the blanking period to zero, e.g. by 'pegontrigger A1 0'.

source See general description.

train Estimate the RMS noise level of the input. This takes a few seconds. Unlike **run** and **cont**, which wait for the source to start, **train** must be executed while the source is already running.

savenoise See **Spikedet**.

loadnoise See **Spikedet**.

info See **Spikedet**.

An off-line version of **Salpa** is available as well; it is called **Posthocartifilt**, and supports most options of the MEABench component through command line switches. The command is self-documenting: type `posthocartifilt --help` for details.

Record

Record is used to record RAW or SPIKE streams to disk. The program can record several streams in parallel. **Record** can optionally respect the triggers on the associated wakeup socket. In this case, only the parts of the stream immediately surrounding the triggers are saved to disk, and an auxiliary file with trigger times is constructed.

These commands are supported:

record *filename* [*comments*] Starts recording to the specified file. The file name is augmented by the type of the data. Optional comments are saved to a description file, if enabled.

multirecord *base-filename* [*comments*] Starts recording to many files, one after the other. The file name is augmented by the start time of each recording, and by the type of the data.

trecord *filename pretrig-ms posttrig-ms* [*comments*] As **record**, but respects trigger information. The window of recording is specified as time before the trigger and time after the trigger, both in ms.

tmultirecord *base-filename pretrig-ms posttrig-ms* [*comments*] This is to **trecord** as **multirecord** is to **record**.

source [*name[/type]* ...] Specifies the sources for recording. **Record** knows the type of most core MEABench streams. If it doesn't know for the stream you name, the type can be specified by appending it to the stream name after a slash. (For example, if *coolsort* is a new stream of type SPIKE, you would say 'source coolsort/spike'.)

describe [0/1] Enables (1) or disables (0) the generation of a description file (file name constructed by augmenting the data file name by '.desc'). Description files contain lots of useful information pertaining to a run and are in human-readable form.

Record can record from several sources simultaneously: just specify them all together as arguments to **source**. For example, 'source raw 60hz spike' would prepare a recording from three sources. If more than one stream of the same type is recorded, the stream names are incorporated in the file names. Recording from several sources does have a few limitations:

- Recording ends immediately when the first stream terminates. A small amount of data from the end of the other streams may be lost.
- For triggered recording, only the first stream will yield a '.trig' file.

To avoid these limitations, it is possible to run several instances of **Record**, and have each of them record from a single stream.

It is possible to terminate a recording session before the source ends by pressing **Ctrl-C** in the **Record** terminal window. This is useful for recording a short segment of RAW data parallel to the beginning of a longer SPIKE data recording.

Replay

Replay replays files recorded by **Record**. Currently, RAW and SPIKE data are supported, and are published as *reraw* and *respike* respectively. Replaying a SPIKE file results in both a RAW and a SPIKE stream, the RAW stream reporting the contexts stored in the SPIKE stream.

These commands are supported:

play [*filename* [*type*]] Plays the given file. Normally, **Replay** automatically detects the type of the data. If it doesn't, help it by specifying it explicitly. Without a file name, plays the last played file again.

slow [*slowdown-factor*] Slows down playback by a given (real valued) factor. Useful for output to **Scope**. Arguments smaller than one cause speed up.

run Alias for **play** without arguments.

source [*filename* [*type*]] Specifies a filename (and optional type) for later playback.

blankout [*period-in-ms* or 0] Enables or disables blanking of electrode channels during a trigger. Blanking is performed by replacing the data by the average of four samples obtained just prior to the stimulus. Blankout works only for triggered files. **Replay** does not detect stimuli itself.

selftrig [0/1] Enables or disables spike detection from an analog channel in the stream. When disabled, triggers stored in the .desc file of a file produced by **Record's** **trecord** command are still reported.

trigchannel [1/2/3] Sets the analog channel on which triggers are detected (if **selftrig** is enabled).

trigthreshold [*digivalue*] Sets the threshold for such trigger detection.

Scope

Scope shows raw data and spikes in a similar format as the graphical parts of my older **Qmeagraph** program, or MultiChannel Systems' **MCRack**. It mostly explains itself. Here are some useful hints:

- Double-clicking on any of the small electrode traces opens a separate window showing that channel.
- Selecting window width, pre-trigger length, or voltage ranges in **Scope** does not affect recording in any way.
- The scrollbar buffer (enabled by clicking the 'Freeze' button) is very useful to home in on some interesting event. Raw data in the scrollbar buffer can be saved to disk using the 'Save' button which appears whenever 'Freeze' is enabled. The scrollbar buffer is about 5 seconds long.
- These known bugs exist in scope:
 - Spike circles at the edges of the electrode traces leave semi-permanent smudges.
 - When scrolling back, older spikes may lose their red marks if there are many detected spikes. This is the result of the SPIKE data shared memory segments being too short.

- When replaying a spike stream, red spike marks sometimes appear out of nowhere. These ghosts are easily recognized, because no context data are plotted around them. This bug has been a mystery so far.

Spikesound

Spikesound makes spikes audible through a PC sound card. It can read of any **SPIKE** stream (e.g., straight from the spike detector, or from **Replay**). The GUI controls are minimal:

Source Selects the MEABench stream to get spikes from. That usually is *spike* to read from **Spikedet**, or *respike* to read from **Replay**.

Play Switches sound on or off

Volume Master volume control.

–ve only Limits sonification to downward spikes. The obvious counterpart can be implemented on request.

Rethreshold It is often useful to set the spike detection threshold fairly low in **Spikedet** and use off-line spike sorting to clean up the data. However, hearing all the near-noise-level spikes is not very appealing. This control allows you to hear only strong spikes. For example, I like to set the **Spikedet** threshold to 4.5σ , and set the **Spikesound** rethreshold factor to 120%.

A1, A2, A3 Enables or disables beeps when a spike is detected on one of the auxillary channels. For example, I like to use A1 and A3 for triggers, and A2 as a 60 Hz lock-in signal. So I may want to hear a sound when a trigger happens, but I don't want to hear the 60 Hz signal.

Customization note: The current version uses output buffers of about 25 ms to improve timing accuracy and reduce lag. If you prefer hearing longer noises, or if your sound card or CPU does not allow you to use such short buffers, you may change the value of `AUDIO_LOG_FRAG` in `spikesound/Audio.H`. The length of the buffer is:

$$\tau = \frac{2^{\text{AUDIO_LOG_FRAG}}}{176.4} \text{ ms.}$$

Flexraster

Flexraster displays raster plots of spike activity in triggered recordings. It currently relies on 'trigger spikes' on channel A1. There are six different ways to create rasters:

Spont Spikes from all channels are combined as blue dots. Each line of the raster is *Pre* plus *Post* ms wide. The raster plot scrolls vertically, showing the most recent interval on top. A new raster line is generated whenever there is a trigger on A1, or every second if there has not been any trigger.

8x8Rec Spikes from each channel are displayed in separate panels. Within each panel, the display is as for Spont.

8x8Stim If the shape of the trigger pulse is used to encode a CR-position, each panel displays all spikes that occurred in response to stimulation on a certain channel. See figure for details.

V Stim and **H Stim** Similar to 8x8Stim, but panels are generated only for those CRs that actually receive stimuli.

Cont Arguably the most useful function, creates a scrollable and scalable raster plot with 60 electrodes stacked vertically and time running continuously left to right. Trigger pulses are indicated by red marks.

