Spatiotemporal dynamics in neocortex : quantification, analysis, models

Lyle Muller

Spatiotemporal dynamics in neocortex: Quantification, analysis, models

Dynamique spatio-temporelle du néocortex: Quantification, analyse, modèles

Lyle E. Muller II

Thèse de Doctorat en Neurosciences Computationnelles

Université Pierre et Marie Curie (UPMC)
École Doctorale Cerveau, Cognition, Comportement

Soutenue le 4 Juin 2014, avec le Jury composé de:

G. Bard Ermentrout (Rapporteur)
Gabriel Peyré (Rapporteur)
Frédéric Chavane (Examinateur)
Michelle Rudolph-Lilith (Examinateur)
Stephane Charpier (Président du Jury)
Alain Destexhe (Directeur de thèse)

Unité de Neurosciences, Information et Complexité (UNIC, UPR-3293),
Centre National de la Recherche Scientifique (CNRS),
Gif-sur-Yvette, FR

Lyle E. Muller II

Spatiotemporal dynamics in neocortex

Quantification, analysis, models

Doctoral Thesis

Unité de Neurosciences, Information et Complexité

Spring 2014

To My Parents

Abstract

This is a very complicated case, Maude. You know,
a lotta ins, lotta outs, lotta what-have-you's.
The Dude

It has only recently been acknowledged to what large extent the internal dynamics of neural networks could play a role in their function. In this respect, synaptic "noise" – that is, the influence of the cortical network on single neurons exerted through the massive recurrent circuity that is the hallmark of neocortex – has recently been shown to have a profound effect on neuronal integrative properties, changing the responses of single neurons across brain states, sometimes within the matter of a few seconds. These internally generated activity states, shaped by and continually shaping the plastic synaptic recurrent connections, then combine with the external inputs to produce a rich repertoire of responses to sensory stimuli in primary cortical regions.

In this thesis, we have focused on the *spatial* aspect of these internal dynamics, specifically the spatial structure of cortical oscillations, spontaneous and stimulus-evoked. Along the way, we have made an extensive review of the literature concerning propagating waves in thalamus and cortex, and studied network models to investigate how waves depend on network state. We have also introduced new tools for the characterization of spatiotemporal activity patterns in noisy multichannel data. The culmination of this work is a demonstration, using voltage-sensitive dye imaging data taken from the awake monkey, that the population response to a small visual stimulus propagates like a wave across a large extent of primary visual cortex during the awake state, a result contradicting a range of previous studies which seemed to suggest that propagating waves disappear in this case. Moving forward, we have begun to investigate the spatiotemporal structure of local field potential and spiking activity in multielectrode recordings taken from the human and monkey in various states of arousal, to address questions prompted by our initial voltage-sensitive dye imaging study in the monkey. In parallel, we have initiated an analysis of the extent to which neural connectivity can be characterized by the "small-world" effect, the main results of which is that neural graphs may in fact reside outside the small-world regime.

The results from these PhD studies thus span the spectrum of scales in neuroscience, from macroscopic activity patterns to microscopic connectivity profiles. It is my sincere hope to expound in these pages a unified theme for these results, and a foundation for further work in neuroscience – a search for structure within the internal architecture of the system under study.

Résumé

Il a récemment été largement reconnu que la dynamique interne des réseaux de neurones pourraient jouer un rôle essentiel dans leur fonction. À cet régard, le "bruit synaptique" – qui représente l'influence du réseau cortical sur les neurones individuels, et qui est une conséquence directe de la circuiterie récurrente massive du néocortex – a récemment été identifié comme un facteur important qui affecte les propriétés intégratives des neurones. Cette activité affecte aussi l'évolution des réponses neuronales en fonction des changements d'états du cerveau, parfois en quelques secondes. Ces états d'activité générés en interne, qui résultent – et eux-même influencent – la plasticité des connexions synaptiques récurrentes, se combinent alors avec les entrées externes pour produire un riche répertoire de réponses aux stimuli sensoriels.

Dans cette thèse, nous nous sommes concentrés sur le aspect *spatial* de ces dynamiques intrinsèques, en particulier la structure spatiale des oscillations corticales, à la fois dans le cas spontané et des réponses évoquées. Nous avons fait un examen approfondi de la littérature concernant la propagation d'ondes dans le thalamus et le cortex, et nous avons proposé un modèle de réseau neuronal pour examiner l'interaction entre les ondes de propagation et l'activité interne du réseau. Nous avons aussi mis en place de nouveaux outils pour la caractérisation de ce type d'activité spatio-temporelle à partir d'enregistrements multicanaux bruités. Le point culminant de ce travail est une démonstration, en utilisant les données d'imagerie par colorants voltage-sensitifs (VSD, "voltage-sensitive dye imaging") obtenues chez le singe éveillé, que la réponse de la population à un stimulus visuel se propage comme une onde sur une grande étendue du cortex visuel primaire. Ce résultat contredit une série d'études précédentes qui semblaient suggérer l'absence d'onde de propagation dans ce cas. Ensuite, nous avons commencé à étudier la structure spatio-temporelle du potentiel de champ local ("local field potential") obtenu à partir d'enregistrements multi-électrodes chez l'homme et le singe, dans divers états cérébraux, pour répondre aux questions suscitées par l'étude initale en imagerie VSD chez le singe. En parallèle, nous avons étudié les caractéristiques de la structure de connectivité de plusieurs systèmes nerveux, en utilisant la théorie des graphes, pour identifier les aspects aléatoires ou structurés ("small-world") de cette connectivité. Le résultat principal est que, contrairement au consensus, la structure de connectivité est beaucoup plus proche d'une connectivité aléatoire.

Les résultats de ces études de doctorat couvrent ainsi un grand spectre d'échelles en neurosciences, de modèles d'activité macroscopiques à des profils de connectivité microscopiques. J'espère sincèrement pouvoir exposer dans ces pages ces résultats de façon unifiée, dans le but de constituer une base pour la poursuite de ces travaux en neurosciences - une recherche de structure au sein de l'architecture interne du système nerveux central.

Acknowledgements

" 'Macte nova virtute, puer, sic itur ad astra,
dis genite et geniture deos...
...nec te Troia capit.' "
Virgil's Aeneid

It is a strange feeling to begin the terminal document of the PhD. After three and a half years, I feel like a wholly different person than the one who began. It is, then, with deep respect to those who have taught me during this time, and those have gotten me through, that I start.

First, and most importantly, I would like to thank my thesis supervisor **Alain Destexhe**, without whose patient mentoring and thoughtful input, these works would not have been possible. He helped me to shape the course of study in a fluid, steady fashion, making changes when necessary, urging me to stay on course through difficult problems, and finally supporting the work when preparing presentations and writing papers. At the most difficult moments during the PhD, he has remained flexible, and for that he has gained my deep gratitude, and my respect.

Secondly, I would like to thank **Claude Bédard**, whose deep knowledge of physics and mathematics are an inspiration. During my PhD studies, Claude has often given extremely well-informed, spontaneous lectures that have impressed upon me both the wonder and the history of mathematics. More than that, he has given me back an ability to play passable chess, which I had lost at some time.

I would be remiss not to acknowledge the aid and support provided by **Yves Frégnac** during this time. Our discussions on propagating waves during the summer of 2012 shaped the main work of this PhD significantly. For this alone, I am extremely grateful. Further, his support during difficult times has proven to be an invaluable resource to the development of my scientific life.

During the collaboration which is the main focus of this PhD, I have gotten the pleasure to work directly and extensively with **Frédéric Chavane**. This guy is an exemplary scientist, seeing deeply into a problem to catch faults in an argument, knowing how to craft the control that proves the point. Fredo is willing to attack the difficult problems, to find the solutions in the toughest settings. I deeply appreciate the work that we have done together, and sincerely look forward to maintaining this collaboration in the future.

And finally, but maybe also most importantly, is **Michelle Rudolph-Lilith**. Her raw talent and unwavering drive have made this the most important intellectual collaboration of my life. I have learned more from her than any one person during my entire studies, not just about mathematics or physics, but also about toughness. Together, we have started to build an idea from scratch, and to succeed in this does not come easily in science. Nevertheless, through sometimes frustrating comments, and even the doubts of many, she has persevered, staying always on the course until the job was done. Now, with a start made towards a project in graph theory, I can only hope that I have the opportunity to continue this collaboration with Michelle well into the future.

I would also like to thank everyone, at UNIC and beyond, who have provided valuable discussions and friendship during this PhD: **Andrew Davison**, **Domenico Guarino**, **Yann Zerlaut**, **Pierre Yger**, **Daniel Shulz**, **Valerie Ego-Stengel**, **Isabelle Ferezou**, **Romain Brette**, **Omar Ahmed**, **James McFarland**, **Arvind Kumar**, **Boris Gutkin**, **Gabriel Peyré**, **Adrien Peyrache**, **Giacomo Benvenuti**, and **Alexandre Reynaud**. The community around computational neuroscience at UNIC, and in Europe more generally, is a vibrant place to work, and I will always be grateful for this opportunity to **École des Neuroscience de Paris (ENP)** for support from 2009 to 2013 and **Complex-V1** (Agence Nationale de la Recherche grant to Alain Destexhe) for support from 2013 to 2014. Finally, I would like to thank the members of the jury to have kindly agreed to review the thesis. It is an honor for them to examine this work.

And lastly, but certainly not least, are the people in my life that have supported me throughout the PhD studies. I am extremely grateful to my parents **Mary Muller** and **Dr. Lyle Muller** for all the opportunities I have received in my life. Finally, I will always be grateful to **Annemarie de Zwaan**, who has been an amazing partner, companion, and source of inspiration in my life, equally during the tough times and the good.

Contents

Chapter 1

Introduction

Motivation and summary

Since the first microelectrode experiments in the primary sensory cortices (Mountcastle, 1957; Hubel and Wiesel, 1962), neurophysiology has been guided by the concept of the cortical map and the principle of hierarchical organization of brain areas. In this framework, smooth representations of stimulus spaces – whether retinotopy in primary visual cortex, somatotopy in primary somatosensory cortex, or tonotopy in primary auditory cortex – serve to map inputs from sensory structures to foci of activity in the two-dimensional plane of the neocortex. Feature detectors at each point within the cortical map in lower sensory areas extract basic stimulus properties, and pass this information on to higher cortical areas, which in turn extract more general representations to be passed on to yet higher cortical areas.

In recent years, however, research in neuroscience has taken up the hypothesis that recurrent and internally generated activity states (Destexhe et al., 2003; Destexhe, 2011) may interact with incoming sensory input to shape the cortical stimulus response (Marre et al., 2005, 2009). Structured by the interplay of single-neuron dynamics with the massive recurrent connectivity of the neocortex, these internally generated activity states may allow cortical networks to display emergent spatiotemporal patterns (Ermentrout, 1998; Ermentrout and Kleinfeld, 2001), and may explain the resurgence of large-scale patterns of activity with similarities to those repeatedly evoked by sensory stimulation (Han et al., 2008). As such, it is clear that the development of tools and techniques both in theoretical and systems neuroscience have now reached a level of sophistication that can enable the consideration of possibilities for cortical dynamics beyond the predominantly feed-forward, hierarchical model.

In the main work of this PhD thesis, we have focused on the spatiotemporal organization of activity both within and among cortical regions. Inspired by recent developments in the imaging of cortical populations, we began a line of research into the quantification and analysis of the signals resulting from these imaging techniques. This

line of research led us to an in-depth analysis of the spatiotemporal form of the population response in the primary visual cortex of the awake monkey, a question which has important consequences for whether cortical neurons participate independently in the population response, or whether the activity from nearby points in cortex affects the processing of visual stimuli already in this "low-level" visual area. Further inspired by recent experimental results that revealed nearly simultaneous responses in many auditory areas of the anesthetized rat (Reimer et al., 2011), we studied the population response in multiple visual areas and observed a novel large-scale organization of activity by simultaneous propagating waves in the primary and secondary visual areas of the awake monkey. Through both of these results, we may thus begin to envision a model of visual processing that consists of parallel, simultaneous activations at the level of single cortical areas.

In further work, we have begun to characterize the spatiotemporal structure of spontaneous, ongoing activity in multielectrode recordings of the neocortex in the human and the monkey. Here, we characterize the ongoing activity by detecting periods of high spatial organization and studying the resulting selected brief epochs of activity for consistent features. In the future, this approach could contribute to the understanding of the spatiotemporal dynamics and the importance of the ongoing irregular activity that is the hallmark of waking states.

Finally, we have recently completed a study applying concepts from graph theory to characterize the connectivity observed in brain networks. In this work, we consider the extent to which observed neural connectivity can be explained by abstract network models, such as the small-world or the scale-free graph. Though our analysis spans from the microscopic connectivity of the neural network in *C. elegans* to the macroscopic white matter digram of the monkey brain, we obtain in this work some surprisingly consistent results.

Structure and outline of the thesis

Chapters presenting published (or submitted) work begin with a brief introduction to motivate methods and results, followed directly by the reprinted work. In Chapter 4, which presents preliminary findings, however, results are presented in line with the text. Both figures and equations are numbered sequentially within each chapter. In the following paragraphs, we provide a brief summary of the individual chapters in the thesis, and a listing of the relevant presentations and publications for each.

In **Chapter 2**, we present a combined review and modeling study on propagating waves across network states. In the first half of this work, we review the experiments and models on propagating waves in the cortical, thalamic, and thalamocortical sys-

tems, to systematically examine to what extent the spatiotemporal properties of cortical activity are influenced by brain state. In the second half of this work, we study a spiking model displaying different network states, and how these internally generated activity states can affect the processing of external inputs.

Relevant Works:

- Muller LE, Destexhe A (2011) A model of propagating waves in cerebral cortex across network states. Poster presentation at CNS*2011.
- Muller L, Destexhe A (2012) Propagating waves in thalamus, cortex and the thalamocortical system: experiments and models. *J Physiology (Paris)* 106: 222-238.

In **Chapter 3**, we present the work from our collaboration with Alexandre Reynaud and Frédéric Chavane (Institut de Neurosciences de la Timone, Marseille), the main result of which is the detection and analysis of stimulus-evoked propagating waves in the primary visual cortex of the monkey during the awake states. We then go on to analyze features of the waves that are consistent across trials, and isolate the most biologically realistic propagation substrate. Finally, we detect simultaneous propagations curring in secondary visual cortex, and observe that these waves maintain precise a precise phase relationship with those occurring in primary visual cortex.

Relevant Works:

- Muller L, Reynaud A, Chavane F, and Destexhe A (2013) Propagating waves structure spatiotemporal activity in visual cortex of the awake monkey. Oral presentation at CNS*2013.
- Muller L, Reynaud A, Chavane F, and Destexhe A (2014) The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nature Communications*, in press.

Chapter 4 reports preliminary results from analysis of multielectrode array data, in collaboration with Giacomo Benvenuti and Frédéric Chavane (primary visual cortex of the monkey, INT Marseille) and Sydney Cash (middle temporal lobe of the human, Harvard University), where we study spatiotemporal structure in higher frequency oscillatory activity (20 - 50 Hz) during recordings of spontaneous local field potentials. For this, we apply the method for automated, non-parametric wave detection developed in earlier work, to detect transient epochs of propagation on the array.

Relevant Works:

- Muller LE, Benvenuti G, Chavane F, and Destexhe A (2014) Measurement of propagating waves from local field potentials and unit activity in the cortex of human and

monkey. Poster presentation at CNS*2014.

In **Chapter 5**, we make a switch from waves to networks and graph theory, in a collaboration with Michelle Rudolph-Lilith (UNIC, CNRS Gif-sur-Yvette). We present the results of a recently submitted study, which quantifies precisely to what extent neural graphs show the small-world property (Watts and Strogatz, 1998). In this work, we find that the specific null hypothesis used to assess the "small-worldness" of a network critically affects the value obtained, and that after taking the degree distribution into account, all networks studied fall into the "boderline" small-worldness regime. We then make use of analytic expressions for graph clustering and distance metrics, in order to understand this borderline regime, and find that the equivalent small-world graphs in this regime are dominated by their random component. This study thus provides the first quantitative evidence that neural graphs reside far from the optimally small-world regime.

Relevant Works:

- Muller L, Destexhe A, and Rudolph-Lilith M (2014) Neural graphs: small-worlds, after all? *New Journal of Physics*, submitted.

Finally, in **Appendix A**, we provide two articles, outside the main scope of the PhD, which resulted from work on software development for neural simulations, in collaboration with the group of Karlheinz Meier in Heidelberg, which develops neuromorphic hardware for fast analog simulation of spiking neural networks.

Relevant Works:

- Brüderle et al. (2011) A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems. *Biological Cybernetics* 104: 263-296.
- Petrovici et al. (2014) Bridging the gap between software simulation and emulation on neuromorphic hardware. In preparation.

Chapter 2

Propagating waves across network states: a review of experiments and models

The need for a comprehensive review

Multichannel recording techniques in neuroscience have just recently come of age. From multielectrode array (MEA) recordings, which are typically employed to measure both spiking activity and local field potential (LFP) in a two-dimensional cortical patch (Maynard et al., 1997), to calcium imaging, which can simultaneously target the spiking activity of many neurons in a local circuit (Grienberger and Konnerth, 2012), to voltage-sensitive dye (VSD) imaging, whose fluorescence signal is used as a measure of the average membrane potential in small groups of neurons at each of thousands of pixels (Grinvald and Hildesheim, 2004), the possibilities for recording active cortical circuits *in vivo* have matured dramatically in the last two decades, to the point where they may now be used in the cortex of the primate and human during waking states. Along the way, a large number of studies have presented findings *in vitro* and *in vivo* under various states of anesthesia, and to make sense of the vast body of literature employing multichannel recording techniques to study population dynamics, we initiated a comprehensive review of experiments and models.

For this work, our goals were twofold. First, we aimed to clarify the terminology surrounding waves in neuroscientific recordings. This clarification was necessary because a variety of terms were in use within the literature to describe the spatiotemporal dynamics observed in recordings. This plethora of terms was, however, natural because, while in some cases propagations were easy to identify by visual inspection of the recordings (see, for example, Xu et al., 2007; Lubenov and Siapas, 2009), no consensus had yet emerged on how to characterize the noisy spontaneous and stimulus-evoked population responses in settings with low signal-to-noise ratio (SNR), where detailed spatiotemporal dynamics often became difficult to distinguish. For this clarification, then, we emphasized in this work the distinction between *propagating waves* and *standing pulses* in terms of a systematic and consistent organization of response phase across cortical space. It is important to note that while such a distinction is just an application

of simple physical definitions, the key component is the emphasis on a phase-based quantification of spatiotemporal dynamics for multichannel neuroscientific recordings, which was applied through the course of this review, and later served as the basis for our own analyses of experimental data.

Secondly, we aimed to provide a comprehensive perspective on the occurrence of propagating waves across experiments, from the dissociated networks of slice preparations *in vitro*, to anesthetized states *in vivo*, and finally to pharmacologically normal waking states, by taking the known electrophysiological properties of each state into account. We focused primarily on literature concerning the mammalian cortical and thalamocortical systems, for which an assessment of the network state is possible with reasonable certainly, given the details on experimental preparation and anesthesia condition (or, in some cases, even a direct measurement of cortical electroencephalogram [EEG]); we then evaluated the activity patterns that were consistently observed across studies, and took specific note of species, stimulation, and/or anesthetic state in the cases that differed. While the differences among studies ranging from the rodent to the primate were indeed profound, several general observations did nevertheless emerge when comparing the patterns of population activity across network states. The most important points on the observed activity patterns across anesthetized and awake states, and the expectations coming from computational modeling and electrophysiological data on cortical dynamics, are briefly summarized in following sections. Before this, however, it is necessary to specify the relationship of network activity to brain state.

Neural activity across network states

From anesthesia to normal waking behavior, the intracellular activity of the neurons in the neocortex passes through a diverse array of activity regimes, each of which is characteristic of a specific internal network state, which is in turn expected to affect the dynamics induced by extrinsic input.

In the relatively synapse-poor networks *in vitro*, where much of the afferent input has been removed, neurons tend to exhibit low membrane potentials and high input resistances (Paré et al., 1998; Cruikshank et al., 2007; Fig. 2.1a). Because of this hyperpolarized and quiescent network state, the excitatory post-synaptic conductances (EPSCs) caused by external input are expected to be strong, thus allowing a local input to synchronously drive many neurons close to spiking threshold. This effect of the network quiescence can explain many observations of very clear, large-amplitude propagating events evoked by electrical or pharmacological stimulation *in vitro* (Wu et al., 1999; Metherate and Cruikshank, 1999; Sanchez-Vives and McCormick, 2000; Wu et al., 2001; see **Work 1**, section 2.1).

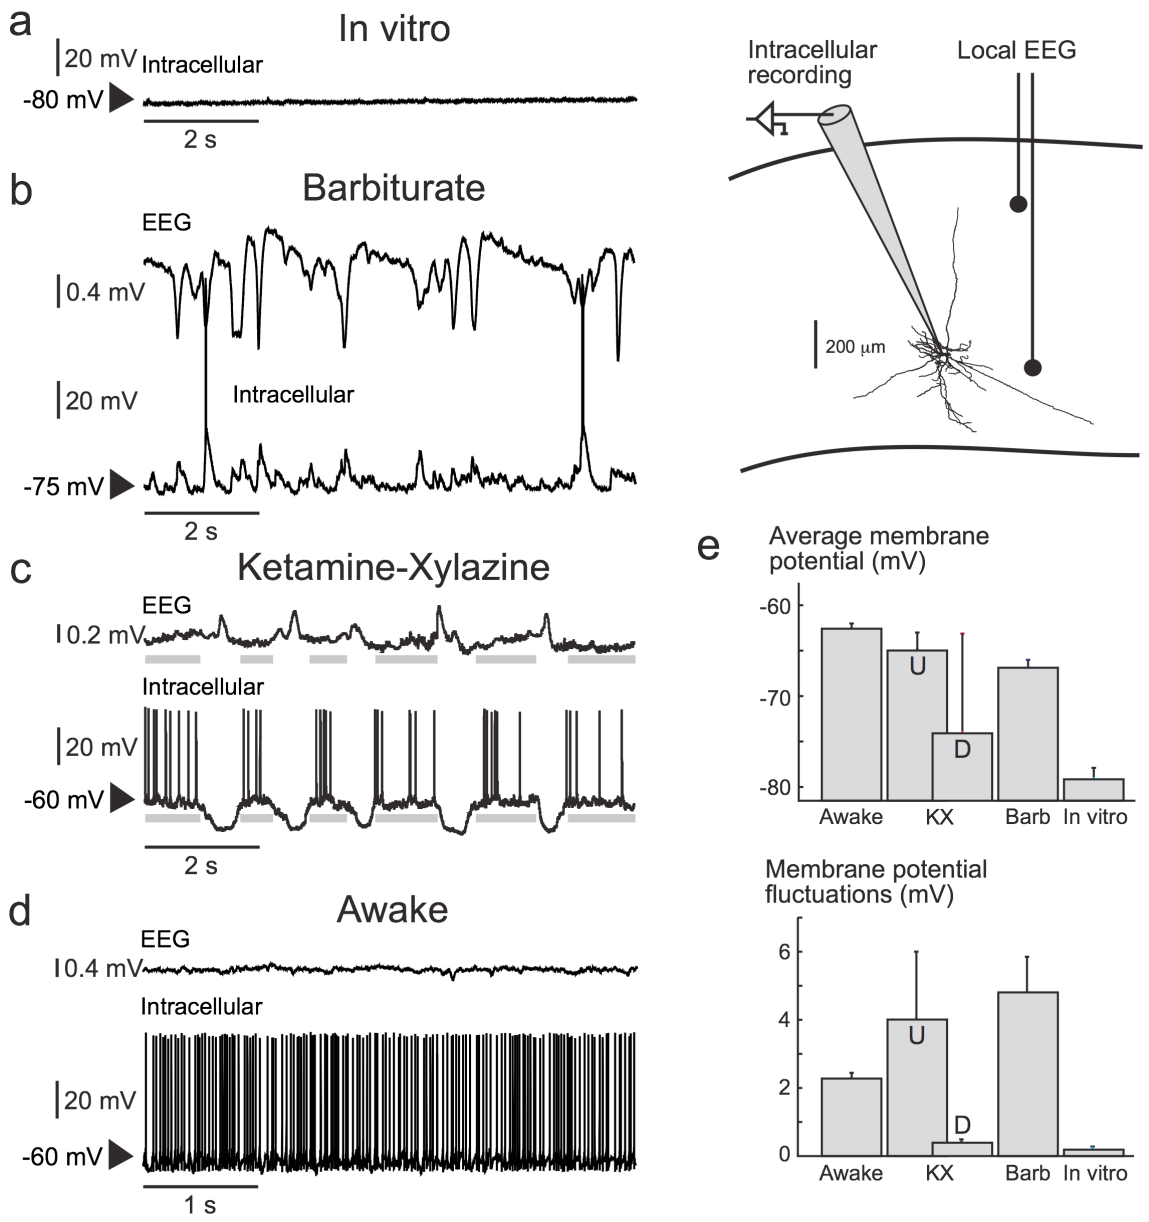

Fig. 2.1 Intracellular and EEG correlates of network state. At top right: schematic representation of simultaneous intracellular and EEG recordings. **(a)** Intracellular membrane potential illustrating the hyperpolarized, quiescent state *in vitro*. **(b)** Intracellular membrane potentials during barbituate anesthesia become slightly more depolarized, but exhibit brief excursions towards spiking threshold. These brief membrane potential excursions are often matched to the slow waves observed in the EEG. **(c)** Under ketamine-xylazine anesthesia, the network displays UP/DOWN-states, where intracellular membrane potentials alternate between periods of transient depolarization and hyperpolarization. **(d)** Desynchronized EEG and intracellular membrane potential displaying high fluctuations and spontaneous spiking. **(e)** Quantification of average membrane potential and its fluctuations across the states of activation displayed in (a-d). Modified from a summary figure presented in Destexhe and Rudolph-Lilith (2012).

Neuronal activation is only slightly higher in the case of barbituate anesthesia (which, for the purposes of comparison, can be considered a "deeply anesthetized" state), where neurons exhibit slightly higher membrane potentials and brief membrane potential excursions toward the spiking threshold (Paré et al., 1998; Fig. 2.1b). These largely quiescent states of anesthesia *in vivo* are expected to show similar large and synchronous responses to external input as cortical slices *in vitro*.

At low doses of ketamine-xylazine anesthesia, the network enters an intermediate level of activation, characterized by intermittent transitions between Up- and Down-states in the network (Fig. 2.1c). During an Up-state, neurons exhibit more depolarized membrane potentials, with the fluctuations and spontaneous discharge characteristic of the waking state. After a transition to a network Down-state, however, neurons exhibit hyperpolarized membrane potentials with reduced fluctuations, similar to deeply anesthetized or *in vitro* networks. In this intermediate, "lightly anesthetized" regime, total network activation remains reduced to some extent, and additionally resides in a somewhat synchronous regime, as reflected by the tight relationship between intracellular and EEG signals (Fig. 2.1c).

Lastly in the spectrum of network activation comes the highly irregular and continually active states expressed in awake animals (Fig. 2.1d). In the awake state, neurons sit close to threshold (Destexhe and Paré, 1999; Destexhe et al., 2003; Poulet and Petersen, 2008), and operate in a "fluctuation-driven" regime (Destexhe et al., 2001; Kuhn et al., 2004; Rudolph et al., 2007), characterized by spontaneous spiking with high irregularity. Background spontaneous spiking in the awake state has been reported to be 6.97 Hz (median) from extracellular recordings in the visual cortex of the monkey (Chen et al., 2009) and 6.1 $\pm$ 6.7 Hz (mean $\pm$ SD) in intracellular recordings of regular-spiking cells in the association cortex of the cat (28.8 $\pm$ 20.4 Hz for fast-spiking cells; Rudolph et al., 2007). By contrast, the median spontaneous background rate reported from extracellular recordings in the anesthetized monkey is about 1 Hz (Ringach et al., 2002; Xing et al., 2004; Smith and Kohn, 2008). While this difference of a few spikes per second in the spontaneous background rate may at first seem to be of little significance to the integrative properties of the network in the awake state, these values become considerable in the context of the massive recurrent circuitry in the neocortex, in which individual neurons receive between 5,000 and 60,000 synaptic inputs (Braitenberg and Schüz, 1998). At minimum, then, a difference of 6 spikes per second in the spontaneous background rate would indicate that a neuron would receive on average 25,000 more incoming spikes per second in the awake state. This large increase in the synaptic bombardment is expected to shape the spatiotemporal dynamics of the awake state considerably in comparison to the anesthetized case. In the next sections, we summarize the experimental observations concerning spontaneous and stimulus-evoked propagating activity in each state, and concerning how this activity is shaped by the network state,

finally with the aim to motivate our subsequent analysis of propagating activity in the visual cortex of the awake monkey at the single-trial level.

Propagating waves in anesthetized states

Strong evidence for propagating waves of activity in the neocortex first came from intracellular recordings *in vivo* (Bringuier et al., 1999), which revealed a progressive increase in response latency with distance from the expected point of thalamocortical input in visual cortex of the anesthetized cat. Since that time, propagating waves have been described in the anesthetized mouse (Ferezou et al., 2006, 2007; Stroh et al., 2013), rat (Xu et al., 2007; Han et al., 2008), ferret (Roland et al., 2006), cat (Benucci et al., 2007; Nauhaus et al., 2009), and monkey (Grinvald et al., 1994; Nauhaus et al., 2009). In both the spontaneous and stimulus-evoked case, these propagating waves are generally large-amplitude, single-cycle events (Ferezou et al., 2007; Xu et al., 2007; Han et al., 2008). These waves generally propagate over a large extent of the whole cortical area, and have been observed to cross over the boundaries between cortical regions (Xu et al., 2007), if not entire cortical hemispheres (Ferezou et al., 2007; Mohajernai et al., 2010; Mohajerani et al., 2013). Because a high signal-to-noise ratio can be achieved with care in these preparations, voltage sensitive dye imaging data can be in some cases presented and analyzed at the single trial level (Ferezou et al., 2007; Xu et al., 2007; Han et al., 2008).

A representative example of experimental observations in the anesthetized state is presented in Figure 2.2. In this experimental recording of bilateral activity in both hemispheres of the mouse (Fig. 2.2a), stimulation at the sensory periphery (here, a stimulus to the C2 whisker) evokes a pulse of activity (red trace, Fig. 2.2b), which then propagates across a wide extent of the system (Fig. 2.2c), next giving way to a more spatially diffuse pattern of activity. The duration of this propagating response typically lasts between 50 and 100 milliseconds. From the expectations about the somewhat suppressed activity state outlined above for anesthetized preparations, it is clear that in this situation, the incoming stimulus-driven spikes probably evoke a large and transient activation, able to propagate across the network relatively unimpeded by ongoing activity, and finally dissipating after coursing across the cortex or interacting with activity evoked at other points.

Do propagations occur during waking states?

As the somewhat suppressed state of anesthesia gives way to the highly irregular activity characteristic of the awake state, however, we may expect incoming spikes from

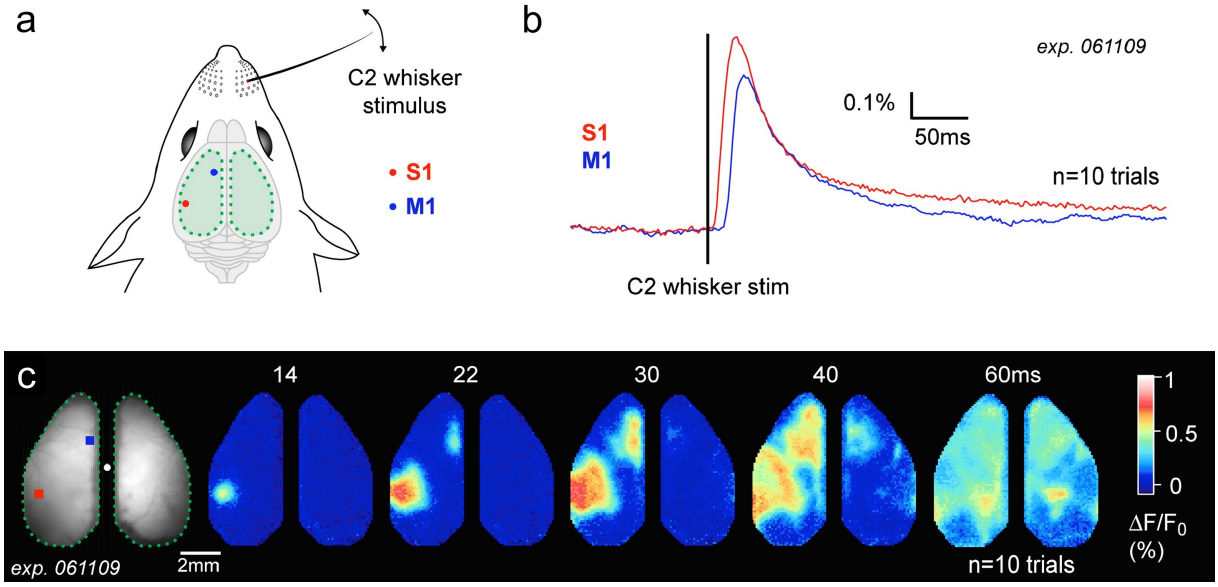

Fig. 2.2 Waves in the anesthetized rodent. (a) Schematic for VSD bilateral hemispheric recordings in the mouse. The green-shaded regions demarcates the craniotomy and imaging zone. The colored dots indicate the points in primary somatosensory cortex (S1, red) and primary motor cortex (M1, blue) whose timecourses are depicted in panel (b). **(b)** Timecourses of population response following brief, naturalistic stimulation of one whisker (C2). Note the single-cycle nature of the response, and the offset between the responses originating in S1 and M1, respectively. **(c)** Spatial pattern of population response to brief whisker deflection. At left is the resting VSD florescence with the points in S1 and M1, and a white dot indicating the location of the bregma. At 14 ms after stimulus presentation, a response begins at the thalamocortical point of input to the whisker barrel C2. Note the clear spatiotemporal organization in both S1 and M1, where activity clearly propagates out from the source at the point of thalamocortical input. Note, furthermore, that this image represents the average response over 10 trials (bottom right). Adapted from Ferezou et al. (2007).

external input be highly dissipated by this ongoing activity regime, where neurons sit close to spiking threshold, correlation coefficients of nearby neurons' membrane potentials are suppressed, and background firing rates are somewhat higher than in anesthesia (Destexhe and Paré, 1999; Destexhe et al., 2003; Poulet and Petersen, 2008). These three points indicate that the driving force for incoming EPSPs will be reduced, leading to smaller membrane potential deflections for the local network activated by a small sensory stimulus, consequently rendering the incoming spike packet unable to evoke a coherent propagation of activity amidst a sea of noise. At the same time, however, it is also possible that the desynchronized activity of the awake state does not immediately dissipate the propagation evoked by incoming sensory input, but rather confines it to a restricted spatial subregion within a single cortical area. The extent to which this dissipation of incoming thalamocortical input occurs in the awake state,

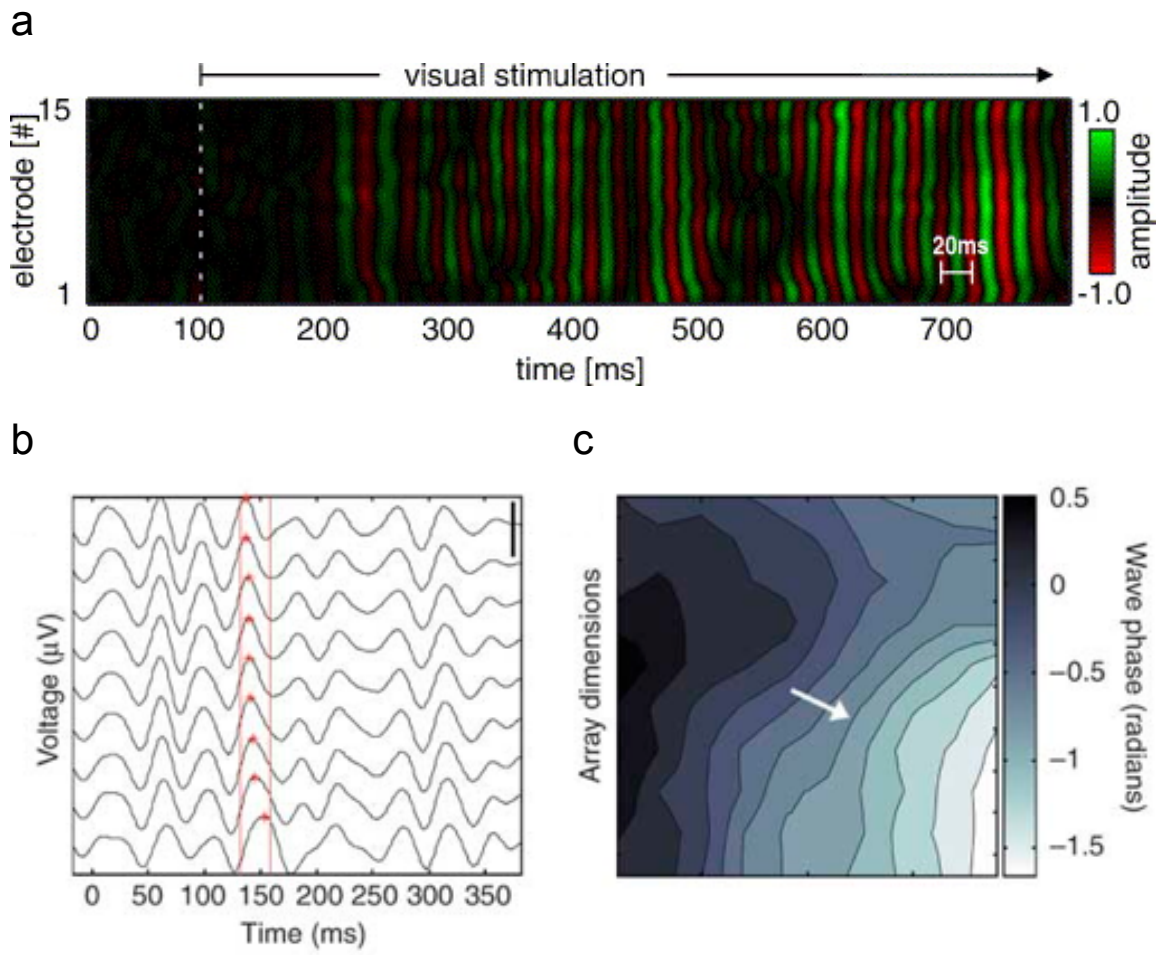

Fig. 2.3 Oscillatory waves in higher frequency bands. (a) Smooth phase organization of gamma-band (30-70 Hz) oscillations in the primary visual cortex of the awake monkey, recorded by a one-dimensional grid of electrodes following sustained visual stimulation. Normalized oscillation amplitude within the passband is presented in falsecolor; note also the 20 ms timescale at the bottom right. Adapted from Gabriel and Eckhorn (2003). **(b)** Beta-band (10-45 Hz) traveling waves observed in MEA recordings of monkey motor cortex. Vertical red lines isolate one oscillation cycle across electrodes, wherein a single phase value is marked with a red asterisk. **(c)** Isophase contours of one oscillation cycle across the MEA, with instantaneous phase plotted in grayscale (color bar, top right). Adapted from Rubino et al. (2006).

then, is the key point that determines whether or not coherent stimulus-evoked propagating activity occurs.