Flexraster is under development. Suggestions for improvement are especially welcome.

Commander

Commander allows you to start and control MEABench programs from within a central shell-like language. This is mainly useful for off-line analysis. You may find an example of its use in *Examples*, above. Commander logs all interaction with the programs it controls to the screen, as well as to an optional log file. Creating log files is highly recommended, because it allows you to keep track of how exactly you processed the data, and to check whether MEABench behaved as you expected it to behave.

Commander supports the following commands:

new *program[id]* [*args*] Starts a new program. An optional *id* may be assigned to the program to disambiguate references to two instances of the same program, or for ease of reference. *Args* are passed to the program unchanged.

tell *program[id command]* [*args*] Sends a command to the named program. *Args* are passed unchanged.

expect *program[id timeout regexp]* Waits until the named program produces output that matches *regexp*. If this does not happen within *timeout* seconds, reports an error.

dieif *program[id regexp]* Looks back at the output captured by the last **expect**, and abort if any line matches *regexp*. (This may be a line that was output long before the line that made **expect** happy.)

dieunless *program[id regexp]* Looks back at the output captured by the last **expect**, and aborts unless some line matches *regexp*. (This may be a line that was output long before the line that made **expect** happy.)

flush *program[id]* Flushes all output from the named program, so that future **expect** commands will not match any output prior to the **flush**.

intr *program[id]* Sends an Interrupt signal to the named program, i.e., simulates pressing **Ctrl-C**.

kill *program[id]* Sends a Term signal to the named program, terminating it. Normally, you would use **close** to achieve the same result more gracefully.

close *program[id]* Closes a running program normally, as if **Ctrl-D** was pressed. If closing fails, the program is terminated as per the **kill** command.

wait *program[id]* Waits for the named program to terminate. Use this if you just sent it a **quit** command. Normally, **close** is an easier way to terminate subprocesses.

log [*logfile*] Writes all future output to the named file, or stops logging if no argument is given.

comment *comments* Writes the specified comments to the log file.

sleep *time-s* Sleeps for the given number of seconds. The subprocesses are not affected.

quit Exit **Commander** after closing all subprocesses.

Unlike other MEABench programs, commander does not present the user with a prompt. You are not really expected to type away at commander from a terminal (although you can). Instead, you would normally prepare a script, and then run commander on it:

```
$ mea commander < myscript.cmdr
```

Any error that happens during script execution (e.g. failure to start a subprocess or failure to read an **expect**-ed string) causes **Commander** to terminate immediately with an error report written to the screen and the log file.

An auxiliary program, **Cmdlog2html** exists to convert log files produced by **Commander** to html format. Log files are human readable, but the html format looks nicer.

Monitor

Monitor is mostly a debugging tool. It displays the status of each of the MEABench servers.

Neurosock and NSsrv

Run **Neurosock** on the machine that contains the physical hardware, and MEABench on any other machine that can connect to the first machine through the internet. It can be run without any arguments for 64-channel MCCards, or as **neurosock -set MC128** for 128-channel MCCards. **NSsrv** is exactly like **Rawsrv**, except that it doesn't record directly from the MEA hardware. Instead, it connects to **Neurosock** on the same or another computer. For 64-channel MCCards, run **NSsrv** without arguments; to use the 2nd half of 128-channel cards, run it like **nssrv -s raw2**. **NSsrv** has the same commands as **Rawsrv**, with one addition:

ip aa.bb.cc.dd Specify the IP address of the computer running **Neurosock**.

Note: Only one instance of **NSsrv** can read from a 64-channel **Neurosock** or half a 128-channel **Neurosock** at a time. It is not possible to specify gain separately for the two halves, and the **gain** command will fail if you try to set gain on one half while the other half is running.

Spikedump

Spikedump converts SPIKE files to human readable form, dropping the context data. It can read from a specified file, or from stdin. It does not presently run off MEABench streams. Output fields are *time* (in seconds), *channel* (hardware order, counting from zero), *height* (digital value) and *width* (in samples).

Doubletxt

Doubletxt combines a RAW stream with a SPIKE stream to construct a file with two context fields per spike. The resulting file can be read with the `loaddblctx.m` Matlab function.

Other Utilities

This documentation is presently incomplete regarding the utilities in the `meabench/utls` and `meabench/perl` subdirectories. These utilities will provide a brief usage message if invoked with a `--usage` argument.

Matlab Functions

A number of Matlab functions included with MEABench can be used to load MEABench data files into Matlab, to perform channel numbering convention conversions, and to perform some common visualization tasks. These functions are installed in `/opt/meabench/matlab` (if you follow the suggested installation procedure). In order to use them, you need to tell Matlab about that directory:

```
>> addpath('/opt/meabench/matlab');
```

The functions are the following:

cleanctxt

`[ctxts, idx] = CLEANCTXT(contexts)` returns cleaned up contexts:

- The first and last 15 values are averaged and used to compute DC offset.
- These two estimates are weighted according to their inverse variance.
- The DC offset is subtracted.
- If any sample in $-1 \dots -0.5$ or $0.5 \dots 1.5$ ms is more than half the peak at 0 ms, the spike is rejected.
- Use `CLEANCTXT(contexts, testidx, relthresh)` to modify this test:
 - testidx* are indices (1...74) of samples to test,
 - relthresh* is a number between 0 and 1.
- Additionally, the area immediately surrounding the peak is tested at the 0.9 level: the spike is also rejected if any sample in $-1 \dots -0.16$ or $0.16 \dots 1.5$ ms has an absolute value more than $0.9 \times$ the absolute peak value. This test is modified on its outer edges by the edges of *testidx*, but cannot be modified independently.

Returns: *ctxts*: the accepted contexts, with DC subtracted.

idx: the index of accepted spikes.

Requirements: *contexts* must be as read from `loadspike`, i.e. $74 \times N$ (or $75 \times N$).

Acknowledgment: The algorithm implemented by this function is due to P. P. Mitra.

cr2hw

`hw = CR2HW(c,r)` converts row and column to hardware channel number. Note that *c*, *r* count from 1 to 8, while *hw* counts from 0.

`hw = CR2HW(cr)` converts combined row and column to hardware channel number. Here, *cr* can be in the range 11...88.

Illegal *c*, *r* values return -1 . If *c*, *r* or *cr* are matrices, the output will also be a matrix. The dimensions of *c*, *r* must agree.

heightscat88

`HEIGHTSCAT88(spks)` plots scatter plots of the (spontaneous activity) spikes in SPKS. It

stacks thin horizontal raster plots for each channel. Within each plot, spikes are positioned based on their height (amplitude).

spks must have been loaded using `loadspike` or `loadspike_noc`.

Column-row numbers are computed from channel numbers using `hw2cr`.

`p = HEIGHTSCAT88(spks)` returns the plot handle.

hist2dar

`mat=hist2dar(X,Y,nx,ny,flag)` returns a matrix suitable for `pcolor` containing the crosstab counts of *X* and *Y* in automatically selected bins: you pick the number of bins, the code determines the edges based on the min and max values in *X* and *Y*.

If optional argument *flag* is present, the result is also plotted.

hist2d

`mat=hist2d(X,Y,x0,dx,x1,y0,dy,y1,flag)` returns a matrix suitable for `pcolor` containing the crosstab counts of *X* and *Y* in the bins edges defined by $x_0:dx:x_1$ resp $y_0:dy:y_1$.