While indications that such coherent propagating evoked activity occurs during waking states had been observed in the awake rodent and cat (Ferezou et al., 2007; Witte et al., 2007; Mohajernai et al., 2010; Mohajerani et al., 2013), and while wave-like organization had been observed in higher-frequency oscillatory activity of awake monkey visual (Fig. 2.3a; Gabriel and Eckhorn, 2003) and motor (Fig. 2.3b,c; Rubino et al., 2006) cortex, the first VSD experiments in the primary visual cortex of the be-

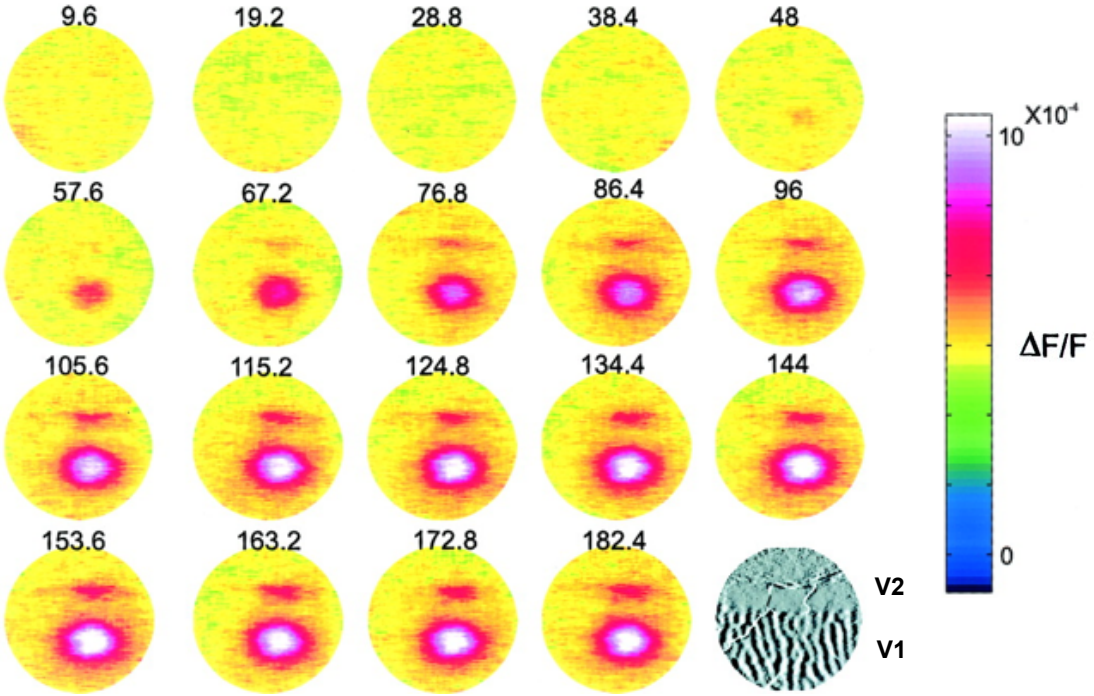

Fig. 2.4 Population response in the awake primate. At bottom right the intrinsic ocular dominance map demarcates the border between primary visual cortex (bottom region, with ocular dominance stripes) and the secondary visual areas (top region, without stripes). In the panels above, VSD activation patterns are depicted for these two areas responding to a small visual stimulus. Time in milliseconds is given above each panel, and the color bar lies to the right. Note that, in contrast to the anesthetized case, the pulse of activity following sensory stimulation appears to remain stationary at the point of thalamocortical input to the system. Note, further, that these data represent the average over 26 trials. Adapted from Slovin et al. (2002).

having monkey seemed to confirm the absence of evoked propagations in the awake state (Slovin et al., 2002; Fig. 2.4). In these VSD experiments, small visual stimuli are generally positioned close to the visual midline, such that the retinotopically evoked activation can be imaged both in the primary visual cortex (V1) and across the border in secondary visual cortex (V2), and in contrast to the large-amplitude sensory-evoked propagating events evoked by sensory stimuli in the anesthetized state (cf. Fig. 2.2), the stimulus-evoked population response appears in this case to be a stationary pulse of activity, remaining in both areas at the point of thalamocortical input to the cortex. This absence of coherent stimulus-evoked propagating activity was observed repeatedly in VSD imaging of the visual cortex in the awake, behaving monkey (Chen et al., 2006; Meirovithz et al., 2010; Ayzenshtat et al., 2010), and it became a generally accepted view that large-amplitude, sensory-evoked propagating waves were predominantly an artifact of the anesthetized state.

However, because the VSD imaging in these experiments requires averaging of the signal over many trials (e.g. over 26 trials in Fig. 2.4), as physiological artifacts are more

difficult to control in awake, behaving animals (Reynaud et al., 2011), and because propagating waves are known to be highly sensitive to trial averaging (Xu et al., 2007; Alexander et al., 2013), we came to develop the hypothesis that propagating waves in the awake state could be present on the single-trial level, on a more restricted spatial scale than under anesthesia, and were being attenuated in the trial-averaged data. To address this point, we formed a collaboration with Frédéric Chavane and Alexandre Reynaud in Marseille, who had just recently developed a denoising technique to allow for the analysis of VSD imaging data at the single trial level (Reynaud et al., 2011). In looking at the data from similar experiments as in Figure 2.4, it was immediately apparent that the population response on the single-trial level had a different spatiotemporal form than in the trial-averaged data (Supplementary Movies 1-4), and indeed, it seemed that propagations could be occurring. However, because of the noise that is inevitable in optical imaging during the awake state (either due to ongoing activity or, more simply, the highest SNR that could realistically be achieved), it was also immediately apparent that simple processing techniques such as smoothing would not be able to make a quantitative determination of whether propagations were in fact occurring in these noisy multichannel data. The development of such a quantitative method, then, and its application to these VSD data taken from the visual cortex of the awake monkey, became the main subject for this thesis work, and is presented in the following chapter.

Work 1: *Propagating waves in thalamus, cortex and the thalamocortical system: experiments and models*

French Summary

Des ondes de propagation d'activité ont été observées chez de nombreuses espèces, dans divers états cérébraux, régions du cerveau, et états d'anesthésie. Dans cet article, nous examinons la littérature expérimentale concernant les ondes de propagation en fonction des conditions expérimentales et des conditions de stimulation. Nous examinons aussi les modèles d'ondes de propagation dans les réseaux thalamiques, corticaux, et le système thalamocortical. Certaines divergences entre des résultats expérimentaux peuvent être expliquées par les differences d'état des réseaux, par exemple entre le sommeil et l'éveil. Nous proposons un modèle produisant divers états d'activité, et nous étudions l'effet des ces états sur la structure spatiale de l'activité auto-entretenue, ainsi que sur l'activité évoquée par une source externe. Cette approche représente une étape vers la compréhension de la façon dont l'activité intrinsèque du cerveau affecte la capacité à traiter et diffuser les signaux afférents.

Journal of Physiology - Paris 106 (2012) 222–238

Contents lists available at SciVerse ScienceDirect

Journal of Physiology - Paris

journal homepage: www.elsevier.com/locate/jphysparis

Review Article

Propagating waves in thalamus, cortex and the thalamocortical system: Experiments and models

Lyle Muller *, Alain Destexhe

Unité de Neurosciences, Information, et Complexité (UNIC), CNRS, Gif-sur-Yvette, France

ARTICLE INFO

Article history:
Available online 2 August 2012

Keywords:
Voltage-sensitive dye
Multi-electrode array
Population dynamics
Propagating waves
Oscillations
Sensory cortices
Spiking neural networks

ABSTRACT

Propagating waves of activity have been recorded in many species, in various brain states, brain areas, and under various stimulation conditions. Here, we review the experimental literature on propagating activity in thalamus and neocortex across various levels of anesthesia and stimulation conditions. We also review computational models of propagating waves in networks of thalamic cells, cortical cells and of the thalamocortical system. Some discrepancies between experiments can be explained by the "network state", which differs vastly between anesthetized and awake conditions. We introduce a network model displaying different states and investigate their effect on the spatial structure of self-sustained and externally driven activity. This approach is a step towards understanding how the intrinsically-generated ongoing activity of the network affects its ability to process and propagate extrinsic input.

Contents

1. Introduction

Recent years have seen an increase in measurements of large-scale spatiotemporal dynamics of neocortical networks, due to improvements in voltage-sensitive dye (VSD) (Shoham et al., 1999) and in multielectrode array (MEA) (Maynard et al., 1997) technologies. With these technological advancements, it is now generally possible to use single-trial imaging to observe the detailed dynamics of cortical circuits, whose trial-to-trial variability may preclude measurement by averaging techniques. Following preliminary evidence from electrophysiological and optical imaging studies *in vitro* (Chagnac-Amitai and Connors, 1989; Langdon and Sur, 1990; Hirsch and Gilbert, 1991; Metherate and Cruikshank, 1999; Sanchez-Vives and McCormick, 2000; Wu et al., 2001; Huang et al., 2004; Pinto et al., 2005), VSD and MEA

* Corresponding author.
E-mail address: muller@inaf.cnrs-gif.fr (L. Muller).

0928-4257/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.jphysparis.2012.06.005

experiments have provided direct observations of spatiotemporally coherent population activity in many cortical areas, anesthetic states, and stimulation conditions. For example, propagating waves have been observed *in vivo* in the visual (Grinvald et al., 1994; Kitano et al., 1994; Slovin et al., 2002; Jancke et al., 2004; Xu et al., 2007; Ahmed et al., 2008; Han et al., 2008; Nauhaus et al., 2009), somatosensory (Derdikman et al., 2003; Petersen et al., 2003; Civillico and Contreras, 2006), auditory (Reimer et al., 2010); and motor (Rubino et al., 2006) cortices under both spontaneous and evoked conditions. These diverse observations suggest that propagating waves could potentially be a general phenomenon in the large-scale dynamics of neocortex. These results, however, have been obtained under a myriad of anesthetic conditions and brain states, and a unified account of the dependence of propagating waves on network state has not yet emerged. It is therefore necessary to first determine the relationship of propagating waves to awake, activated brain states.

In the present paper, we first review the experimental literature from multichannel recording techniques, focusing specifically on propagating activity in thalamic, cortical and thalamocortical networks. In cortical networks, we also emphasize the brain state involved in each study, to assess the functional relevance of propagating waves to the awake brain. The critical factor in determining a network's responsiveness to perturbations is its conductance state (Destexhe and Paré, 1999; Destexhe et al., 2003), as the conductance state determines the average membrane potential throughout the network and the driving force on a single input, given a fixed conductance change at the synapse. Changes in global brain state, such as anesthesia or arousal (from sleep to wake), affect the spatiotemporal dynamics of cortical networks via changes in conductance state. Thus, we will analyze changes in brain state within the framework of conductance-based effects at the network level. We will focus mainly on experimental results from primary visual cortex and refer to the results from other brain areas (auditory, somatosensory, motor) for purposes of comparison. We also review both experiments and models of propagating waves in thalamic, cortical and thalamocortical networks. Finally, we present preliminary results from a computational study using network models of nonlinear adaptive exponential integrate and fire (AdEx) neurons (Brette and Gerstner, 2005). Neuronal adaptation has previously been shown to be critical for modeling the transition between UP/DOWN and AI states (Destexhe, 2009); here, we study the contributions of neuronal adaptation to low-frequency activity (1–4 Hz) and excitatory/inhibitory (E/I) interactions to high-frequency activity (20–80 Hz) as parallel factors determining the network state. Moreover, because the adaptation variable in the AdEx model has a straightforward interpretation in terms of specific membrane conductances (K^+), which are also those affected by many anesthetic drugs (Sanchez-Vives and McCormick, 2000; Franks, 2008; Destexhe, 2009), a connection among basic pharmacology, brain state, and spatiotemporal network dynamics becomes possible.

2. Propagating waves in different networks and network states

2.1. In vitro

The possibility of coherent propagating activity was first raised by VSD, MEA, and intracellular studies *in vitro*. Though some studies have used disinhibited slices to study the spatial component of epileptiform activity (Chagnac-Amitai and Connors, 1989; Huang et al., 2004; Pinto et al., 2005), many studies have observed propagating activity in pharmacologically normal slices (Langdon and Sur, 1990; Hirsch and Gilbert, 1991; Metherate and Cruikshank, 1999, Wu et al., 2001). It is well known that neurons *in vitro*, which lack a large fraction of synaptic input, have low membrane potentials and high input resistances (Cruikshank et al., 2007) compared to those measured *in vivo* (Steriade et al., 2001), similar to a neocortical DOWN state (Steriade et al., 1993; Destexhe et al., 2003). Propagating activity *in vitro* is typically initiated by means of electrical stimulation, either directly to cortical areas (Wu et al., 1999, 2001; Buonomano, 2003; Pinto et al., 2005) or to thalamocortical afferents (Metherate and Cruikshank, 1999), although in at least one study activity was stimulated using local application of glutamate (Sanchez-Vives and McCormick, 2000). While these stimuli certainly have different statistics from those induced by the sensory stimulation delivered *in vivo*, these artifically induced depolarizations may serve as a basic pulse perturbation. Because of the quiescent state of these neuronal networks, the driving force on the EPSPs evoked by these stimuli will be strong and synchronously drive many neurons close to spiking threshold. Interestingly, recent evidence *in vitro* also suggests a critical role for the infragranular layers (and interlayer interactions) in supporting the horizontal spread of activity across the cortex (Wester and Contreras, 2012).

The pharmacological dependence of propagating activity has been well-characterized by *in vitro* studies and has been localized to individual receptor classes. Several studies have shown that the non-NMDA glutamatergic ionotropic receptor antagonists CNQX[1] and DNQX[2] block horizontal propagation (Fukuda et al., 1998; Metherate and Cruikshank, 1999; Sanchez-Vives and McCormick, 2000; Wu et al., 2001; Pinto et al., 2005), specifically implicating polysynaptic fast glutamatergic transmission in sustaining propagating activity. NMDA-mediated conductances have also been shown to play a role in horizontal propagation, albeit to a lesser extent, with some studies showing a clear dependence of the generation of propagating activity on these receptors (Metherate and Cruikshank, 1999; Wu et al., 2001) and others showing only a modulatory effect (Fig. 1A) (Fukuda et al., 1998; Sanchez-Vives and McCormick, 2000). Blockage of $GABA_A$ receptors has a dramatic affect on neuronal activity, transforming normal horizontal propagation into epileptiform activity (Wu et al., 2001; Pinto et al., 2005), speeding up the propagation (from 11 ± 6 mm/s to 125 ± 24 mm/s – Fig. 1B), and focusing activity around the wavefront.

Combined with intracellular studies of horizontal axonal conduction, which falls in the range of 100–500 mm/s for unmyelinated intracortical fibers across species and cortical areas (Bringuier et al., 1999; Hirsch and Gilbert, 1991; González-Burgos et al., 2000; Murakoshi et al., 1993; Telfeian and Connors, 2003), VSD and MEA studies *in vitro* captured the first estimates of the speed of horizontal propagation of population events across the surface of the cortex (Fukuda et al., 1998). This estimation of propagation speed is in general agreement for the speeds observed *in vivo* (Grinvald et al., 1994; Jancke et al., 2004; Nauhaus et al., 2009), though some studies have reported values one order of magnitude lower in the anesthetized rat (Xu et al., 2007; Han et al., 2008) and slice preparations (Sanchez-Vives and McCormick, 2000; Wester and Contreras, 2012). While the cause for this discrepancy is unclear, possible sources include the vast differences in brain state for individual experiments, species-specific differences between rodents and other mammals, or differing techniques for measuring propagation speed (e.g. center of mass methods (Xu et al., 2007; Han et al., 2008), latency analysis (Tanifuji et al., 1994), and the offset of maximum correlation (Sanchez-Vives and McCormick, 2000)).

[1] 6-Cyano-7-nitroquinoxaline.
[2] 6,7-Dinitroquinoxaline-2,3-diome.

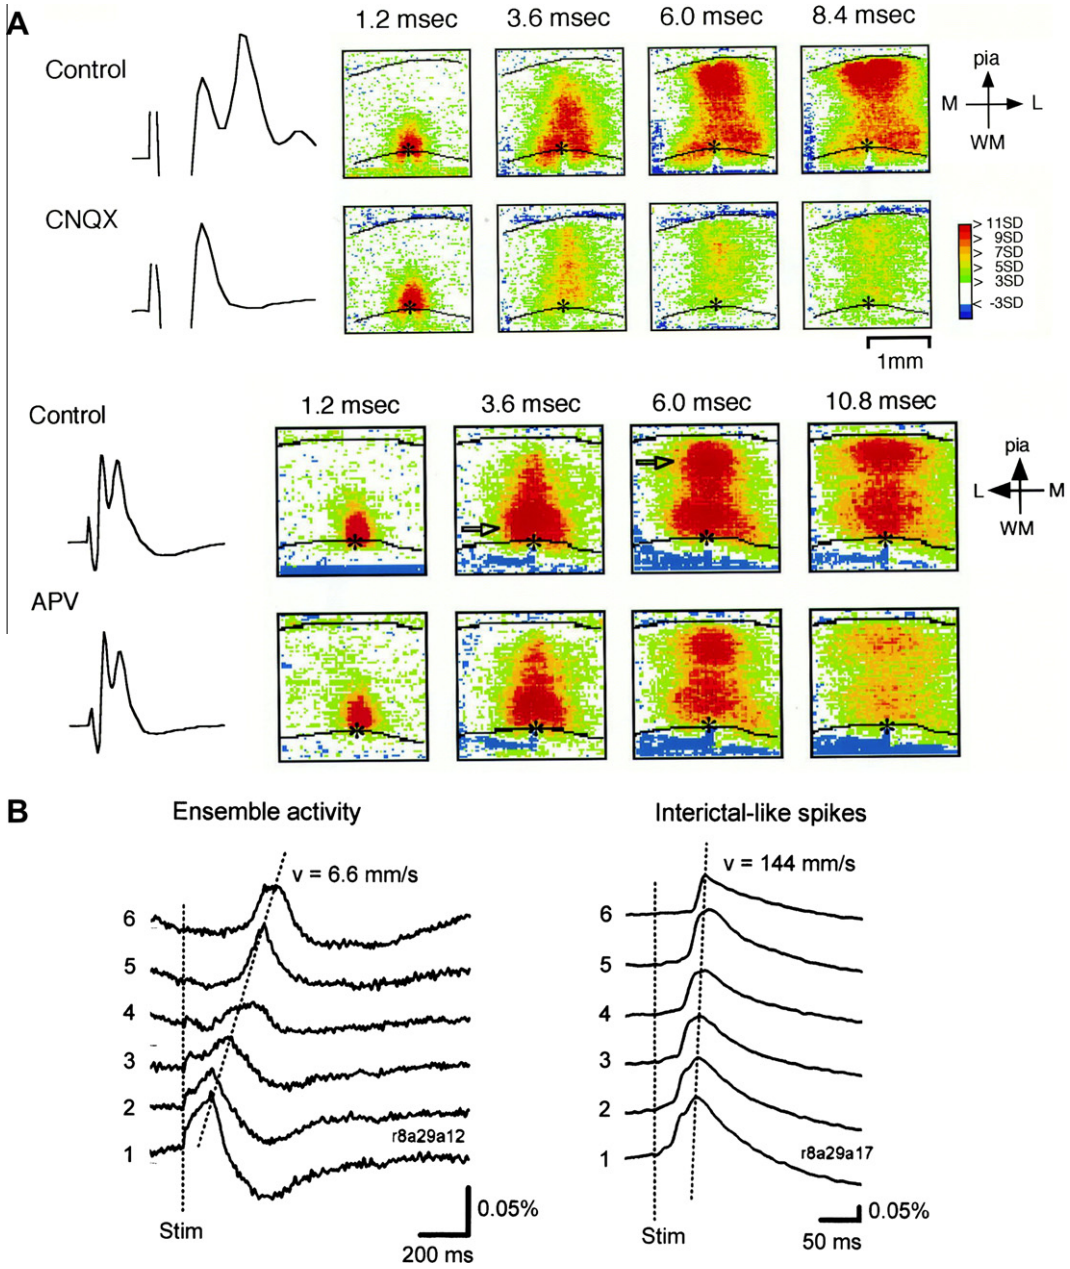

Fig. 1. Propagating activity in cortical slices. (A) Two different pharmacological conditions illustrating clear dependence of horizontal spread of activity on fast glutamatergic synaptic transmission mediated by AMPA receptors (CNQX, bath application) and partial dependence on NMDA-mediated conductances (APV, bath application). Data taken from coronal slices of rat primary visual cortex. Adapted from Fukuda et al. (1998). (B) Measurement of propagation speed with the center of mass method, across normal (ensemble activity) and pharmacologically disinhibited (interictal-like spikes) network states *in vitro*. Adapted from Wu et al. (2001).

Besides cerebral cortex, propagating waves have also been observed *in vitro* in the thalamus, in particular in slices of the visual thalamus of ferrets (Kim et al., 1995) (Fig. 2). Slices comprising the lateral geniculate nucleus (LGN) and the associated reticular sector (perigeniculate nucleus, PGN), as well as their synaptic connections, were shown to display spontaneous spindle waves (von Krosigk et al., 1993). Using multisite extracellular recordings, it was shown that all spindle waves initiate at one side of the slice (sometimes in two sites simultaneously), and propagate in the ventro-dorsal axis. Two types of waves were found, either spindle waves in control conditions, or slow oscillatory waves after blocking $GABA_A$ receptors. These two types of oscillations are compared in Fig. 2B, and both propagate, with the slow oscillation traveling at a lower speed (Kim et al., 1995).

Note that thalamic waves are different from the "monosynaptic" propagating waves of cerebral cortex, which mostly involve a single type of synapse, the excitatory interactions between pyramidal cells. In contrast, thalamic waves rely on the mutual interaction between thalamic relay and reticular nuclei (LGN and PGN here). As we will see in the modeling section, this oscillatory wave activity results from the mutual recruitment of LGN and PGN neurons, which are excitatory and inhibitory, respectively. Thus, thalamic waves involve two types of synapses, and in this sense, they are "polysynaptic" waves. Another consequence is that thalamic waves are difficult to observe because they require slices containing both nuclei and their interconnections, which explains why only one study so far succeeded to obtain this propagating activity *in vitro* (Kim et al., 1995).

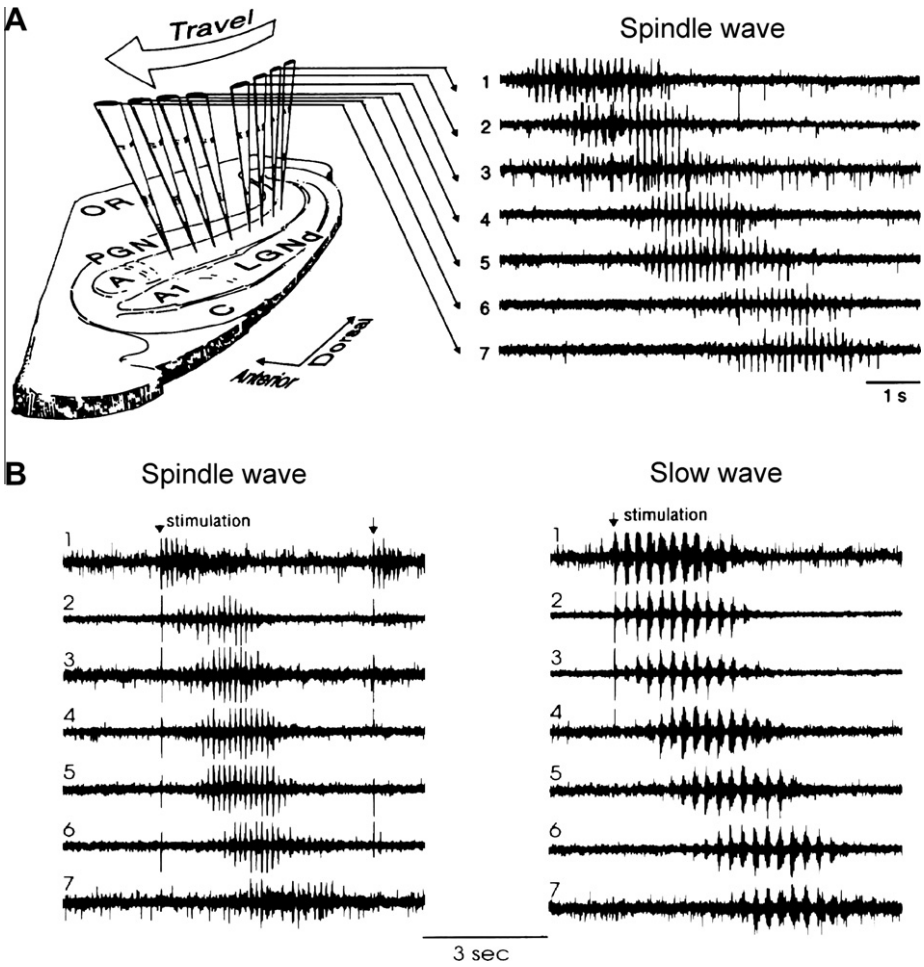

Fig. 2. Propagating waves in thalamic slices. (A) Propagating spindle waves in multisite extracellular recordings in ferret thalamic slices. Spontaneous spindle waves propagate in the ventral–dorsal axis of the slice. (B) Comparison of a spindle wave evoked by electric stimulation (left) with a slow synchronized oscillation obtained after blocking $GABA_A$ receptors (right). In the latter case, the propagation velocity was lower. Modified from Kim et al. (1995).

2.2. In vivo – anesthetized recordings

With multichannel recordings under anesthesia, large, low-frequency spreading depolarization has been observed in the visual cortex of the rat (Xu et al., 2007; Han et al., 2008),[3,4] cat (Jancke et al., 2004; Nauhaus et al., 2009)[5,6] and macaque (Grinvald et al., 1994; Nauhaus et al., 2009),[7,8] lasting hundreds of milliseconds and spreading over the entire cortical area. It is important to distinguish here between propagating waves, which are defined as oscillations propagating across the network, and spreading depolarizations, with a center of mass that remains stationary throughout the response and with a level of activation that decays monotonically with distance. While some studies under anesthetized conditions have observed large, low-frequency propagating waves of activity that spread across the cortical area (Grinvald et al., 1994; Ferezou et al., 2006; Xu et al., 2007; Han et al., 2008), other studies have observed spreading depolarizations (Derdikman et al., 2003; Jancke et al., 2004); the observation of propagating waves in the VSD signal may be dependent on the specific anesthetic level achieved in each study, with more deeply anesthetized preparations favoring large propagating waves, in the case of visual cortex. Additionally, Reimer et al. (2010)[9] have observed propagating waves in the rat auditory cortex on the timescale of 50 ms and traveling across multiple auditory fields under three different anesthetic conditions, and Kral et al. (2009) have observed propagating waves on a similar timescale in the auditory cortex of the cat.

Under anesthesia, these waves generally appear as large single-cycle events in the VSD signal (Fig. 3A) (Ferezou et al., 2006; Xu et al., 2007; Han et al., 2008; Reimer et al., 2010), display repeatable patterns across trials (Jancke et al., 2004; Han et al., 2008; Reimer et al., 2010), and have a dependence on stimulus history (Fig. 3B), both on short (Jancke et al., 2004) and long (Han et al., 2008) timescales. It is interesting to note the dependence on stimulus history in light of the NMDA-mediated conductances that have a modulatory effect on horizontal propagation mentioned above (Fukuda et al., 1998; Metherate and Cruikshank, 1999; Sanchez-Vives and McCormick, 2000; Wu et al., 2001), as a possible indication that organization of synaptic weights on the microscopic level of neuronal activity could have a clear effect on the spatiotemporal dynamics of population events on the macroscopic level of whole cortical areas. The wave generation is generally localized to the point of thalamocortical input (Fig. 3A, left panels), leading to retinotopically (Grinvald et al., 1994; Xu et al., 2007;

[3] Isoflurane.
[4] Pentobarbitol.
[5] Halothane.
[6] Ketamine-xylazine, sodium penthotal, fentanyl, N_2O.
[7] Sodium pentothal.
[8] Sufentanil, propofol.
[9] Either isoflurane, isoflurane + N_20, or ketamine.

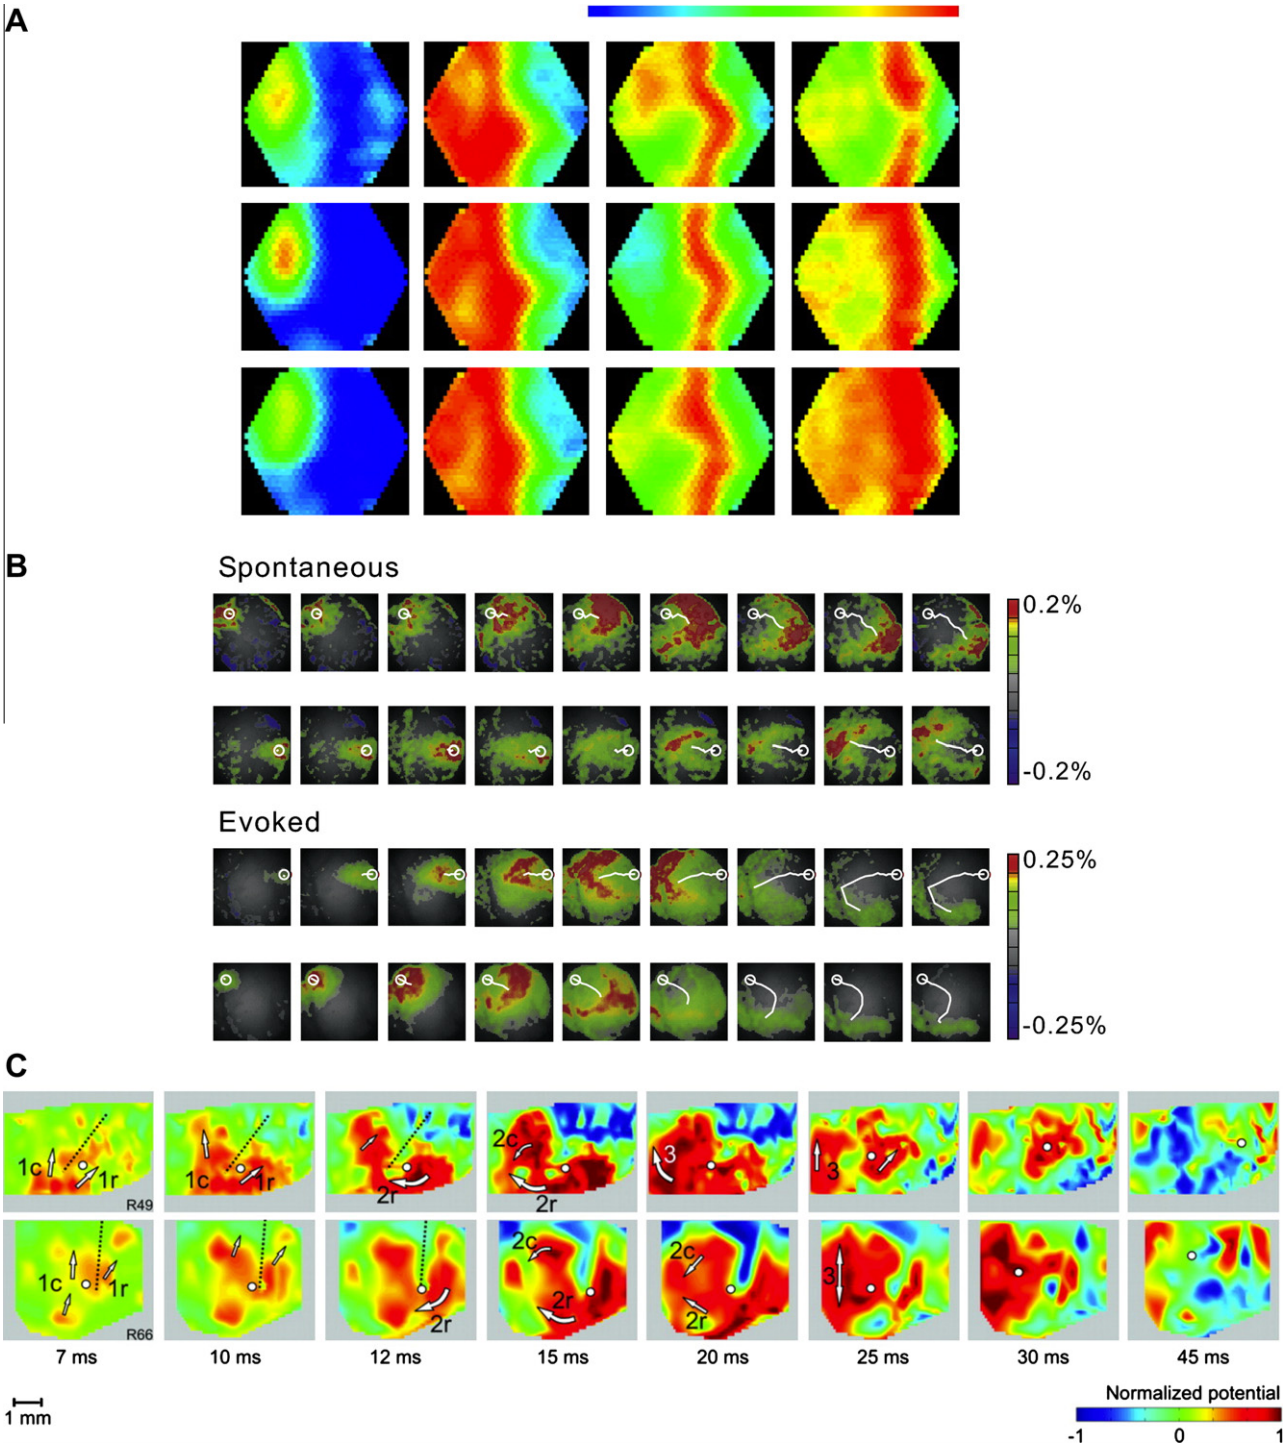

Fig. 3. Propagating waves in cerebral cortex *in vivo* under anesthesia. (A) Three VSDI runs from the same animal demonstrating trial-variability of the spatiotemporal dynamics of propagating waves. The four panels in turn depict initiation, propagation, compression, and reflection of stimulus-evoked waves in the primary visual cortex. The color bar is located above the panels. Adapted from Xu et al. (2007). (B) Spatiotemporal dynamics of propagation compared across spontaneous and evoked conditions with VSD. In this study (Han et al., 2008), spontaneous cortical waves were found to become more similar to the waves evoked by a given stimulus after repeated presentations. Color bars lie to the right of the panels. Open circles indicate the point of generation, and white lines indicate the trajectory of the center of mass during the trial. Adapted from Han et al. (2008). (C) Propagating waves observed with MEA recordings in the auditory cortex of the rat. Interestingly, in this study, onset latencies across the four participating cortical areas were found to be of similar magnitude, in contrast to the classic hierarchical sequential model of cortical processing.

Han et al., 2008; Chavane et al., 2011) or tonotopically (Harel et al., 2000; Versnel et al., 2002; Ojima et al., 2005; Song et al., 2006) organized spatiotemporal activity. Interestingly, it has been recently observed that the onset latency of the propagating waves is similar across auditory cortical areas in the anesthetized rat (Reimer et al., 2010), a result which challenges the notion of feed-forward hierarchy in cortical computations (Fig. 3C).

These waves, however, have been recorded under a myriad of anesthetic conditions, ranging from an almost totally quiescent state, to the UP/DOWN state, and – increasingly – to the desynchronized EEG of the awake state. While it is certain that brain state will have a large effect on the spatiotemporal dynamics of these networks (Destexhe and Paré, 1999; Steriade et al., 2001; Léger et al., 2005), no unified account of the change in propagating

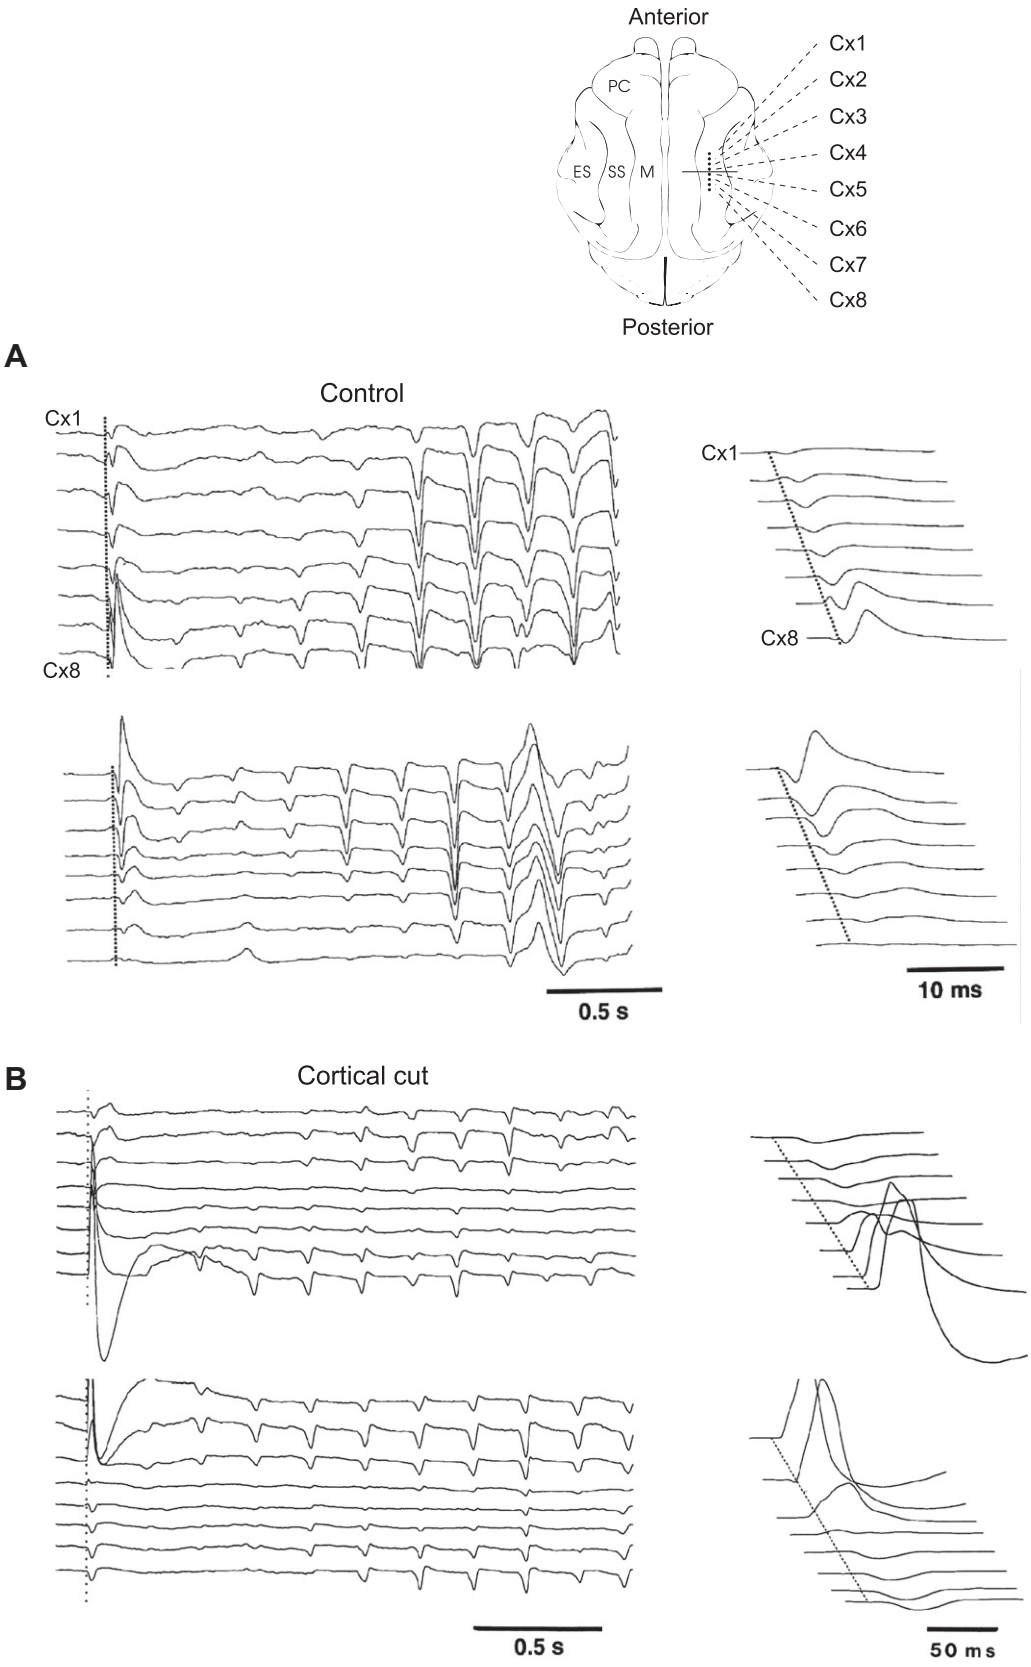

Fig. 4. Propagating waves in cat association cortex *in vivo*. Top: scheme of the recording configuration with eight extracellular electrodes inserted in suprasylvian (SS) cortex in cats anesthetized with barbiturates. (A) Electrical stimulation in the posterior (top) or anterior (bottom) side resulted in propagating spindle waves. (B) Same experiment after a deep cut in cerebral cortex (horizontal bar in top scheme). Although direct corticocortical connections were cut, the propagation still occurred. Modified from Contreras et al. (1997).

activity as a function of brain state has emerged. Above threshold levels, general anesthetics induce either UP/DOWN or totally quiescent brain states (Franks, 2008), which are characterized by increased adaptation conductances and external input that is either decreased or totally absent. In the neocortical DOWN state, membrane potentials are relatively hyperpolarized and background firing rates are low (Steriade et al., 1993; Destexhe et al., 2003; Destexhe et al., 2003; Léger et al., 2005), reflecting many features of the synapse-poor networks *in vitro*. Because the driving force on synaptic inputs during the DOWN state is high, incoming stimulus-driven spikes to the network will have a large depolarizing effect (Destexhe and Paré, 1999; Léger et al., 2005), causing a localized membrane potential deflection that can spread easily through the quiescent network, unimpeded by ongoing activity.