If optional argument *flag* is present, the result is also plotted. `[mat,xx,yy]=hist2d(...)` returns *x* and *y* arrays as well, so you can call `pcolor(xx,yy,mat)`.

hw2cr

`[c,r] = HW2CR(hw)` converts the hardware channel number *hw* (0...63) to column and row numbers (1...8).

`cr = HW2CR(hw)` converts the hardware channel number *hw* (0...63) to a 1+row+10×col format (11...88). If *hw* is a matrix, the result will also be a matrix.

Illegal hardware numbers result in $c = 0, r = 0$.

loaddesc

`d = LOADDESC(fn)` reads the description file *fn* or *fn.desc* and returns a structure with values for each line read.

The fields are named from the label in the description file.

All values are converted to double. Original strings are stored in the field *label_str*.

Repeated keys are stored in a cell array.

loadraw

`y=LOADRAW(fn)` reads the raw MEA datafile *fn* and stores the result in *y*.

`y=LOADRAW(fn,range)` reads the raw MEA datafile *fn* and converts the digital values to voltages by multiplying by *range*/2048.

Range values 0,1,2,3 are interpreted specially:

Range value	Electrode range (μV)	Auxiliary range (mV)
0	3410	4092
1	1205	1446
2	683	819.6
3	341	409.2

“electrode range” is applied to channels 0...59, “auxiliary range” is applied to channels 60...63. Note that channel *hw* is stored in the (*hw*+1)-th row of the output.

loadspike

`y=LOADSPIKE(fn)` loads spikes from given filename into structure *y* with members

time (1×N) (in samples)
channel (1×N)
height (1×N)
width (1×N)
context (75×N)
thresh (1×N)

`y=LOADSPIKE(fn,range,freq_khz)` converts times to seconds and width to milliseconds using the specified frequency, and the height and context data to microvolts by multiplying by *range*/2048.

As a special case, *range* = 0...3 is interpreted as a MultiChannel Systems gain setting, as for `loadraw`. In this case, the frequency is set to 25 kHz unless specified.

loadspike_noc

`y=LOADSPIKE_NOC(fn)` loads spikes from given filename into structure *y* with members as for `loadspike`, except that *context* is not loaded.

`y=LOADSPIKE_NOC(fn,range,freq_khz)` converts times to seconds and width to milliseconds using the specified frequency, and the height and context data to microvolts by multiplying by *range*/2048.

As a special case, *range* = 0...3 is interpreted as a MultiChannel Systems gain setting, as for `loadraw`. In this case, the frequency is set to 25 kHz unless specified.

randscat88

`RANDSCAT88(spks)` plots scatter plots of the (spontaneous activity) spikes in *spks*. It stacks thin horizontal raster plots for each channel. Within each plot, spikes are randomly positioned vertically for clarity.

spks must have been loaded using `loadspike` or `loadspike_noc`.

Column-row numbers are computed from channel numbers using `hw2cr`.

timeclust

`[t0,cnt,dt] = TIMECLUST(tms_s,bin_s,thr_mean,thr_abs)`

Given a set of times *tms_s* and a bin size *bin_s* (both nominally in seconds), find the locations, volumes, and widths of peaks in the time distribution.

A peak is (primitively) defined as a contiguous area of bins exceeding *thr_abs* and exceeding *thr_mean* times the mean bin count. Either *thr_mean* or *thr_abs* may be left unspecified, in which case *thr_abs* defaults to 2.

File Formats

The current version of MEABench uses two binary file formats: RAW, and SPIKE. All other files are plain text.

Raw Files

Raw files contain electrode voltage values in digitized form. The scaling factor between digital values and microvolts is not stored in the `.raw` file, but in an accompanying `.raw.desc` file. The number of channels in a raw file is currently fixed to 64, and not noted in the `.raw.desc` file.

This will be changed in a future release. File format is: one 16-bit integer for every channel, 64 channels per scan, repeated for the length of the file. Thus, in (almost) C notation, a raw file would be defined as:

```
typedef short int scan[64];
typedef scan rawfile[];
```

Spike Files

Spike files contain information about detected spikes. The time of detection (in samples) as well as the channel number, height (in digital units), and width (in samples) are stored in a structure that also contains a limited amount of ‘context’, or sample values immediately surrounding the spike. The BandFilt and AdaFilt spike detectors also record the threshold used for detecting that particular spike in the structure. In (almost) C notation, a spike file would be defined as:

```
typedef struct {
    long long time; // 64 bits of time (in sample periods from start of file)
    short channel; // a channel number, 0-63
    short height; // height from baseline, may be negative (in digital units)
    short width; // the width of the spike (in samples)
    short context[74]; // 24 samples before, and 49 samples after peak
    short threshold; // detection threshold used for this spike
} spike;
typedef spike spikefile[];
```

As for raw files, conversion factors may be found in a description file (.spike.desc).

Hints and Tips

This section contains solutions to some common problems and provides some hopefully helpful hints.

None of the Programs Will Run

You may find that none of the programs will run, and complain like this:

```
/opt/meabench/bin/replay: error in loading shared libraries: libmea.so: cannot open shared object
file: No such file or directory
```

This means that they cannot find the common libraries *libmea.so* or *libmeagui.so*. This problem may be fixed by typing

```
export LD_LIBRARY_PATH=/opt/meabench/lib:$LD_LIBRARY_PATH
```

before running the command. Similarly, if the perl programs complain about missing libraries, try:

```
export PERL5LIB=/opt/meabench/libexec:$PERL5LIB
```

(If you installed MEABench somewhere else, you will know how to modify the above lines.)
If you use **Mea** this problem is less likely to show.

Client X Keeps Saying ‘Waiting for START from Y’ and Won’t Run

Most clients synchronize the start of their run with the start of a run of the server they are connected to. To make this work smoothly, start the client off (typically by typing **run**) *before* starting the server. Some clients allow you to disable this synchronization behavior, e.g. through a **wait** command.

My Recordings Are Truncated or Empty

If you press **Ctrl-C** in **Record**, it stops immediately without saving any buffered data. This is usually not what you want. You can avoid this problem by always ending recording sessions by pressing **Ctrl-C** in **Rawsrv** or **NSsrv** instead.

Shared Memory Problems

You may find that a server program refuses to run, complaining like this:

```
Error from ShmSrv: Segment exists, please delete using 'ipcrm shm 3417399'
```

This normally means that the server crashed on a previous run, leaving behind the shared memory segment used for its output stream. Executing the suggested command often solves the problem, but you may be rewarded by the following when you rerun the program:

```
Error from ShmSrv: Segment exists, and cannot be accessed. Any lingering clients? Before deleting  
the segment using ipcrm (see man page) these may have to be stopped.
```

This means that some client is still connected to the (defunct) crashed server. Quitting and restarting such clients is a pretty sure way of solving the problem. If quitting the client is undesirable, most clients can be convinced to renegotiate the connection to their server using the command **source**. Once the clients have been cleared, the segment normally goes away spontaneously. If the problem persists, you may have to resort to **Ipcs** and **Ipcrm**. See the linux man pages for details.

```
Error from ShmSrv: creat
```

This exceedingly unhelpful error message may mean that you don’t have a `.meabench` subdirectory in your home directory. This is where MEABench tries to link all its shared memory segments and wakeup sockets. To make this problem go away for ever, create the directory using `mkdir ~/.meabench`.