Finally, propagating waves have also been observed as arising from thalamocortical interactions under barbiturate anesthesia *in vivo*. In the parietal cortex of cats anesthetized with light doses of barbiturates, it was shown that electric stimulation of cerebral cortex can result in triggering spindle oscillations, with propagating properties (Contreras et al., 1997). This propagating wave is illustrated in Fig. 4 and arose from stimulation on either side of the electrode array (Fig. 4A). Interestingly, these propagating oscillatory waves still occur following a deep cut across the whole extent of the cortical gyrus (Fig. 4B). Thus, in contrast to the cortical propagating waves reviewed above, these propagating waves do not critically depend on intracortical connections, and are thought to arise from thalamocortical interactions (see Section 3.1). This emphasizes again the polysynaptic nature of propagating waves involving the thalamus, as we will further see in the modeling section.

2.3. In vivo – awake recordings

Considering the above points on stimulus integration in the neocortical DOWN state, it would be reasonable to expect that the large, low-frequency propagating waves observed during anesthesia would give way to localized bumps of activity in the awake or lightly anesthetized case, where neurons sit a few millivolts below threshold, correlation coefficients of nearby neurons' membrane potentials are suppressed, and background firing rates are somewhat higher than in anesthesia (Destexhe and Paré, 1999; Destexhe et al., 2003; Poulet and Petersen, 2008). In the activated network state, the driving force on synaptic inputs is suppressed, and incoming stimulus-driven spikes will induce smaller EPSPs, causing a smaller local membrane potential deflection when compared to similar stimulation during neocortical DOWNs. Additionally, because of ongoing desynchronized activity in the network, the synchronous depolarization of nearby neurons may be unlikely to spread across the network as a large propagating wave.

The first VSDI results in awake monkeys seem to confirm these expectations (Slovin et al., 2002; Chen et al., 2006; Meirovithz et al., 2010; Ayzenshtat et al., 2010). In these studies, visual stimuli elicit localized peaks of activity whose center of mass remains stationary throughout the response, with the level of activation decaying monotonically with distance. This peak is localized to the point of thalamocortical input, and appears at a latency of 50–100 ms after stimulus onset (Fig. 5A) (Slovin et al., 2002; Chen et al., 2006). While this stationary central peak is a clear feature repeatedly observed in the cortical population response of the awake animal, it is important to note that these first VSDI studies in awake behaving monkeys have presented only trial-averaged data, presumably because of noise constraints, and that the trial-variable spatiotemporal dynamics of propagating waves could be missed under such averaging conditions (Xu et al., 2007). It is important to note additionally that propagating waves of activity have been observed in MEA recordings of the awake cat auditory cortex (Witte et al., 2007), where neurons tend to display properties similar to the low-conductance state of the classic neocortical DOWN across anesthetized and awake conditions (DeWeese and Zador, 2006), and that preliminary evidence for the existence of traveling waves in the spontaneous background activity of the awake state has been presented recently in MEA recordings of the monkey primary visual cortex (Nauhaus et al., 2012). Furthermore, with the advent of new techniques in VSDI denoising that improve the signal-to-noise ratio on the level of single trials (Reynaud et al., 2011), it may presently become possible to detect spontaneous and stimulus-evoked propagating events in VSD recordings from the awake, behaving monkey (Muller et al., in preparation). Finally, at least one study has observed a trend contrary to the expectations for the awake state in the barrel cortex – with depolarization spreading farther and lasting longer in awake mice compared to the anesthetized state (Fig. 5C) (Ferezou et al., 2006). This study, however, also showed stationary spreading depolarization in the somatosensory cortex of the mouse under awake conditions with active perception. From the experimental observations of propagating activity and conductance state across cortical areas, we can see the large extent to which the nature of propagating waves will be shaped by network state, in a manner that depends on stimuli and specific cortical area.

Additionally, the lack of large, low-frequency propagating activity in the awake animal does not preclude the possibility of phase gradients in the high-frequency (20–80 Hz) oscillations classically observed *in vivo* (Gray and Singer, 1989). Indeed, traveling high-frequency oscillations have been observed in MEA recordings of the primary visual cortex (Gabriel and Eckhorn, 2003) and in motor cortex of the awake monkey (Rubino et al., 2006) and human (Takahashi et al., 2011) (Fig. 5B). VSD recordings under similar experimental conditions do not show these high-frequency oscillations (Slovin et al., 2002; Chen et al., 2006; Meirovithz et al., 2010; Ayzenshtat et al., 2010), probably because of the low-pass filtering inherent in the VSD signal (Wu et al., 2001) or stimulus averaging due to noise constraints (Xu et al., 2007). In contrast to the dynamics of the low-frequency propagating waves observed during anesthesia (Sanchez-Vives and McCormick, 2000; Destexhe, 2009), these high-frequency traveling oscillations are likely driven by the interaction between the excitatory and inhibitory populations in the network, as previously proposed (Brunel and Hakim, 1999; Brunel and Wang, 2003). Thus, to reconcile the observations of traveling oscillations in sensory cortices in awake behaving monkeys and the VSDI result of stationary peaks of activation, it must be possible that propagating activity exists throughout the spectrum of network activation, at low frequencies in anesthetized or sleep states and at high frequencies in activated cortical states. These high frequency propagating events may be localized to the central peak of activation, as measured by the VSD signal, or travel farther across the two-dimensional cortical sheet. Interestingly, this hypothesis of large, low-frequency propagating activity in the anesthetized case shifting to smaller, high-frequency propagating activity in the awake case is consistent with observations of increased decay of spatial correlation across the two brain states (Destexhe et al., 1999).

3. Network simulations

3.1. Thalamocortical interactions

Traveling spindle waves were simulated by computational models of networks of interconnected thalamocortical (TC) and thalamic reticular (RE) cells (respectively called LGN and PGN cells in the visual thalamus) in one dimension (see scheme in Fig. 6A), in two independent modeling studies (Destexhe et al., 1996; Golomb

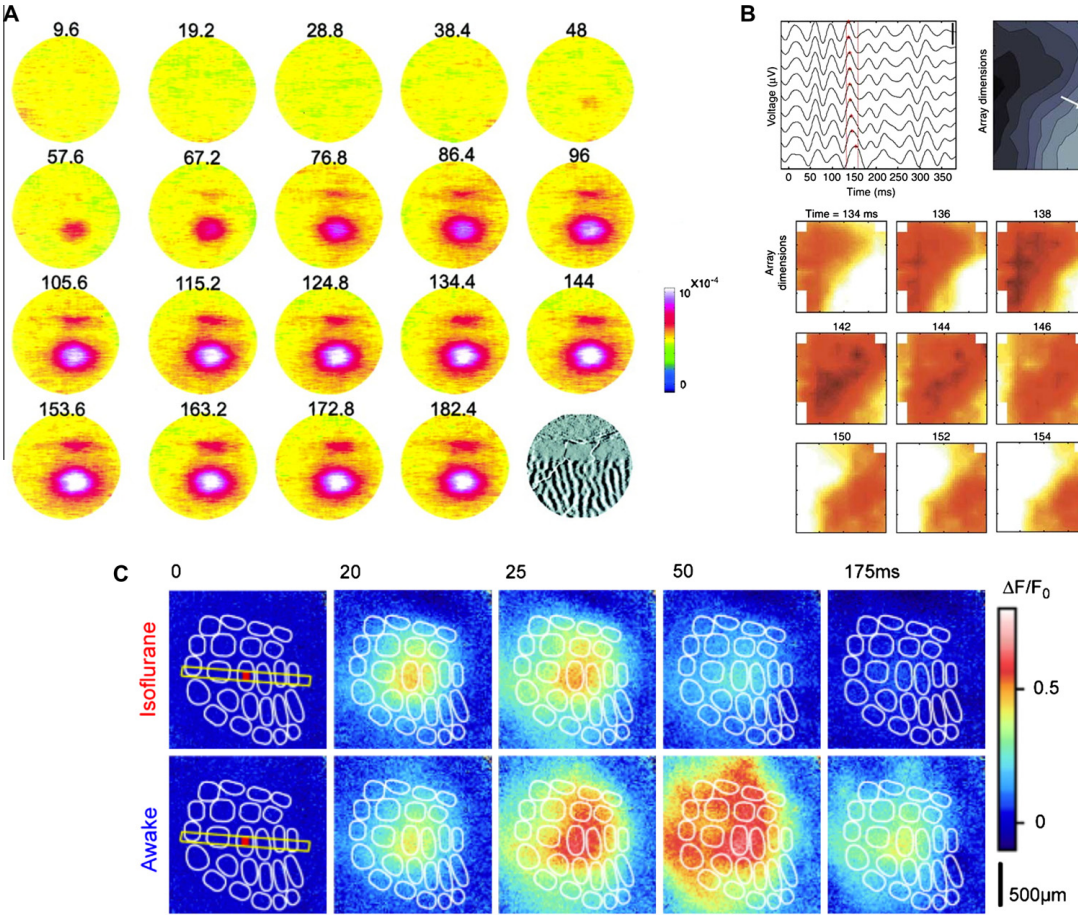

Fig. 5. Propagating waves in the cerebral cortex of awake animals. (A) VSD activation patterns in the primary visual cortex (bottom activity) and secondary visual cortex (top activity) of the macaque in response to a small retinotopic visual stimulus. The V1/V2 border is depicted in the intrinsic ocular-dominance map in the last panel. Time in milliseconds is given above each panel, and the color bar lies to the right. Adapted from Slovin et al. (2002). (B) Beta-band propagating waves observed in MEA recordings of monkey motor cortex (bottom). The top left panel depicts the phase-based method used for measuring propagation speed (not shown) and direction (indicated by white arrow, top right). Isophase contours are depicted in the top right panel, with instantaneous phase across the recording array plotted in grayscale (color bar, top right). Adapted from Rubino et al. (2006) (C) Comparison across isoflurane anesthesia and awake states in VSD recordings of the mouse somatosensory cortex, illustrating the surprising finding that propagating activity spreads farther and lasts longer in awake mice. Color bar is to the right, with spatial scale indicated under the color bar. Adapted from Ferezou et al. (2006).

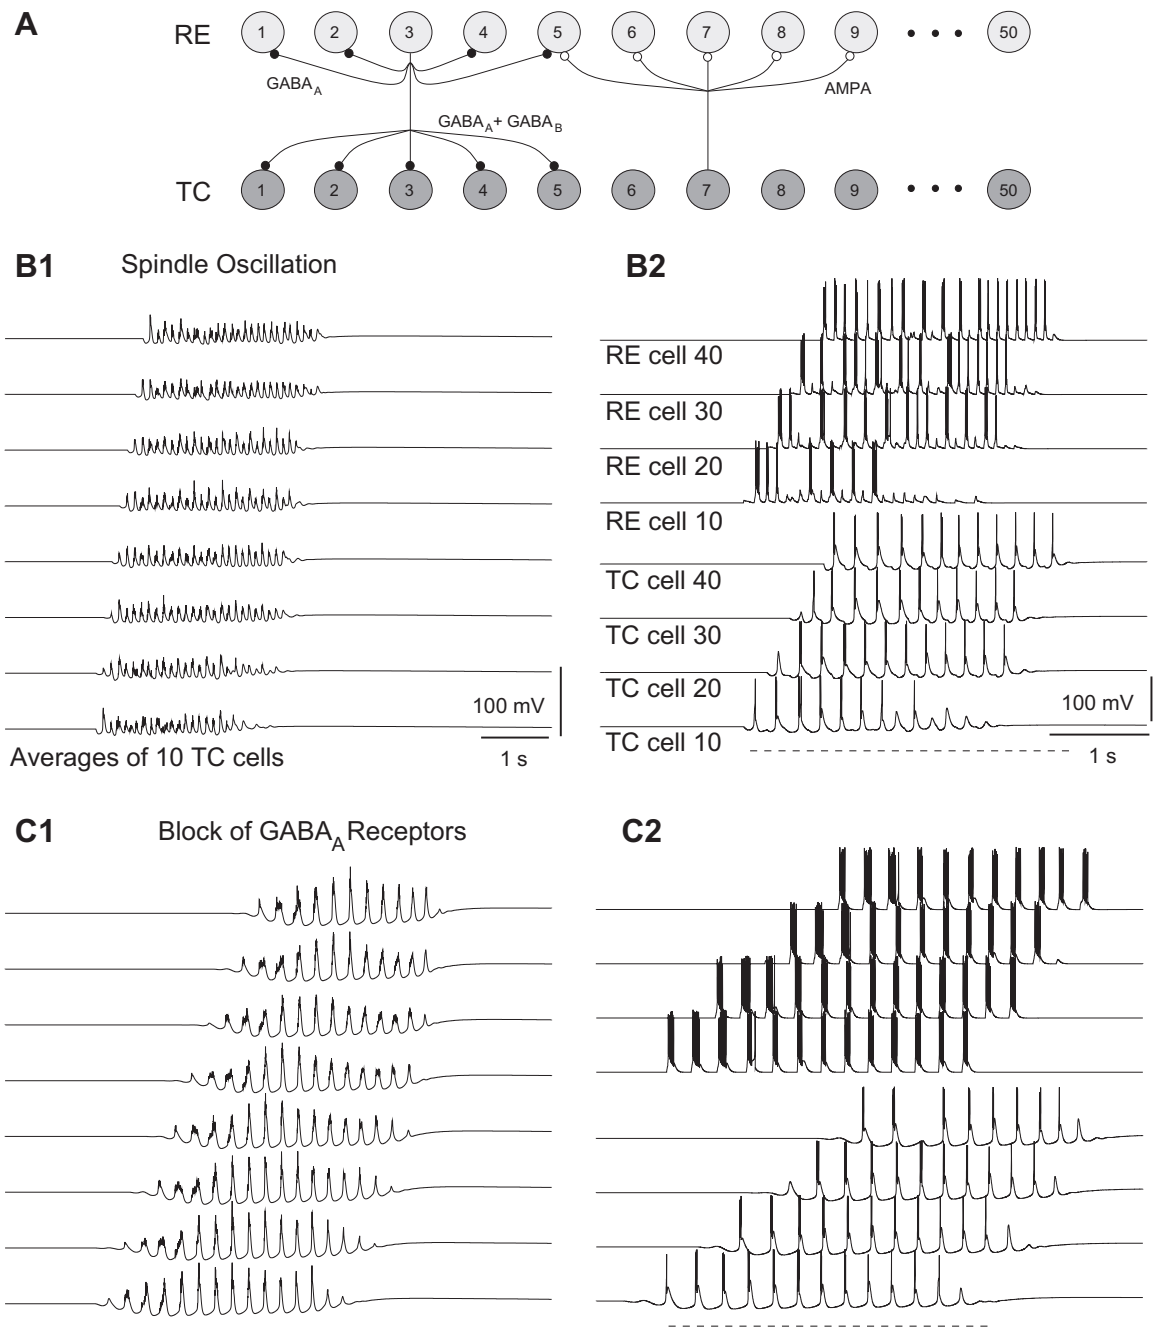

Fig. 6. Network model of propagating waves in thalamic slices. (A) Scheme of the network model, with two types of neurons, thalamocortical (TC) and reticular (RE), interconnected as indicated. (B) Spindle oscillations occurring as an interaction between TC and RE cells. Due to the local connectivity (schematized in A), the oscillation propagated. (C) Slow oscillation following block of $GABA_A$ receptors. The oscillation had a slower frequency and propagated with a lower velocity. Modified from Destexhe et al. (1996).

et al., 1996). These thalamic networks used Hodgkin-Huxley type models of TC and RE neurons, and assumed that there was a topographic connectivity between TC and RE layers, consistent with anatomical data (see details in Destexhe et al. (1996)). Under these conditions, the models generated traveling waves consistent with *in vitro* data (Fig. 6B). Suppression of $GABA_A$ receptors led to slow oscillations around 3 Hz, which propagated at a lower velocity (Fig. 6C), as also observed experimentally.

These models emphasize the fact that the propagating waves observed in thalamic slices (Kim et al., 1995) are the result of polysynaptic interactions involving TC and RE cells. In these models, the spindle oscillation arises from the mutual interaction between TC and RE cells, and the propagating activity involves the simultaneous recruitment of entire populations of neurons, and has therefore a discontinuous character, which gives the appearance of a "lurching wave" (Destexhe et al., 1996; Golomb et al., 1996; Rinzel et al., 1998). This type of wave contrasts with the propagating waves observed in cerebral cortex, which are nearly continuous because they are based on interactions within a relatively homogeneous population.

To investigate propagating waves in the thalamocortical system, a thalamocortical network model was developed by combining the previous model of thalamic slices with a model of deep cortical layers with pyramidal (PY) cells and interneurons (IN)

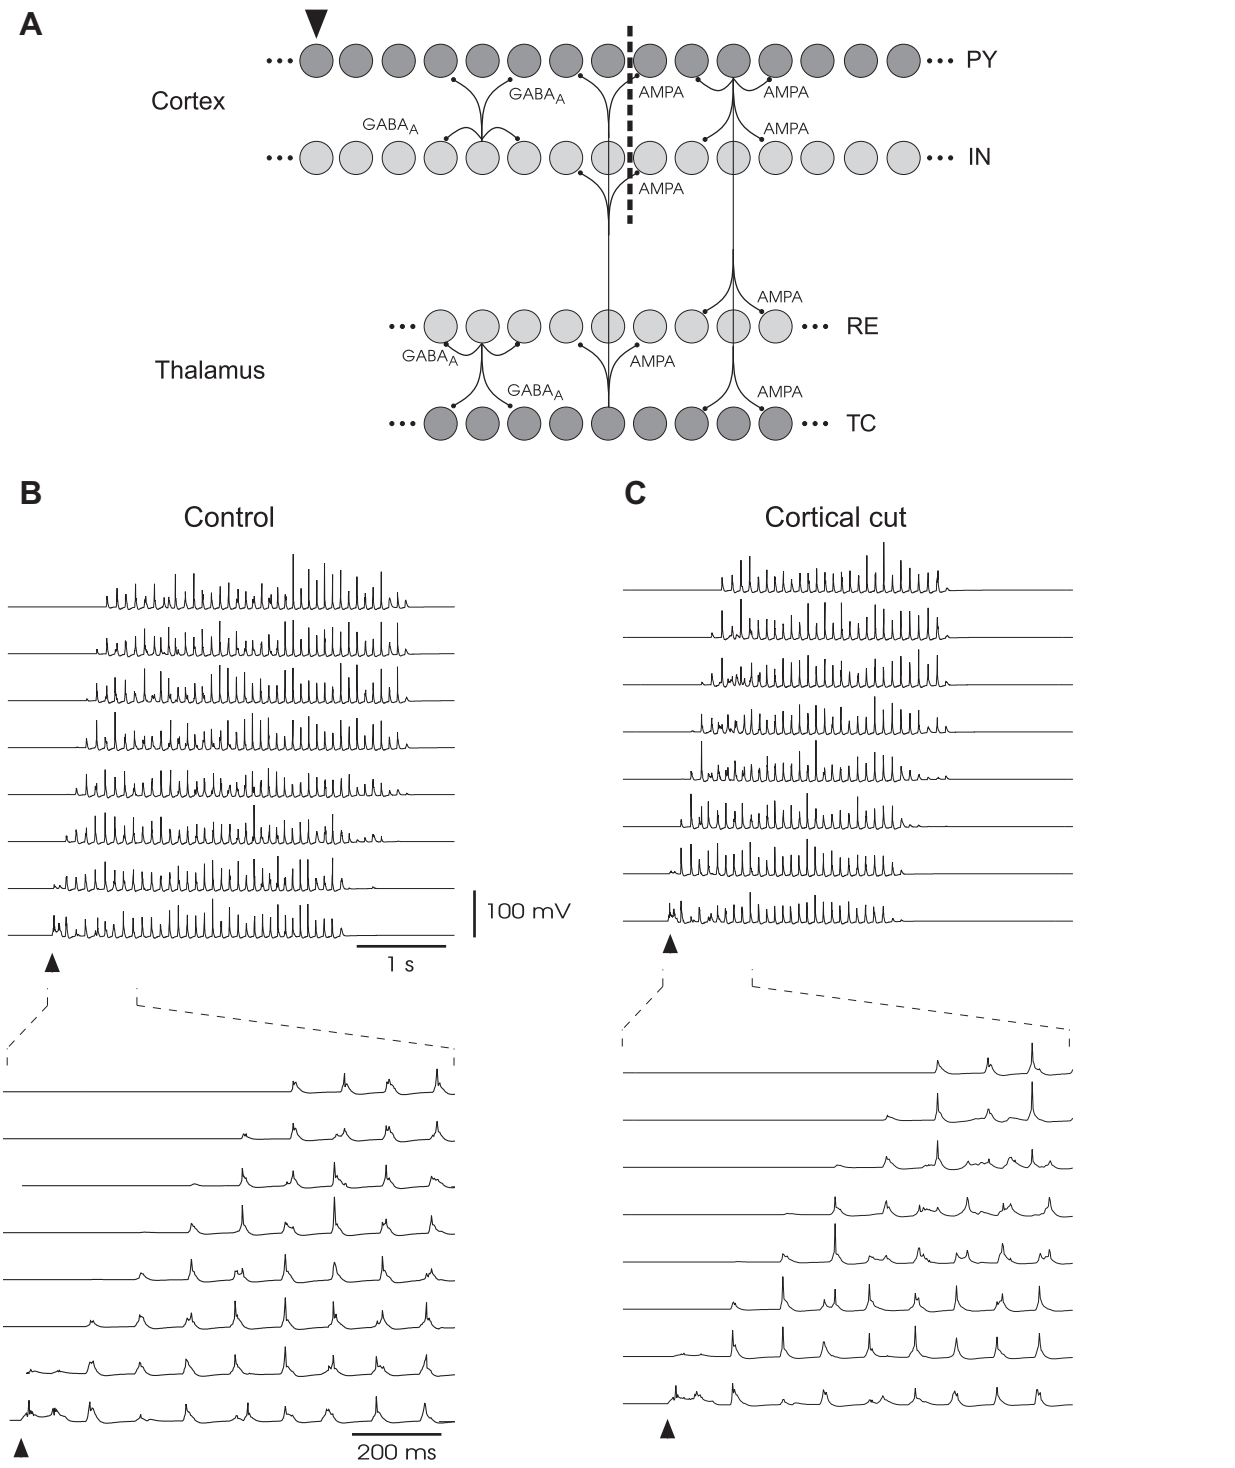

Fig. 7. Network model of propagating waves in the thalamocortical system. (A) Scheme of the network model, with four layer of neurons, cortical pyramidal (PY) cells and interneurons (IN), and thalamic thalamocortical (TC) and reticular (RE) cells. The connectivity between layers was topographic and mediated by glutamate AMPA receptors and $GABA_A$ receptors, as indicated. (B) Control spindle oscillation in the network, and its propagating properties at eight equidistant sites (magnification at the bottom). (C) Same simulation after cutting cortico-cortical connections (dotted line in A). As in experiments, the propagation resisted cutting cortical connections, and involved thalamo-cortical loops. Modified from Destexhe et al. (1998).

(Fig. 7A; see details in Destexhe et al. (1998)). This model could reproduce all experiments on propagating activity, under the assumption that the cortex recruited TC cells primarily through inhibition. This "inhibitory dominant" cortical feedback to the thalamus is consistently observed experimentally, and can explain propagation through thalamocortical loops (Destexhe et al., 1998). In these conditions, the model was capable of reproducing the experiments on evoked traveling waves *in vivo*, as well as their resistance to cutting corticocortical connections (Fig. 7B and C). The same model can also explain the genesis of absence-type of epileptic seizures when the excitability of cortical neurons is too high (Destexhe, 1998). Consistent with this model, propagating activity has indeed been observed in the thalamus of decorticated cats *in vivo* (Contreras et al., 1996).

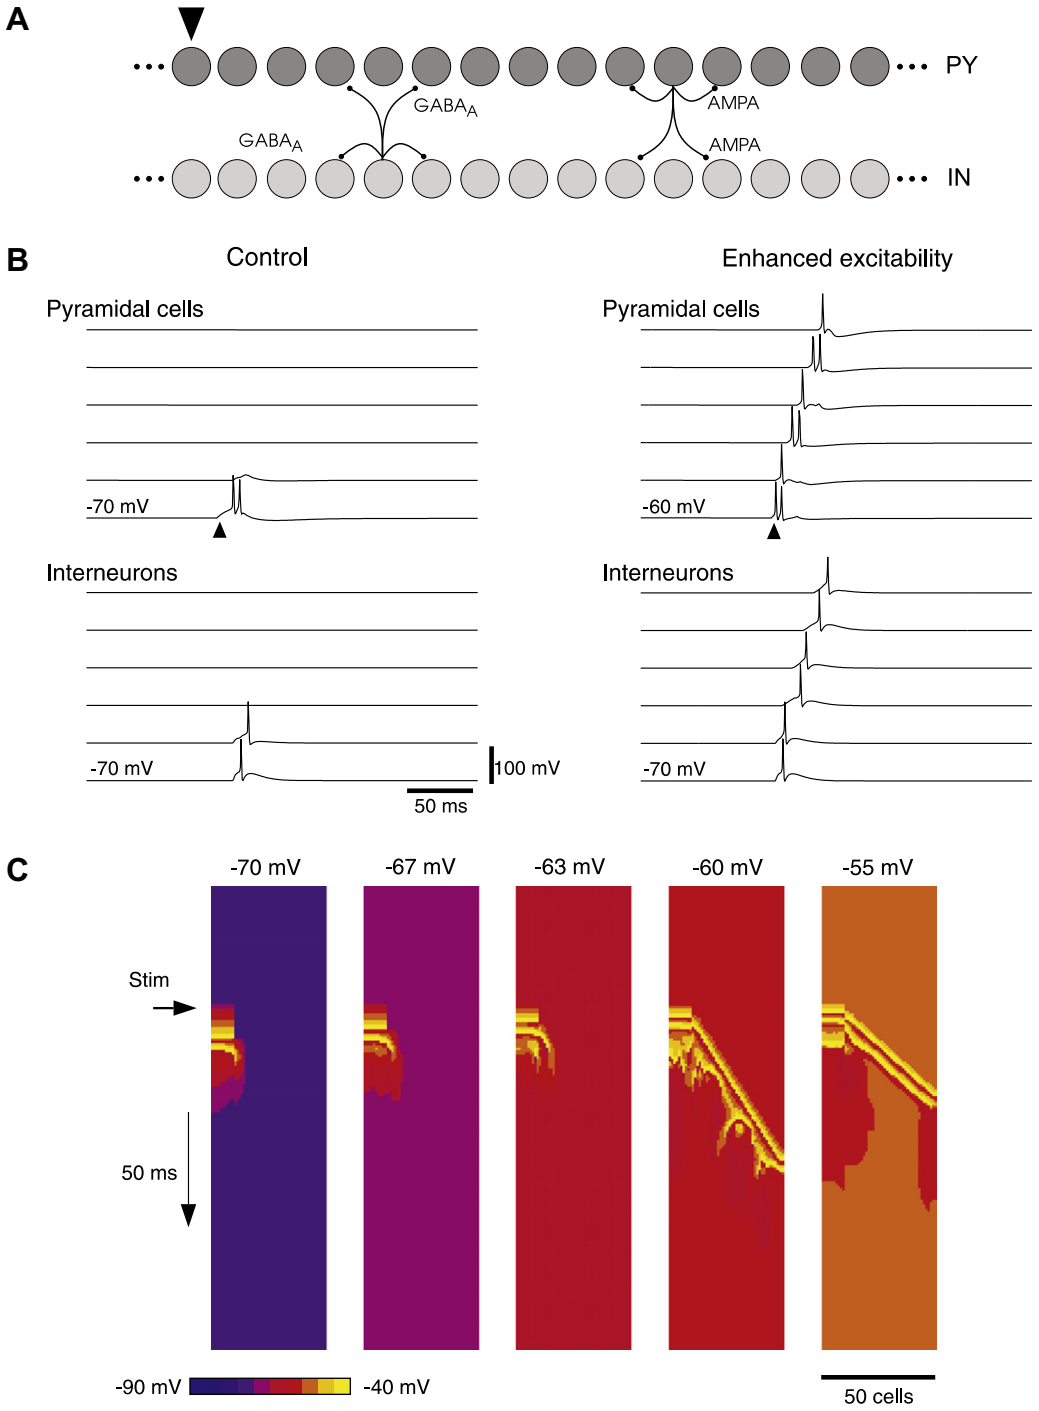

Fig. 8. Network model of propagating waves in cerebral cortex. (A) Scheme of the 1-dimensional network of 50 PY and 50 IN cells, with local axon collateral connections mediated by glutamate AMPA receptors and $GABA_A$ receptors as indicated. (B) Examples of cell activities in the network. In control conditions (left), the resting V_m was of −70 mV and the network did not sustain propagation. With enhanced excitability (right), the more depolarized V_m at −60 mV for PY cells led the network to sustain propagation. (C) Snapshots of activity showing the dependence of propagating activity on the V_m level of PY cells. Modified from Destexhe et al. (1999).

This again emphasizes the polysynaptic aspects of thalamic and thalamocortical propagating waves. In the thalamocortical system, the successive recruitment of larger portions of cortical tissue during the wave depends on several successive synapses:

PY → RE → TC → PY.

The added axonal divergence of each projection is responsible for a marked "lurching" character of these propagating waves, as in the case of the thalamus.

To show the prominent role of cortical excitability, a one-dimensional model was constructed, with PY and IN cells identical to the thalamocortical model above, and with topographical connectivity (Destexhe et al., 1999) (Fig. 8A). This simple cortical model generated monosynaptic propagating waves, in a manner that was critically dependent on the resting membrane potential. For a given connection strength, a propagating wave occurred only at some level (−60 mV in Fig. 8B and C), and the velocity of the wave was also dependent on the resting level of PY cells.

Table 1
Model parameters.

Simulation parameters		
Parameter class	Parameter	Value (units)
Cell	g_L	16.67 nS
	V_r	−70.0 mV
	V_{reset}	−70.0 mV
	V_{thresh}	−50.0 mV
	T_{refrac}	5.0 ms
	Δ_T	2.5 mV
	τ_w	600.0 ms
	C_m	250.0 pF
	a	1.0 ns
	b (exc)	5.0/155.0 pA
	b (inh)	0.0 nA
Synaptic	E_e	0.0 mV
	E_i	−80.0 mV
	τ_e	2.0/5.0 ms
	τ_i	5.0 ms
Network	L	1.0/2.0/4.0 mm
	v_c	0.2 m/s
	K	400
	K_{exc}	320
	K_{inh}	80
	σ_c	0.2 mm
	ν_{ext}	0–25 Hz
ODE	dt	0.1 ms
	T	2–10 s

In contrast to the "lurching" waves of the thalamus and thalamocortical system seen above, these cortical waves were continuous and were highly dependent on proximity of the mean V_m to spike threshold. Thus, we can infer from these simulations that the genesis of such monosynaptic propagating waves by cortical networks will be sensitive to the excitability of cortical cells, which could be affected by neuromodulatory factors. Propagating wave activity will also be sensitive to any factor that changes the mean V_m of the neurons, such as for example the general "state" of activity of cerebral cortex (Destexhe et al., 2003), as investigated in the next section.

3.2. Corticocortical interactions

To investigate the dependence of propagating waves on "network state", the simulations shown above must be augmented with the ability to display spontaneous activity, to match the high levels of spontaneous activity observed *in vivo* (reviewed in Destexhe et al. (2003)). To address this point, we considered network models displaying more realistic activity states, where input to the system is fundamentally shaped by background activity. For this purpose, we considered a small-scale topographic network ($N = 8000$, with $N_e = 6400$ and $N_i = 1600$). Individual neurons followed the bidimensional adaptive exponential integrate-and-fire (AdEx) model (Brette and Gerstner, 2005), described by the equations:

$$C_m \frac{dV}{dt} = -g_L(V - V_r) + g_L \Delta_T e^{\frac{V - V_T}{\Delta_T}} - w + g_e(t)(E_e - V) + g_i(t)(E_i - V) \quad (1)$$

$$\tau_w \frac{dw}{dt} = a(V - V_r) - w \quad (2)$$

where C_m is the membrane capacitance, g_L is the leak conductance, V_r is the resting membrane potential, Δ_T is the slope factor, V_T is the threshold voltage, g_e and g_i are the time-dependent synaptic conductances, E_e and E_i are the reversal potentials for excitatory and inhibitory synaptic transmission, and w is the adaptation variable. The level of adaptation in this model is specified by parameters a (subthreshold) and b (spike frequency adaptation). At a membrane potential high above threshold, where the exponential nonlinearity dominates, the membrane potential is reset to V_r, and the adaptation variable w is incremented by the spike frequency adaptation parameter b. To ensure against residual synchronization in the network, small Gaussian heterogeneities were incorporated in the C_m, g_L, and V_T parameters. Membrane potentials were randomly initialized from a uniform distribution between V_r and V_T. Additional simulations with randomly initialized w variables in the pyramidal neurons were run to ensure that the results reported were not the result of initial transients. Glutamatergic fast synaptic transmission via AMPA receptors plays a central role in the generation of horizontal propagation *in vitro*, while NMDA receptors have only a modulatory effect.[10] Thus, we chose to consider only fast synaptic conductances, with time constants in the range of 2–5 ms. Conductance changes at individual synapses follow an alpha-function. Parameter values for the AdEx neurons were adapted from Destexhe (2009). A listing of parameter values not explicitly discussed in the text is given in Table 1. All simulations were performed using the PyNN interface (Davison et al., 2008) to the NEST simulation environment (Gewaltig and Diesmann, 2007).

In order to study the spatial aspects of activity propagation, we embedded the network on a one-dimensional ring, with distance-dependent connection probabilities and delays as in previous work (Yger et al., 2011). Connection probabilities between neurons decay as a Gaussian with distance (ℓ_{ij}), following the spatial profile:

$$p_{ij} = e^{-\frac{\ell_{ij}^2}{2\sigma_c^2}} \quad (3)$$

where σ_c is the standard deviation – in this context, the spatial spread – of the Gaussian profile (Mehring et al., 2003). The axonal conduction delay between two neurons (d_{ij}) grows linearly with distance, following the relation:

$$d_{ij} = d_{syn} + \frac{\ell_{ij}}{v} \quad (4)$$

where d_{syn} (0.3 ms) is the minimum synaptic delay, and v is the axonal conduction velocity (0.2 m/s) (Bringuier et al., 1999). Each neuron receives $K_e = 320$ excitatory and $K_i = 80$ inhibitory recurrent synapses from the network drawn from the Gaussian spatial profile, giving 5% connectivity in the case of 8000 neurons. Additionally, we consider that the neuron receives an equivalent number of excitatory synaptic inputs from outside the network (Morrison et al., 2007; Kumar et al., 2008), each firing with a rate ranging from 0 to 25 Hz.

To characterize the results of the network simulations, we considered the mean firing rate, the mean pairwise cross-correlation, and the peak frequency of population oscillations. Mean firing rate is computed as the mean spike count over neurons during the simulation interval; however, because of heterogeneity in the network, some neurons are silent and do not spike during this interval, and these are not included in the mean firing rate estimate (Kumar et al., 2008). Mean pairwise cross-correlation is computed as the normalized zero-lag correlation coefficient averaged across 500 randomly chosen neuron pairs (Aertsen et al., 1989). To characterize the spatiotemporal patterns of subthreshold activity, we recorded the mean membrane potential ($\langle V_m \rangle$) across the network in an array of 10 detectors at intervals of 1 ms.

It has been previously shown that the adaptation present in this neuron model can facilitate the transition from activated, asynchronous irregular (AI) states to UP/DOWNs (Destexhe, 2009); here, we have extended this result to understand the spatial organization of high and low frequency activity across network states. To illustrate the dependence of the spatial structure of activity on the intensity of external Poisson input, we studied these

[10] see Section 2.1 – *in vitro*.

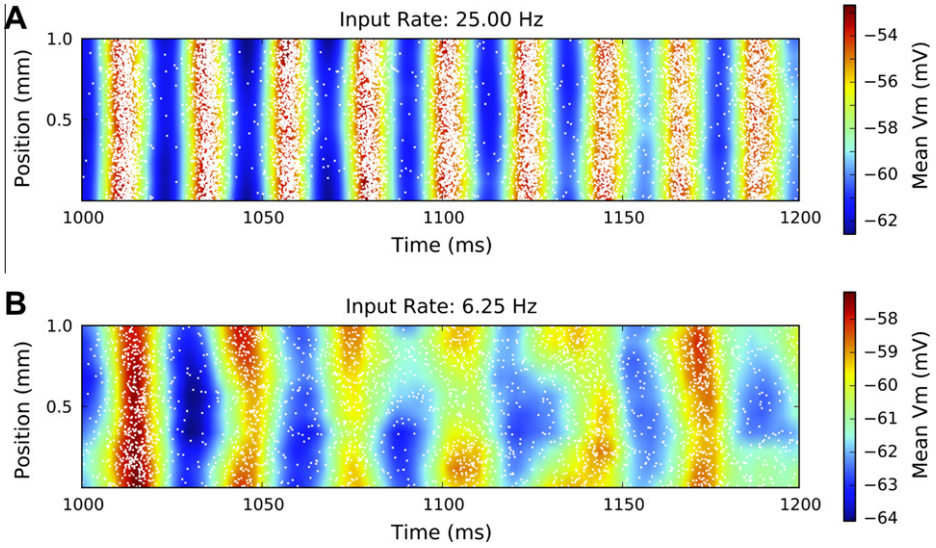

Fig. 9. Short-timescale spatial structure of network population activity driven by external input. In both panels, mean membrane potential in 0.1 mm bins is plotted in false color (color bars at right), and single-unit spikes are plotted in white. (A) Coherent population oscillations. With high external input (individual synapses firing at 25 Hz) and low neuronal adaptation (b – 0.005 nA), the network displays highly synchronized oscillations, with irregular dynamics at the single-neuron level. (B) Traveling oscillations. With lower external input (6.25 Hz), both oscillation frequency and coherence are decreased, revealing clear phase gradients on the level of individual cycles.