Abbreviating Commands

All programs that provide the user with a command line accept non-ambiguous abbreviations for commands. For example, in most programs, **run** can be abbreviated to **r**.

Passing Commands at Run Time

All programs that provide the user with a command line can also process commands passed at startup time. For example, instead of entering into the following dialog:

```
$ spikedet
spikedet> type 3
Type is BandFit-25
spikedet> threshold 4
Threshold is 4
spikedet> source 60hz
Source is 60hz
spikedet>
```

you may also say:

```
$ spikedet ty:3 thr:4 so:60hz
Type is BandFit-25
Threshold is 4
Source is 60hz
spikedet>
```

This is especially useful if you find yourself quitting and restarting a program frequently, since with the second scheme, the shell history facilities can be used.

At start up time, commands are separated from their arguments by a colon rather than a space, and arguments are separated from each other by a comma rather than a space.

Interrupting Long Commands

All command line programs respond to **Ctrl-C** by returning to their prompt. If they don't respond immediately, press **Ctrl-C** again after a second or two. Pressing **Ctrl-C** twice in quick succession kills the program forcefully, quite possibly leaving a shared memory segment behind as explained above. Pressing **Ctrl-C** while the program is waiting for user input does not cause the program to exit. All programs support a **quit** command to exit cleanly, but **Ctrl-D** (end of input) is also a fine way to exit them.

Debugging Information

Most programs support two additional commands to aid debugging. These are:

dbx [0/1] Enables (1) or disables (0) debugging output. Debugging output varies wildly from component to component and from release to release. Mainly of use for developers, who can sprinkle their code with `sdbx(...)` calls to track bugs down.

clients Prints a list of all clients currently connected to this server program. Mainly intended for debugging, the output format is not very user friendly.

Contacting the Author

For further information, or to report bugs, please contact the author, Daniel Wagenaar. My address is in flux, so please refer to the web site below for current information.

Suggestions for improvement are always welcome, but I cannot guarantee I will have time to implement them. You too can contribute by sending me your bug-fixing and feature-adding patches by e-mail. Please use ‘diff -C2’ against the latest public release version.

The latest versions of MEABench and this documentation can be found at <http://www.its.caltech.edu/~wagenaar/meabench.html>. Major new releases of MEABench will be announced on the meausers mailing list. Please refer to the web site for details.

Reporting Bugs

When reporting bugs, please include the following information:

- The version of MEABench you are using.
- The version of g++ you are using (type `g++ --version` to get this information).
- The version of Qt you are using (type `moc -v` to get this information).
- The output of the ‘configure’ script.
- If compilation failed: the output of the ‘make’ command (not just the lines containing the error).
- For runtime errors: Which programs you are running, which commands you have executed within those programs, and a description of what you are trying to do as well as any relevant screen output. A succinct set of conditions that reproducibly produces the error is very much appreciated.

Bibliography

- Abeles M, 1989: *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press.
- Amit D J, 1989: *Modeling brain function: The world of attractor neural networks*. Cambridge University Press.
- Bai Q and Wise K D, 2001: Single-unit neural recording with active microelectrode arrays. *IEEE Trans Biomed Eng* **48** (8): 911–920.
- Bakkum D J, Shkolnik A C, Ben-Ary G, Gamblen P, DeMarse T B, and Potter S M, 2004: Removing some ‘A’ from AI: Embodied cultured networks. In F Iida, L Steels, and R Pfeifer (eds.), *Embodied Artificial Intelligence*. Springer-Verlag. In press.
- Barabanov M, 1997: *A Linux-based real-time operating system*. Ph.D. thesis, New Mexico Institute of Mining and Technology.
- Beggs J M and Plenz D, 2003: Neuronal avalanches in neocortical circuits. *J Neurosci* **23** (35): 11167–11177.
- Ben-Ari Y, 2001: Developing networks play a similar melody. *Trends Neurosci* **24** (6): 353–360.
- Ben-Menachem E, Manon-Espaillat R, Ristanovic R, Wilder B J, Stefan H, Mirza W, Tarver W B, and Wernicke J F, 1994: Vagus nerve stimulation for treatment of partial seizures: 1. A controlled study of effect on seizures. First International Vagus Nerve Stimulation Study Group. *Epilepsia* **35**: 616–626.
- Bi G Q and Poo M M, 2001: Synaptic modification by correlated activity: Hebb’s postulate revisited. *Annu Rev Neurosci* **24**: 139–166.
- Bove M, Grattarola M, Tedesco M, and Verreschi G, 1994: Characterization of growth and electrical activity of nerve cells cultured on microelectronic substrates—

- towards hybrid neuro-electronic devices. *J Mater Sci-Mater Med* **5**: 684–687.
- Bove M, Grattarola M, and Verreschi G, 1997: In vitro 2-D networks of neurons characterized by processing the signals recorded with a planar microtransducer array. *IEEE Trans Biomed Eng* **44**: 964–977.
- Bragin A, Hetke J, Wilson C L, Anderson D J, Engel J, and Buzsaki G, 2000: Multiple site silicon-based probes for chronic recordings in freely moving rats: implantation, recording and histological verification. *J Neurosci Methods* **98** (1): 77–82.
- Branner A and Normann R A, 2000: A multielectrode array for intrafascicular recording and stimulation in sciatic nerve of cats. *Brain Res Bull* **51** (4): 293–306.
- Brewer G J and Cotman C W, 1989: Survival and growth of hippocampal-neurons in defined medium at low-density—advantages of a sandwich culture technique or low oxygen. *Brain Research* **494**: 65–74.
- Brewer G J, Torricelli J R, Evege E K, and Price P J, 1993: Optimized survival of hippocampal neurons in B27-supplemented Neurobasal(tm), a new serum-free medium combination. *J Neurosci Res* **35** (5): 567–576.
- Buitenweg J R, Rutten W L C, Willems W P A, and Van Nieuwkastele J W, 1998: Measurement of sealing resistance of cell-electrode interfaces in neuronal cultures using impedance spectroscopy. *Med Biol Eng Comput* **36** (5): 630–637.
- Buonomano D V and Merzenich M M, 1998: Cortical plasticity: From synapses to maps. *Annu Rev Neurosci* **21**: 149–186.
- Canepari M, Bove M, Maeda E, Cappello M, and Kawana A, 1997: Experimental analysis of neuronal dynamics in cultured cortical networks and transitions between different patterns of activity. *Biol Cybern* **77** (2): 153–162.
- Chao Z C, Bakkum D J, Wagenaar D A, and Potter S M, 2005: Effects of random external background stimulation on network synaptic stability after tetanization: a modeling study. *Neuroinformatics*: In press.
- Chiu C Y and Weliky M, 2001: Spontaneous activity in developing ferret visual cortex in vivo. *J Neurosci* **21** (22): 8906–8914.

- Contreras D, Destexhe A, Sejnowski T J, and Steriade M, 1997: Spatiotemporal patterns of spindle oscillations in cortex and thalamus. *J Neurosci* **17** (3): 1179–1196.
- Cooper I S, Amin I, and Gilman S, 1973: The effect of chronic cerebellar stimulation upon epilepsy in man. *Trans Am Neurol Assoc* **98**: 192–196.
- Cooper I S, Amin I, Riklan M, Waltz J M, and Poon T P, 1976: Chronic cerebellar stimulation in epilepsy. Clinical and anatomical studies. *Arch Neurol* **33**: 559–570.
- Cooper I S, Amin I, Upton A, Riklan M, Watkins S, and McLellan L, 1977: Safety and efficacy of chronic stimulation. *Neurosurgery* **1**: 203–205.
- Corner M A and Ramakers G J A, 1992: Spontaneous firing as an epigenetic factor in brain-development—physiological consequences of chronic tetrodotoxin and picrotoxin exposure on cultured rat neocortex neurons. *Dev Brain Res* **65** (1): 57–64.
- Corner M A, Van Pelt J, Wolters P S, Baker R E, and Nuytinck R H, 2002: Physiological effects of sustained blockade of excitatory synaptic transmission on spontaneously active developing neuronal networks—an inquiry into the reciprocal linkage between intrinsic biorhythms and neuroplasticity in early ontogeny. *Neurosci Biobehav Rev* **26** (2): 127–185.
- Cossart R, Aronov D, and Yuste R, 2003: Attractor dynamics of network UP states in the neocortex. *Nature* **423** (6937): 283–288.
- Curtis A S, Breckenridge L, Connolly P, Dow J A, Wilkinson C D, and Wilson R, 1992: Making real neural nets: design criteria. *Med Biol Eng Comput* **30**. CE33–6.
- Darbon P, Scicluna L, Tscherter A, and Streit J, 2002: Mechanisms controlling bursting activity induced by disinhibition in spinal cord networks. *Eur J Neurosci* **15** (4): 671–683.
- Davis R, 2000: Cerebellar stimulation for cerebral palsy spasticity, function, and seizures. *Arch Med Res* **31**: 290–299.
- DeAngelis G C, Cumming B G, and Newsome W T, 1998: Cortical area MT and the perception of stereoscopic depth. *Nature* **394** (6694): 677–680.
- DeMarse T B, Wagenaar D A, Blau A W, and Potter S M, 2001: The neurally con-

- trolled animat: Biological brains acting with simulated bodies. *Auton Robot* **11**: 305–310.
- Dempster A P, Laird N M, and Rubin D B, 1977: Maximum likelihood from incomplete data via EM algorithm. *J R Stat Soc Ser B-Methodol* **39** (1): 1–38.
- Dichter M A, 1978: Rat cortical neurons in cell culture: Culture methods, cell morphology, electrophysiology, and synapse formation. *Brain Res* **149**: 279–293.
- Droge M H, Gross G W, Hightower M H, and Czisny L E, 1986: Multielectrode analysis of coordinated, multisite, rhythmic bursting in cultured CNS monolayer networks. *J Neurosci* **6**: 1583–92.
- Echevarria D and Albus K, 2000: Activity-dependent development of spontaneous bioelectric activity in organotypic cultures of rat occipital cortex. *Dev Brain Res* **123** (2): 151–164.
- Egert U, Knott T, Schwarz C, Nawrot M, Brandt A, Rotter S, and Diesmann M, 2002: MEA-Tools: an open source toolbox for the analysis of multi-electrode data with MATLAB. *J Neurosci Methods* **117** (1): 33–42.
- Egert U, Schlosshauer B, Fennrich S, Nisch W, Fejtl M, Knott T, Muller T, and Hammerle H, 1998: A novel organotypic long-term culture of the rat hippocampus on substrate-integrated multielectrode arrays. *Brain Research Protocols* **2**: 229–242.
- Emery D G, Lucas J H, and Gross G W, 1991: Contributions of sodium and chloride to ultrastructural damage after dendrotomy. *Exp Brain Res* **86**: 60–72.
- Erickson J, Tooker A, Tai Y C, and Pine J, 2005: Caged neuron multi-electrode array: a new generation of neurochip. *Proc Soc for Neurosci* Submitted.
- Eytan D, Brenner N, and Marom S, 2003: Selective adaptation in networks of cortical neurons. *J Neurosci* **23** (28): 9349–9356.
- Fisher R, 1922: On the mathematical foundations of theoretical statistics. *Philosophical Trans of the Royal Society of London Series A* **222**: 309–368.
- Fisher R S, Krauss G L, Ramsay E, Laxer K, and Gates J, 1997: Assessment of vagus nerve stimulation for epilepsy: Report of the Therapeutics and Technology Asses-

- ment Subcommittee of the American Academy of Neurology. *Neurology* **49**: 293–297.
- Franaszczuk P J, Kudela P, and Bergey G K, 2003: External excitatory stimuli can terminate bursting in neural network models. *Epilepsy Res* **53** (1–2): 65–80.
- Fu Y X, Djupsund K, Gao H F, Hayden B, Shen K, and Dan Y, 2002: Temporal specificity in the cortical plasticity of visual space representation. *Science* **296** (5575): 1999–2003.
- Furshpan E J and Potter D D, 1989: Seizure-like activity and cellular-damage in rat hippocampal neurons in cell-culture. *Neuron* **3** (2): 199–207.
- Gray C M, Maldonado P E, Wilson M, and McNaughton B, 1995: Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods* **63** (1–2): 43–54.
- Gross G W, 1979: Simultaneous single unit recording in vitro with a photoetched laser deinsulated gold multimicroelectrode surface. *IEEE Trans Biomed Eng* **26** (5): 273–279.
- Gross G W, Harsch A, Rhoades B K, and Gopel W, 1997: Odor, drug and toxin analysis with neuronal networks in vitro: Extracellular array recording of network responses. *Biosens Bioelectron* **12** (5): 373–393.
- Gross G W and Kowalski J M, 1999: Origins of activity patterns in self-organizing neuronal networks in vitro. *J Intell Mater Syst Struct* **10** (7): 558–564.
- Gross G W, Rhoades B K, Reust D L, and Schwalm F U, 1993: Stimulation of monolayer networks in culture through thin-film indium-tin oxide recording electrodes. *J Neurosci Methods* **50** (2): 131–143.
- Grumet A E, 1999: *Electrical stimulation parameters for an epi-retinal prosthesis*. Ph.D. thesis, Massachusetts Institute of Technology.
- Grumet A E, Wyatt Jr J L, and Rizzo III J F, 2000: Multi-electrode stimulation and recording in the isolated retina. *J Neurosci Methods* **101** (1): 31–42.
- Habets A M M C, Van Dongen A M J, Van Huizen F, and Corner M A, 1987: Spon-