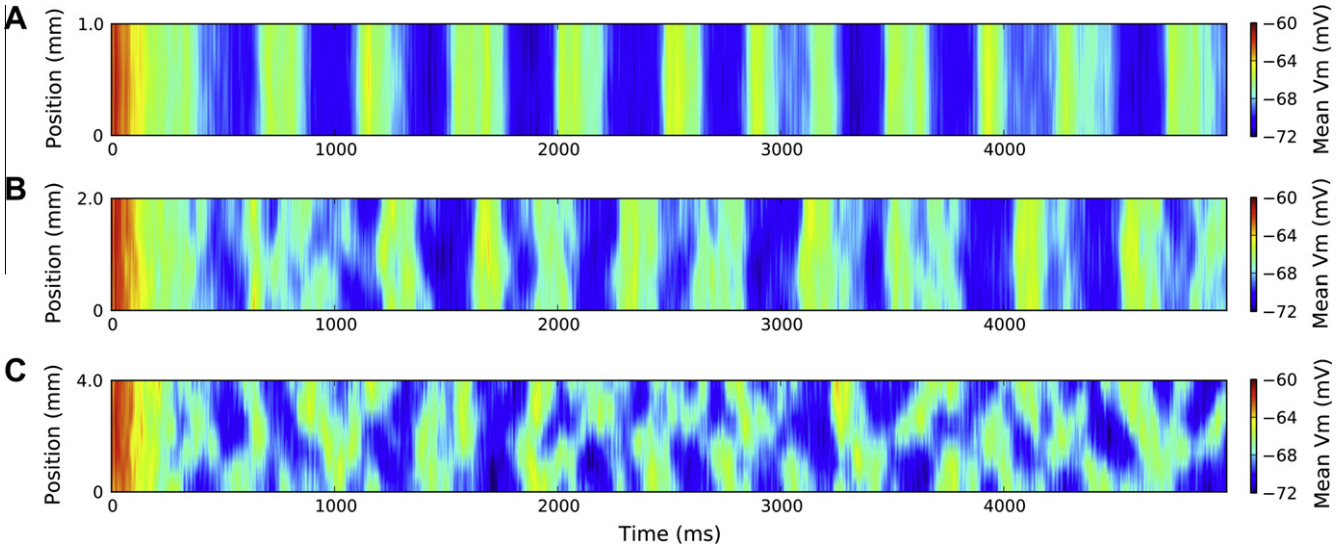

Fig. 10. Scaling of spatial structure in low-frequency oscillations with total network length. The three panels present the results of simulations of coupled one-dimensional ring networks with local connections and conduction delays. Only one layer is depicted here, across three different spatial scales (L = 1, 2, 4 mm). It is important to note the difference in timescales between this and the preceding Fig. (5 s versus 200 ms in Fig. 9). (A) L = 1 mm. In this case, the low-frequency population activity is highly synchronized, with little phase shift across space. (B) L = 2 mm. In this intermediate regime, the scale of the network begins to have a clear effect on the population activity, structuring the oscillations with a gradual phase shift. (C) L = 4 mm. After a second doubling of the network length, the spatial structure of the population activity changes qualitatively, with periods of propagations alternating with periods of incoherent, patchy activity.

small-scale locally connected ring networks in an oscillatory regime (τ_e – 2 ms, τ_i – 5 ms, g_e – 0.8 ns, g – 4) and varied the level of external input across the range of biologically relevant values. Given high external input (v_{ext} – 25 Hz) and low adaptation (b – 5 pA), the network displays coherent, high frequency oscillations (with a peak at 46 Hz in the power spectrum of multi-unit activity), low single-unit firing rate (11.52 Hz) and low average cross correlation (0.081 $\langle CC \rangle$) at the single-neuron level, consistent with previous theoretical work (Brunel and Hakim, 1999; Brunel and Wang, 2003) (Fig. 9A). With a lower external drive to the system (v_{ext} – 6.25 Hz), however, network population events have a lower frequency (30 Hz) and become less coherent, producing oscillations that are not fully synchronized, but have a gradual temporal (i.e. phase) shift across space (Fig. 9B), while the firing of individual neurons remains sparse (6.33 Hz) and uncorrelated (0.011 $\langle CC \rangle$).

To probe the spatial structure of self-sustained activity of the type discussed in Destexhe (2009), we studied two-layer ring networks in a self-sustained activity regime (τ_e – 5 ms, τ_i – 5 ms, g_e – 5 ns, g_i – 70 ns) – one layer with weak adaptation (b – 5 pA), the other with stronger adaptation (b – 155 pA) – displaying bistable dynamics resulting from the transient self-sustained activity of one layer being continually reignited by activity in the second layer.[11] It is important to emphasize here that the membrane potential bistability in this model is not solely driven by strong adaptation

[11] Activity was initiated by a short excitatory pulse of external Poisson input to a local section of one layer.

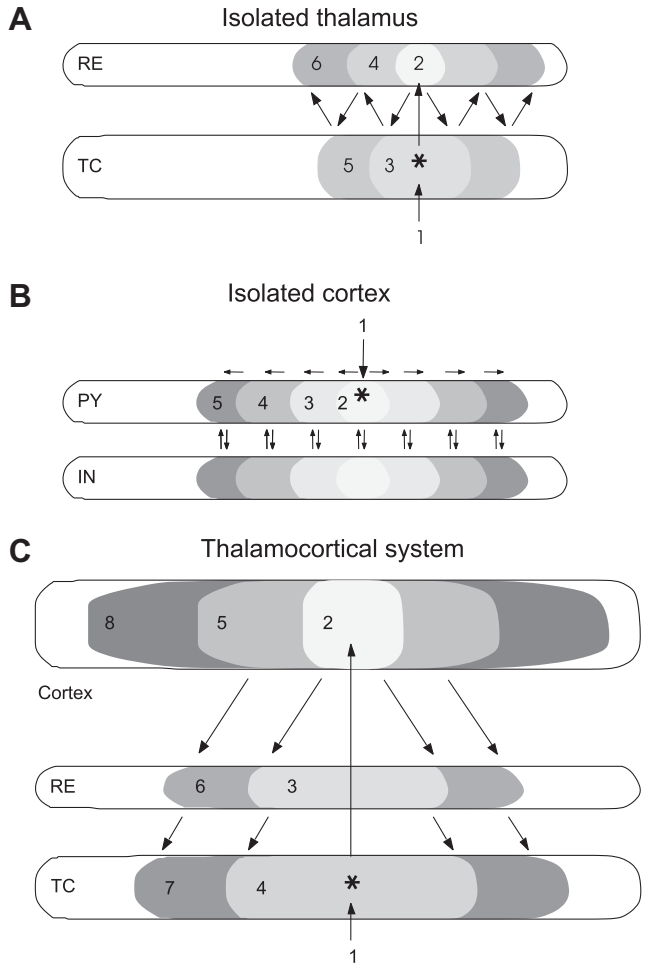

Fig. 11. Different propagation mechanisms in the thalamocortical system. (A) Propagation in the isolated thalamus. The numbers (1), (2), (3),... indicate successive populations of neurons recruited in the propagating wave. In this case, the propagation occurs through successive recruitments between the two layers. (B) Propagation in isolated cortex. In this case, propagation occurs through the local axon collaterals within cortex (horizontal arrows), also recruiting the interneurons (vertical arrows). (C) Propagating waves in the thalamocortical system. The successive recruitments occur through thalamocortical loops. Modified from Destexhe and Sejnowski (2003).

currents, but rather results from an interplay between the self-sustained activity regime at the network level and the adaptation currents at the single neuron level (Destexhe, 2009). To visualize the propagation of the low-frequency oscillations produced in this activity regime, we increased the size of the simulated neural population (N = 112,500, with N_e = 90,000 and N_i = 22,500), with the number of incoming connections per neuron unchanged except for a sparse interlayer excitatory projection (32 synapses per neuron), and systematically changed the length on which the network was wrapped from 1 to 4 mm. Fig. 10 illustrates the change in the spatial structure of population activity with increasing distances in the network. With the original network length (L = 1 mm), the low-frequency population activity is highly synchronized, with very little dependence of the signal on space. After doubling the network length (L = 2 mm), however, the population activity gained a notable shift across space, much greater than that observed in the first case (L = 1 mm). After a second doubling of the network length (L = 4 mm), the population activity again changes qualitatively, with the appearance of incoherent, patchy local activity. This sharp dependence of the structure in population activity on the topology in which the network is embedded during these oscillatory regimes is quite interesting; further considerations, such as the size of the intermediate regime observed here (L = 2 mm), during which coherent yet not fully synchronized population activity is displayed by the network, will be the subject of future study.

Previous studies have begun to address the effect of the topographic nature of cortical networks on their spatiotemporal dynamics (Mehring et al., 2003; Roxin et al., 2005; Kumar et al., 2008; Kilpatrick and Bressloff, 2010; Yger et al., 2011); here, we aim to begin a connection between spiking neural networks with realistic background activity and the experimental results from VSD and MEA recordings taken across brain states. Moreover, whereas topographic models of synfire chain propagation (Mehring et al., 2003; Kumar et al., 2008) or firing rate models with conduction delays (Roxin et al., 2005) focus specifically on propagation of suprathreshold activity, we have focused here on coherent propagation of subthreshold activity, which may be a more robust phenomenon in recurrent networks, specifically in the ongoing activity regime (cf. dissociation between the spatiotemporal patterns in spikes and $\langle V_m \rangle$, Fig. 9B). Additionally, several studies have modeled results from recordings of propagating activity in cortical slices, establishing a detailed connection between propagation and membrane-level conductances, specifically relating their results to the paroxysmal activity state (Golomb and Amitai, 1997; Golomb, 1998; Golomb and Ermentrout, 2002). Here, we aim to study propagation in the context of the activity states observed *in vivo* under normal awake and anesthetized conditions; however, it would be interesting to explore these states in future work. Furthermore, while we have focused predominantly in this work on the activity dynamics of spiking neural networks, it is important to note the substantial literature on propagating waves in neural fields (Ermentrout, 1998; Bressloff, 1999; Goulet and Ermentrout, 2011; Kilpatrick and Ermentrout, 2012) and in networks of coupled oscillators (Kopell and Ermentrout, 1986; Ermentrout and Kleinfeld, 2001), which could provide strong insights into the spatial structure of activity dynamics in spiking neural networks in the future.

Figs. 9 and 10 provide a succinct illustration of possible states for the locally connected networks with conduction delays studied here, as a function of both external input and neuronal adaptation. This model illustrates the transition from low-frequency dominated network states commonly observed during anesthesia and slow-wave (SWS) sleep to high-frequency dominated states observed in waking, activated cortex as an interplay between network-level parameters (e.g. external input rate) and single-unit parameters (level of adaptation). Furthermore, the existence of low-frequency and high-frequency activity need not be mutually exclusive in the model and the two may coexist, consistent with experimental observations of brief high-frequency epochs occurring between the slow-wave complexes of SWS (Destexhe et al., 1999). Finally, it should be noted that while only single points in rather large parameter spaces are presented here as illustrative examples for the purposes of this preliminary study, they serve as a summary of the results of large-scale parameter scans that have been made as part of a forthcoming exhaustive numerical study.

4. Conclusions

4.1. Different "propagating modes" in the thalamocortical system

Fig. 11 summarizes the different types of propagating mechanisms reviewed here. In thalamic slices or in the isolated thalamus, the propagation occurs through mutual recruitment of TC and RE cells, and is therefore polysynaptic. In Fig. 11A, initiation of activity in TC cells (1) recruits a larger population of RE cells (2) due to

divergence of connectivity, which recruits a larger population of TC cells (3), etc. The propagation occurs through successive recruitments between the two layers, as there are no collateral connections in this system. This gives a discontinuous or "lurching" appearance to the propagating wave. Interestingly, because the polysynaptic recruitment involves GABAergic inhibition and inhibitory-rebound sequences, this type of propagating wave is particularly sensitive to altering of GABA receptors. With $GABA_A$ antagonists, the propagating wave is slower, because the wave then depends on the sole activation of the slower type of ($GABA_B$-mediated) inhibition.

In cerebral cortex (Fig. 11B), the propagation occurs through the abundant excitatory collateral connections, and is therefore essentially monosynaptic. In the example of Fig. 11B, activity initiates in PY cells (1), recruits a larger population of PY cells (2) due to local axon collaterals, which itself recruits an even larger population (3), etc. The interneurons are recruited in parallel by the same collaterals. In this case, because propagation speed depends on excitability and proximity of the V_m to spike threshold, blocking $GABA_A$ inhibition is expected to increase excitability and thus also increase propagation velocity (Wu et al., 2001), which is the opposite effect as in the thalamus.

In the thalamocortical system, polysynaptic propagating waves can also occur through mutual recruitment of thalamus and cortex. This is illustrated in Fig. 11C, where initiation in TC cells (1) recruits a larger population in cortex (2) which recruits larger and larger populations in the thalamus (3 and 4), resulting in a more extended recruited population in in cortex (5), etc. This scheme occurs entirely through thalamocortical loops, and is not critically dependent on corticocortical connections. Note that in the case of the thalamic and thalamocortical propagating activities, the propagation is dependent on an oscillation (generated in the thalamus). The models reviewed here reproduce all these interactions.

4.2. Propagating waves depend on network state

Understanding how the topographic nature of cortical networks influences their dynamics and computational properties is critical for a complete description of stimulus integration and localization in the brain. We have reviewed here the available experimental literature from multichannel recordings to assess the functional significance of propagating waves to the waking, activated state of cortical networks and introduced a computational model whose purpose is to account for the range of spatially structured activity generated from the balance of internal and external drives across brain states. Consistent with this approach, we can account for the transition from long- to short-timescale activity as an interplay between single-neuron and network-level properties.

At this time, several key experimental questions stand ready to be answered, such as the explicit dependence of propagating waves on the brain state of the animal and how these spatiotemporal patterns of activation interact with the cortical maps observed after averaging many trials in electrophysiological or optical recording. Additionally, the spatiotemporal interaction of multiple propagating events across the cortical sheet is an important point to consider in the context of the complexity of naturalistic visual or auditory stimuli.

While we have focused here specifically on the potassium-mediated adaptation conductances affected by anesthesia, it is important to note that many neuromodulatory substances modify conductances at the single-neuron level, and may cause determinable effects at the macroscopic level. The transition from large, low-frequency propagating waves to smaller, high-frequency propagating activity analyzed here is one example of analyzing the spatiotemporal dynamics and stimulus integration of a topographic neural network by considering the network conductance state. Certainly, it is possible to uncover more effects of this type when considering ascending modulatory pathways such as the dopaminergic or noradrenergic systems (Bear et al., 2002), and this is a possibility for future work. Moreover, the influence of top-down inputs mediated by processes such as attention during active perception may be an important consideration when understanding the state-dependent shift in propagating activity (Krupa et al., 2004; Ferezou et al., 2006; Poulet and Petersen, 2008).

Acknowledgments

The authors are grateful to an anonymous reviewer for helpful comments. This work has been supported by the CNRS, the Agence Nationale de la Recherche (ANR HR-Cortex, Complex-V1) and the European Community (FET Grants FACETS FP6-015879, BRAINSCALES FP7-269921). L.M. is a PhD fellow from École des Neurosciences de Paris (ENP).

References

Aertsen, A.M.H.J., Gerstein, G.L., Habib, M.K., Palm, G., 1989. Dynamics of neuronal firing correlation: modulation of "Effective Connectivity". J. Neurophys. 61, 900–917.

Ahmed, B., Hanazawa, A., Undeman, C., Eriksson, D., Valentiniene, S., Roland, P.E., 2008. Cortical dynamics subserving visual apparent motion. Cereb. Cortex 18, 2796–2810.

Ayzenshtat, I., Meirovithz, E., Edelman, H., Werner-Reiss, U., Bienenstock, E., Abeles, M., Slovin, H., 2010. Precise spatiotemporal patterns among visual cortical areas and their relation to visual stimulus processing. J. Neurosci. 30, 11232–11245.

Bear, M.F., Connors, B.W., Paradiso, M.A., 2002. Neuroscience: Exploring the Brain. Lippincott, Williams & Wilkins, Baltimore, MA.

Bressloff, P.C., 1999. Synaptically generated wave propagation in excitable neural media. Phys. Rev. Lett. 82, 2979–2982.

Brette, R., Gerstner, W., 2005. Adaptive exponential integrate- and-fire model as an effective description of neuronal activity. J. Neurophysiol. 94, 3637–3642.

Bringuier, V., Chavane, F., Glaeser, L., Fregnac, Y., 1999. Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. Science 283, 695–699.

Brunel, N., Hakim, V., 1999. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. Neural Comput. 11, 1621–1671.

Brunel, N., Wang, X.J., 2003. What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation–inhibition balance. J. Neurophysiol. 90, 415–430.

Buonomano, D., 2003. Timing of neural responses in cortical organotypic slices. Proc. Natl. Acad. Sci. 100, 4897–4902.

Chagnac-Amitai, Y., Connors, B.W., 1989. Horizontal spread of synchronized activity in neocortex and its control by GABA-mediated inhibition. J. Neurophysiol. 61, 747–758.

Chavane, F., Sharon, D., Jancke, D., Marre, O., Frégnac, Y., Grinvald, A., 2011. Lateral spread of orientation selectivity in V1 is controlled by intracortical cooperativity. Front. Sys. Neurosci., 10.3389/fnsys.2011.00004.

Chen, Y., Geisler, W.S., Seidemann, E., 2006. Optimal decoding of correlated neural population responses in the primate visual cortex. Nat. Neurosci. 9, 1412–1420.

Civillico, E.F., Contreras, D., 2006. Integration of evoked responses in supragranular cortex studied with optical recordings in vivo. J. Neurophysiol. 96, 336–351.

Contreras, D., Destexhe, A., Sejnowski, T.J., Steriade, M., 1996. Control of spatiotemporal coherence of a thalamic oscillation by corticothalamic feedback. Science 274, 771–774.

Contreras, D., Destexhe, A., Sejnowsky, T.J., Steriade, M., 1997. Spatiotemporal patterns of spindle oscillations in cortex and thalamus. J. Neurosci. 17, 1179–1196.

Cruikshank, S.J., Lewis, T.J., Connors, B.W., 2007. Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. Nat. Neurosci. 10, 462–468.

Davison, A.P., Brüderle, D., Eppler, J.M., Kremkow, J., Muller, E., Pecevski, D.A., Perrinet, L., Yger, P., 2008. PyNN: a common interface for neuronal network simulators. Front. Neuroinform., doi:10.3389/neuro.11.011.2008.

Derdikman, D., Hildesheim, R., Ahissar, E., Arieli, A., Grinvald, A., 2003. Imaging spatiotemporal dynamics of surround inhibition in the barrels somatosensory cortex. J. Neurosci. 23, 3100–3105.

Destexhe, A., 1998. Spike-and-wave oscillations based on the properties of $GABA_B$ receptors. J. Neurosci. 18, 9099–9111.

Destexhe, A., 2009. Self-sustained asynchronous irregular states and up/down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. J. Comput. Neurosci. 27, 493–506.

Destexhe, A., Paré, D., 1999. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. J. Neurophys. 81, 1531–1547.
Destexhe, A., Sejnowski, T., 2003. Interactions between membrane conductances underlying thalamocortical slow-wave oscillations. Physiol. Rev. 83, 1401–1453.
Destexhe, A., Bal, T., McCormick, D.A., Sejnowski, T.J., 1996. Ionic mechanisms underlying synchronized oscillations and propagating waves in a model of ferret thalamic slices. J. Neurophys. 76, 2049–2070.
Destexhe, A., Contreras, D., Steriade, M., 1998. Mechanisms underlying the synchronizing action of corticothalamic feedback through inhibition of thalamic relay cells. J. Neurophys. 79, 999–1016.
Destexhe, A., Contreras, D., Steriade, M., 1999. Cortically-induced coherence of a thalamic-generated oscillation. Neuroscience 92, 427–443.
Destexhe, A., Contreras, D., Steriade, M., 1999. Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. J. Neurosci. 19, 4595–4608.
Destexhe, A., Rudolph, M., Paré, D., 2003. The high-conductance state of neocortical neurons in vivo. Nat. Rev. Neurosci. 4, 739–751.
DeWeese, M.R., Zador, A.M., 2006. Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. J. Neurosci. 26, 12206–12218.
Ermentrout, E., 1998. The analysis of synaptically generated traveling waves. J. Comput. Neurosci. 5, 191–208.
Ermentrout, G.B., Kleinfeld, D., 2001. Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. Neuron 29, 33–44.
Ferezou, I., Bolea, S., Petersen, C.C.H., 2006. Visualizing the cortical representation of whisker touch: voltage-sensitive dye imaging in freely moving mice. Neuron 50, 617–629.
Franks, N.P., 2008. General anaesthesia: from molecular targets to neuronal pathways of sleep and arousal. Nat. Rev. Neurosci. 9, 370–386.
Fukuda, M., Hata, Y., Ohshima, M., Tsumoto, T., 1998. Role of NMDA receptors in the propagation of excitation in rat visual cortex as studied by optical imaging. Neurosci. Res. 31, 9–21.
Gabriel, A., Eckhorn, R., 2003. A multi-channel correlation method detects traveling γ-waves in monkey visual cortex. J. Neurosci. Methods 131, 171–184.
Gewaltig, M.O., Diesmann, M., 2007. NEST (neural simulation tool). Scholarpedia 2 (4), 1430.
Golomb, D., 1998. Models of neuronal transient synchrony during propagation of activity through neocortical circuitry. J. Neurophysiol. 79, 1–12.
Golomb, D., Amitai, Y., 1997. Propagating neuronal discharges in neocortical slices: computational and experimental study. J. Neurophysiol. 78, 1199–1211.
Golomb, D., Ermentrout, G.B., 2002. Slow excitation supports propagation of slow pulses in networks of excitatory and inhibitory populations. Phys. Rev. E 65, 061911.
Golomb, D., Wang, X.J., Rinzel, J., 1996. Propagation of spindle waves in a thalamic slice model. J. Neurophys. 75, 750–769.
González-Burgos, G., Barrionuevo, G., Lewis, D.A., 2000. Horizontal synaptic connections in monkey prefrontal cortex: an in vitro electrophysiological study. Cereb. Cortex 10, 82–92.
Goulet, J., Ermentrout, G.B., 2011. The mechanisms for compression and reflection of cortical waves. Biol. Cybern. 105, 253–268.
Gray, C.M., Singer, W., 1989. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. PNAS 86, 1698–1702.
Grinvald, A., Lieke, E.E., Frostig, R.D., Hildesheim, R., 1994. Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. J. Neurosci. 14, 2545–2568.
Han, F., Caporale, N., Dan, Y., 2008. Reverberation of recent visual experience in spontaneous cortical waves. Neuron 60, 321–327.
Harel, N., Mori, N., Sawada, S., Mount, R.J., Harrison, R.V., 2000. Three distinct auditory areas of cortex (AI, AII, and AAF) defined by optical imaging of intrinsic signals. NeuroImage 11, 302–312.
Hirsch, J.A., Gilbert, C.D., 1991. Synaptic physiology of horizontal connections in the cat's visual cortex. J. Neurosci. 11, 1800–1809.
Huang, X., Troy, W.C., Yang, Q., Ma, H., Laing, C.R., Schiff, S.J., Wu, J.Y., 2004. Spiral waves in disinhibited mammalian neocortex. J. Neurosci. 24, 9897–9902.
Jancke, D., Chavane, F., Naaman, S., Grinvald, A., 2004. Imaging cortical correlates of illusion in early visual cortex. Nature 428, 423–426.
Kilpatrick, Z.P., Bressloff, P.C., 2010. Spatially structured oscillations in a two-dimensional excitatory neuronal network with synaptic depression. J. Comput. Neurosci. 28, 193–209.
Kilpatrick, Z.P., Ermentrout, B., 2012. Response of traveling waves to transient inputs in neural fields. Phys. Rev. E 85, 021910.
Kim, U., Bal, T., McCormick, D.A., 1995. Spindle waves are propagating synchronized oscillations in the ferret LGNd in vitro. J. Neurophys. 74, 1301–1323.
Kitano, M., Niiyana, K., Kasamatsu, T., Sutter, E.E., Norcia, A.M., 1994. Retinotopic and nonretinotopic field potentials in cat visual cortex. Vis. Neurosci. 11, 953–977.
Kopell, N., Ermentrout, G.B., 1986. Symmetry and phaselocking in chains of weakly coupled oscillators. Commun. Pure Appl. Math. 39, 623–660.
Kral, A., Tillein, J., Hubka, P., Schiemann, D., Heid, S., Hartmann, R., Engel, A.K., 2009. Spatiotemporal patterns of cortical activity with bilateral cochlear implants in congenital deafness. J. Neurosci. 29, 811–827.
Krupa, D.J., Wiest, M.C., Shuler, M.G., Laubach, M., Nicolelis, M.A.L., 2004. Layer-specific somatosensory cortical activation during active tactile discrimination. Science 304, 1989–1992.
Kumar, A., Rotter, S., Aertsen, A., 2008. Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. J. Neurosci. 28, 5268–5280.
Langdon, R.B., Sur, M., 1990. Components of field potentials evoked by white matter stimulation in isolated slices of primary visual cortex: spatial distributions and synaptic order. J. Neurophysiol. 64, 1484–1501.
Léger, J.F., Stern, E.A., Aertsen, A., Heck, D., 2005. Synaptic Integration in rat frontal cortex shaped by network activity. J. Neurophysiol. 93, 281–293.
Maynard, E.M., Nordhausen, C.T., Normann, R.A., 1997. The Utah intracortical electrode array: a recording structure for potential brain–computer interfaces. Electroenceph. Clin. Neurophys. 102, 228–239.
Mehring, C., Hehl, U., Kubo, M., Diesmann, M., Aertsen, A., 2003. Activity dynamics and propagation of synchronous spiking in locally connected random networks. Biol. Cybern. 88, 395–408.
Meirovithz, E., Ayzenshtat1, I., Bonneh, Y.S., Itzhack, R., Werner-Reiss, U., Slovin, H., 2010. Population response to contextual influences in the primary visual cortex. Cereb. Cortex 20, 1295–1304.
Metherate, R., Cruikshank, S.J., 1999. Thalamocortical inputs trigger a propagating envelope of gamma-band activity in auditory cortex in vitro. Exp. Brain Res. 126, 160–174.
Morrison, A., Aertsen, A., Diesmann, M., 2007. Spike-time dependent plasticity in balanced recurrent networks. Neural Comput. 19, 1437–1467.
Muller, L., Reynaud, A., Chavane, F., Destexhe, A., in preparation. Single-trial voltage sensitive dye imaging reveals propagating waves in awake behaving monkey.
Murakoshi, T., Guo, J.Z., Ichinose, T., 1993. Electrophysiological identification of horizontal synaptic connections in rat visual cortex in vitro. Neurosci. Lett. 163, 211–214.
Nauhaus, I., Busse, L., Carandini, M., Ringach, D.L., 2009. Stimulus contrast modulates functional connectivity in visual cortex. Nat. Neurosci. 12, 70–76.
Nauhaus, I., Busse, L., Ringrach, D.L., Carandini, M., 2012. Robustness of traveling waves in ongoing activity of visual cortex. J. Neurosci. 32, 3088–3094.
Ojima, H., Takayanagi, M., Potapov, D., Homma, R., 2005. Isofrequency band-like zones of activation revealed by optical imaging of intrinsic signals in the cat primary auditory cortex. Cereb. Cortex 15, 1497–1509.
Petersen, C.C.H., Grinvald, A., Sakmann, B., 2003. Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured in vivo by voltage-sensitive dye imaging combined with whole-cell voltage recordings and neuron reconstructions. J. Neurosci. 23, 1298–1309.
Pinto, D.J., Patrick, S.L., Huang, W.C., Connors, B.W., 2005. Initiation, propagation, and termination of epileptiform activity in rodent neocortex in vitro involve distinct mechanisms. J. Neurosci. 25, 8131–8140.
Poulet, J.F.A., Petersen, C.C.H., 2008. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. Nature 454, 881–885.
Reimer, A., Hubka, P., Engel, A.K., Kral, A., 2010. Fast propagating waves within the rodent auditory cortex. Cereb. Cortex. http://dx.doi.org/10.1093/cercor/bhq073.
Reynaud, A., Takerkart, S., Masson, G.S., Chavane, F., 2011. Linear model decomposition for voltage-sensitive dye imaging signals: application in awake behaving monkey. NeuroImage 54, 1196–1210.
Rinzel, J., Terman, D., Wang, X.J., Ermentrout, B., 1998. Propagating activity patterns in large-scale inhibitory neuronal networks. Science 279, 1351–1355.
Roxin, A., Brunel, N., Hansel, D., 2005. he role of delays in shaping spatio-temporal dynamics of neuronal activity in large networks. Phys. Rev. Lett. 94, 238103.
Rubino, D., Robbins, K.A., Hatsopoulos, N.G., 2006. Propagating waves mediate information transfer in the motor cortex. Nat. Neurosci. 9, 1549–1557.
Sanchez-Vives, M.V., McCormick, D.A., 2000. Cellular and network mechanisms of rhythmic recurrent activity in neocortex. Nat. Neurosci. 3, 1027–1034.
Shoham, D., Glaser, D.E., Arieli, A., Kenet, T., Wijnbergen, C., Toledo, Y., Hildesheim, R., Grinvald, A., 1999. Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes. Neuron 24, 791–802.
Slovin, H., Arieli, A., Hildesheim, R., Grinvald, A., 2002. Long-term voltage-sensitive dye imaging reveals cortical dynamics in behaving monkeys. J. Neurophysiol. 88, 3421–3438.
Song, W.J., Kawaguchi, H., Totoki, S., Inoue, Y., Katura, T., Maeda, S., Inagaki, S., Shirasawa, H., Nishimura, M., 2006. Cortical intrinsic circuits can support activity propagation through an isofrequency strip of the guinea pig primary auditory cortex. Cereb. Cortex 16, 718–729.
Steriade, M., Nunez, A., Amzica, F., 1993. A novel slow (<1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. J. Neurosci. 13, 3252–3265.
Steriade, M., Timofeev, I., Grenier, F., 2001. Natural waking and sleep states: a view from inside neocortical neurons. J. Neurophysiol. 85, 1969–1985.
Takahashi, K., Saleh, M., Penn, R.D., Hatsopoulos, N.G., 2011. Propagating waves in human motor cortex. Front. Hum. Neurosci., 5, doi: 10.3389/fnhum.2011.00040.
Tanifuji, M., Sugiyama, T., Murase, K., 1994. Horizontal propagation of excitation in rat visual cortical slices revealed by optical imaging. Science 266, 1057–1059.
Telfeian, A.E., Connors, B.W., 2003. Widely integrative properties of layer 5 pyramidal cells support a role for processing of extralaminar synaptic inputs in rat neocortex. Neurosci. Lett. 343, 121–124.
Versnel, H., Mossop, J.E., Mrsic-Flogel, T.D., Ahmed, B., Moore, D.R., 2002. Optical imaging of intrinsic signals in ferret auditory cortex: responses to narrowband sound stimuli. J. Neurophysiol. 88, 1545–1558.
von Krosigk, M., Bal, T., McCormick, D.A., 1993. Cellular mechanisms of a synchronized oscillation in the thalamus. Science 261, 361–364.
Wester, J.C., Contreras, D., 2012. Columnar interactions determine horizontal propagation of recurrent network activity in neocortex. J. Neurosci. 32, 5454–5471.

Witte, R.S., Rousche, P.J., Kipke, D.R., 2007. Fast wave propagation in auditory cortex of an awake cat using a chronic microelectrode array. J. Neural Eng. 4, 68–78.
Wu, J.Y., Guan, L., Tsau, Y., 1999. Propagating activation during oscillations and evoked responses in neocortical slices. J. Neurosci. 19, 5005–5015.
Wu, J.Y., Guan, L., Bai, L., Yang, Q., 2001. Spatiotemporal properties of an evoked population activity in rat sensory cortical slices. J. Neurophysiol. 86, 2461–2474.
Xu, W., Huang, X., Takagaki, K., Wu, J.Y., 2007. Compression and reflection of visually evoked cortical waves. Neuron 55, 119–129.
Yger, P., El Boustani, S., Destexhe, A., Frégnac, Y., 2011. Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons. J. Comput. Neurosci. 31, 229–245.

Chapter 3

Propagating waves in the visual cortex of the awake monkey

[Work in collaboration with Alexandre Reynaud and Frédéric Chavane]

The analytic signal

As outlined in Chapter 2, to quantify the spatiotemporal dynamics observed in noisy multichannel recordings, we adopted a phase-based approach (Prechtl et al., 1997; Mitra and Pesaran, 1999), and work with the instantaneous phase of the response at each pixel on the recording array via the analytic signal (following Rubino et al., 2006). Before specifying how this framework is utilized for detection of propagating waves, we will first recall a few introductory points in order to establish the theoretical background. These points are to be found in the general literature on this topic (Papoulis, 1977; Johansson, 1999; Marple, 1999; Feldman, 2011), yet we repeat them here for the sake of introduction to our particular application of the analytic signal.

Let $t \in \mathbb{R}$, and let f be a real-valued function on t that can be represented by an inverse Fourier transform. We may then write

$$\begin{aligned} f(t) =& \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega) e^{i\omega t} d\omega \\ =& \frac{1}{2\pi} \left[\int_{-\infty}^{0} F(\omega) e^{i\omega t} d\omega + \int_{0}^{\infty} F(\omega) e^{i\omega t} d\omega \right] \\ =& \frac{1}{2\pi} \left[\int_{0}^{\infty} F(-\omega) e^{-i\omega t} d\omega + \int_{0}^{\infty} F(\omega) e^{i\omega t} d\omega \right] \\ =& \frac{1}{2\pi} \int_{0}^{\infty} \left(F^*(\omega) e^{-i\omega t} + F(\omega) e^{i\omega t} \right) d\omega, \end{aligned} \tag{3.1}$$

where the last line of Equation 3.1 shows that the positive half of the frequency spectrum is sufficient to represent a real signal (Johansson, 1999). With this point in mind,

we may define a continuous-time analytic signal $z(t)$ using only the positive half of the frequency domain (Papoulis, 1977; Marple, 1999),

$$Z(\omega) = \begin{cases} 2\mathcal{F}(\omega), & \omega > 0 \\ 0, & \omega \leq 0 \end{cases} \tag{3.2}$$

which we may simply rewrite as

$$Z(\omega) = 2\Theta(\omega)F(\omega) \tag{3.3}$$

where Θ is the Heaviside step function, and we take $\Theta(0)$ to be 0. Taking the inverse Fourier transform of this equation,

$$\begin{aligned} \mathcal{F}^{-1}\{Z(\omega)\} =& \mathcal{F}^{-1}\{2\Theta(\omega)F(\omega)\} \\ z(t) =& f(t) * \left[\delta(t) + \frac{i}{\pi t}\right] \\ =& f(t) + i\left[f(t) * \frac{1}{\pi t}\right] \\ =& f(t) + i\hat{f}(t) \end{aligned} \tag{3.4}$$

where, in evaluating the first step in Equation 3.4, we have used the Fourier transform pair $\delta(t) + (i/\pi t) \leftrightarrow 2\Theta(\omega)$ (Papoulis, 1977, section 7-4). The last line in Equation 3.4 is now in the form of the analytic signal (Gabor, 1946; Feldman, 2011):

$$z(t) = f(t) + iH[f(t)], \tag{3.5}$$

(where H is the Hilbert transform), but with a term $\hat{f}(t)$ parallel to $H[f(t)]$. This parallel term $\hat{f}(t)$, however, is simply the convolution of the original signal $f(t)$ with $(1/\pi t)$, which in turn is exactly the definition of the Hilbert transform (Feldman, 2011).

Thus, by starting with the desired properties of the transformed quantity (Equation 3.2), and performing the required inverse transformation, we recover the definition of the Hilbert transform in the temporal domain. In this way, we see that the properties of the Hilbert transform allow one to rewrite a real signal as a complex exponential $\left[A(t)e^{i\phi(t)}\right]$, thereby separating the amplitude and phase of the real signal into separately accessible components, respectively termed **instantaneous amplitude** $A(t)$ and **instantaneous phase** $\phi(t)$.

It is important to note at this point, however, that such a representation does not come for free, and this transform must be carefully applied in any particular context. With this in mind, in applying the Hilbert transform throughout this analysis, we used several checks to ensure the validity of the analytic signal representation, and to ensure that the signal of interest was well-represented under the numerical implementation of the transform. Signal reconstruction error was verified to be negligible (Goswami and Hoefel, 2004). Continuity of the instantaneous phase was assessed via the distribution of angular acceleration values (rad/s^2), to ensure against the presence of "phase slips", which can occur when the total signal bandwidth becomes large relative to the pass-band of interest. Variation in instantaneous frequency across channels was assessed, and verified to have no systematic organization across space.

Stationary pulses, propagating waves, and their analytic signals

Now, we want to make a distinction between propagating waves and stationary pulses of activity, using their analytic signal forms. To do this, we can introduce two mathematical models to formalize the spatiotemporal profile of the cortical population response in each case. With this motivation, then, we first introduce a simple, space-time separable model, which we term a "stationary" or "Gaussian pulse". We write this model as a function of space and time,

$$f(r,t) = g(r)h(t) \tag{3.6}$$

where r represents the norm of the distance vector from the origin, $g(r)$ represents a Gaussian profile in space $\left[g(r) = A\exp(-r^2/2\sigma^2)\right]$ and $h(t)$ is a sinusoidal response timecourse $\left[h(t) = (1+\cos\omega t)/2\right]$, with amplitude A, spatial spread σ, and angular frequency ω. Note that this equation is defined for arbitrary t, but here, we focus on one cycle of response (from 0 to $2\pi/\omega$ seconds).

Given this definition of the Gaussian pulse, we may then take the Hilbert transform of $f(r,t)$,

$$H\left[f(r,t)\right] = H\left[g(r)h(t)\right] = g(r)H\left[h(t)\right] \tag{3.7}$$

and evaluate $H[h(t)]$,

$$\begin{aligned} H[h(t)] =& \frac{1}{2}\left[sgn(\omega)\sin(\omega t)\right] \\ =& \frac{1}{2}\sin(\omega t) \end{aligned} \tag{3.8}$$

Now, forming the analytic signal $z(r,t)$,

$$\begin{aligned} z(r,t) =& f(r,t) + i\hat{f}(r,t) \\ =& \frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \Big[1 + \cos(\omega t)\Big] + i\frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \sin(\omega t) \\ =& \frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \Big[1 + \cos(\omega t) + i\sin(\omega t)\Big] \\ =& \frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \Big[1 + e^{i\omega t}\Big] \end{aligned} \tag{3.9}$$

where, in this representation, we have phasors rotating in the complex plane, each with identical phase angle, and with amplitude modulated by position in space.

Similarly, if we repeat this calculation for a sinusoidal target wave expanding out from a point source, again with $g(r) = A\exp(-r^2/2\sigma^2)$ and now $h(r,t) = \big(1 + \cos(\omega t - \lambda r)\big)/2$, the analytic signal becomes

$$\begin{aligned} z(r,t) =& \frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \Big[1 + \cos(\omega t - \lambda r)\Big] + i\frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \sin(\omega t - \lambda r) \\ =& \frac{A}{2} e^{-\frac{r^2}{2\sigma^2}} \Big[1 + e^{i(\omega t - \lambda r)}\Big] \end{aligned} \tag{3.10}$$

where we have phasors rotating in the complex plane, but with a linear phase shift with distance from the origin.

Equations 3.9 and 3.10 prove most useful when visualizing their temporal progression in the complex plane. In Fig. 3.1a, we plot the Gaussian pulse in the amplitude domain (left) and its analytic signal in the complex plane (right). In the complex plane, the Gaussian pulse appears as a series of concentric rings (illustrating the amplitude dependence across space), with identical phase angles at each point in time (illustrating the independence of phase across space). For the target wave (Fig. 3.1b, left), however, these concentric rings experience a clockwise rotation across space (illustrating the dependence of phase on space, Fig. 3.1b, right), so that the smallest-amplitude (and most distal) oscillation visualized experiences a rotation of almost one half-cycle.