- taneous neuronal firing patterns in fetal rat cortical networks during development in vitro: a quantitative analysis. *Exp Brain Res* **69**: 43.
- Hammond E J, Uthman B M, Reid S A, Wilder B J, and Ramsay R E, 1990: Vagus nerve stimulation in humans: Neurophysiological studies and electrophysiological monitoring. *Epilepsia* **31:S2**: 51–59.
- Handforth A, DeGiorgio C M, Schachter S C, Uthman B M, Naritoku D K, Tecoma E S, Henry T R, Collins S D, Vaughn B V, Gilmartin R C, Labar D R, Morris G L r, Salinsky M C, Osorio I, Ristanovic R K, Labiner D M, Jones J C, Murphy J V, Ney G C, and Wheless J W, 1998: Vagus nerve stimulation therapy for partial-onset seizures: a randomized active-control trial. *Neurology* **51**: 48–55.
- Harsch A, Konno K, Takayama H, Kawai N, and Robinson H, 1998: Effects of alphapompilidotoxin on synchronized firing in networks of rat cortical neurons. *Neurosci Lett* **252**: 49–52.
- Harsch A and Robinson H P C, 2000: Postsynaptic variability of firing in rat cortical neurons: the roles of input synchronization and synaptic NMDA receptor conductance. *J Neurosci* **20** (16): 6181–6192.
- Hastie T and Loader C, 1993: Local regression: Automatic kernel carpentry. *Stat Sci* **8** (2): 120–143.
- Hebb D O, 1949: *The organization of behavior: a neuropsychological theory*. John Wiley and Sons, New York.
- Heck D, 1995: Investigating dynamic aspects of brain-function in slice preparations—spatiotemporal stimulus patterns generated with an easy-to-build multielectrode array. *J Neurosci Methods* **58** (1–2): 81–87.
- Heuschkel M O, Fejtl M, Raggenbass M, Bertrand D, and Renaud P, 2002: A three-dimensional multi-electrode array for multi-site stimulation and recording in acute brain slices. *J Neurosci Methods* **114** (2): 135–148.
- Honma S, Shirakawa T, Katsuno Y, Namihira M, and Honma K, 1998: Circadian periods of single suprachiasmatic neurons in rats. *Neurosci Lett* **250**: 157–160.
- Horowitz P and Hill W, 1996: *The art of electronics*. 2nd ed, Cambridge University

- press.
- Iasemidis L D, 2003: Epileptic seizure prediction and control. *IEEE Trans Biomed Eng* **50** (5): 549–558.
- Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, and Yuste R, 2004: Synfire chains and cortical songs: Temporal modules of cortical activity. *Science* **304** (5670): 559–564.
- Jackson L B, 1996: *Digital filters and signal processing, with MATLAB exercises*. Kluwer, Boston.
- Jimbo Y, Kasai N, Torimitsu K, Tateno T, and Robinson H P C, 2003: A system for MEA-based multisite stimulation. *IEEE Trans Biomed Eng* **50** (2): 241–248.
- Jimbo Y and Kawana A, 1992: Electrical stimulation and recording from cultured neurons using a planar electrode array. *Bioelectrochem Bioenerg* **29** (2): 193–204.
- Jimbo Y, Kawana A, Parodi P, and Torre V, 2000: The dynamics of a neuronal culture of dissociated cortical neurons of neonatal rats. *Biol Cybern* **83** (1): 1–20.
- Jimbo Y, Robinson H P C, and Kawana A, 1998: Strengthening of synchronized activity by tetanic stimulation in cortical cultures: Application of planar electrode arrays. *IEEE Trans Biomed Eng* **45** (11): 1297–1304.
- Jimbo Y, Tateno T, and Robinson H P C, 1999: Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys J* **76** (2): 670–678.
- Kamioka H, Maeda E, Jimbo Y, Robinson H P C, and Kawana A, 1996: Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. *Neurosci Lett* **206** (2–3): 109–112.
- Keefer E W, Norton S J, Boyle N A J, Talesa V, and Gross G W, 2001: Acute toxicity screening of novel AChE inhibitors using neuronal networks on microelectrode arrays. *Neurotoxicology* **22**: 3–12.
- Kellinghaus C, Loddenkemper T, Moddel G, Tergau F, Luders J, Ludemann P, Nair D R, and Luders H O, 2003: Electric brain stimulation and epilepsy. *Nervenarzt*

- 74** (8): 664–676.
- Kim K H and Kim S J, 2000: Neural spike sorting under nearly 0-dB signal-to-noise ratio using nonlinear energy operator and artificial neural-network classifier. *IEEE Trans Biomed Eng* **47** (10): 1406–1411.
- Koo B, 2001: EEG changes with vagus nerve stimulation. *J Clin Neurophysiol* **18**: 434–441.
- Kovacs G T A, 1994: Introduction to the theory, design, and modeling of thin film microdevices for neural interfaces. In D A Stenger and T M McKenna (eds.), *Enabling technologies for cultured neural networks*, pp. 121–162. Academic Press.
- Krahe R and Gabbiani F, 2004: Burst firing in sensory systems. *Nature Rev Neurosci* **5** (1): 1–11.
- Latham P E, Richmond B J, Nirenberg S, and Nelson P G, 2000: Intrinsic dynamics in neuronal networks. II. Experiment. *J Neurophysiol* **83** (2): 828–835.
- Leinekugel X, Khazipov R, Cannon R, Hirase H, Ben-Ari Y, and Buzsaki G, 2002: Correlated bursts of activity in the neonatal hippocampus in vivo. *Science* **296** (5575): 2049–2052.
- Leonardo A and Fee M S, 2005: Ensemble coding of vocal control in birdsong. *J Neurosci* **25** (3): 652–661.
- Lesser R P, Kim S H, Beyderman L, Miglioretti D L, Webber W R, Bare M, Cysyk B, Krauss G, and Gordon B, 1999: Brief bursts of pulse stimulation terminate afterdischarges caused by cortical stimulation. *Neurology* **53**: 2073–2081.
- Lilly J C, Hughes J R, Alvord E C, and Galkin T W, 1955: Brief, noninjurious electric waveform for stimulation of the brain. *Science* **121** (3144): 468–469.
- Litt B and Echauz J, 2002: Prediction of epileptic seizures. *Lancet Neurol* **1** (1): 22–30.
- Liu C, Weaver D R, Strogatz S H, and Reppert S M, 1997: Cellular construction of a circadian clock: Period determination in the suprachiasmatic nuclei. *Cell* **91**: 855–860.

- Lopes da Silva F, Blanes W, Kalitzin S N, Parra J, Suffczynski P, and Velis D N, 2003: Dynamical diseases of brain systems: Different routes to epileptic seizures. *IEEE Trans Biomed Eng* **50** (5): 540–548.
- Lüders H O, Lesser R P, Dinner D S, Morris H H, Wyllie E, and Godoy J, 1988: Localization of cortical function: New information from extra-operative monitoring in patients with epilepsy. *Epilepsia* **29**: S56–S65.
- Maeda E, Kuroda Y, Robinson H P C, and Kawana A, 1998: Modification of parallel activity elicited by propagating bursts in developing networks of rat cortical neurones. *Eur J Neurosci* **10**: 488–496.
- Maeda E, Robinson H P C, and Kawana A, 1995: The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J Neurosci* **15** (10): 6834–6845.
- Magill P J, Sharott A, Bolam J P, and Brown P, 2004: Brain state-dependency of coherent oscillatory activity in the cerebral cortex and basal ganglia of the rat. *J Neurophysiol* **92**: 2122–2136.
- Maher M P, Dvorak-Carbone H, Pine J, Wright J A, and Tai Y C, 1999a: Microstructures for studies of cultured neural networks. *Med Biol Eng Comput* **37** (1): 110–118.
- Maher M P, Pine J, Wright J, and Tai Y C, 1999b: The neurochip: a new multi-electrode device for stimulating and recording from cultured neurons. *J Neurosci Methods* **87** (1): 45–56.
- Malinow R and Malenka R C, 2002: AMPA receptor trafficking and synaptic plasticity. *Annu Rev Neurosci* **25**: 103–126.
- McAdams E T, Lacknermeier A, McLaughlin J A, and Macken D, 1995: The linear and non-linear electrical properties of the electrode–electrolyte interface. *Biosens Bioelectr* **10**: 67–74.
- McCreery D B, Agnew W F, Yuen T G H, and Bullara L, 1990: Charge density and charge per phase as cofactors in neural injury induced by electrical stimulation. *IEEE Trans Biomed Eng* **37** (10): 996–1001.

- McIntyre C C and Grill W M, 1999: Excitation of central nervous system neurons by nonuniform electric fields. *Biophys J* **76** (2): 878–888.
- McIntyre C C and Grill W M, 2001: Finite element analysis of the current-density and electric field generated by metal microelectrodes. *Ann Biomed Eng* **29** (3): 227–235.
- Meister M, Pine J, and Baylor D A, 1994: Multi-neuronal signals from the retina—acquisition and analysis. *J Neurosci Methods* **51** (1): 95–106.
- Meister M, Wong R O L, Baylor D A, and Shatz C J, 1991: Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science* **252** (5008): 939–943.
- Morefield S I, Keefer E W, Chapman K D, and Gross G W, 2000: Drug evaluations using neuronal networks cultured on microelectrode arrays. *Biosens Bioelectron* **15** (7–8): 383–396.
- Motamedi G K, Lesser R P, Miglioretti D L, Mizuno-Matsumoto Y, Gordon B, Webber W R, Jackson D C, Sepkuty J P, and Crone N E, 2002: Optimizing parameters for terminating cortical afterdischarges with pulse stimulation. *Epilepsia* **43**: 1441.
- Mukai Y, Shiina T, and Jimbo Y, 2003: Continuous monitoring of developmental activity changes in cultured cortical networks. *Electr Eng Jpn* **145** (4): 28–37.
- Murphy T H, Blatter L A, Wier W G, and Baraban J M, 1992: Spontaneous synchronous synaptic calcium transients in cultured cortical neurons. *J Neurosci* **12**: 4834–45.
- Nádasdy Z, Hirase H, Czurkó A, Csicsvari J, and Buzsáki G, 1999: Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci* **19** (21): 9497–9507.
- Nakanishi K and Kukita F, 1998: Functional synapses in synchronized bursting of neocortical neurons in culture. *Brain Res* **795** (1–2): 137–146.
- Nicolelis M A L, Ghazanfar A A, Faggin B M, Votaw S, and Oliveira L M O, 1997: Reconstructing the engram: Simultaneous, multisite, many single neuron recordings. *Neuron* **18** (4): 529–537.

- Nicolelis M A L, Ghazanfar A A, Stambaugh C R, Oliveira L M O, Laubach M, Chapin J K, Nelson R J, and Kaas J H, 1998: Simultaneous encoding of tactile information by three primate cortical areas. *Nat Neurosci* **1** (7): 621–630.
- Novak J L and Wheeler B C, 1988: Multisite hippocampal slice recording and stimulation using a 32 element microelectrode array. *J Neurosci Methods* **23** (2): 149–159.
- Obaid A L, Farries M A, Kisley M A, Sakai T, and Salzberg B M, 1996: Simultaneous optical recording of electrical activity with single cell resolution from several ganglia in a mammalian plexus. *J Gen Physiol* **108**: 27a–28a.
- Okajima Y, Tsubahara A, Kondo K, Chino N, Noda Y, and Tomita Y, 1995: A new method of estimating the distribution of muscle fiber conduction velocities. *Electromyogr Mot Control-Electroencephalogr Clin Neurophysiol* **97** (6): 310–317.
- O’Keefe D T, Lyons G M, Donnelly A E, and Byrne C A, 2001: Stimulus artifact removal using a software-based two-stage peak detection algorithm. *J Neurosci Methods* **109** (2): 137–145.
- Opitz T, De Lima A D, and Voigt T, 2002: Spontaneous development of synchronous oscillatory activity during maturation of cortical networks in vitro. *J Neurophysiol* **88** (5): 2196–2206.
- Pancrazio J J, Bey Jr P P, Loloe A, Manne S R, Chao H C, Howard L L, Gosney W M, Borkholder D A, Kovacs G T A, Manos P, Cuttino D S, and Stenger D A, 1998: Description and demonstration of a CMOS amplifier-based system with measurement and stimulation capability for bioelectrical signal transduction. *Biosens Bioelectron* **13** (9): 971–979.
- Penfield W and Jasper H, 1954: *Epilepsy and the functional anatomy of the human brain*. Little Brown and co, Boston.
- Pine J, 1980: Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J Neurosci Methods* **2** (1): 19–31.
- Poirazi P and Mel B W, 2001: Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron* **29** (3): 779–796.

- Potter S M, 2001: Distributed processing in cultured neuronal networks. In M A L Nicolelis (ed.), *Progress in Brain Research*, vol. 130, pp. 49–62. Elsevier Science.
- Potter S M and DeMarse T B, 2001: A new approach to neural cell culture for long-term studies. *J Neurosci Methods* **110** (1–2): 17–24.
- Potter S M, Wagenaar D A, and DeMarse T B, 2006: Closing the loop: Stimulation feedback systems for embodied MEA cultures. In M Taketani and M Baudry (eds.), *Advances in network electrophysiology: Using multi-electrode arrays*. Springer, New York. In press.
- Prut Y, Vaadia E, Bergman H, Haalman I, Slovin H, and Abeles M, 1998: Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *J Neurophysiol* **79** (6): 2857–2874.
- Quiroga R Q, Nadasdy Z, and Ben-Shaul Y, 2004: Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* **16** (8): 1661–1687.
- Rattay F, 1999: The basic mechanism for the electrical stimulation of the nervous system. *Neuroscience* **89** (2): 335–346.
- Regehr W G, Pine J, Cohan C S, Mischke M D, and Tank D W, 1989: Sealing cultured invertebrate neurons to embedded dish electrodes facilitates long-term stimulation and recording. *J Neurosci Methods* **30** (2): 91–106.
- Regehr W G, Pine J, and Rutledge D B, 1988: A long-term in vitro silicon-based microelectrode–neuron connection. *IEEE Trans Biomed Eng* **35** (12): 1023–1032.
- Rhoades B K, Weil J C, Kowalski J M, and Gross G W, 1996: Distribution-free graphical and statistical analysis of serial dependence in neuronal spike trains. *J Neurosci Methods* **64**: 25–37.
- Rissanen J, 1978: Modeling by shortest data description. *Automatica* **14** (5): 465–471.
- Ruaro M E, Bonifazi P, and Torre V, 2005: Toward the neurocomputer: Image processing and pattern recognition with neuronal cultures. *IEEE Trans Biomed Eng* **52** (3): 371–383.