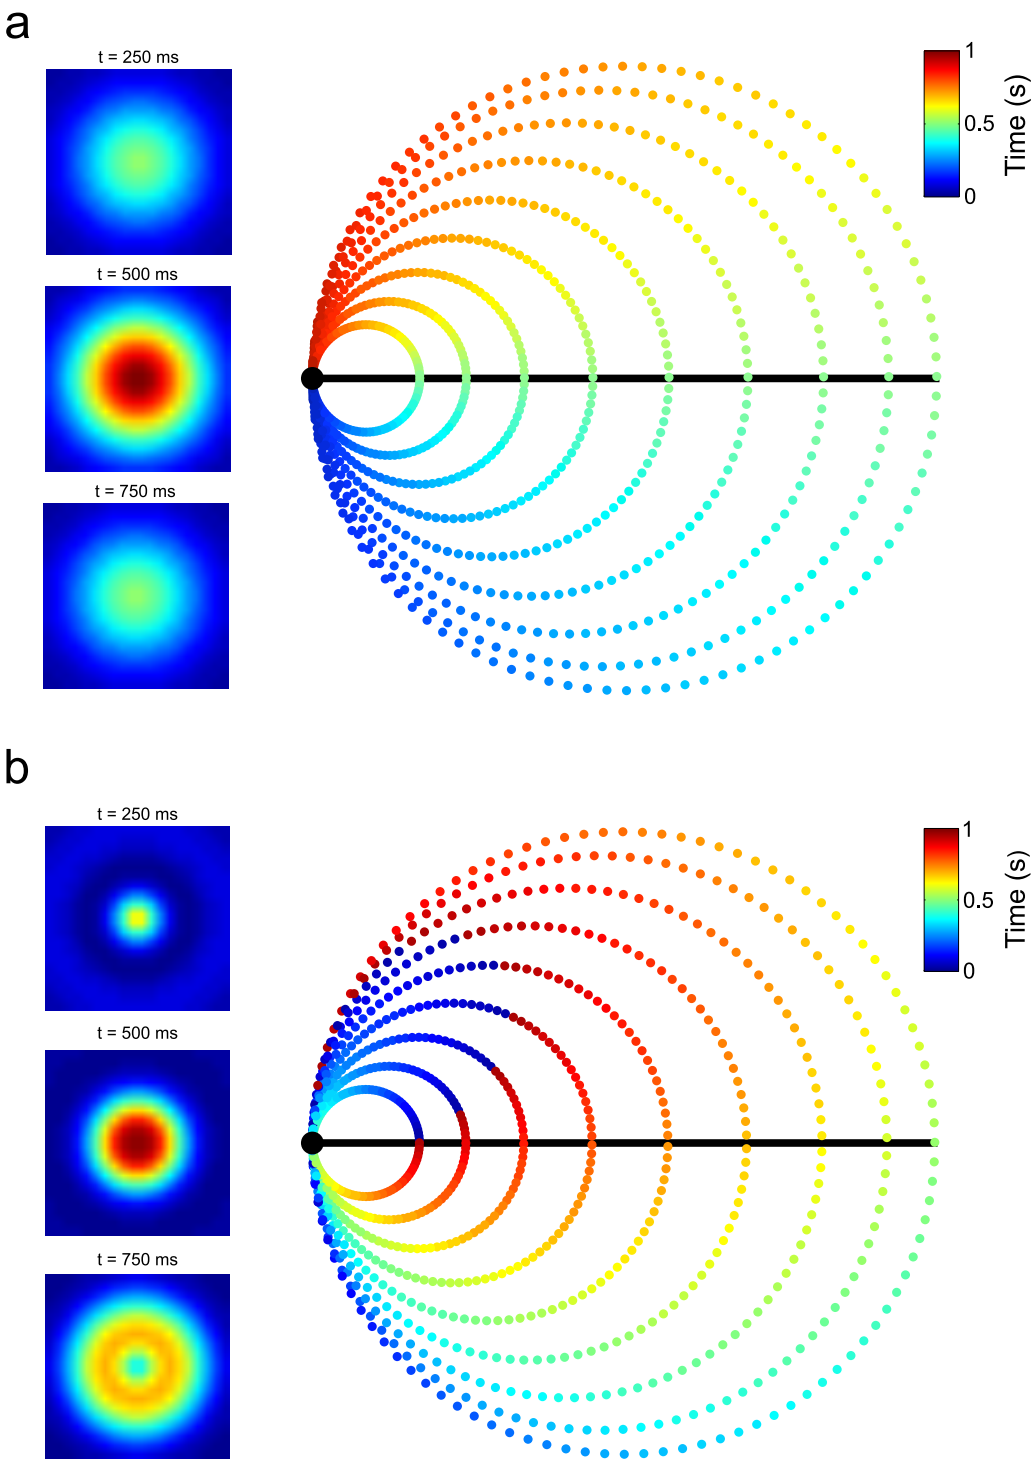

Fig. 3.1 Gaussian pulse and target wave. The Gaussian pulse and target wave examples for the spatiotemporal form of the population response are plotted in the amplitude domain (left), and the analytic signals from the central column as points in the complex plane (right). The real axis appears as a thick black line, with the origin plotted as a black point at the left. Time during the response is indicated by the color of each point in the complex plane (colorbar at top right). Note that a change in the direction of rotation has been introduced for the purposes of visualization. In this representation, the temporal offset in the two responses can be immediately identified by the color of the points at a given phase crossing (e.g. at 2π).

"Phase latency"

These visualizations in the complex plane can furthermore help to explain the numerical algorithm we use to distinguish between stationary pulses and propagating waves in noisy multichannel data. For this, we calculate a quantity we term **phase latency**, which represents a sort of integrated phase progression, in reference to absolute time, at each pixel in the data. To calculate the phase latency at a given pixel, we select a point in time within the oscillation cycle of interest (Fig. 3.2, green points, which represent a starting time at 500 ms), and calculate the latency in absolute time to a given phase crossing for the channel (Fig. 3.2, larger points falling on the horizontal, where the phase value crosses 2π). When the phase crossing of interest falls between two

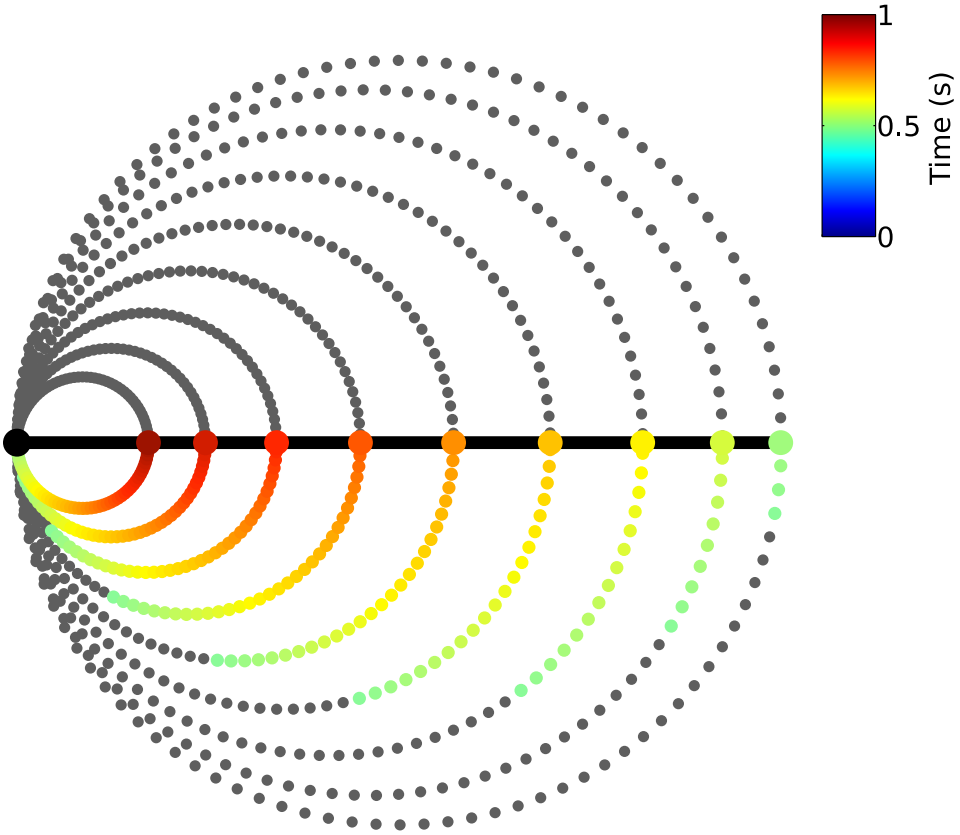

Fig. 3.2 Phase latency illustration. The analytic signal of the sinusoidal target wave depicted in Fig. 3.1b is used to demonstrate the calculation of phase latency from a complex-valued timeseries. The horizontal black line is again the real axis, with the origin plotted as a thick black point at left. The calculation of phase latency begins at a chosen time point (at 500 ms, represented by the first colored green dots), and extends until the phase of the timeseries crosses a given phase (here 2π). The extracted time for the phase crossing in the example above is given here by the color of the larger dots falling on the real axis.

discrete samples, the precise time for the crossing is calculated by linear interpolation between the two points by their instantaneous frequency ($d\phi/dt$).

Here, further, we calculate instantaneous frequency through a complex multiplication, using a method from Feldman (2011). To briefly explain, let x_n be a discretely-sampled real signal, and let $X_n = x_n + i\hat{x}_n$ be its discrete-time analytic signal representation. The product of X_n and its complex conjugate at the next sample X^*_{n+1} produces a new complex sequence whose angle is equal to the instantaneous frequency of the original:

$$\Delta\phi_n = \arg(X_n X^*_{n+1}). \tag{3.11}$$

The benefits of this method are firstly that we may calculate the instantaneous frequency without phase unwrapping and secondly that it is much less sensitive to noise than numerical differentiation of the signal instantaneous phase. After much experimentation, then, we have found that by working with this "integrated" phase quantity and calculating instantaneous frequency though complex multiplication, we are able to deal with the noise in these multichannel data quite effectively, so that a clear distinction can be made between true propagating waves and noisy, synchronous pulses. Precisely how this wave detection is done is explained in the next section.

Detecting arbitrarily shaped waves in noisy multichannel data

Up to this point, we have focused on the calculation of phase latency at the level of single pixels. We next quantify spatial patterns in the phase latency in several steps. First, we calculate the phase latency at each channel (thereby producing a "phase latency map" across space); then, we estimate the wave source from the minimum of the smoothed phase latency maps. Note that if the phase latency map contains only noise, the estimated wave source will be more or less randomly distributed across space. Next, after calculating the Euclidean distance matrix from the source point, we can estimate the correlation of phase latency with linear distance. Note that the correlation with linear distance is simply a first approximation, and could be refined in the future. Finally, by extracting the propagation speed from the slope of the phase latency with distance, and by assessing the strength of the correlation coefficient of this relationship, we can make a definite statement on whether a wave is occurring at this particular time point. The application of this method to quantifying the spatiotemporal form of the stimulus-evoked population response in visual cortex of the awake monkey is the subject of **Work 2:** Muller et al. (2014); following this, we then apply the method to detect epochs of phase organization across space in higher-frequency oscillatory activity in MEA data taken from the monkey and human, in **Chapter 4**.

Work 2: *The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave.*

French Summary

Les ondes d'activité se propagent dans de nombreaux milieux excitables, et elles ont été observées récemment dans des systèmes neuronaux, de la rétine au néocortex. Alors que les ondes de propagation d'activité sont clairement présentes sous anesthésie, leur occurence chez le sujet éveillé est controversée. Une explication possible est que ces ondes sont systématiquement supprimées dans l'analyse des données moyennées, en raison de leur variabilité. Nous présentons dans ce travail une nouvelle méthode pour détecter des ondes de propagation dans des enregistrements multicanaux bruités. En appliquant cette méthode à des données d'imagerie voltage sensitifs (VSD, "voltage-sensitive dye") chez le singe, nous montrons que la réponse évoquée dans la population neuronale du cortex visuel primaire se propage comme une onde, avec une dynamique consistante et fiable d'un essai à l'autre. Un modèle de réseau de neurones indique que cette fiabilité est la signature du réseau de fibres horizontales des couches superficielles du neocortex. Des ondes d'activité ayant des propriétés similaires se propagent indépendamment dans le cortex visuel secondaire, mais entretiennent des relations de phase précises avec les ondes du cortex visuel primaire. Ces résultats montrent que, en réponse à un stimulus visuel, des ondes de propagation d'activité sont systematiquement évoquées dans plusieures régions visuelles du cortex, générant une trame spatio-temporelle cohérente pour les interactions neuronales.

nature COMMUNICATIONS

ARTICLE

Received 24 Sep 2013 | Accepted 17 Mar 2014 | Published 28 Apr 2014

DOI: 10.1038/ncomms4675 OPEN

The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave

Lyle Muller[1,*], Alexandre Reynaud[2,*], Frédéric Chavane[2] & Alain Destexhe[1]

Propagating waves occur in many excitable media and were recently found in neural systems from retina to neocortex. While propagating waves are clearly present under anaesthesia, whether they also appear during awake and conscious states remains unclear. One possibility is that these waves are systematically missed in trial-averaged data, due to variability. Here we present a method for detecting propagating waves in noisy multichannel recordings. Applying this method to single-trial voltage-sensitive dye imaging data, we show that the stimulus-evoked population response in primary visual cortex of the awake monkey propagates as a travelling wave, with consistent dynamics across trials. A network model suggests that this reliability is the hallmark of the horizontal fibre network of superficial cortical layers. Propagating waves with similar properties occur independently in secondary visual cortex, but maintain precise phase relations with the waves in primary visual cortex. These results show that, in response to a visual stimulus, propagating waves are systematically evoked in several visual areas, generating a consistent spatiotemporal frame for further neuronal interactions.

[1] Unité des Neurosciences, Information et Complexité (UNIC), UPR-3293, CNRS, 1 Avenue de la Terrasse, Gif-sur-Yvette 91198, France. [2] Institut de Neurosciences de la Timone (INT), CNRS and Aix-Marseille Université, UMR 7289, Campus Santé Timone, 27 boulevard Jean Moulin, Marseille 13005, France. * These authors contributed equally to this work. Correspondence and requests for materials should be addressed to F.C. (email: frederic.chavane@univ-amu.fr) or to A.D. (email: destexhe@unic.cnrs-gif.fr).

Propagating waves of cortical activity have been a long-standing source of interest in neuroscience, from macroscopic scales spanning multiple cortical areas[1] to microscopic scales involving single neurons[2]. Since the recent introduction of multichannel recording techniques capable of resolving the dynamics of single cortical areas[3,4], the spatiotemporal structure of activity in these mesoscopic cortical regions has become a prime subject of study. With these methods, stimulus-evoked propagating waves have been observed in several sensory cortices of the anaesthetised animal[5–14].

In awake preparations, however, the results are still unclear. In recent advances, multielectrode array recordings have uncovered travelling high-frequency gamma- and beta-band oscillations propagating across the primary visual[15] and motor[16] cortices of the monkey, and fast wave propagation in the auditory cortex of the awake cat[17]; further, voltage-sensitive dye (VSD) imaging studies have observed propagating waves spanning large parts of the cerebral cortex in the awake rodent[18–20]. Nevertheless, the central question here still remains: does the stimulus-evoked population response in primary sensory cortex during waking states remain stationary at the point of thalamocortical input, or propagate across some extent of the cortical area? This question has serious implications for models of cortical coding[21]. While VSD imaging in the awake monkey is an ideal method to provide an answer, with its inherently high spatial and temporal resolution[22], the first VSD imaging studies in the primary visual cortex have reported a spread of activity in the stimulus-evoked population response not clearly consistent with a stationary pulse or propagating wave[23–26]. The results of these first studies in the awake monkey were based on trial-averaged data, however, and given the well-known sensitivity of propagating waves *in vivo* to trial averaging[10,27], it is possible that propagating waves on the single-trial level were attenuated by the averaging procedure.

Here, to rigorously quantify the spatiotemporal dynamics underlying the population response in the awake monkey, we designed a phase-based method that works on unsmoothed data at the single-trial level, and provides a quantitative means to detect propagating waves in noisy multichannel data. We used data denoised for single-trial analysis[28], and sought to distinguish unambiguously between a stationary population response (termed here a 'Gaussian pulse') and a spatiotemporally structured propagating wave. With this methods, we demonstrate that the stimulus-evoked population response is a propagating wave that travels across a large section of visual cortex in the awake monkey. On no-stimulus 'blank' trials, intermittent spontaneous propagations also occur. The evoked propagation occurs during all tested visual stimulation paradigms, and shows an intrinsic reliability: while wave onset is variable across trials, the phase variability of the waves does not increase during their course across cortical space, indicating a propagation substrate with a highly consistent propagation speed. The simplest mechanism accounting for this is the horizontal fibre network of the superficial cortical layers; using a topographic spiking network model, we confirm horizontal fibres can mediate such reliability under very general conditions. A direct consequence of this observation is that if two cortical areas are activated by a visual stimulus with short latency difference (such as primary and secondary visual cortex), then the propagations induced in these areas should be correlated in their course across cortical space. Indeed, we observe a significant correlation of stimulus-evoked waves in those two areas. Such a consistent spatiotemporal organization across neocortical circuits may be an important feature of the network-level computations in visual cortex during awake, 'activated' cortical states.

Results

VSD Imaging on the Single-Trial and Trial-Averaged Level. For the experiments analysed in this study, the primary and secondary visual areas (V1 and V2) of the macaque monkey were imaged under spontaneous and stimulus-evoked conditions. The recording chamber was positioned over an area spanning portions of these cortical areas (18 mm diameter, Fig. 1a; see Methods). Stimulation with a short (50 ms), small Gaussian blob evoked a fast response, which peaked within 50–100 ms after stimulus onset (Fig. 1b, mean ± s.d.). Noting the temporal characteristics, we selected a frequency band encompassing this fast response (5–20 Hz) for all further analysis.

Figure 1c illustrates the difference between the single-trial and trial-averaged level in these data. Smoothed surface plots of the stimulus-evoked VSD responses in three trials are plotted alongside the trial-averaged data (animations available online as Supplementary Movies 1–4). The form of the response is clearly different between the single-trial and trial-averaged data. A detailed inspection of the surface plots on the single-trial level shows that activity first emerges proximally (near the vertical reference line) and is followed in later frames by the development of stronger activity at more distal locations, suggesting the presence of a propagating wave. In contrast, on the trial-averaged level, distal responses are heavily attenuated as a result of the averaging, resulting in a spatiotemporal pattern similar to a Gaussian pulse (see Introduction). Nevertheless, while propagations seem to be present on the single-trial level, the phase relationship between proximal and distal sites cannot be clearly established in these amplitude-domain, smoothed surface plots, where the detailed spatiotemporal dynamics remain obscured. For this reason, we need more sophisticated methods to characterize the 'phase latency' of individual channels in the unsmoothed data.

Phase latency method. To quantify the spatiotemporal properties of visual responses, we compute instantaneous phase at each pixel via the analytic signal framework[16,29,30], wherein a real-valued, narrowband time series on channel i $\left(V_t^i\right)$ can be written as a complex exponential:

$$V_t^i + jH[V_t^i] = A_t^i e^{j\phi_t^i} \quad (1)$$

where H denotes the Hilbert transform, A_t^i the instantaneous amplitude of channel i at time t, ϕ_t^i denotes the instantaneous phase of the same, and j is the imaginary unit. To illustrate this concept, Fig. 2a depicts the representation of a damped oscillation in the real, imaginary and complex planes via the analytic signal, in order to provide an example of the transformation of a real signal into a complex phasor.

In this work, we specifically consider the latency in absolute time to a given phase crossing as the phasor at each channel rotates in the complex plane ('phase latency', Fig. 2b). The calculation of phase latency starts from a selected point in the time series (first black point, Fig. 2b), and consists of finding the two samples on either side of a given phase crossing (2π throughout). Linear interpolation based on the instantaneous frequency of the signal between these two points returns the exact time for the crossing within the interval. This calculation is made independently for all channels in the bandpass-filtered, spatially unsmoothed data. By analysing spatial patterns in the phase latency across channels, we can unambiguously distinguish between stationary pulses and propagating waves in VSD imaging data.

To clarify this distinction, we focused on two idealized models for the form of the population response: a spatiotemporally separable, stationary 'Gaussian pulse' and a spatiotemporally

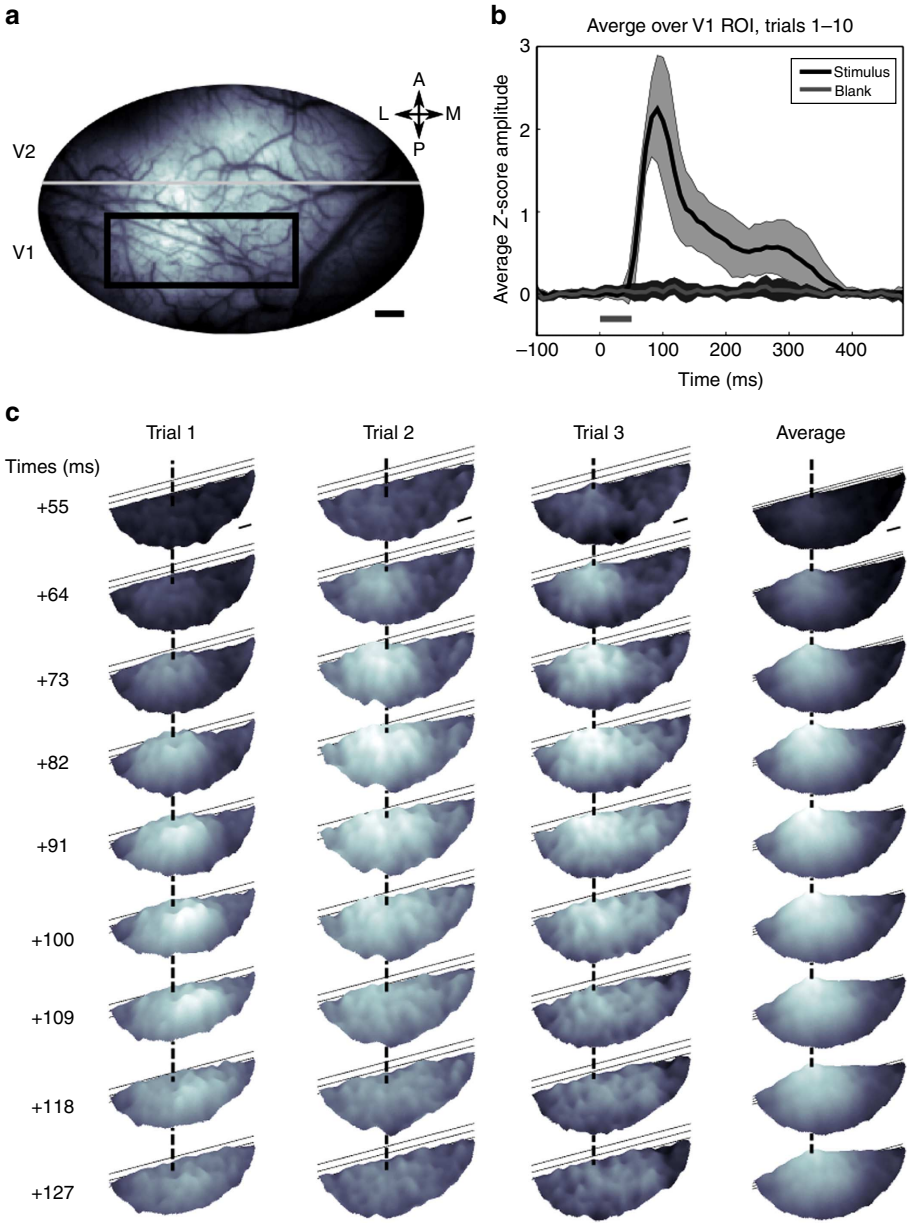

Figure 1 | Imaged cortical area and stimulus response time course. (**a**) Vasculature pattern. Black rectangle indicates V1 ROI for response in Figs 3,4. Black bars indicate 1 mm spatial scale throughout this work (here, bottom right). (**b**) Average time course in the V1 ROI for stimulus (black) and blank (grey) conditions, averaged over first 10 trials (mean ± s.d.). The grey bar from 0–50 ms marks stimulus presentation. (**c**) Surface plots of the response, for three single trials and the trial average. Individual frames are smoothed with a Gaussian (0.23 mm s.d.), and scaled to the maximum *Z*-score in the trial. Horizontal black lines indicate *Z*-scores of 0, 1 and 2. Dashed vertical lines provide a visual reference across surface plots.

inseparable propagating wave. Note that here we use 'pulse' as a purely temporal concept, indicating a lack of phase organization across space, whereas a 'propagating wave' indicates this spatiotemporal organization, for both single- and multi-cycle events. We write the Gaussian pulse as a separable function of space and time:

$$f(x,t) = g(x)h(t) \tag{2}$$

where $g(x)$ represents a Gaussian profile in space $\left(g(x) = Ae^{-x^2/2\sigma^2}\right)$ and $h(t)$ is a sinusoidal response time course $(h(t) = 1/2[1+\cos(\omega t)])$, with amplitude A, spatial spread σ, and angular frequency ω, over one cycle of response. Given this definition of the Gaussian pulse, its analytic signal representation $z(x, t)$ becomes:

$$z(x,t) = \frac{A}{2}e^{-\frac{x^2}{2\sigma^2}}[1+e^{j\omega t}] \tag{3}$$

which has the properties discussed above—phasors rotating in the complex plane, each with identical phase angle, whose amplitudes are modulated by a Gaussian in space (see Methods—Analytic Signal Representation of the Gaussian Pulse). Clearly, because of the spatiotemporal separability of this model, the phase-latency map calculated for a representative Gaussian pulse is flat, exhibiting no organization across space (Fig. 2c,d). By contrast, the phase-latency map calculated for an expanding target wave $(\sin(\lambda r - \omega t)$, where r is the norm of the distance vector $\vec{r}$ from the source) captures the phase offset across space, producing a

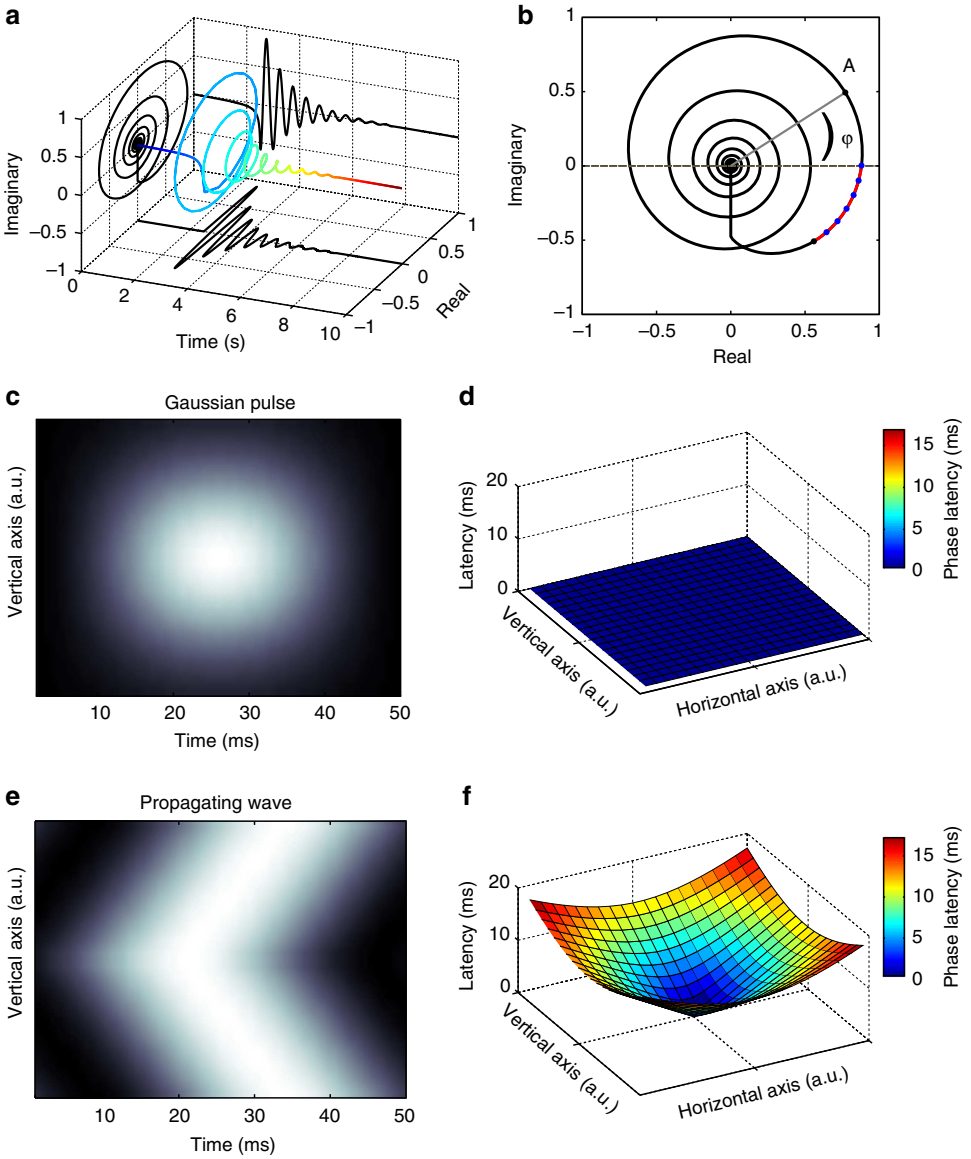

Figure 2 | Phase latency method. (**a**) The real, imaginary, and complex plane projections of the analytic signal (coloured by heat in time) for a damped oscillation make explicit the decomposition of a real signal into a complex phasor[30]. (**b**) The complex plane projection in (**a**) is used in this work to analyse instantaneous amplitude (*A*, complex modulus) and phase (ϕ, complex angle) in the real signal. The red curve in (**b**) illustrates the interval over which an example phase latency calculation takes place. (**c**) A vertical section through time of a representative Gaussian pulse. Note the spatial axes have arbitrary units (a.u.). (**d**) Phase-latency map for the Gaussian pulse response depicted in (**c**). Note that the minimum latency has been subtracted from all maps, to aid visualization. (**e**) A section through time of a representative target wave. (**f**) The phase-latency map for this example captures the spatiotemporal organization in these data.

basin centred on the wave source (Fig. 2e,f). It is important to emphasize that, in contrast, a latency analysis in the amplitude domain cannot distinguish between Gaussian pulses and propagating waves (Supplementary Fig. 1). Several tests were made with various signals and sources of noise to ensure that the phase latency method is robust against false-positive propagation detections (Supplementary Fig. 2).

Phase-latency maps reveal spontaneous and evoked waves. Phase-latency maps were calculated pixel by pixel starting from +72.7 ms after stimulus onset, then smoothed for purposes of visualization (Methods). Individual phase-latency maps at the single-trial level show a clear increase in latency with distance from the source of the basin (Fig. 3a), in the range expected for a propagation mediated by the horizontal fibre network[6,31,32]. While no spatiotemporal organization was observed in the blank trial data at the corresponding temporal point, we did observe several instances of spontaneous propagating waves, with similar propagation speed as the evoked, emerging at various spatial and temporal points during the no-stimulus imaging sessions (Fig. 3b). Note that the size, but not the position, of the ROI in Fig. 3b corresponds to that in Fig. 3a.

We next addressed whether these propagations occurred consistently across trials. For this, individual phase-latency maps were calculated at the same temporal point for trials 1–10, smoothed and then averaged. This average phase-latency map (Fig. 4a, top) again shows a clear increase in latency with distance

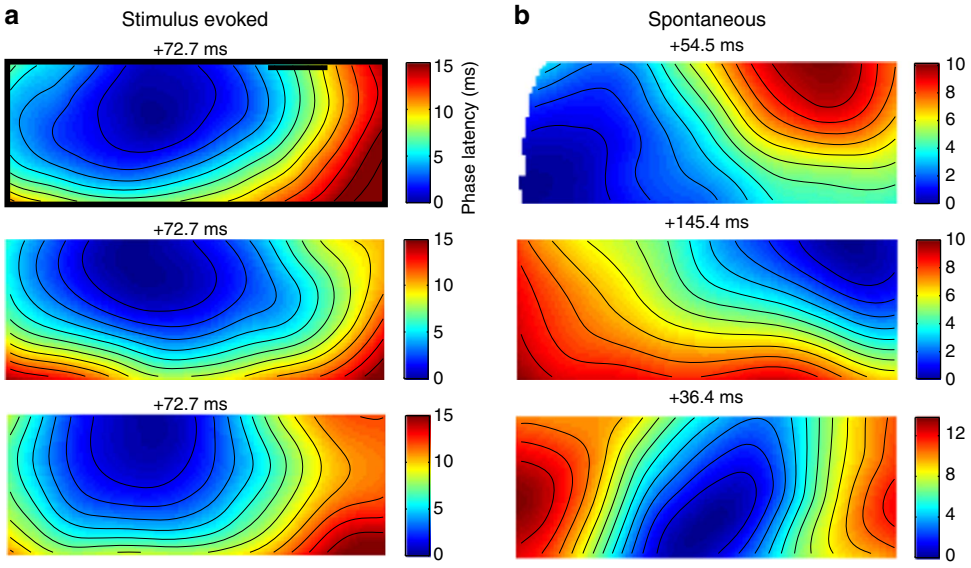

Figure 3 | Single-trial evoked and spontaneous waves in awake monkey. (**a**) Single-trial phase-latency maps for V1 ROI in the 50 ms stimulus presentation condition, from trials 1, 3 and 10. These maps are calculated at +72.7 ms after stimulus onset. Note black box in top panel corresponds to V1 ROI in Fig. 1a. (**b**) Phase-latency maps for spontaneous waves observed during no-stimulus, blank conditions. Note the varying colour axis and temporal points for the spontaneous maps.

from the source, consistent with the single-trial maps. No spatiotemporal organization is observed in the blank trial data at the corresponding temporal point, which produces a flat phase-latency map (Fig. 4a, bottom). This trial-average over phase-latency maps is much less sensitive to noise than the trial-average over the signal, first because the waves are extracted in the individual trials before the average, and second because this measure is insensitive to amplitude fluctuations (*cf.* Fig. 1c).

Nevertheless, these smoothed phase-latency maps do not provide a quantitative criterion for the distinction between a significant propagation and background noise; for this, we use the original, unsmoothed phase-latency maps to test for a linear relationship between phase latency and distance (Methods). We extracted the propagation speed from the slope of this relationship across 40 trials in the 50-ms stimulus presentation condition (Fig. 4b). Propagation speeds followed a positively skewed distribution, ranging from 0.25–1.35 m s^{-1}, with a median speed of 0.57 m s^{-1} and a median absolute deviation of 0.18 m s^{-1}, in close correspondence to the observed range of horizontal conduction speeds in monkey V1 (ref. 32). The correlation coefficient (ρ_d) for stimulus (black points) and blank (grey points) conditions is given in the inset. Positive detection of a propagating wave for each trial at the given time point was determined by significance of the correlation coefficient ($\alpha = 0.01$, one-tailed t-test, H_0: $\rho_d = 0$, H_1: $\rho_d > 0$, with Bonferroni correction), coupled with a criterion for the propagation speed determined from the slope of the regressor line to be within the biophysically plausible range of horizontal conduction speeds (0.05–0.8 m s^{-1}, shaded area in Fig. 4b)[6,31,32]. Under these dual detection criteria, we found stimulus-evoked propagations over the imaging array in 32 of 40 trials (80.0% detection rate, red points in Fig. 4b, inset). In contrast, for the blank trial data, no propagations were detected at the chosen temporal point.

We analysed additional stimulus presentation conditions in two animals (Fig. 4c, 10- and 100-ms stimulus in monkey WA, and 600-ms stimulus in monkey GR). The phase-latency maps for each condition (also calculated at +72.7 ms) contained similar modulations across space for the 10-, 100- and 600-ms stimuli (80.0, 52.6 and 20% detection rate, respectively). No major differences were observed in the features of the waves between animals and stimulus presentation durations; rather, whether waves could be detected on the single-trial level was related to the signal–to-noise ratio achieved during the individual imaging session. Note that shifts of phase latency minima reflect different visuospatial coordinates for stimuli in each condition, and were verified to correspond to the expected point of thalamocortical input. As previously, no propagations were detected in the matched blank trial data.

As a control on the results derived from the phase latency method, we divided the V1 ROI into successive regions by their measured phase latency (six regions total in the smoothed maps, in 2-ms bins), and averaged all channels across trials and within the region in the unfiltered data from the 50-ms stimulus presentation condition (Fig. 4d). Despite the fact that the process of trial-averaging dramatically attenuates the spatiotemporal frequency components observed on the single-trial level (*cf.* Fig. 1), by averaging over trials and region defined by phase latency calculated in the bandpass data, we are able to recover a trial-consistent phase offset even in the unprocessed data (Fig. 4d, black dots). This control analysis confirms that the propagating response is well characterized by the phase information in the 5–20 Hz frequency band, and rules out artifacts in the phase representation or bandpass filtering of the data. Further, it is important to note that additional checks confirmed the 5–20 Hz frequency band chosen here captures the full waveform of the fast stimulus-evoked population response well: as an example, repeating the analysis in Fig. 4b with a lowpass filter at 27.5 Hz (that is, the lower half of the available frequency spectrum) results in a similar number of wave detections (34 of 40 trials in the 50-ms stimulus condition), while a highpass filter at 27.5 Hz results in zero wave detections.

Phase gradient. Having confirmed the existence of stimulus-evoked propagating waves of activity in these VSD imaging data, we then examined the response wavevectors taken from the spatial gradient of phase at each time point. Following previous

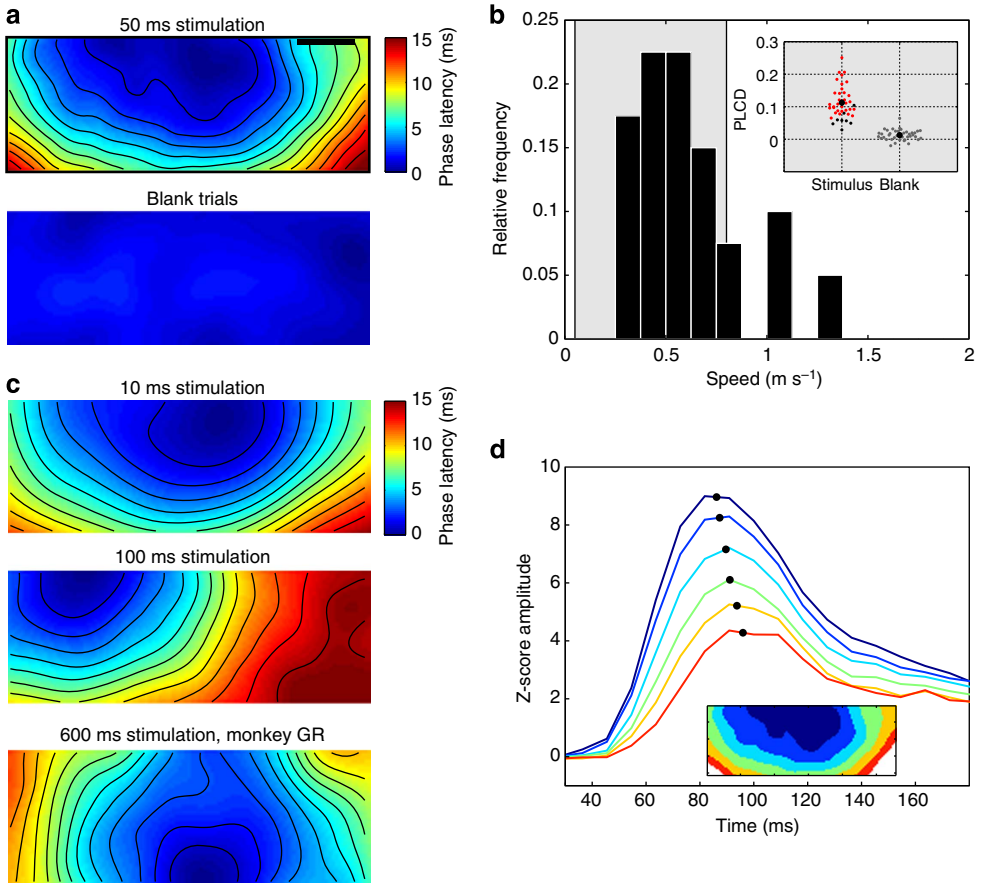

Figure 4 | Stimulus-evoked propagating waves in V1 of the awake monkey. (**a**) Average phase-latency map for V1 ROI for stimulus (top) and blank (bottom) data. (**b**) Propagation speeds extracted from the slope of the phase latency with distance in the unsmoothed maps, 50 ms stimulus presentation condition. (Shaded area) biophysically plausible range of horizontal conduction speeds (0.05–0.8 m s^{-1}). (Inset) phase latency correlation with distance (PLCD), stimulus (black) and blank (grey) conditions. Trials with positive detections are marked in red. Mean correlation coefficient ± s.e.m. plotted with larger black dots (Fisher *z*-transform). Note that error bars are obscured behind the dots. (**c**) Additional phase-latency maps across stimulus presentation length (10 and 100 ms, top and middle), and across animals (600-ms stimulus, monkey GR). Note only significant trials are included in this representation. Stimulus coordinates were 3.5°, 3° and 5° below the horizontal meridian, respectively, and 3° left of the vertical in the 600 ms condition. (**d**) ROIs of increasing phase latency (bottom inset, concentric regions coded by heat) show a temporal offset with distance (time courses, matching colour code). Interpolated maxima plotted directly on time courses (black dots).

work[16,33,34], the wavevector $\vec{k}_t$ at time t is:

$$\vec{k}_t = -\nabla\phi_t \quad (4)$$

Figure 5a depicts the phase gradient at +72.7 ms, averaged over the first ten trials of the session for the stimulus (top) and blank (bottom) conditions. In this wavevector plot over all V1 channels, a 'pinwheel' of colour hues in a region of high phase gradient magnitude marks the source of propagation during the response epoch (Fig. 5a, top). In contrast, the wavevector plot for the blank trial data has low-phase gradient magnitude and no such spatial organization (Fig. 5a, bottom). To estimate the spatial extent of the propagations in these data, we calculated the contour at 2.32 standard deviation (s.d.) values away from the mean of the phase gradient magnitude in the blank trial data (attained by <1% of values in the blank), and delimited a contiguous, approximately elliptical region on the wavevector map in the stimulus condition, with a long axis of 7.5 mm and a short axis of 3.7 mm (calculated via normalized second central moments of the equivalent ellipse), compatible with known anatomy[35,36] and reported spatial spread[23,26]. Using a magnification factor of 0.35° mm^{-1} at an eccentricity of 3.5° (refs 37,38) and an anisotropy index of 1.5 along the inferior vertical meridian[39], this cortical area corresponds to an ellipse of 2° on the horizontal axis and 2.7° on the vertical axis of the visual field. In looking at the phase gradient wavevector length time series averaged over the V1 ROI for the first 10 trials (mean ± s.e.m., Fig. 5b), a sharp phasic increase occurs during the stimulus response epoch, in comparison with the phase gradient time course from the blank trial data (here, high pass filtered with 5 Hz cutoff frequency). These additional measurements based on the spatial gradient of phase clarify the spatial extent of the propagations observed in the data, and corroborate their appearance in a consistent manner across trials.