- Rubini A and Corbet J, 2001: *Linux device drivers, 2nd ed.* O'Reilly.
- Salzman C D, Britten K H, and Newsome W T, 1990: Cortical microstimulation influences perceptual judgments of motion direction. *Nature* **346** (6280): 174–177.
- Schachter S C, 2002: Vagus nerve stimulation therapy summary: Five years after FDA approval. *Neurology* **59**: S15–S20.
- Segal M M, Baugman R W, Jones K A, and Huettnner J E, 1998: Mass cultures and microislands of neurons from postnatal rat brain. In G Banker and K Goslin (eds.), *Culturing nerve cells*, pp. 309–338. MIT Press, Cambridge.
- Segev R, Benveniste M, Hulata E, Cohen N, Palevski A, Kapon E, Shapira Y, and Ben-Jacob E, 2002: Long-term behavior of lithographically prepared in vitro neuronal networks. *Phys Rev Lett* **88** (11). Art. no.-118102.
- Segev R, Shapira Y, Benveniste M, and Ben-Jacob E, 2001: Observations and modeling of synchronized bursting in two-dimensional neural networks. *Phys Rev E* **6401** (1). Art. no.-011920.
- Shahaf G and Marom S, 2001: Learning in networks of cortical neurons. *J Neurosci* **21** (22): 8782–8788.
- Shepherd R K, Matsushima J, Millard R E, and Clark G M, 1991: Cochlear pathology following chronic electrical stimulation using non-charge balanced stimuli. *Acta Oto-Laryngol* **111** (5): 848–860.
- Shulz R, Lüders H O, Tuxhorn I, Ebner A, Holthausen H, Hoppe M, Noachtar S, Pannek H, May T, and Wolf P, 1997: Localization of epileptic auras induced on stimulation by subdural electrodes. *Epilepsia* **38**: 1321–1329.
- Sigworth F J and Neher E, 1980: Single Na⁺ channel currents observed in cultured rat muscle cells. *Nature* **287** (5781): 447–449.
- Slutzky M W, Cvitanovic P, and Mogul D J, 2003: Manipulating epileptiform bursting in the rat hippocampus using chaos control and adaptive techniques. *IEEE Trans Biomed Eng* **50** (5): 559–570.
- Staubli U and Lynch G, 1987: Stable hippocampal long-term potentiation elicited by

- theta pattern stimulation. *Brain Res* **435** (1–2): 227–234.
- Stopfer M, Jayaraman V, and Laurent G, 2003: Intensity versus identity coding in an olfactory system. *Neuron* **39** (6): 991–1004.
- Streit J, Tschertter A, Heuschke M O, and Renaud P, 2001: The generation of rhythmic activity in dissociated cultures of rat spinal cord. *Eur J Neurosci* **14** (2): 191–202.
- Tehovnik E J, 1996: Electrical stimulation of neural tissue to evoke behavioral responses. *J Neurosci Methods* **65** (1): 1–17.
- Thomas Jr C A, Springer P A, Loeb G E, Berwald-Netter Y, and Okun L M, 1972: A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exp Cell Res* **74**: 61–66.
- Tschertter A, Heuschkel M O, Renaud P, and Streit J, 2001: Spatiotemporal characterization of rhythmic activity in rat spinal cord slice cultures. *Eur J Neurosci* **14** (2): 179–190.
- Turrigiano G G, 1999: Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci* **22** (5): 221–227.
- Ueda N, Nakano R, Ghahramani Z, and Hinton G E, 2000: SMEM algorithm for mixture models. *Neural Comput* **12** (9): 2109–2128.
- Van Pelt J, Corner M A, Wolters P S, Rutten W L C, and Ramakers G J A, 2004a: Long-term stability and developmental changes in spontaneous network burst firing patterns in dissociated rat cerebral cortex cell cultures on multielectrode arrays. *Neurosci Lett* **361** (1–3): 86–89.
- Van Pelt J, Wolters P S, Corner M A, Rutten W L C, and Ramakers G J A, 2004b: Long-term characterization of firing dynamics of spontaneous bursts in cultured neural networks. *IEEE Trans Biomed Eng* **51** (11): 2051–2062.
- Van Staveren G W, Buitenvweg J R, Marani E, and Rutten W L C, 2005: The effect of training of cultured neuronal networks, can they learn? In *Proc 2nd Intl IEEE EMBS Conf on Neural Eng*, pp. 328–331.
- Velasco F, Velasco M, Velasco A L, Menez D, and Rocha L, 2001: Electrical stimulation

- for epilepsy: Stimulation of hippocampal foci. *Stereotact Funct Neurosurg* **77** (1–4): 223–227.
- Velasco M, Velasco F, Velasco A L, Boleaga B, Jimenez F, Brito F, and Marquez I, 2000: Subacute electrical stimulation of the hippocampus blocks intractable temporal lobe seizures and paroxysmal EEG activities. *Epilepsia* **41**: 158–169.
- Voigt T, Baier H, and deLima A D, 1997: Synchronization of neuronal activity promotes survival of individual rat neocortical neurons in early development. *Eur J Neurosci* **9** (5): 990–999.
- Vonck K, Boon P, Achten E, De Reuck J, and Caemaert J, 2002: Long-term amygdalo-hippocampal map stimulation for refractory temporal lobe epilepsy. *Ann Neurol* **52**: 556–565.
- Wagenaar D A, DeMarse T B, and Potter S M, 2005a: MEABench: A toolset for multi-electrode data acquisition and on-line analysis. In *Proc 2nd Intl IEEE EMBS Conf on Neural Eng*, pp. 518–521.
- Wagenaar D A, Madhavan R, Pine J, and Potter S M, 2005b: Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *J Neurosci* **25** (3): 680–688.
- Wagenaar D A, Pine J, and Potter S M, 2004: Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J Neurosci Methods* **138** (1–2): 27–37.
- Wagenaar D A and Potter S M, 2002: Real-time multi-channel stimulus artifact suppression by local curve fitting. *J Neurosci Methods* **120** (2): 113–120.
- Wagenaar D A and Potter S M, 2004: A versatile all-channel stimulator for electrode arrays, with real-time control. *J Neural Eng* **1**: 39–44.
- Wang S S H and Augustine G J, 1995: Confocal imaging and local photolysis of caged compounds—dual probes of synaptic function. *Neuron* **15** (4): 755–760.
- Wang X J, 2001: Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* **24** (8): 455–463.

- Weiland J D, Anderson D J, and Humayun M S, 2002: In vitro electrical properties for iridium oxide versus titanium nitride stimulating electrodes. *IEEE Trans Biomed Eng* **49** (12): 1574–1579.
- Weliky M, 1999: Recording and manipulating the in vivo correlational structure of neuronal activity during visual cortical development. *J Neurobiol* **41** (1): 25–32.
- Wheless J W, 2001: Vagus nerve stimulation. In E Wyllie (ed.), *The treatment of epilepsy—principles and practice*, pp. 1007–1018. Lippincott, Williams & Wilkins, Boston.
- Wilkinson C D W, 1993: Research on information-processing by neural networks cultured on substrates. *Jpn J Appl Phys Part 1 - Regul Pap Short Notes Rev Pap* **32** (12): 6210–6212.
- Wilson H R and Cowan J D, 1973: Mathematical theory of functional dynamics of cortical and thalamic nervous-tissue. *Kybernetik* **13** (2): 55–80.
- Wise K D and Angell J B, 1975: Low-capacitance multielectrode probe for use in extracellular neurophysiology. *IEEE Trans Biomed Eng* **22** (3): 212–219.
- Wong R O, Meister M, and Shatz C J, 1993: Transient period of correlated bursting activity during development of the mammalian retina. *Neuron* **11** (5): 923–938.
- Zeck G and Fromherz P, 2001: Noninvasive neuroelectronic interfacing with synaptically connected snail neurons immobilized on a semiconductor chip. *Proc Natl Acad Sci U S A* **98** (18): 10457–10462.
- Zhang L I and Poo M M, 2001: Electrical activity and development of neural circuits. *Nat Neurosci* **4**: S1207–S1214.
- Zhu G, Okada M, Murakami T, Kamata A, Kawata Y, Wada K, and Kaneko S, 2000: Dysfunction of M-channel enhances propagation of neuronal excitability in rat hippocampus monitored by multielectrode dish and microdialysis systems. *Neurosci Lett* **294** (1): 53–57.