Trial variability and network model. We next analysed the variability of these propagations in time. During an oscillatory epoch, the phase distribution over channels will in general be peaked and have a well-defined mean direction (MD)[40]. The distribution of MDs (+64 ms, 5–20 Hz frequency band) across 40 trials spans more than a quarter-cycle in the 50-ms stimulus condition (Fig. 6a, black arrows); given the mean instantaneous frequency at this temporal point (11.26 Hz), this span implies a variability on the order of tens of milliseconds (29.08 ms). Note that these results were verified at several time points throughout the stimulus response.

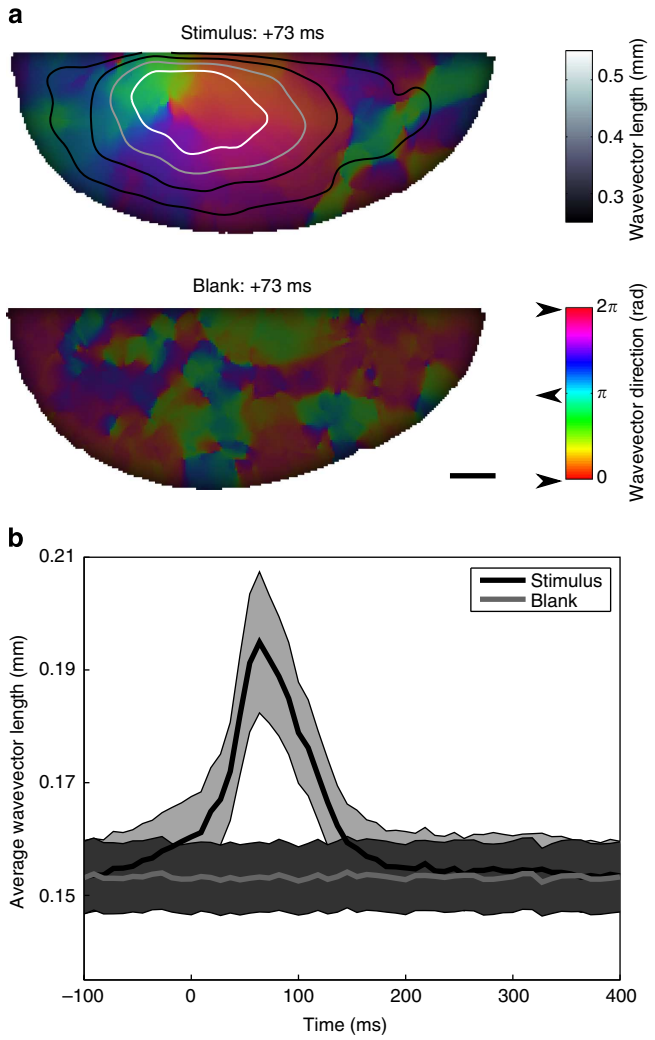

Figure 5 | Phase gradient magnitude and direction plot. (**a**) Wavevector plot of the phase gradient in the stimulus (top) and blank trials (bottom), averaged over the first 10 trials of the 50-ms stimulus condition. Propagation direction is represented by pixel hue (colour bar, bottom right), with wavevector length represented by pixel luminance and grayscale contours (colour bar, top right). (**b**) Phase gradient averaged over V1 ROI, for the stimulus (black) and blank (grey) conditions, mean ± s.e.m.

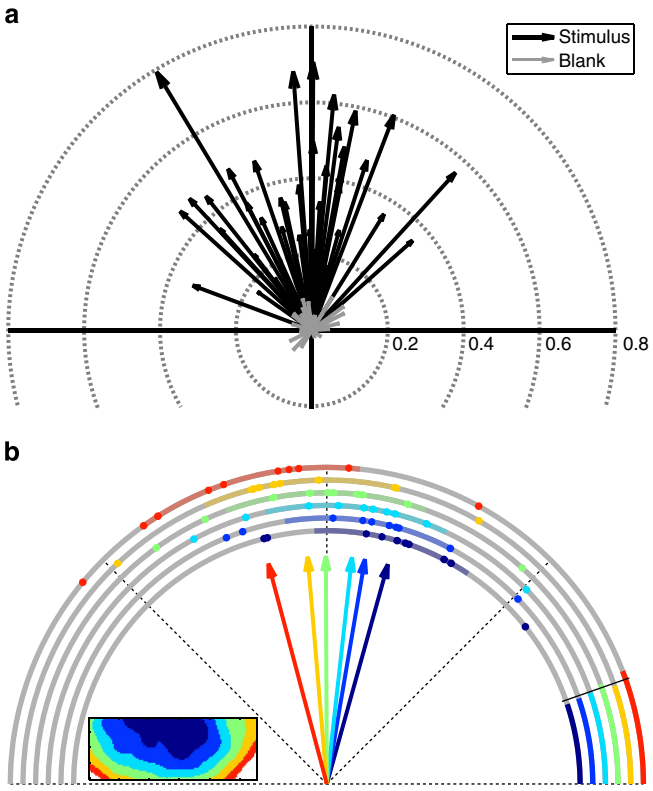

Figure 6 | Variability analysis. (**a**) Mean direction of phase distribution at +64 ms, scaled by length of the resultant vector, for 50-ms stimulus (black arrows) and blank (grey) conditions. (**b**) Mean direction across individual phase-latency regions (inset at bottom left, as defined in Fig. 4d), first 10 trials of the stimulus condition (coloured points). (Coloured arrows) mean direction of the cross-trial distribution. (Shaded regions) mean direction ± angular deviation. (Coloured regions, bottom right) angular deviations, with line through mean of the deviations for visualization (black line).

By dividing this analysis over-phase latency regions (as in Fig. 4d), we studied this variability across space (Fig. 6b). Interestingly, while the mean directions of the cross-trial MD distribution in each region experience a 30.4-degree phase shift across space (reflecting the underlying propagating wave dynamics), the angular deviations differ by no more than 2.5 degrees (Fig. 6b, lower right). In other words, the cross-trial phase variability of the waves does not increase with distance from their source. This strongly suggests that the intrinsic dynamics of the propagations are invariant across trials and space, and that the observed trial variability in these propagations can be explained by simple temporal shifts of the propagating wave from trial to trial.

To investigate possible mechanisms explaining these observations, we studied a one-dimensional topographic spiking network model composed of excitatory and inhibitory neurons operating in a balanced regime, with local connections and linear conduction delays (Methods), and with length L (7.0 mm, matched to the spatial extent of the wave propagation, see Results–Phase Gradient). Stimulation of the middle 100 excitatory cells in the network with a Gaussian spike packet (centre time μ_{sp}; Fig. 7a) ignites a wave of subthreshold activity, which propagates along a rigid spatiotemporal path, determined by the fixed axonal conduction delays and rise time of the postsynaptic potential (Fig. 7b). We divided the network into six regions based on distance from the stimulation zone (coloured regions; Fig. 7a,b) and used the analytic signal representation for the response waveform, as with the experimental data (Methods).

The most likely mechanism underlying the trial-invariant dynamics observed in our experiments is a predominantly monosynaptic propagation carried by the horizontal fibre network of the superficial cortical layers, whose axonal conduction velocities are solely a function of the axonal diameter[41], and are therefore fixed from trial to trial. We used this spiking network model to further understand how this mechanism could produce the phase distributions observed experimentally in two ways: by adding variability in the timing of the input to the system, or by adding variability in the axonal conduction velocity. The rationale for the choice in parameters studied is this: in observing the constant phase variability across distance (as in Fig. 6b), we postulate that this is due to a static underlying mechanism, observed under various sources of noise. If this were not the case, however, and the spatiotemporal dynamics observed in our analysis were mediated by some other process, such as a gain-control mechanism[25], we would expect variability in the

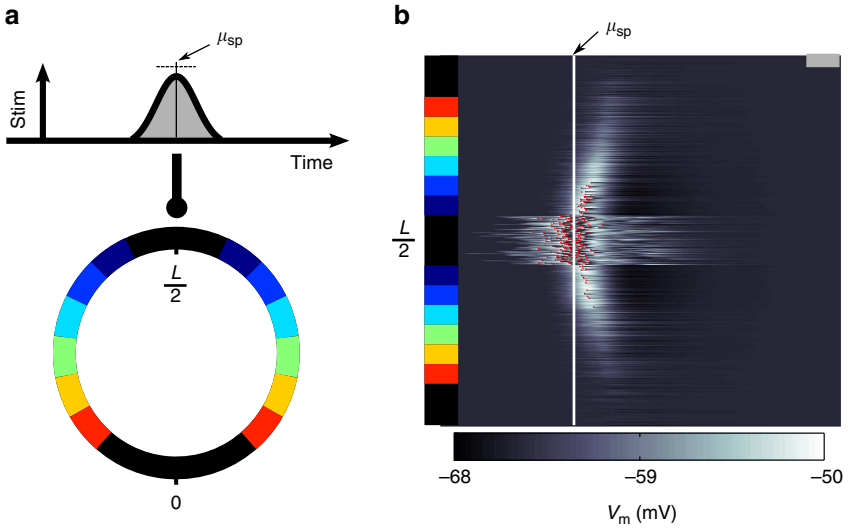

Figure 7 | A spiking network model of the trial-invariant response. (**a**) Schematic of the network model and stimulation paradigm. The network is composed of excitatory and inhibitory cells operating in a balanced regime, and is wrapped onto a one-dimensional ring of length L with distance-dependent connections. A Gaussian spike packet is delivered to the network with a centre time μ_{sp}. Six regions with increasing distance from the edge of the stimulation zone are selected for analysis of their averaged response waveform. (**b**) Example of a stimulus giving rise to a wave of activity. The membrane potential (V_m) across all cells in the excitatory population is plotted in grayscale, with analysis regions depicted at left. Spikes are depicted on the membrane potential time courses by red dots.

apparent conduction speed, as well, which would then manifest as an increase in angular deviation with distance.

To study these two sources of variability, the stimulation paradigm is repeated over ten trials, and the cross-trial phase distribution is plotted as in Fig. 6b. With no additional source of variability in the system, all cross-trial phase distributions have zero angular deviation, as the response phase is totally consistent from trial to trial (Fig. 8a). If the centre time of the Gaussian spike packet (μ_{sp}) is drawn from a uniform distribution ranging from 95 to 105 ms, the angular deviations become nonzero, but remain constant across analysis regions (input jitter; Fig. 8b), as in the experimental data. If the axonal conduction speed for the network in each trial is drawn from a uniform distribution (from 0.2–0.5 ms^{-1}), however, the angular deviations then increase markedly with distance from the source (speed variability; Fig. 8c), in stark contrast to the experimental data. The rate of the increase in angular deviation with distance from the stimulation zone is directly proportional to the range of the uniform distribution from which the axonal conduction speed for the network is drawn at each trial (Fig. 8d, lines of increasing brightness). Note that additional control simulations were performed to test that the combination of these two sources of variability, rather than being mutually obscuring, produces an even larger increase in variability with distance than either source in isolation.

This network model reproduces the cross-trial phase distributions observed in the experimental data under very general conditions. Moreover, the simplicity of the mechanism studied does not preclude its uniqueness: other mechanisms, whose apparent conduction speeds are affected by the ongoing dynamics in the network, do not show this non-increase in phase variability across space. In this way, we demonstrate not only that the propagation must be mediated by the horizontal fibre network of the superficial cortical layers, but also that this mechanism of propagation generation provides a unique physical substrate for delivering a precisely timed event with trial-invariant dynamics (e.g. under temporal variability in the input), preserving across trials temporal patterns in the population response over several millimetres of primary visual cortex.

Dual, independent propagations in V1 and V2. The horizontal propagation mechanism furthermore directly implies that propagations in V1 and V2 should be synchronized in their course across cortical space, as these areas are retinotopically activated by a visual stimulus in a fixed, rapid temporal sequence[42,43]. Specifically, because these areas are activated with comparable response latencies (points labelled **A** in Fig. 9a) and the horizontal fibre network in each area will carry the propagations with similar speeds, the propagations should reach points in the cortex with similar distance from the source at the same time (points labelled **B** in Fig. 9a). In accordance with this expectation, we did in fact observe dual stimulus-evoked propagating waves travelling across V1 and V2, respectively (Fig. 9b). Next, to verify that the propagations observed in this region are also likely mediated by the horizontal fibre network, we performed the trial variability analysis in V2, just as in Fig. 6b, and indeed observed that the cross-trial phase variability similarly does not increase with distance (Fig. 9c).

To summarize the situation observed here, our results suggest that the retinotopic feedforward stream rapidly activates primary and secondary visual cortex in succession, generating two independent waves propagating along the horizontal network in each area. These two waves, which both express consistent dynamics indicative of the horizontal network (Figs 6b and 9c) and similar spatiotemporal dyamics (Figs 4a,c and 9b), are thus expected to exhibit phase correlations across time for each pair of pixels in individual trials. Interestingly, because of the smooth retinotopic organization in V1 and V2, this correlation across cortical space will also correspond to correlation in retinotopic space.

To quantify this effect systematically, we considered spatial sections of V1 and V2, aligned by flipping V2 vertically across the border, and computed the circular correlation coefficient (ρ_{cc}) of the instantaneous phase between each pair of aligned pixels at each point in time (Fig. 10a; see Methods). Note that this calculation is performed in the temporally filtered but spatially unsmoothed data, to preclude any smoothing artifact from introducing spurious correlation. In the average over the first ten

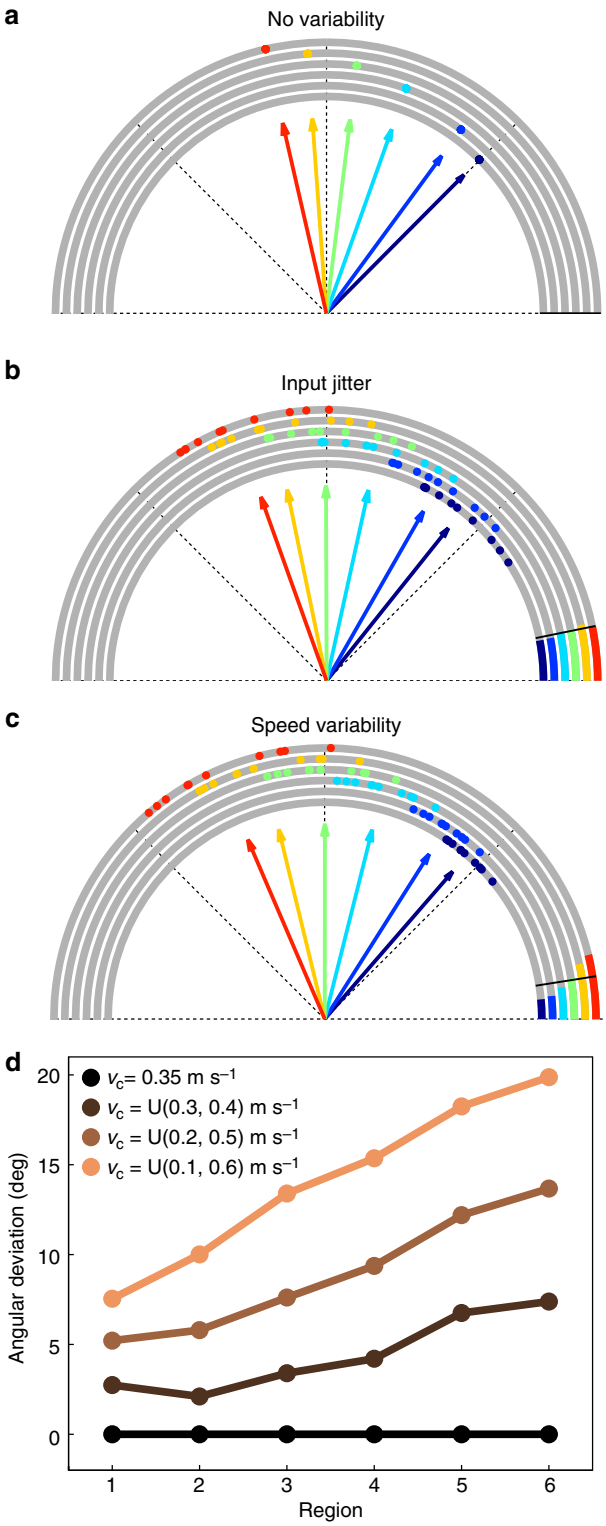

Figure 8 | Input jitter explains the observed phase distribution, while speed variability does not. (**a**) Cross-trial phase distribution over analysis regions, plotted as in Fig. 6b, over 10 iterations. The cross-trial phase distributions have zero angular deviation, as no variability is present in the system. (**b**) Cross-trial phase distribution under 10-ms jitter in the centre time of the Gaussian spike packet input. The angular deviation remains constant across analysis regions. (**c**) Phase distribution under speed variability (uniform distribution ranging from 0.2 to 0.5 m s^{-1}). The angular deviation increases markedly across analysis regions. (**d**) The angular deviation across regions is plotted for several uniform distributions of v_c, with a range expanding around the centre speed 0.35 m s^{-1}.

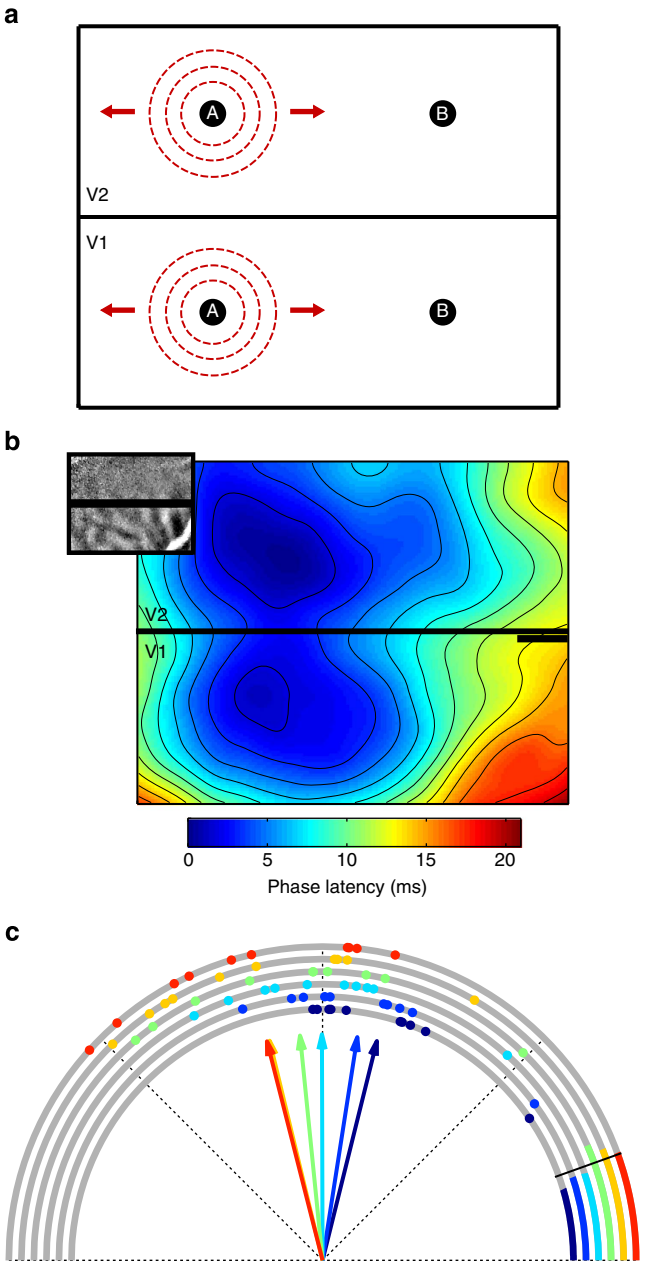

Figure 9 | Dual propagating waves in V1 and V2. (**a**) Schematic of dual stimulus-evoked propagations in V1 and V2. Simultaneous input at *A* evokes a propagating wave. Because the waves are carried by horizontal fibres (with nearly constant propagation velocity), the propagations will arrive in a synchronized fashion at *B*. (**b**) Phase-latency map at V1/V2 border, 50-ms stimulus condition, average over 10 trials. (Inset) ocular dominance map. (**c**) Cross-trial phase distribution calculated over analysis regions, as in Fig. 6b.

trials of the 50-ms stimulus presentation condition, a sharp, transient increase in the correlation coefficient between V1 and V2 is evident (black line, mean ± s.e.m., Fig. 10a), indicating that a similar pattern of phase develops in each cortical area during the wave propagation epoch. No such transient increases in the interarea correlation coefficient were observed during the matched blank trials (dark grey, Fig. 10a). Interestingly, shuffling the spatial position of each time course in the imaging array removes this transient correlation (light grey, Fig. 10a), indicating that it is indeed due to specific spatial patterns shared between the

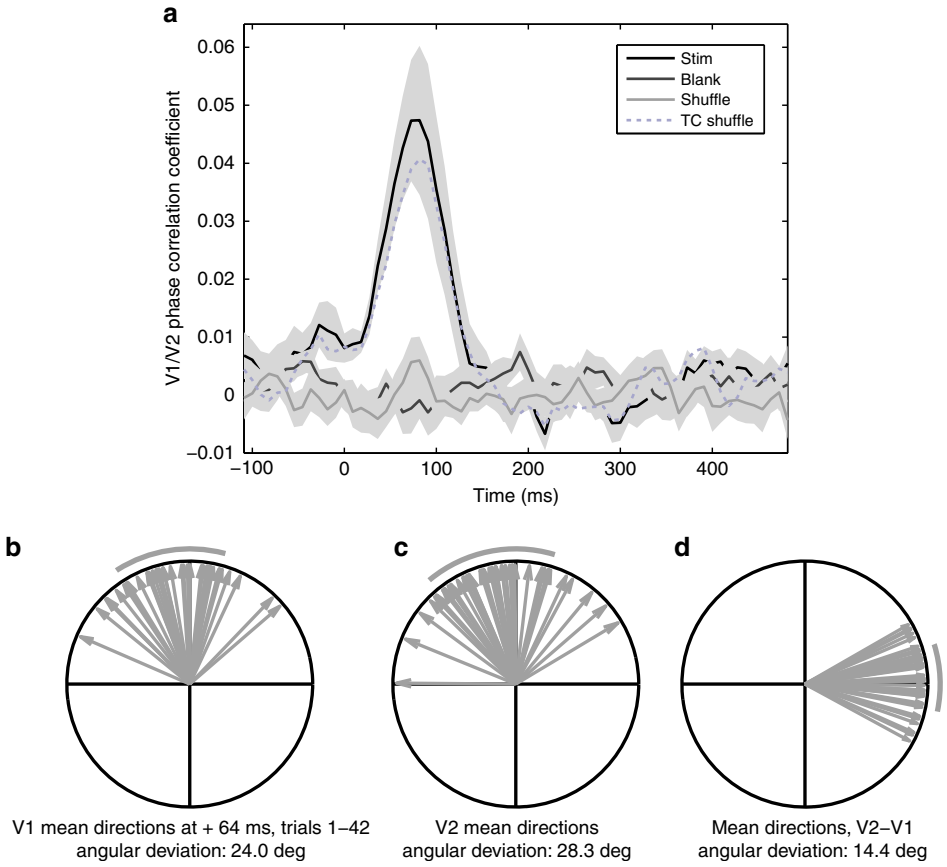

Figure 10 | Propagating waves in V1 and V2 are correlated but independent. (**a**) Interarea phase correlations, mean ± s.e.m. for stimulus (black), blank (dark grey), shuffled stimulus (light grey) and thalmocortical input control (dashed light blue) cases. (**b**) Mean direction distribution in V1 at + 64 ms for the 50 ms stimulus condition. (Grey curve) angular deviation of the cross-trial MD distribution. (**c**) V2 MD distribution and angular deviation. (**d**) The angular difference of the MD distributions in V1 and V2 exhibits an angular deviation that is moderately decreased, but non-vanishing.

two areas. Even more interestingly, when only the subset of time courses corresponding to the points nearest the thalamocortical input is shuffled (light blue, dashed, Fig. 10a), the transient increase remains nearly unchanged from the stimulus case, indicating that the spatial pattern of the phase responsible for this transient increase is diffuse, spread through a large extent of the individual cortical regions. Such diffuse, correlated phase patterns indicate a novel large-scale organization of the dynamics in these first two regions of the visual hierarchy, which presumably extends to areas higher in the visual hierarchy, and which could serve to organize computations in the visual system across space and time.

In light of this result, it is natural to ask further whether these dual, correlated waves in neighbouring cortical regions are in fact independent, or part of one continuous wave. To answer this question, we analysed to what extent the variability in the mean direction of the phase distribution in V2 is explained by the variability in V1 across trials by taking the circular difference of the MDs in each area (Fig. 10b–d). Both MD distributions in V1 and V2 have similar angular deviations (24.0 and 28.3 degrees, respectively), and similar phase values overall, leading to a distribution of the circular differences centred at zero radians (Fig. 10d), illustrating again the synchrony of these dual waves. Importantly, however, the angular deviation of this distribution of the differences between MDs is only moderately decreased relative to the angular deviation of the individual MD distributions in V1 and V2. This indicates that while part of the variability in V2 is explained by that of V1, the waves are not perfectly correlated, consistent with the interpretation that each wave is generated independently by the retinotopic feedforward volley, and then mediated by the horizontal fibre network of each cortical area in a predominantly monosynaptic fashion. Finally, it is important to specify that throughout this analysis, we did not observe a phase resetting, or second oscillation cycle, that would indicate the presence of a wave from either region crossing the border between the two. Thus, the propagating waves observed in both cortical areas imaged here seem to respect an anatomical border, in contrast to the wave dynamics observed in the visual cortex of the anaesthetised rat[10]. Whether this difference is due to variation among species or to cortical state will be an interesting topic for future research.

Discussion

While the existence of spontaneous and stimulus-evoked propagating waves during waking states in primary sensory cortices has been highly debated in recent years (see Sato *et al.*[44]; Muller and Destexhe[45] for review), our analysis shows for the first time that these waves occur systematically in the awake, behaving monkey, and that these waves are consistently evoked on the great majority of trials. Furthermore, by analysing the temporal variability in these propagations across trials, we observe that the intrinsic neocortical dynamics are invariant apart from a variable onset time for each propagation. Using a network model, we confirm that the most likely propagation mechanism—a predominantly monosynaptic propagation mediated by the horizontal fibre network of the superficial cortical layers—produces

trial-invariant dynamics as observed in the experimental data. In addition, we also find that the retinotopic feedforward stream, which activates primary and secondary visual cortex in rapid succession[42,43], with short, fixed latency differences[46], evokes a wave in V2 which then propagates nearly simultaneously with the wave in V1. These dual propagating waves induce in the two cortical areas correlated spatial patterns of phase, which establish a consistent spatiotemporal frame for further neuronal interactions, possibly including those in regions downstream in the visual hierarchy.

Previous VSD imaging studies in primate visual cortex worked primarily on the level of trial-averaged data[5,23,25,26], where the transient dynamics of the single-trial VSD response are heavily attenuated (see Fig. 1), and either did not specifically study the spatiotemporal form of the population response[5,23], or observed synchronous dynamics early in the response (attributed to a gain-control mechanism)[25]. Here, working always on the single-trial level, we observe a consistent offset across space throughout the response (Fig. 4d) and a trial-invariant dynamics indicative of a propagation mechanism that remains static from trial to trial (Figs 6–8). Together, these results conclusively demonstrate that the fast component of the stimulus-evoked VSD response is indeed a propagating wave mediated by the horizontal fibre network of the superficial cortical layers.

To observe and characterize propagating waves in a primary sensory cortex of the awake animal, we have introduced a phase-based method representing the first step towards automated, nonparametric identification of arbitrarily shaped propagating waves in noisy multichannel data. Similar approaches have either restricted detection possibilities to plane wave phenomena[16], or used a template matching approach in the amplitude domain[47], which precludes the generality of the method and hinders its robustness to noise. Furthermore, this method is able to capture the transient, trial-variable single-cycle propagations observed in these data, which is difficult with methods based on the phase of Fourier components[9], or spike-triggered averaging[48]. It is not to be claimed, however, that this method offers a fully automated approach at this stage; for example, the results in this study were fully and meticulously checked by visual inspection. With a complete characterization of the method on well-controlled surrogate data, however, fully automated detection of propagations in noisy data could soon become possible.

It is thought that the majority of the VSD signal is due to dendritic activity of superficial neurons, with minor contributions due to spiking activity and to deep layers[49]. The waves detected in the VSD signal in this study thus correspond to transient depolarization of excitatory and inhibitory neurons travelling across the cortex, carried primarily by the horizontal fibre network of the superficial cortical layers[35,36]. Such transient depolarizations will affect the processing of future stimuli at distant points in visual cortex[6,7,26], as ascending input will combine with the travelling depolarization at a correctly timed interval, changing the spiking response. Furthermore, because neurons during waking, 'activated' cortical states sit a few millivolts below threshold[50–52], operating in a fluctuation-driven regime[53–55], it is reasonable to expect that such input fluctuations will transiently change the spiking probability in the local circuit. Such possibilities will be the subject of future experimental and theoretical work. Finally, note that the phase latency correlations with distance observed in these data are significant, but not overly strong (see Fig. 4b). Whether this is due to the noise inherent in VSD imaging on the single-trial level, or to the noisy, asynchronous-irregular type background activity that is the hallmark of the awake state[51], remains to be addressed in future imaging studies, with continually improved signal-to-noise ratios.

The horizontal fibre network has previously been implicated in active computational roles[6,7,26,56–59]. By proving here that the population response in awake monkey V1 and V2 is indeed a propagating wave, we implicate the horizontal fibre network in a much more specific function. Previous work has discussed the possibility that waves propagating across cortical maps could label sensory inputs with a unique phase[21], and in this work, we have provided a conclusive experimental demonstration of the plausibility for such a computational role. In particular, the discovery that independent propagating waves in V1 and V2 maintain precise phase relations shows that, at least for these two areas, this coordinated activity will influence, and perhaps actively participate in, cortical computations. In light of these results, we suggest that the existence of propagating waves in awake, conscious states expands the basic function of cortical maps, as these spatiotemporal dynamics will animate these maps with well-defined phase relations both within and across regions.

Methods

Surgical preparation. Experiments were conducted on two male rhesus monkeys (*macaca mulatta*, aged 5 and 8 years old). The monkeys were chronically implanted with a head-holder and a recording chamber located above the cortical areas V1 and V2 of the right hemisphere. The dura was surgically removed over a surface corresponding to the recording aperture (18 mm diameter) and a silicon-made artificial dura was inserted under aseptic conditions. Before each recording session, the cortex was stained with the VSD RH-1691 (Optical Imaging). For this, the optical chamber was opened, the artificial dura removed and the cortical surface was cleaned under strictly sterile conditions[60]. The dye solution was prepared in artificial cerebrospinal fluid at a concentration of 0.2 mg ml^{-1}, and filtered through a 0.2-μm filter. The recording chamber was filled with this solution and closed for 3 h. The chamber was then rinsed thoroughly with filtered artificial cerebrospinal fluid to wash off any supernatant dye, the artificial dura was placed back in position and the chamber was filled with transparent agar before closing with a transparent cover.

During the recordings, the cortex was illuminated at 630 nm to excite the dye for 600–1,000 ms. The electromagnetic coil technique was used to record eye movements[61]. Briefly, a search coil (Skalar BNV, Holland) was inserted below the ocular sclera, and the subject was exposed to two magnetic fields alternating at different frequencies. The demodulation (DNI) of the induced voltage in the coil at these two frequencies then allows simultaneous measurement of horizontal and vertical eye movements.

Experimental protocols have been approved by the Marseille Ethical Committee in Neuroscience (approval #A10/01/13, official national registration #71-French Ministry of Research). All procedures complied with the French and European regulations for animal research, as well as the guidelines from the Society for Neuroscience.

VSD imaging protocol. Experimental controls and online eye position monitoring were performed by a PC running the REX software (NEI-NIH) with the QNX operating system[62]. Optical signals were recorded with a Dalstar camera (512 × 512 pixels, 110 Hz frame rate) driven by the Imager 3001 system (Optical Imaging Ltd.). Both online behavioural control and image acquisition were heartbeat regulated. Heartbeat was detected with a pulse oximeter (Nonin 8600V). The visual stimuli were computed online using VSG2/5 libraries and were displayed on a 22 inch CRT monitor at a resolution of 1,024 × 768 pixels. Refresh rate was set to 100 Hz. Viewing distance was 57 cm. Luminance values were linearized by means of a look-up table.

Visual stimuli. During a single trial, the monkey had to fixate on a central red dot for 1–2 s. The animal's gaze was constrained in a window of 2° × 2°. Stimuli were presented during fixation, and a reward (water drop) was given after the trial if the monkey maintained fixation during the acquisition period. We later verified that the s.d. of eye position during stimulation was 0.32° and 0.28° in the horizontal and vertical axes, respectively. Each trial ran for 510–1,100 ms: 100-ms delay, 10–600-ms stimulation period, 400-ms blank period. Stimuli were local Gaussian blobs with a s.d. of 0.5° in space. Stimuli were presented at 0.5° on the left of the vertical meridian and 3.5° below the horizontal, except where noted (Fig. 4c). Four different durations were used: 10, 50, 100 ms (monkey WA), and 600 ms (monkey GR).

Linear-model denoising method. Stacks of images were stored on hard drives for offline analysis. Preprocessing was carried out with Matlab R2009a (The MathWorks) using the Optimization, Statistics and Signal Processing toolboxes. Data were preprocessed using a linear model-based denoising method as described previously[28]. A physically motivated set of basis vectors was designed by

characterizing each source of signal and noise in a series of control experiments. For individual trials, the raw signal was decomposed across these basis vectors. Data were denoised by removing the components that are not linked to the evoked response in each individual trial, including physiological artifacts, environmental noise, and dye bleaching. Data were *Z*-scored via standard methods for further analysis.

Temporal filtering. Temporal filtering was carried out with an 8*th*-order digital Butterworth filter, forward–reverse in time to prevent phase distortion (see MATLAB function *filtfilt*). The cutoff frequencies were 5 and 20 Hz, unless otherwise noted. All results were checked with multiple cutoff frequencies to ensure against parametric sensitivity.

Phase-latency maps. In the analytic signal framework, a real-valued, narrowband time series is transformed into a phasor rotating in the complex plane. After selecting an appropriate temporal point within the oscillation cycle and before the phase crossing threshold (+ 72.7 ms for all panels in Fig. 4), we look for the temporal latency to a given phase crossing (throughout this work, $0/2\pi$) at each channel in the unsmoothed, temporally filtered data. Results were checked with several phase crossing values to ensure robustness. Exact times for between-sample $0/2\pi$ phase crossings were calculated by linear interpolation based on the instantaneous frequency of the signal between the two nearest samples. Because the instantaneous frequencies $(d\phi/dt)$ of individual channels may vary across the tested epoch, tracking phase crossing latency (in absolute time) is the most direct option to precisely compare phase offsets. All phase-latency maps are smoothed via a nonparametric, automated algorithm before averaging across trials for purposes of visualization[63]. For further calculations on the phase-latency maps, however, no smoothing is used.

Throughout the analysis, several checks were made to ensure the validity of the analytic signal representation. Reconstruction error for this representation was verified to be negligible[64]. Continuity of the phase variable was assessed via the distribution of angular acceleration values $\left(\frac{rad}{s^2}\right)$, to ensure against the presence of 'phase slips', or sudden discontinuities.

Phase derivatives and instantaneous frequency. The instantaneous frequency of an analytic signal is the time derivative of phase $\left(\frac{d\phi}{dt}\right)$, and is normally calculated by numerical differentiation of the signal instantaneous phase. In this work, all phase derivatives were recast as complex multiplications. Let x_n be a discretely sampled real signal, and let $X_n = x_n + j\hat{x}_n$ be its analytic representation. The product of X_n and its complex conjugate at the next sample X^*_{n+1} produces a new complex sequence whose angle is equal to the instantaneous frequency of the original[30]:

$$\Delta\phi_n = \arg(X_n X^*_{n+1}) \tag{5}$$

In this way, we can estimate the signal instantaneous frequency without phase unwrapping.

Phase latency correlation with distance. With the raw phase latencies calculated, we then find the position of the minimum of the smoothed phase-latency map, and calculate the Euclidean distance of each pixel from this point. We then compute the linear correlation coefficient of phase latency with distance (ρ_d), as part of the assessment of a significant propagation on the imaging array at the given temporal point.

Phase gradient. To compute coherent direction maps of the phase gradient, it was necessary to consider only the low spatial frequency components of the signal, by multiplication of the signal at each time slice with a Gaussian in spatial frequency domain (with 0.53 cyc mm^{-1} s.d.); however, for the phase gradient magnitude, no spatial smoothing was used. For convenience in measuring signal against the high-frequency spatial noise, we use the reciprocal of this measure (pixel brightness and contour plot, Fig. 6a), rendering the higher spatial frequencies as short wavevector lengths.

Spiking network model. The network model considered here is composed of leaky integrate-and-fire (IF) neurons (75% excitatory, 25% inhibitory) with conductance-based synapses. The network is topographically arranged in one dimension, with length L, and wrapped onto the circle, to avoid boundary effects. As in previous work[45,65], the network has local connections defined by a Gaussian spatial profile:

$$p_{ij} = Ae^{-\frac{\ell_{ij}^2}{2\sigma_c^2}} \tag{6}$$

where A scales the distribution to the connection probability (p_c), ℓ_{ij} is the distance between individual neurons, and σ_c is the standard deviation (or spatial spread) of the connectivity profile. The axonal conduction delay between two neurons grows linearly with distance:

$$d_{ij} = d_{syn} + \frac{\ell_{ij}}{v_c} \tag{7}$$

where d_{syn} is the minimum synaptic delay, and v_c is the axonal conduction velocity[6]. For these simulations, we used the PyNN interface[66] to the NEST simulation environment[67], with a standard IF neuron model and conductance-based alpha-function synaptic inputs (*IF_cond_alpha*). The relevant parameters for the cells and the network are given in Supplementary Table 1. All other parameters for the model neurons may be found in the PyNN documentation (http://www.neuralensemble.org/docs/PyNN/standardmodels.html).

To deliver a stimulus to the middle 100 excitatory neurons in the network, we used the Gaussian pulse packet module in NEST (*pulsepacket_generator*), which distributes N_{sp} spikes to each target neuron, following a Gaussian distribution in time, with parameters μ_{sp} and σ_{sp}. As mentioned in the main text, the pulse packet stimulus ignites a wave of subthreshold activity, that is, only 10–20% of spikes in the network occur outside the stimulation zone.

To analyse the results of the simulations, we selected six equal, symmetric regions based on the distance from the edge of the stimulation zone, average the V_m waveforms over these regions, lowpass filter (4th-order Butterworth, 100 Hz cutoff frequency), and put the resulting waveform into its analytic signal representation. The cross-trial phase distributions were then constructed from the signal phase at the temporal point + 115 ms into the simulation. To facilitate comparison between model and experiment, the measured phase values were then rotated by $\pi/2$ and multiplied by a complex exponential term to have the same frequency sign as in the experimental data. Note that relative phase values were unchanged by this procedure. Additional checks were performed at each step to ensure the resulting signal is free from discontinuities (for example, the spike reset of the IF neuron), and is well-captured by the final analytic signal representation, so that the comparison between phase variables in the experimental data and the model is well controlled. Finally, additional analysis was performed on the phase distribution of the synaptic conductance variables, which have no such reset discontinuity, to confirm the results of the phase distribution analysis in the model.

Circular correlation coefficient. To quantify the phase synchronization between neighbouring cortical areas in our VSD data, we computed the pixel-by-pixel circular correlation coefficient[40,68,69]. For two angular sequences α_i and β_i of length N with angular means $\bar{\alpha}$ and $\bar{\beta}$, the circular correlation coefficient (ρ_{cc}) is:

$$\rho_{cc} = \frac{\sum_i \sin(\alpha_i - \bar{\alpha})\sin(\beta_i - \bar{\beta})}{\sqrt{\sum_i \sin^2(\alpha_i - \bar{\alpha})\sin^2(\beta_i - \bar{\beta})}} \tag{8}$$

To compute the circular correlation coefficient at each time point in these data, we first selected rectangular regions of interest with an equal number of channels. After putting the time courses in each area into their analytic signal representation, we flipped the V2 map vertically and computed the circular correlation coefficient on the phase angles of the aligned pixel pairs in each region. This computation thus estimates the correlation between the spatial pattern of phase in each area. To exclude the regions receiving feedforward input for the thalamocortical input control (light blue dashed line, Fig. 10a), we selected a region of low-phase latency in the V1 trial-averaged phase-latency map (from + 35 to + 40 ms), and shuffled the channels within in this region only (15.5% of the V1 ROI). In all cases, image sequences were confirmed to be well registered by manual inspection of both the amplitude and phase maps of the system. Note finally that all means and standard errors of the mean (s.e.m.) were obtained via Fisher *z*-transform.

Analytic signal representation of the gaussian pulse. As a control on the spatiotemporal form of the population response in visual cortex, we used a simple model of space timeseparable dynamics, termed 'Gaussian pulse' in the text. Again, we write this as a separable function of space and time,

$$f(x,t) = g(x)h(t) \tag{9}$$

where $g(x)$ represents a Gaussian profile in space,

$$g(x) = Ae^{-\frac{x^2}{2\sigma^2}} \tag{10}$$

and $h(t)$ is a sinusoidal response time course,

$$h(t) = \frac{1}{2}[1 + \cos(\omega t)] \tag{11}$$

with amplitude A, spatial spread σ, and angular frequency ω. We may then take the Hilbert transform of $f(x,t)$,

$$H[f(x,t)] = H[g(x)h(t)] \tag{12}$$

$$= g(x)H[h(t)] \tag{13}$$

and evaluate $H[h(t)]$,

$$H[h(t)] = H\left[\frac{1}{2}\cos(\omega t) + \frac{1}{2}\right] \quad (14)$$

$$= H\left[\frac{1}{2}\cos(\omega t)\right] + H\left[\frac{1}{2}\right] \quad (15)$$

$$= H\left[\frac{1}{2}\cos(\omega t)\right] \quad (16)$$

$$= \frac{1}{2}[sgn(\omega)\sin(\omega t)] \quad (17)$$

$$= \frac{1}{2}\sin(\omega t) \quad (18)$$

Now, forming the analytic signal $z(x, t)$,

$$z(x,t) = f(x,t) + j\hat{f}(x,t) \quad (19)$$

$$= \frac{A}{2}e^{-\frac{x^2}{2\sigma^2}}[1+\cos(\omega t)] + j\frac{A}{2}e^{-\frac{x^2}{2\sigma^2}}\sin(\omega t) \quad (20)$$

$$= \frac{A}{2}e^{-\frac{x^2}{2\sigma^2}}[1+\cos(\omega t) + j\sin(\omega t)] \quad (21)$$

$$= \frac{A}{2}e^{-\frac{x^2}{2\sigma^2}}\left[1+e^{j\omega t}\right] \quad (22)$$

which has the form of equation (3) in the main text. Note that an arbitrary phase shift can be added directly into the complex exponential term in equation (22), although this was not written explicitly here.

References

1. Ribary, U. *et al.* Magnetic field tomography of coherent thalamocortical 40-Hz oscillations in humans. *Proc. Natl Acad. Sci. USA* **88,** 11037–11041 (1991).
2. Chavane, F. *et al.* The visual cortical association field: a Gestalt concept or a psychophysiological entity? *J. Physiol. Paris* **94,** 333–342 (2000).
3. Maynard, E., Nordhausen, C. & Normann, R. The Utah intracortical electrode array: a recording structure for potential brain-computer interfaces. *Electroencephalogr. Clin. Neurophysiol.* **102,** 228–239 (1997).
4. Grinvald, A. & Hildesheim, R. VSDI: a new era in functional imaging of cortical dynamics. *Nat. Rev. Neurosci.* **5,** 874–885 (2004).
5. Grinvald, A., Lieke, E., Frostig, R. & Hilesheim, R. Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *J. Neurosci.* **14,** 2545–2568 (1994).
6. Bringuier, V., Chavane, F., Glaeser, L. & Frégnac, Y. Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. *Science* **283,** 695–699 (1999).
7. Jancke, D., Chavane, F., Naaman, S. & Grinvald, A. Imaging cortical correlates of illusion in early visual cortex. *Nature* **428,** 423–426 (2004).
8. Roland, P. *et al.* Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. *Proc. Natl Acad. Sci. USA* **103,** 12586–12591 (2006).
9. Benucci, A., Frazor, R. & Carandini, M. Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron* **55,** 103–117 (2007).
10. Xu, W., Huang, X., Takagaki, K. & Wu, J. Compression and reflection of visually evoked cortical waves. *Neuron* **55,** 119–129 (2007).
11. Han, F., Caporale, N. & Dan, Y. Reverberation of recent visual experience in spontaneous cortical waves. *Neuron* **60,** 321–327 (2008).
12. Huang, X. *et al.* Spiral wave dynamics in neocortex. *Neuron* **68,** 978–990 (2010).
13. Gao, X. *et al.* Interactions between two propagating waves in rat visual cortex. *Neuroscience* **216,** 57–59 (2012).
14. Stroh, A. *et al.* Making waves: initiation and propagation of corticothalamic Ca2+ waves *in vivo*. *Neuron* **77,** 1136–1150 (2013).
15. Gabriel, A. & Eckhorn, R. A multi-channel correlation method detects traveling γ-waves in monkey visual cortex. *J. Neurosci. Methods* **131,** 171–184 (2003).
16. Rubino, D., Robbins, K. & Hatsopoulos, N. Propagating waves mediate information transfer in the motor cortex. *Nat. Neurosci.* **9,** 1549–1557 (2006).
17. Witte, R., Rousche, P. & Kipke, D. Fast wave propagation in auditory cortex of an awake cat using a chronic microelectrode array. *J. Neural. Eng.* **4,** 68–78 (2007).
18. Ferezou, I. *et al.* Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* **56,** 907–923 (2007).
19. Mohajerani, M., McVea, D., Fingas, M. & Murphy, T. Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *J. Neurosci.* **30,** 3745–3751 (2010).
20. Mohajerani, M. *et al.* Spontaneous cortical activity alternates between motifs defined by regional axonal projections. *Nat. Neurosci.* **16,** 1426–1435 (2013).
21. Ermentrout, G. & Kleinfeld, D. Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. *Neuron* **29,** 33–44 (2001).
22. Shoham, D. *et al.* Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes. *Neuron* **24,** 791–802 (1999).
23. Slovin, H., Arieli, A., Hildesheim, R. & Grinvald, A. Long-term voltage-sensitive dye imaging reveals cortical dynamics in behaving monkeys. *J. Neurophysiol.* **88,** 3421–3428 (2002).
24. Chen, Y., Geisler, W. & Seidemann, E. Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat. Neurosci.* **9,** 1412–1420 (2006).
25. Sit, Y., Chen, Y., Geisler, W., Miikkulainen, R. & Seidemann, E. Complex dynamics of V1 population responses explained by a simple gain-control model. *Neuron* **64,** 943–956 (2009).
26. Reynaud, A., Masson, G. & Chavane, F. Dynamics of local input normalization result from balanced short- and long-range interactions in area V1. *J. Neurosci.* **32,** 12558–12569 (2012).
27. Alexander, D. M. *et al.* Traveling waves and trial averaging: The nature of single-trial and averaged brain responses in large-scale cortical signals. *Neuroimage* **73,** 95–112 (2013).
28. Reynaud, A., Takerkart, S., Masson, G. & Chavane, F. Linear model decomposition for voltage-sensitive dye imaging signals: application in awake behaving monkey. *Neuroimage* **54,** 1196–1210 (2011).
29. Gabor, D. Theory of communication. Part 1: the analysis of information. *J. Inst. Elec. Eng. –Part 3* **93,** 429–441 (1946).
30. Feldman, M. *Hilbert Transform Applications in Mechanical Vibration* 1st edn (Wiley, 2011).
31. González-Burgos, G., Barrionuevo, G. & Lewis, D. Horizontal synaptic connections in monkey prefrontal cortex: an *in vitro* electrophysiological study. *Cereb. Cortex.* **10,** 82–92 (2000).
32. Girard, P., Hupé, J. & Bullier, J. Feedforward and feedback connections between areas V1 and V2 of the monkey have similar rapid conduction velocities. *J. Neurophysiol.* **85,** 1328–1331 (2001).
33. Prechtl, J., Cohen, L., Pesaran, B., Mitra, P. & Kleinfeld, D. Visual stimuli induce waves of electrical activity in turtle cortex. *Proc. Natl Acad. Sci. USA* **94,** 7621–7626 (1997).
34. Mitra, P. & Pesaran, B. Analysis of dynamic brain imaging data. *Biophys. J.* **76,** 691–708 (1999).
35. Stettler, D., Das, A., Bennett, J. & Gilbert, C. Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* **36,** 739–750 (2002).
36. Angelucci, A. *et al.* Circuits for local and global signal integration in primary visual cortex. *J. Neurosci.* **22,** 8633–8646 (2002).
37. Dow, B., Snyder, A., Vautin, R. & Bauer, R. Magnification factor and receptive field size in foveal striate cortex of the monkey. *Exp. Brain. Res.* **44,** 213–228 (1981).
38. Tehovnik, E. J. & Slocum, W. M. What delay fields tell us about striate cortex. *J. Neurophysiol.* **98,** 559–576 (2007).
39. Van Essen, D. C., Newsome, W. T. & Maunsell, J. H. The visual field representation in striate cortex of the macaque monkey: asymmetries, anisotropies, and individual variability. *Vis. Res.* **24,** 429–448 (1984).
40. Fisher, N. *Statistical Analysis of Circular Data* 1st edn (Cambridge University Press, 1993).
41. Waxman, S. & Bennett, M. Relative conduction velocities of small myelinated and non-myelinated fibres in the central nervous system. *Nat. New. Biol.* **238,** 217–219 (1972).
42. Kennedy, H. & Bullier, J. A double-labeling investigation of the afferent connectivity to cortical areas V1 and V2 of the macaque monkey. *J. Neurosci.* **5,** 2815–2830 (1985).
43. Ayzenshtat, I. *et al.* Precise spatiotemporal patterns among visual cortical areas and their relation to visual stimulus processing. *J. Neurosci.* **30,** 11232–11245 (2010).
44. Sato, T., Nauhaus, I. & Carandini, M. Traveling waves in visual cortex. *Neuron* **75,** 218–229 (2012).
45. Muller, L. & Destexhe, A. Propagating waves in thalamus, cortex and the thalamocortical system: experiments and models. *J. Physiol. Paris* **106,** 222–238 (2012).
46. Bullier, J. Integrated model of visual processing. *Brain. Res. Rev.* **36,** 96–107 (2001).
47. Takagaki, K., Zhang, C., Wu, J. & Ohl, F. Flow detection of propagating waves with temporospatial correlation of activity. *J. Neurosci. Methods* **200,** 207–218 (2011).
48. Nauhaus, I., Busse, L., Carandini, M. & Ringrach, D. Stimulus contrast modulates functional connectivity in visual cortex. *Nat. Neurosci.* **12,** 70–76 (2009).
49. Chemla, S. & Chavane, F. A biophysical cortical column model to study the multi-component origin of the VSDI signal. *Neuroimage* **53,** 420–438 (2010).

50. Destexhe, A. & Paré, D. Impact of network activity on the integrative properties of neocortical pyramidal neurons *in vivo*. *J. Neurophysiol.* **81,** 1531–1547 (1999).
51. Destexhe, A., Rudolph, M. & Paré, D. The high-conductance state of neocortical neurons *in vivo*. *Nat. Rev. Neurosci.* **4,** 739–751 (2003).
52. Poulet, J. & Petersen, C. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* **454,** 881–885 (2008).
53. Destexhe, A., Rudolph, M., Fellous, J. & Sejnowski, T. Fluctuating synaptic conductances recreate *in-vivo*-like activity in neocortical neurons. *Neuroscience* **107,** 13–24 (2001).
54. Kuhn, A., Aertsen, A. & Rotter, S. Neuronal integration of synaptic input in the fluctuation-driven regime. *J. Neurosci.* **24,** 2345–2356 (2004).
55. Rudolph, M., Pospischil, M., Timofeev, I. & Destexhe, A. Inhibition determines membrane potential dynamics and controls action potential generation in awake and sleeping cat cortex. *J. Neurosci.* **27,** 5280–5290 (2007).
56. Gilbert, C. Horizontal integration and cortical dynamics. *Neuron* **9,** 1–13 (1992).
57. Grossberg, S., Mingolla, E. & Ross, W. Visual brain and visual perception: how does the cortex do perceptual grouping? *Trends Neurosci.* **20,** 106–111 (1997).
58. Gilbert, C. Adult cortical dynamics. *Physiol. Rev.* **78,** 467–485 (1998).
59. Raizada, R. & Grossberg, S. Towards a theory of the laminar architecture of cerebral cortex: computational clues from the visual system. *Cereb. Cortex* **13,** 100–113 (2003).
60. Arieli, A., Grinvald, A. & Slovin, H. Dural substitute for long-term imaging of cortical activity in behaving monkeys and its clinical implications. *J. Neurosci. Methods* **114,** 119–133 (2002).
61. Robinson, D. A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans. Biomed. Eng.* **10,** 137–145 (1963).
62. Hays, A., Richmond, B. & Optican, L. A UNIX-based multiple process system for real-time data acquisition and control. *WESCON Conf. Proc.* **2,** 1–10 (1982).
63. Garcia, D. Robust smoothing of gridded data in one and higher dimensions with missing values. *Comput. Stat. Data Anal.* **54,** 1167–1178 (2010).
64. Goswami, J. & Hoefel, A. Algorithms for estimating instantaneous frequency. *Signal Proc.* **84,** 1423–1427 (2004).
65. Yger, P., Boustani, S. E., Destexhe, A. & Frégnac, Y. Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons. *J. Comput. Neurosci.* **31,** 229–245 (2011).
66. Davison, A. *et al.* PyNN: a common interface for neuronal network simulators. *Front. Neuroinform.* **2,** 1–10 (2009).
67. Gewaltig, M. & Diesmann, M. NEST (Neural Simulation Tool). *Scholarpedia* **2,** 1430 (2007).
68. Jammalakadaka, S. & Sengupta, A. *Topics in Circular Statistics* 1st edn (World Scientific, Singapore, 2001).
69. Berens, P. CircStat: a Matlab toolbox for circular statistics. *J. Stat. Softw.* **31** (2009).

Acknowledgements

We are grateful to T. Movshon and Y. Frégnac for helpful comments on the manuscript; I. Balansard for veterinarian assistance; S. Takerkart for help with the linear model; A. De Moya, X. De Giovanni, M. Martin and J. Baurberg for excellent technical support; and Q. Montardy for help with some experiments. This work was supported by CNRS, ANR (Complex-V1) and the European Community (FET grants FACETS FP6-015879 and BrainScaleS FP7-269921). L.M. is a PhD fellow from École des Neurosciences de Paris (ENP). A.R. was supported by FACETS and Fondation des Aveugles et Handicapés Visuels de France.

Author contributions

L.M., A.R., F.C. and A.D. designed the study; A.R. and F.C. designed and performed experiments; A.R. performed linear-model denoising; analyses and modelling were initially designed by L.M. and A.D., discussed by L.M., F.C. and A.D., and performed by L.M. All authors discussed the results and contributed to writing the paper. Experiments and preprocessing were performed at INT. Analysis and modelling were performed at UNIC. F.C. and A.D. jointly supervised this work.

Additional information

Supplementary Information accompanies this paper at http://www.nature.com/naturecommunications

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at http://npg.nature.com/reprintsandpermissions/

How to cite this article: Muller, L. *et al.* The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nat. Commun.* 5:3675 doi: 10.1038/ncomms4675 (2014).

Supplementary Information

Supplementary Movie 1. Single Trial VSD Response, Trial 1. Animated version of the VSD response presented in Figure 1c (left). As previously, individual frames are smoothed with a Gaussian (0.23 mm standard deviation), and scaled to the maximum z-score in the trial. Black lines indicate z-scores of 0, 1, and 2. Time after stimulus onset is indicated above the surface plot in all frames. One millimeter spatial scale bar depicted in red at bottom right.

Supplementary Movie 2. Single Trial VSD Response, Trial 2. Animated version of the corresponding surface plot presented in Figure 1c.

Supplementary Movie 3. Single Trial VSD Response, Trial 3. Animated version of the corresponding surface plot presented in Figure 1c.

Supplementary Movie 4. Trial-Averaged VSD Response. Animated version of the VSD response averaged over all trials in the 50 ms stimulus presentation condition, as in Figure 1c (right). Note the markedly different form of the response in comparison to that in the single-trials.

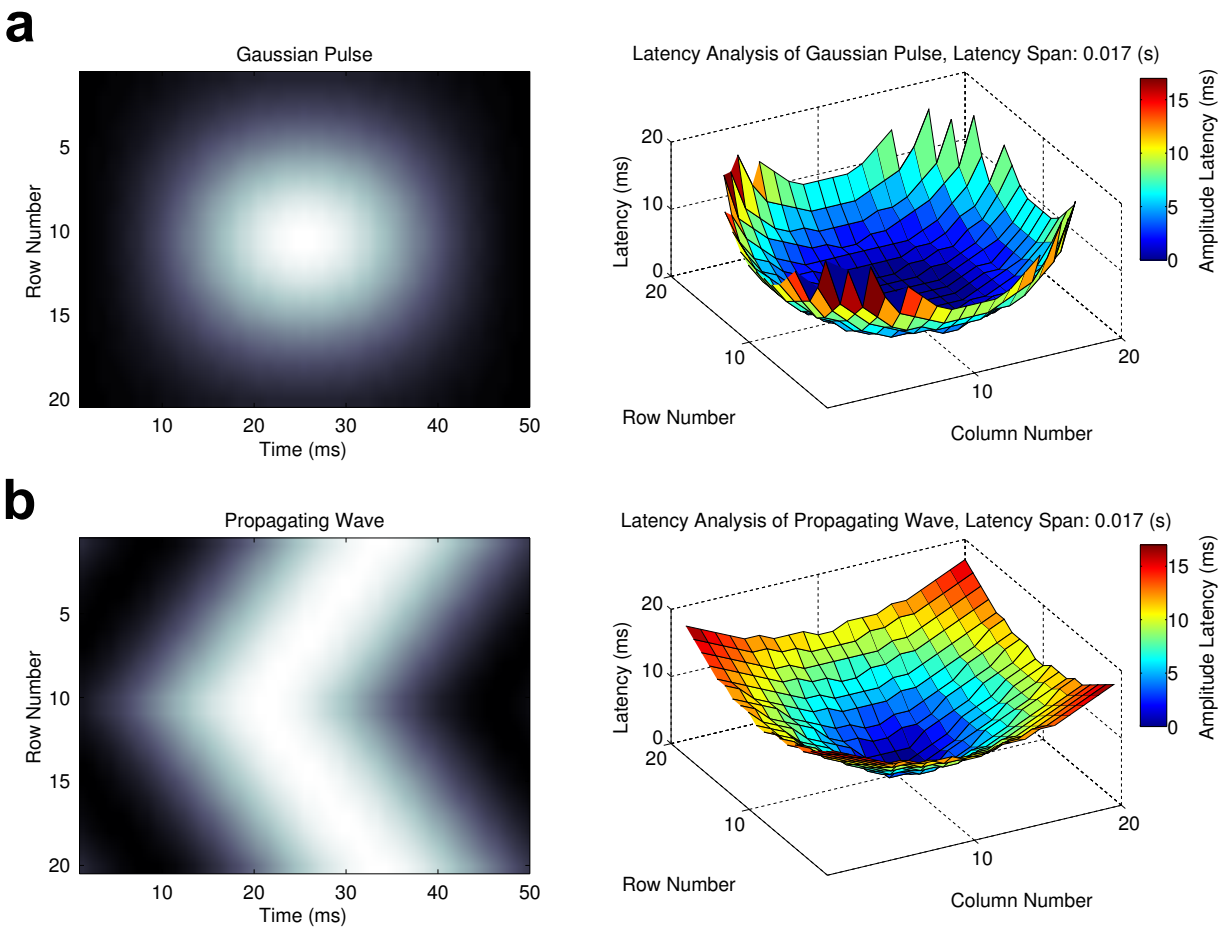

Supplementary Figure 1. Latency Analysis in the Amplitude Domain. Latency analysis in the amplitude domain is unable to distinguish between spatiotemporally separable and inseparable forms. Threshold-crossing latency analysis was conducted on a Gaussian pulse **(a)** and propagating wave **(b)**. Vertical sections of the data through time illustrate the temporal form of the response in both the spatiotemporally separable (Gaussian pulse) and inseparable (propagating wave) case (top panels). Even though the two synthetic datasets display markedly different spatiotemporal behavior, a latency analysis in the amplitude domain registers an increase in latency for each case, confounding the differences between the two. The correlation of latency with distance is strong in both cases ($\sim$ .8 - .9). Note that the shape of the latency basin recovered from this analysis depends on the form of the envelope (as illustrated by the difference in latency basin shape for the two models here), and that it is straightforward to specify a spatiotemporally separable model with exactly the same amplitude latency basin as the spatiotemporally inseparable propagating wave model in (b).

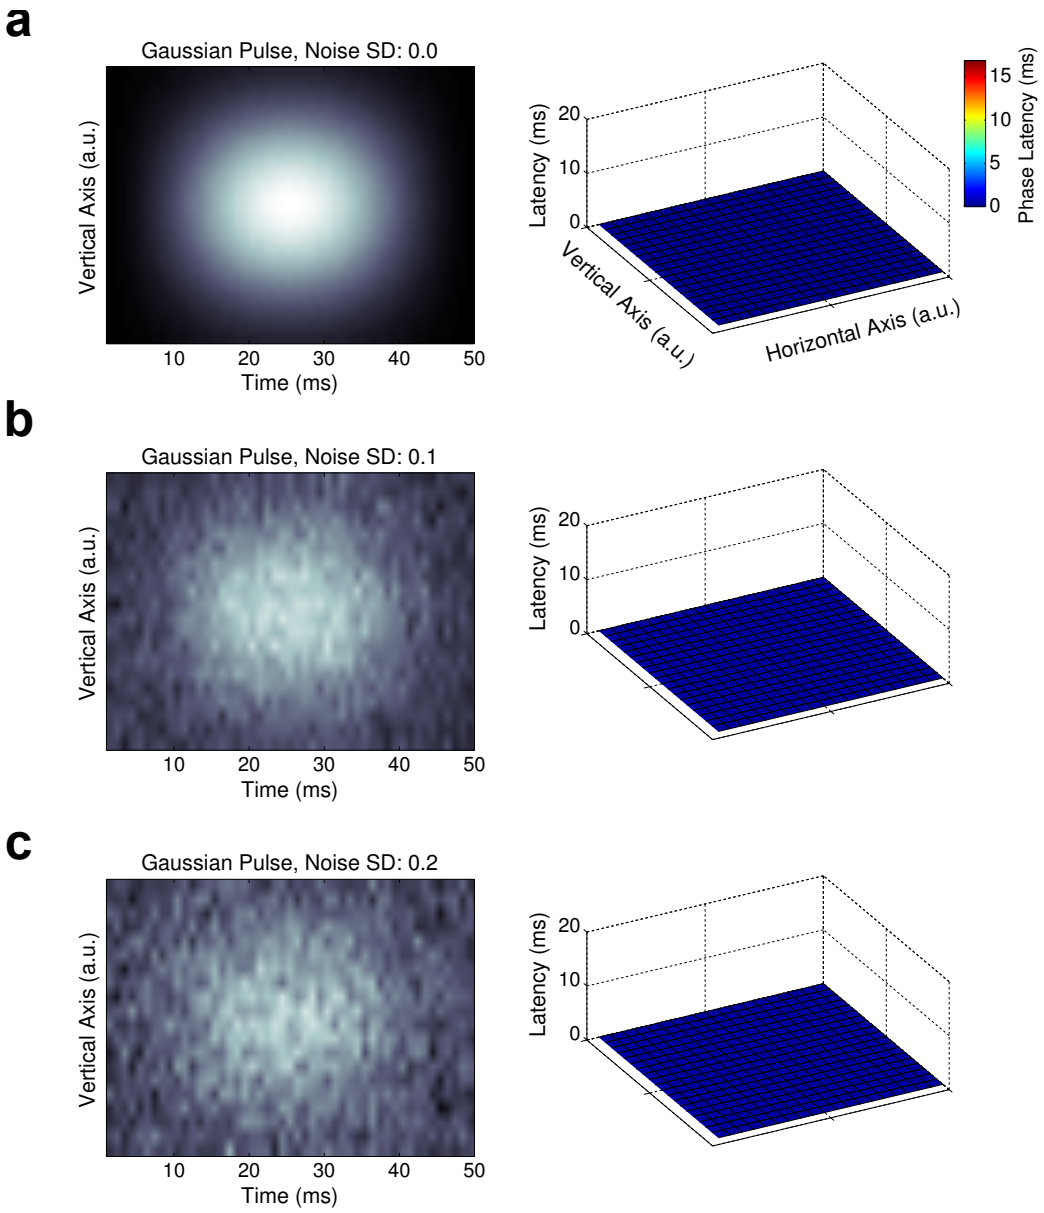

Supplementary Figure 2. Phase Latency Analysis Under Noise. To test against the possibility of false positive detections using the phase latency analysis, we started with a Gaussian pulse (with unitary amplitude) and added increasing levels of noise (SD 0, 0.1, 0.2, in panels **a**, **b**, and **c**, left column). The fact that increasing levels of noise have no effect on the phase latency maps extracted from the timecourses illustrates the robustness and applicability of the method for use in noisy multichannel data.

Parameter Class	Parameter	Value (Units)
Cell	τ_m	20.0 ms
	V_{rest}	-65.0 mV
	V_{reset}	-65.0 mV
	V_{thresh}	-50.0 mV
Synaptic	E_e	0.0 mV
	E_i	-70.0 mV
	g_e	1.8 nS
	g_i	18 nS
	τ_e	5.0 ms
	τ_i	5.0 ms
Network	N	1000
	N_e	750
	N_i	250
	p_c	0.1
	L	7.0 mm
	v_c	0.35 m/s
	K	100
	K_{exc}	75
	K_{inh}	25
	σ_c	1.0 mm
ODE	dt	0.1 ms
	T	300 ms

Supplementary Table 1. Parameters for Spiking Network Simulations. Parameters used for the simulations presented in Figures 7 and 8. Values for parameters not listed here can be found in the PyNN documentation (*IF_cond_alpha*).

Chapter 4

Investigation of spatial structure in LFP activity across the frequency spectrum: a multielectrode array study in human and monkey.

[Work in collaboration with Giacomo Benvenuti and Frédéric Chavane]

Utah array recordings

Recording activity using two-dimensional high-density MEAs offers the unique opportunity to study the spatiotemporal dynamics of the neocortex both in the human and monkey. In a recent paper employing MEA data in the human (Peyrache et al., 2012), extracellularly recorded units were categorized as excitatory or inhibitory during sleep, based on putative monosynaptic connections. This study also provided evidence for distinctive network dynamics for excitatory and inhibitory neurons during spontaneous activity in drowsiness and sleep.

In this work, we use the same database of human recordings, paired with MEA recordings from the primary visual cortex of the monkey (data from INT Marseille), to make an analysis of the spatiotemporal dynamics of the neocortex in both settings. In the data taken in the human (an example of which is shown in Fig. 4.1), the Utah array is positioned in human temporal cortex, and recordings are obtained at multiple scales. Electrocorticogram (ECoG) recordings are obtained using a grid of subdural electrodes (green traces, Fig. 4.1), while the 96 electrodes of the Utah array provide recording of the LFP (black traces, Fig. 4.1), as well as neuronal units on most electrodes, which could be discriminated into excitatory and inhibitory cells (blue and red dots, respectively).

Application of the method to multielectrode arrays

Having studied in **Work 2** the spatiotemporal dynamics of the "principal" component of the stimulus-evoked population response, we next aimed in further analyses to study the spatial organization of higher-frequency oscillations. As noted in Chapter 1, the spatial organization of high-frequency, stimulus-induced oscillatory activity

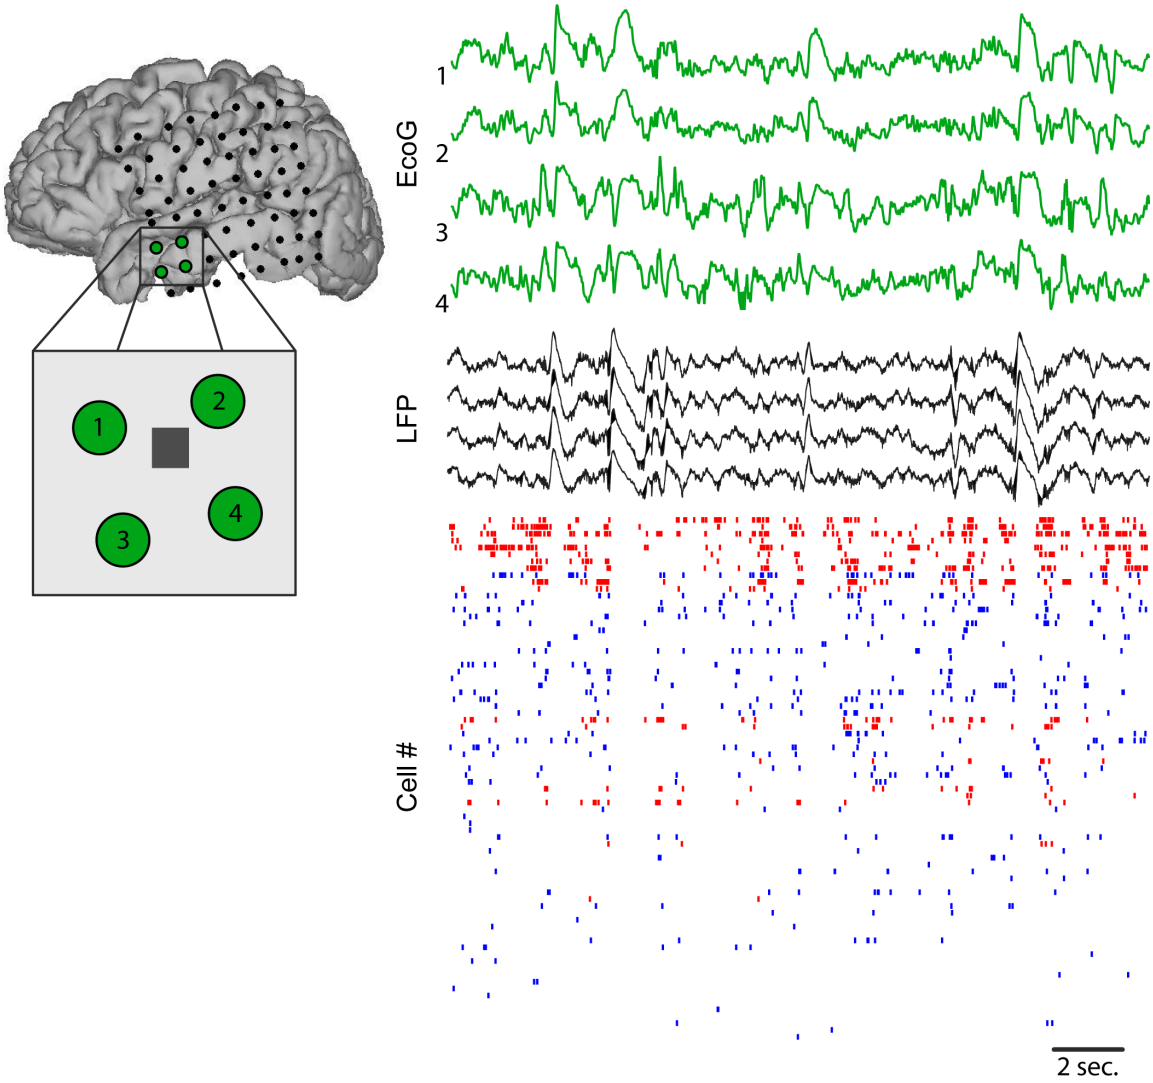

Fig. 4.1 Human recordings with Utah arrays. Left: Localization of subdural electrodes (green) and the Utah electrode array (dark gray square). Right: Recordings obtained at each scale. The green traces show electrocorticogram (ECoG) of the four closest contacts to the microelectrode (numbered from 1 to 4). The black traces show LFP. The raster plot shows the pooled firing of inhibitory (red) and excitatory (blue) cells for this period of slow-wave sleep. Modified from Peyrache et al. (2012).

has been studied previously in Gabriel and Eckhorn (2003) and Rubino et al. (2006), not to mention Roelfsema et al. (1997). In this analysis, however, we use our method to detect epochs of phase organization in the spontaneous ongoing activity of LFPs in data recorded by Utah arrays implanted in the middle temporal gyrus of the human and the primary visual cortex of the monkey.

These LFP data present an interesting application of our method, which was initially applied to VSD imaging data, because the analysis of MEA recordings represents quite a different signal processing regime. While VSD has a high spatial resolution (about 40

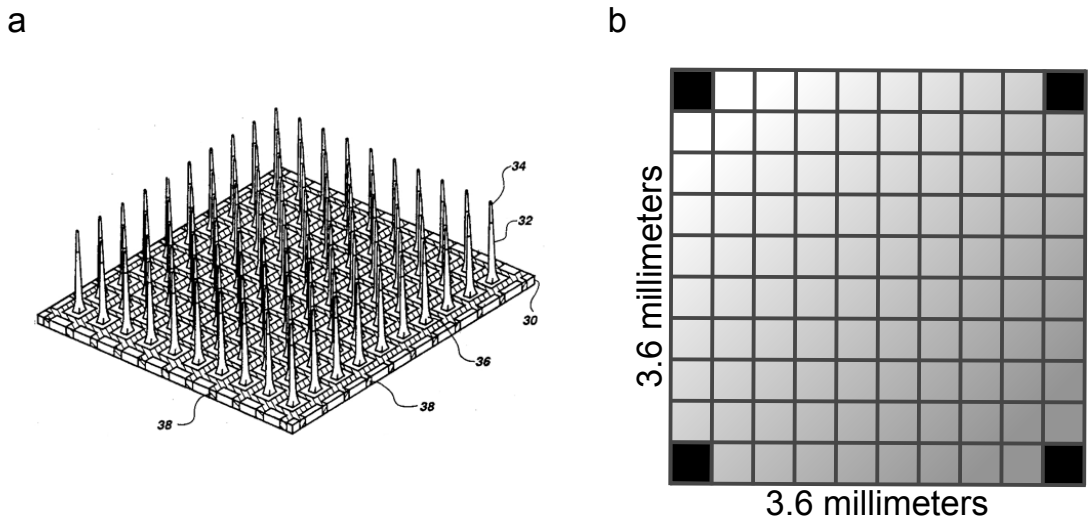

Fig. 4.2 The Utah multielectrode array. (a) Schematic representation of the Utah array, illustrating penetrating electrodes organized in a two-dimensional grid. Adapted from Campbell et al. (1993). **(b)** Further illustration depicting channels used to record LFP (in gray).

μm per pixel) but a relatively low temporal resolution (from 100 to 500 Hz), the 10x10 electrode Utah array has an interelectrode spacing of 400 μm, but a relatively high sampling rate in time (at 1 kHz or above). On the one hand, such a temporal resolution is quite advantageous for our analysis, which at the first step makes a precise quantification of the temporal dynamics on the level of single channels; on the other hand, however, the recording array covers a smaller spatial area (Fig. 4.2b; cf. 18 mm diameter for the recording aperture in **Work 2**), which admits the possibility of capturing the spatial extent waves observed on the array only in part. We thus adapted the questions we asked with the phase latency calculations to these long recordings of spontaneous LFP, to extract brief epochs of wave propagation on the array, where the spatial pattern of phase exhibits a significant level of organization (Fig. 4.3).

Extracting wave propagation epochs

To extract wave propagations in spontaneous background activity, as opposed to the stimulus-evoked case considered in **Work 2**, we apply the same method as in the previous work, and calculate the phase latency correlation with distance (ρ_d) on the array at each time sample.[1] This leads to a time series in ρ_d (Figure 4.4). By taking into account the full distribution of correlation values observed across the array during the

[1] Note that, as in **Work 2**, we use an 8th-order Butterworth filter, digital forward-reverse in time, to analyze the data within the selected passband.

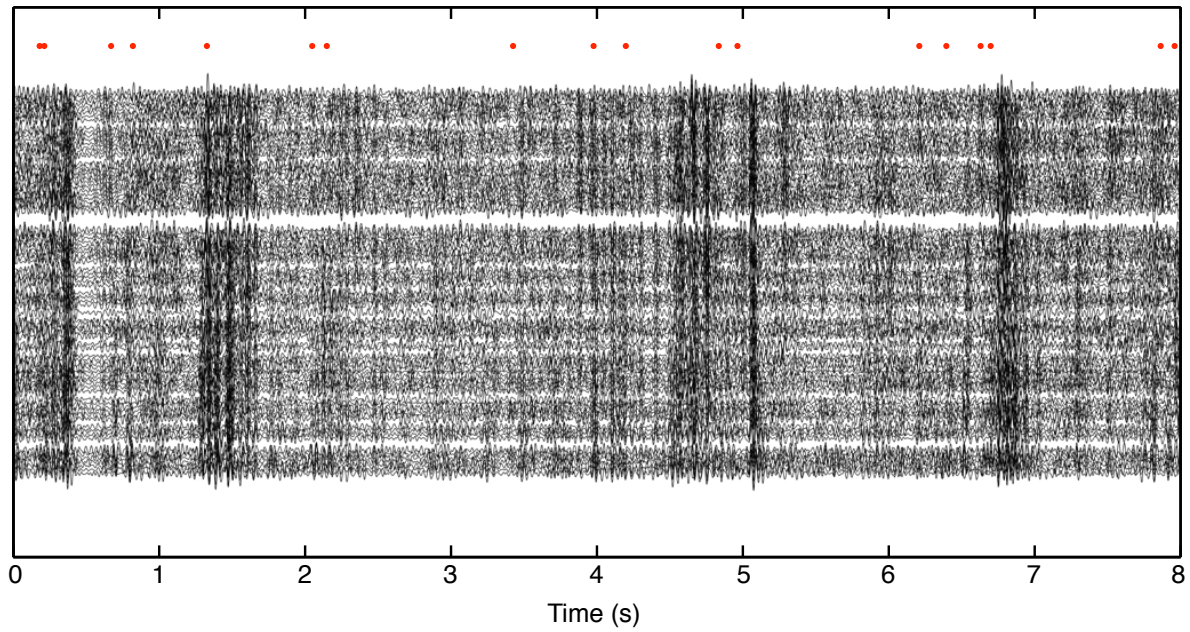

Fig. 4.3 Wave propagation epochs. Eight seconds of spontaneous LFP activity in the 20 - 50 Hz frequency band is displayed across the 96 recording channels of the MEA. Eighteen detected epochs of wave propagation are marked by red dots above the timecourses.

analysis period (Fig. 4.4, right), we can calculate the mean correlation coefficient of the distribution (horizontal black line, Fig. 4.4), its standard deviation, and the correlation value achieved by only 1% of values in the equivalent normal distribution ($\mu + 2.32\sigma$; dotted red line, Fig. 4.4). Phase latency correlation values crossing this threshold then mark epochs of wave propagation, as these values represent periods where the spatial organization of phase across the array is highest within the analyzed recording session. Note further that this method represents a quite conservative approach to detecting waves in the spontaneous background activity. For the first analysis within this spontaneous background activity, however, we feel that a more careful approach is judicious, while more exploratory techniques can be applied in later analyses.

Quantifying epochs – single events and population statistics

With this criterion for the critical threshold on ρ_d, we then analyzed 80 seconds of spontaneous LFP activity from Utah array data taken in the human and the monkey.[2] Within the spontaneous background activity, we detected 162 individual wave propagation epochs in the human, and 97 epochs in the monkey. These propagation epochs

[2] Note that the data from the primate serves as an important control in this analysis, because the human subjects in these recordings suffered from epilepsy (though, in the subject analyzed, the activity on the array did not show signs of epileptiform activity, and thus can be considered effectively normal).

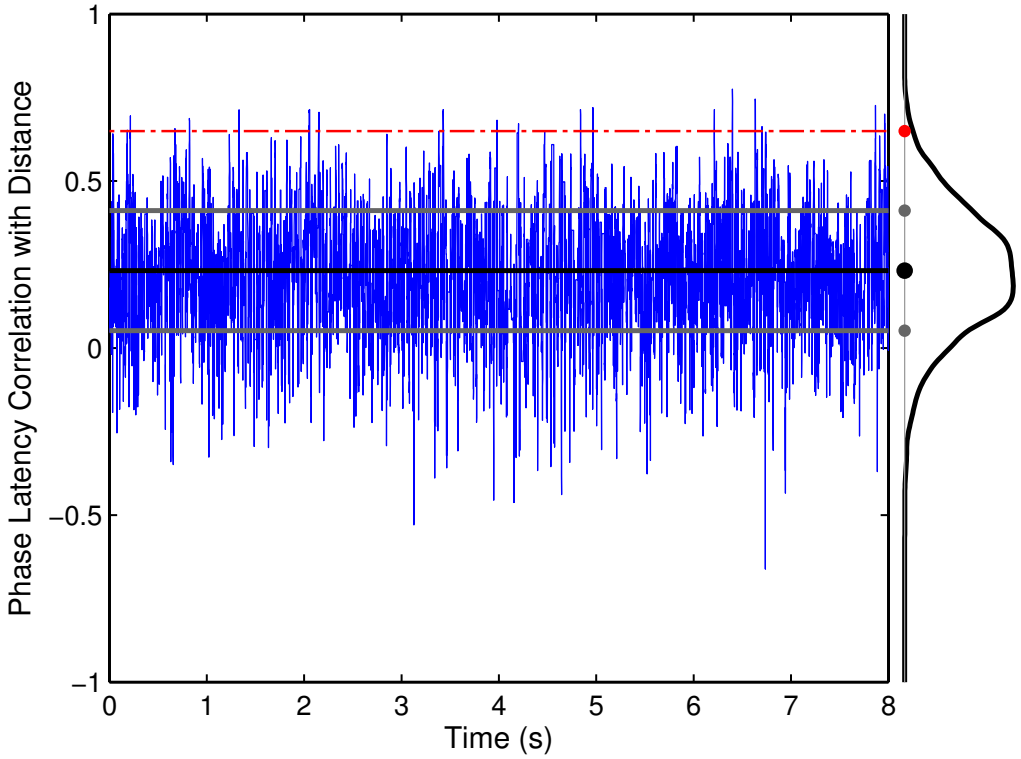

Fig. 4.4 Phase latency correlation time series. The phase latency correlation with distance (ρ_d), calculated at each sample during 8 seconds of spontaneous activity, is given as a time series (blue trace). The cumulative distribution of ρ_d is given at right, along with its mean (black dot), standard deviation (gray), and critical 1% threshold level (red). Wave propagation epochs were extracted from these threshold-crossing ρ_d values.

were distributed throughout the analyzed session, at an average rate of 2.03 and 1.21 events/s, respectively. It is important to note that the recordings in monkey visual cortex occur with the subject sitting in a darkened room, thus reducing the possibility in this case of the measured cortical activity being affected by uncontrolled visual stimulation, or by voluntary saccades.

In quantifying the wave propagation epochs, it is interesting first to consider the spectrum of amplitudes achieved during these periods of transient, high spatial organization on the array. The propagation epochs detected in the human showed a right-skewed distribution of oscillation amplitudes ranging from 0.04 (in z-score units) to 4.87 standard deviations above the mean, with a median of 1.21 $\pm$ 0.78 median absolute deviation. The propagation epochs detected in the monkey showed a similar distribution overall, with a median of 1.26 $\pm$ 0.44, but with a somewhat reduced range, extending from 0.18 to only 2.49 z-score units. This reduced amount of relatively high-amplitude epochs of phase organization in the monkey may be due to a different levels of noise in the two recording environments; further analyses will focus on this specific comparison of the amplitudes during extracted epochs. In general, however, it is inter-

esting to note that by focusing on a novel measure in these data – that is, the strength of the spatial organization of phase across the array – we can now look for consistencies in these detected phenomena, which would not be accessible when analyzing only oscillation epochs with a z-score threshold of 3 or 4 standard deviations above the mean.

We began characterizing the properties of these wave propagation epochs at the level of single events (see Supplementary Movie 5 for one example). For each event, we compute the phase latency map across space, and the relationship of phase latency with distance (Fig. 4.5a). Note that in this panel, the value of the phase latency at the point closest to the wave source has been subtracted out, to give a representation of relative phase latency at both a point within a detected propagation epoch (Fig. 4.5a, black dots) and one at which no wave has been detected (gray dots). This comparison makes the difference between the spatiotemporal structure of the activity in each case clear: within the propagation epoch the value of ρ_d is high (0.69) and the extracted propagation speed is low (124.9 mm/s), falling again into the range of the axonal conduction velocity for the unmyelinated horizontal fiber network (Bringuier et al., 1999; González-Burgos et al., 2000; Girard et al., 2001), as observed in **Work 2**, while outside the epoch, ρ_d is low (0.15) and the propagation speed is higher by almost an order of magnitude (1,225 mm/s). It it important to note that though these much faster propagation speeds observed outside the wave detection epochs could in principle be mediated by other fiber classes, such as top-down projections (Bullier, 2001) or thalamocortical arborizations (Metherate and Cruikshank, 1999), the observation of a high propagation speed in tandem with a low value of ρ_d indicates rather that the spatial pattern at point in question is more consistent with that produced by filtered noise (or, more simply, by the baseline level of spatial autocorrelation in the data). Finally, from the smoothed version of the phase latency map (Fig. 4.5a, inset), we can estimate the overall direction of phase flow and the spatial extent of the wave propagation. We note here, however, that the thresholding criterion used in this analysis probably does not take waves with spatial extent smaller than the recording array into account, because in this case, not all channels on the array will be able to contribute to the magnitude of the phase latency correlation with distance.

Moving next from single events to population-level statistics, we quantified the distribution of propagation speeds across all the detected wave epochs. To visualize this distribution, we constructed the kernel density estimate for the observed speed values, under the constraint of positive support (black line, Fig. 4.5b). Across the 162 epochs detected in the human MEA data, the distribution of propagation speeds shows a clear peak between 100 - 300 mm/s, consistently within the conduction velocities of the horizontal fiber network. In addition, this distribution of propagation speeds is reproduced across the 97 wave propagation epochs detected in the monkey (Fig. 4.6). It is important to emphasize here that though the criterion for propagation detection only

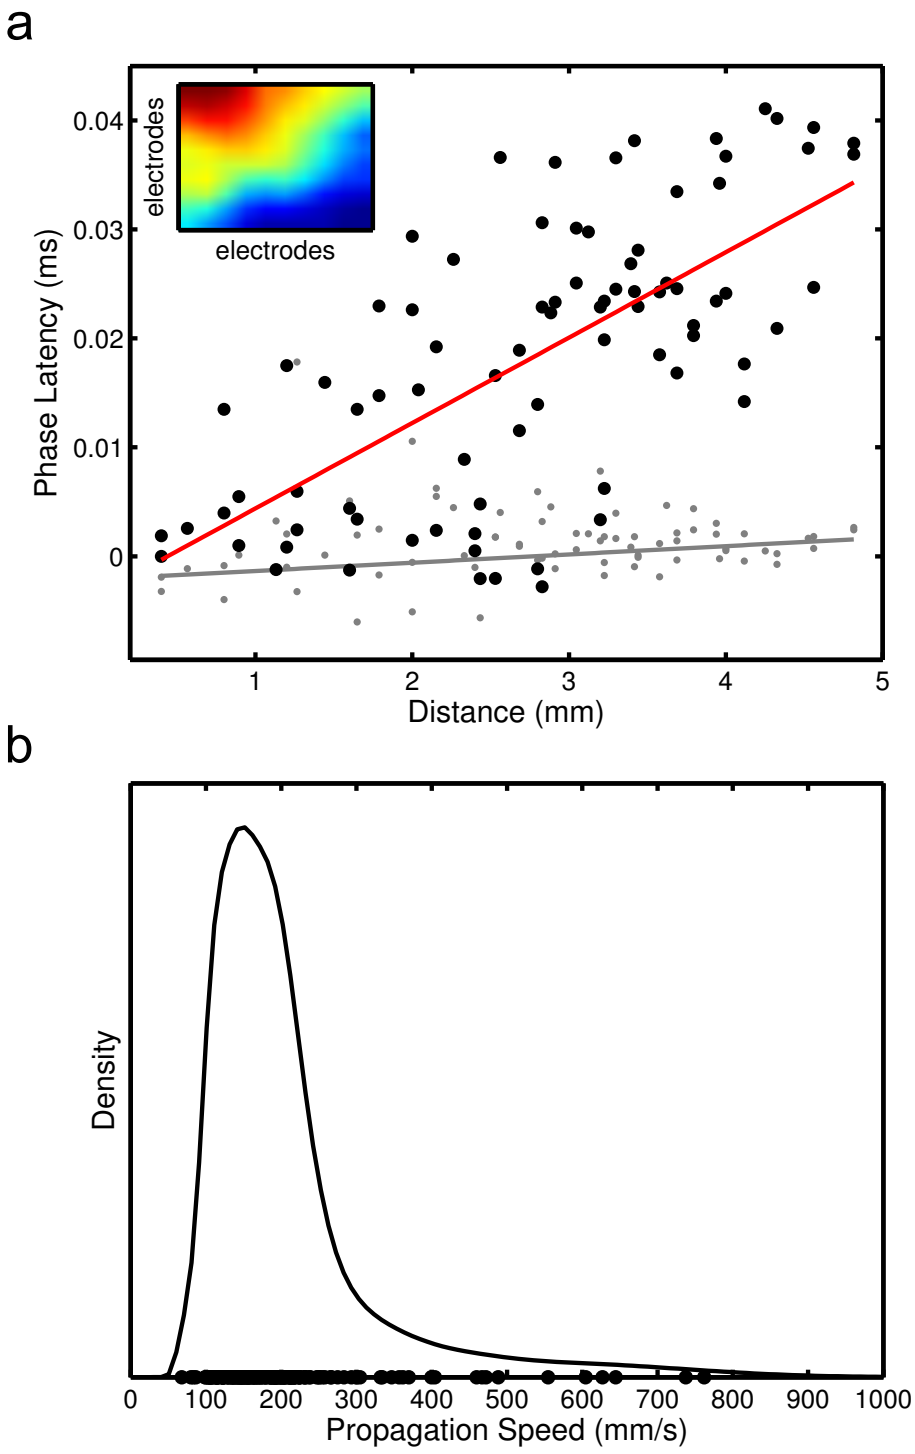

Fig. 4.5 Epoch statistics at the single-event and population level. (a) Phase latency with distance from the source point, within a detected wave propagation epoch (black dots) and at a randomly selected sample (gray dots). The smoothed phase latency map for the wave propagation epoch is displayed as an inset (top left). Propagation speed corresponds to a linear fit of the relationship of phase latency with distance for the within epoch (red line) and control (gray line) case. **(b)** The distribution of propagation speeds over the 162 detected epochs (black line), constructed by the kernel density estimate. Values for individual epochs are given at the bottom of the plot (black dots).

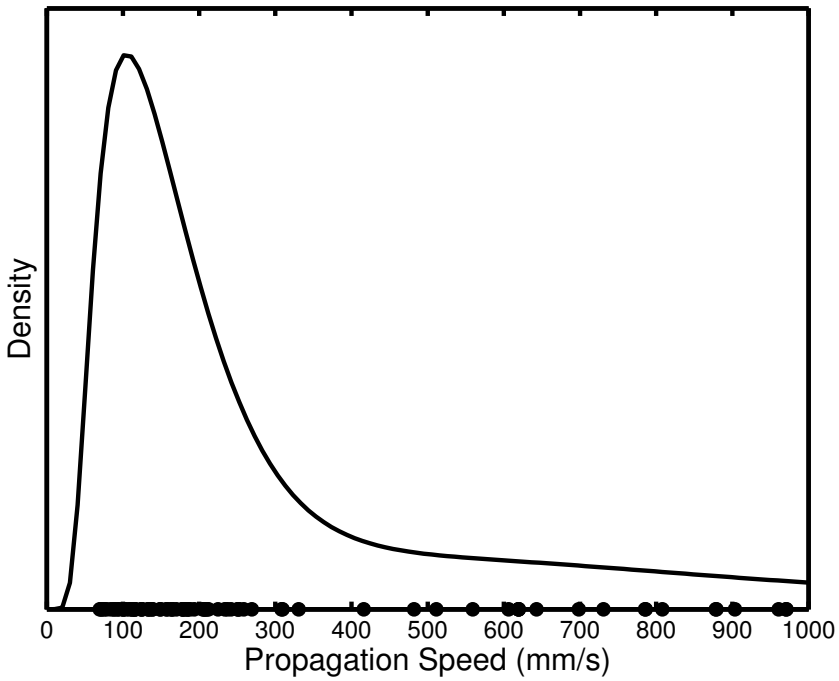

Fig. 4.6 Speed density in the monkey. The distribution of propagation speeds over the 97 detected wave epochs, constructed as in 4.5b, with individual speed values (black dots) and the kernel density estimate (black line).

captures the periods of highest spatial organization measured on the array, that these detected epochs then display a characteristic propagation speed is quite interesting. These observations indicate that, at least at a preliminary level, we can state that when ongoing activity enters into periods with significant spatial patterns of phase on the array, the horizontal fiber network may be playing an important role in organizing these spatiotemporal dynamics.

To make this statement more definitive, however, requires additional work, both experimental and theoretical, aimed at characterizing the exact relationship between the conduction speed of the propagation substrate and the distribution of propagation speeds observed in the LFP. In further experimental analyses, these results can be confirmed through the analysis of multiunit spiking activity (MUA), which has a more directly understood link with the dynamics in the biological network, and the phase relationship of MUA to the LFP waves. In future theoretical work, the relationship between synaptic delays in spiking networks and the spatiotemporal structure of network activity is an interesting problem that could fruitfully be addressed by adapting established methods, such as the mean-field formalism for the temporal dynamics of spiking networks studied in Brunel and Hakim (1999) and Brunel and Wang (2003), and the integro-differential approach studied in Pinto and Ermentrout (2001), Roxin et al. (2005), and Goulet and Ermentrout (2011).

Chapter 5

Brain networks: small-worlds, after all?

[Work in collaboration with Michelle Rudolph-Lilith]

The small-world network model

The canonical model of the small-world network was proposed only recently (Watts and Strogatz, 1998), but since that time, has spread through the imaginations of the broader scientific community like wildfire. Starting with regular ring graphs perturbed by the random rewiring of a few connections (Fig. 5.2), small-world networks display an intriguing emergent behavior, posed between the total order of regular graphs and the total disorder of random graphs. Within the small-world regime, these networks retain the high clustering of connections of the ring graph, while gaining the property of a short distance between any two nodes from the random component. Below, we will briefly recall the details of the construction of Watts-Strogatz networks and the mathematical definitions for the network properties of interest.

The construction of the small-world network model proceeds as follows. We start with a ring graph of N_N nodes, with each node connected to each of its neighbors within a distance k on the circular lattice (Fig. 5.2, top). Note that this will result in a total of $2kN_N$ edges in the final graph. Next, we select by turns each of the edges in the graph, and reconnect it to a randomly selected target node ("rewire") with a probability q. At small values of q, only a few edges will be rewired (Fig. 5.2, middle). At $q = 1$, then, the graph becomes fully randomized (Fig. 5.2, bottom).

To quantify the notions of clustering and distance in graphs, a panoply of measures have been employed. Watts and Strogatz gave the definition for the clustering coefficient of a single node (Watts and Strogatz, 1998; Newman, 2003)

$$C_i = \frac{\text{number of triangles connected to node } i}{\text{number of triangles possible at node } i}, \tag{5.1}$$

where for nodes with degree 0 or 1, for which this fraction evaluates to $0/0$, the value of C_i is set to 0. The clustering coefficient for the network is then the average over the local clustering coefficients $C = (1/N_N) \sum_i C_i$. The notion of distance between nodes in the

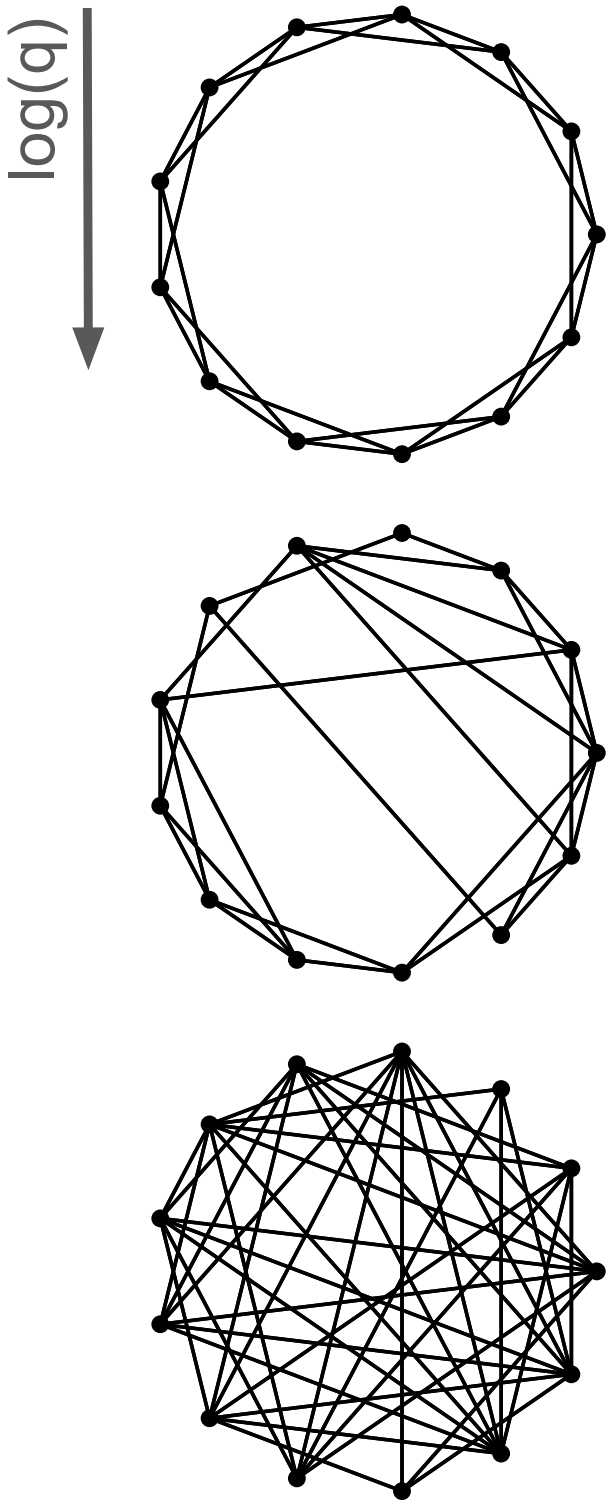

Fig. 5.1 The Watts-Strogatz network model. The progression of rewiring in the Watts-Strogatz model from a regular ring to a random graph, as a function of the logarithm of the rewiring parameter q. The graph starts as a regular network ($q = 0$, top), with each node connected to k neighbors in each direction. As the rewiring parameter increases ($q = 0.1$, middle), progressively more edges from the ring graph are "rewired" to random nodes in the network. Finally, at $q = 1$ (bottom), the network becomes fully randomized.

network is typically quantified by the average shortest path length between nodes in the network (Boccaletti et al., 2006)

$$L = \frac{1}{N_N(N_N-1)} \sum_{\substack{i,j=1 \\ i \neq j}}^{N_N} d_{ij} \tag{5.2}$$

where d_{ij} is the N_NxN_N matrix of all shortest paths between nodes in the network, and is computed numerically (e.g. by breadth first-search or Dijkstra's algorithm). In this definition, the contribution of the diagonal elements in the matrix is excluded (hence the term $N_N[N_N - 1]$); the handling of this point varies between authors, however because most set the distance from a node to itself to 0, adding these elements only constitutes a correction factor of order N_N^{-1} (Newman, 2003). Note that later, we will adopt the definition for the distance of a node with a self-loop (i.e. a connection to itself) to 1, which leads to a more consistent usage in practice.

Finally, because one is often not simply interested in raw value of these quantities in the context of assessing whether or not a given network exhibits the small-world property, and further because the comparison of these individual quantities in a given real-world graph to their equivalents in a random graph may not always be clear, a ratio-based index for the "small-worldness" of a network was recently proposed (Humphries and Gurney, 2008)

$$S = \frac{C^{obs}}{C^{rand}} \frac{L^{rand}}{L^{obs}} \tag{5.3}$$

where obs and $rand$ stand for the values observed in the given graph and its random counterpart, respectively. A visualization of the dependence of S on q is given in Figure 5.2 for a representative Watts-Strogatz graph.

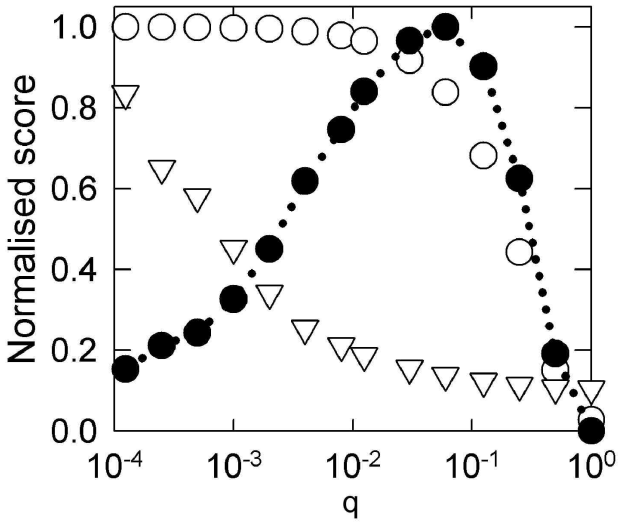

Fig. 5.2 The small-worldness index. The clustering coefficient (open circles), average path length (open triangles), and small-worldness index (filled circles) averaged over 20 realizations of a Watts-Strogatz graph with $N_N = 1000$ and $k = 10$, and plotted as a function of q. Note the logarithmic scale for q, and that all measures are presented in normalized form. Note finally maximal small-worldness occurs before 10% rewiring of edges in the graph. Adapted from Humphries and Gurney (2008)

Because of the many varied definitions for the quantities in this ratio, there are many ways to measure S in practice. The comprehensive characterization of S across several neural graphs, along with the how the choice of the random graph null hypothesis affects this index, is the starting point for the analysis that is presented in **Work 2**. From that point, we go on to apply theoretical expressions for the clustering coefficient in small-world and random graphs, and for the average geodesic graph distance in random graphs, which allows us to construct bounds on the maximum and minimum small-worldness of a Watts-Strogatz graph matched to the number of nodes and edges in a particular neural graph as a function in the rewiring parameter q. Because no good approximation for the average distance in a small-world graph exists, we cannot directly construct an analytical expression for S in a Watts-Strogatz graph matched to the neural data; here instead, we assess this last term numerically, thus computing the small-worldness index by semi-analytic means. The results from this work have interesting implications concerning the applicability of the small-world network paradigm in explaining the connectivity of the biological networks in the nervous system.

Work 3: *Brain networks: small-worlds, after all?*

French Summary

Le concept de connectivité «petit-monde» (“small-world”) introduit par Watts et Strogatz a joué un rôle central dans la théorie des réseaux, en particulier dans l’analyse des réseaux complexes du système nerveux. De l’échelle cellulaire à l’échelle des régions cérébrales, de nombreuses analyses ont révélé des caractéristiques de réseaux de type «petits-monde». Dans ce travail, nous revisitons la quantification du concept de petit-monde dans des réseaux neuronaux. Nous trouvons que les réseaux neuronaux tombent dans la région frontière entre les réseaux petit-monde et réseaux aléatoires, spécialement lorsque les séquences de dégrés de connectivité sont prises en compte. Nous appliquons des expressions de la théorie des graphes introduites récemment pour les mesures de distance et d’aggrégation (“clustering”), pour étudier cette région frontiére. Nous dérivons des limites théoriques pour l’indice de petit-monde minimal et maximal pour un réseau, et par des moyens semi-analytiques, nous pouvons étudier l’indice de petit-monde pour un réseau donné. En utilisant cette nouvelle approche, nous trouvons que des réseaux artificiels équivalents aux réseaux naturels sont dominés par leur composante aléatoire. Ces résultats fournissent des indications que la connectivité des réseaux neuronaux se situe entre une connectivité aléatoire et une connectivité de type petit-monde, mais très proche de la connectivité aléatoire.

Brain networks: small-worlds, after all?

Lyle Muller, Alain Destexhe and Michelle Rudolph-Lilith

Unité des Neurosciences, Information et Complexité (UNIC)
Centre National de la Recherche Scientifique (CNRS)
1 Avenue de la Terrasse, Gif-sur-Yvette, FR

E-mail: rudolph@unic.cnrs-gif.fr

Abstract. Since its introduction, the "small-world" effect has played a central role in network science, particularly in the analysis of the complex networks of the nervous system. From the cellular level to that of interconnected cortical regions, many analyses have revealed small-world properties in the networks of the brain. In this work, we revisit the quantification of small-worldness in neural graphs. We find that neural graphs fall into the "borderline" regime of small-worldness, residing close to that of a random graph, especially when the degree sequence of the network is taken into account. We then apply recently introducted analytical expressions for clustering and distance measures, to study this borderline small-worldness regime. We derive theoretical bounds for the minimal and maximal small-worldness index for a given graph, and by semi-analytical means, study the small-worldness index itself. With this approach, we find that graphs with small-worldness equivalent to that observed in experimental data are dominated by their random component. These results provide the first thorough analysis suggesting that neural graphs may reside far away from the optimally small-world regime.

1. Introduction

The small-world network model of Watts and Strogatz is quite possibly the pinnacle of network science. In their seminal work [1], the authors showed that there exists an intermediate state between regular and random graphs, where high clustering of network connections and low distance between any two nodes coexist, and that this intermediate "small-world" state is reached by random rewiring of only a few connections in a regular ring graph. Because of these dual properties stemming from the structure of a regular graph combined with the efficiency of a random graph, the pertinence of the small-world network model to real-world systems, both natural and technological, was immediately clear.

Since that time, indications of the small-world property have been uncovered in a vast array of real-world datasets. From the World-Wide Web [2] to the electric power grid [3], social networks [4], and various neural systems [5], network connectivity has often been observed to exhibit clustering exceeding that of a random graph, while simultaneously maintaining a low average number of links between nodes. While in many networks the small-world property may be strikingly clear in comparison to a random graph, however, in some cases the difference may not be particularly great, especially in the case of small systems. For this reason, a ratio-based index that quantifies the "small-worldness" of a graph was recently introduced [6], where values beyond unity indicate the presence of the small-world property.

Immediately upon the introduction of this measure, it was apparent that, while systems such as the electric power grid or the co-authorship network of mathematicians could exhibit a small-worldness index ranging from 100 even to 10,000 and above, certain other systems fall into a "borderline" category, defined by small-worldness indices ranging between 1 and 3 [6]. Among these borderline cases is the neural connectivity diagram of the macaque cortex, and further, since the introduction of this index, several neural graphs have been reported to fall within this regime [7, 8], including the area-level connectivity diagram of the human cerebral cortex [9].

Here, motivated by interest in the small-world property and the meaning of this borderline regime, we present a comprehensive analysis of small-worldness in neural graphs. The connectivity diagrams used for this analysis span in scale from the microscopic neural connectivity of *C. elegans* to the macroscopic connectivity of cortical regions in the monkey brain, and in size from a few tens to several hundred nodes. We quantify clustering and distance in the neural graphs using several available measures, in order to consider all relevant formulations for the small-worldness index in this analysis. Within this multifaceted approach, a consistent result nevertheless emerges: neural graphs fall in general into the borderline small-worldness regime, particularly when the degree sequence of the network is taken into account.

To investiate the implications of this borderline regime, we employ a recently introduced exact expression for the clustering coefficient of small-world and random graphs [10] and an excellent analytical approximation for the average geodesic graph

distance [11]. With these expressions, we derive theoretical bounds for the maximal and minimal small-worldness index of graph models constrained by neural data. Furthermore, through a semi-analytical approach, we study the behavior of the small-worldness index directly. To reproduce the small-worldness observed in the data, we find that corresponding graph models require a dominant random rewiring, in all cases one order of magnitude larger than the small rewiring first noted to bring a regular graph into the optimally small-world regime [1]. Finally, we discuss these results and their implications for the study of network structure in the nervous system.

2. Graph theory preliminaries

Mathematically, a (relational) graph or network is a pair $G = (\mathcal{N}, \mathfrak{a})$, where $\mathcal{N}$ denotes a set of N_N nodes (vertices) and $\mathfrak{a} : \mathcal{N} \times \mathcal{N} \longrightarrow \{0, 1\}$ an adjacency relation which assigns to each pair of nodes (i, j) a number $a_{ij} \in \{0, 1\}$. If $a_{ij} = 1$, the nodes i and j are said to be adjacent (connected) to each other, and the adjacency relation defines an edge between both nodes. For $a_{ij} = 0$, no link (edge) exists between nodes i and j. If the relation $\mathfrak{a}$ is commutative, i.e. $\forall(i, j) : a_{ij} = a_{ji}$, the graph is said to be undirected. A graph is directed, if at least one node pair (i, j) exists for which $a_{ij} \neq a_{ji}$. For such digraphs, edges are endowed with a direction pointing from a source node to a target node. If $a_{ii} = 1$, node i is said to be self-looped. A graph with at least one self-looped node is called self-looped graph.

The adjacency relations a_{ij} form a matrix, which provides a complete description of a relational graph's structure. This adjacency matrix defines a whole host of graph theoretical measures, with which a given graph can be further characterized (e.g, see [12, 13] for comprehensive reviews). The connectedness (or connectance) Co, a measure of relative graph connectivity, of a not self-looped graph is defined as

$$Co = \frac{1}{N_N(N_N - 1)} \sum_{i,j=1}^{N_N} a_{ij}, \tag{1}$$

with $0 \leq Co \leq 1$. The node in- and out-degrees for digraphs are defined as

$$a_i^{\{in,out\}} = \sum_{j=1}^{N_N} a_{\{ji,ij\}}, \tag{2}$$

and the average node degree as

$$\langle a_i \rangle = \frac{1}{N_N} \sum_{i=1}^{N_N} a_i^{in} = \frac{1}{N_N} \sum_{i=1}^{N_N} a_i^{out}. \tag{3}$$

A graph is comprised of a number of connected components, that is pairwise disjoint subgraphs in which each node can reach and is reachable by any other node in the subgraph. More specifically, a connected component is defined as a subgraph consisting of a set of nodes from which all other nodes in the subgraph can be reached, and which can be reached from all other nodes, by following existing edges. In digraphs, the

so defined connected components are also called strongly connected components. The analysis in this paper will will restrict to the strongly connected components.

3. Analysis methods

In this paper, we study structural aspects of a number of publicly available biological neural graphs used in the literature in the past two decades. These data include areal connectivity graphs of the cat cortex (CC1 and CC2), the macaque brain (MB1), macaque cortex (MC1 and MC2), macaque visual cortex (MVC1 and MVC2) and macaque neocortex (MNC1), as well as the neuronal connectivity graph of the nematode *C. elegans* (CE1, CE2 and CE3). References for each data source are provided in Supplementary Table 1. Note that some of these graphs have experience slight modifications and refinement since their first investigation. Throughout this study, we analyze the giant strongly connected component of the network (see Supplement, Connected components).

3.1. Numerical tools

Numerical analysis was performed using the custom software *cygraph* and *Mathematica*. A *cygraph* binary (Mac OSX), all graph data, and analysis protocols are available for download [14].

3.2. Network models

The assessment of the small-world properties of the the investigated biological neural graphs requires their comparison with equivalent random graphs [6, 15]. To that end, we constructed surrogate randomized networks with the same number of nodes (N_N), number of edges (N_E), and node in- and out-degrees (a_i^{in} and a_i^{out}, respectively) as in the neural data using the well-known edge swapping algorithm (MS) [16]. We next confirmed all results in the study using an algorithm for sampling from the ensemble of exact degree-matched Erdős-Rényi (EDM) random graphs introduced in [17, 18]. Finally, to quantitatively understand the effect of the network degree distribution on the small-worldness index, we constructed Erdős-Rényi (ER) random graphs with the same number of nodes and edges as in the neural data, without taking the degree-distribution into account.

In order to compare the obtained results with a classical model of small-worlds, we constructed Watts-Strogatz (WS) graph models [1] with the same number of nodes and initial node degree to match the number of edges in the network. To construct WS digraphs, we first constructed appropriately matched undirected WS graphs, after which the direction of each edge was randomized with a probability of 0.5. When considering the MS, EDM, ER, and WS network models, 1,000 random realizations were used at each point in parameter space to ensure statistical stability.

4. Results

In the following sections, we first present results of the geodesic distance and clustering analysis of the considered neural graphs and degree-preserved randomized counterparts. These results are then used to quantitatively assess the small-worldness index (section 4.3). We conclude in section 4.4 by applying analytical expressions for these network measures to derive theoretical bounds for these quantities and conducting a semi-analytical calculation for the small-worldness of Watts-Strogatz graphs matched to the neural data, expressed as a function of the rewiring parameter.

4.1. Geodesic distances

First, we performed a classical geodesic distance analysis. In the past, various measures of characteristic distance in relational graphs were proposed (for a review of various measures, see [12, 13]), all based on the geodesic distance d_{ij} between two nodes i and j. Here, d_{ij} is defined as the number of edges in the shortest path connecting two nodes, and the set of d_{ij} forms a matrix called the (geodesic) distance matrix. Typically, unconnected nodes are assigned a geodesic distance of infinity, but as we restrict our analysis to the (strongly) connected components of each graph, no unconnected pair of nodes can occur. The diagonal elements d_{ii} of the distance matrix are called geodesic loops, or closed geodesics. In looped graphs, the closed geodesics of a self-looped node i (i.e. a node for which $a_{ii} = 1$) is the shortest nontrivial path which connects the node with itself. This definition contrasts the commonly used notion, in which d_{ii} of self-looped nodes is set to 0. The analysis presented in this paper will include closed geodesics of all nodes.

The geodesic node distance matrix d_{ij} is a $N_N \times N_N$ matrix from which a number of measures can be conceived. First, the geodesic node in- and out-distance degree $d_i^{\{in,out\}}$ is defined for directed graphs as

$$d_i^{\{in,out\}} = \sum_{j=1}^{N_N} d_{\{ji,ij\}} \tag{4}$$

and, due to the symmetry of the geodesic distance matrix, for undirected graphs by

$$d_i = d_i^{in} = d_i^{out} = \sum_{j=1}^{N_N} d_{ij}. \tag{5}$$

If closed geodesics are excluded, the sum in the above equations runs over $j \neq i$.

Secondly, the maximal geodesic distance between a given node i and any other node in the graph,

$$\varepsilon_i = \max_j \{d_{ij}\}, \tag{6}$$

defines the geodesic node eccentricities. In the case closed geodesics are excluded from the analysis, the maximum is evaluated over all $d_{ij}, j \neq i$ for a given i. With these

measures, one has

$$D = \begin{cases} \sum\limits_{i,j=1}^{N_N} d_{ij} & \text{closed geodesics included} \\ \sum\limits_{\substack{i,j=1 \\ i\neq j}}^{N_N} d_{ij} & \text{closed geodesics excluded} \end{cases} \tag{7}$$

$$r = \min_i\{\varepsilon_i\} \tag{8}$$

$$d = \begin{cases} \max\limits_{i,j}\{d_{ij}\} & \text{closed geodesics included} \\ \max\limits_{\substack{i,j=1 \\ i\neq j}}\{d_{ij}\} & \text{closed geodesics excluded} \end{cases} \tag{9}$$

as the geodesic graph distance, geodesic graph radius and diameter, respectively. Furthermore,

$$\langle d\rangle = \frac{D}{N_C} \tag{10}$$

$$\langle d_i\rangle = \frac{D}{N_N} \tag{11}$$

define the average geodesic graph and node distance, respectively, where N_C denotes the number of connected node pairs in a given graph. In the case of (strongly) connected components, $N_C = N_N^2$ in the case closed geodesics are included, and $N_C = N_N(N_N - 1)$ if closed geodesics are not considered.

Finally, the characteristic geodesic path length is given by the median of the means of the shortest geodesic distances connecting each node to all other nodes, i.e.

$$L = \operatorname*{median}_i \left\{ \frac{1}{N_N} d_i \right\}, \tag{12}$$

and the geodesic connectivity length is defined by the harmonic mean of geodesic node distances

$$\ell = N_C \left\{ \sum_{i,j=1}^{N_N} \frac{1}{d_{ij}} \right\}^{-1}. \tag{13}$$

If closed geodesics are excluded from the analysis, the sum in the last equation runs over all i, j with $i \neq j$. Results of the application of the above measures to the considered biological neural graphs are summarized in table 1.

In accordance with expectation, the various geodesic distance measures will decrease for increasing connectedness of a graph. This is illustrated in figure 1, which shows the characteristic geodesic path length, geodesic connectivity length and average geodesic graph distance as function of the graph connectedness Co (figures 1(a)–(c), respectively) for both the neural graphs (green dots) and their degree-preserved randomized counterparts (black dots). A rough theoretical estimate for the average graph distance in ER random graphs is known [19], and given by

$$\langle d\rangle_{ER}^{theor} = \frac{\ln(N_N)}{\ln(\langle a_i\rangle)}, \tag{14}$$

geodesic distance analysis								
	N_C	r	d	D	$\langle d \rangle$	$\langle d_i \rangle$	L	ℓ
CC1	9025	2	4	17359	1.9234	182.7263	1.8842	1.6894
CC2	2704	2	4	4899	1.8118	94.2115	1.8077	1.5814
CE1	57121	4	14	227938	3.9904	953.7155	3.1674	3.1299
CE3	75076	4	7	215363	2.8686	785.9964	2.8376	2.5700
MB1	123201	3	6	322355	2.6165	918.3903	2.5385	2.3513
MC1	4900	3	5	11296	2.3053	161.3714	2.2143	1.9756
MC2	7225	3	5	13453	1.8620	158.2706	1.8000	1.5794
MNC1	2209	3	4	4537	2.0539	96.5319	2.0851	1.7551
MVC1	900	2	4	1553	1.7256	51.7667	1.7333	1.5128
MVC2	900	2	4	1553	1.7256	51.7667	1.7333	1.5128

Table 1: Geodesic distance analysis across the neural graphs. Shown are values for the number of connected node pairs N_C, the geodesic graph distance D, graph radius r and diameter d, (equations (7)–(9)), average geodesic graph and node distance ($\langle d \rangle$ and $\langle d_i \rangle$, respectively; equations (10) and (11)), as well as the characteristic path length L, equation (12), and geodesic connectivity length ℓ, equation (13).

where $\langle a_i \rangle$ denotes the average node degree. Excluding for simplicity self-loops, the average node degree of a classical ER random graph is $\langle a_i \rangle = Co(N_N - 1)$, which yields

$$\langle d \rangle_{ER}^{theor} = \frac{\ln(N_N)}{\ln(Co(N_N - 1))} \tag{15}$$

for the expected average geodesic graph distance. This theoretical prediction qualitatively matches with the results found for the biological neural graphs when considering $\langle d \rangle$ as a function of connectedness, but generally underestimate the observed values (figure 1(c), compare dots with upward triangles). Recently, a more accurate theoretical prediction of $\langle d \rangle$ for classical ER random graphs was provided, which takes the form [11]

$$\langle d \rangle_{ER}^{theor} = \frac{\ln(N_N) - \gamma}{\ln(Co \cdot N_N)} + \frac{1}{2}, \tag{16}$$

where γ denotes the Euler-Mascheroni constant. Comparing the numerical results of equation (16) with the average geodesic graph distance for the corresponding biological neural graphs shows an excellent agreement both qualitatively and quantitatively (figure 1(c), compare dots with downward triangles). Moreover, this agreement provides further support for the recent finding that the structural makeup of small neural graphs is dominated by a classical ER random component constrained only by a non-scale-free node-degree distribution [20, 21].

To further quantify the comparison between the observed and randomized graphs,

we defined the following ratios [6]

$$\lambda_L = \frac{L}{L_{MS}}, \tag{17}$$

$$\lambda_{\langle d\rangle} = \frac{\langle d\rangle}{\langle d\rangle_{MS}}, \tag{18}$$

$$\lambda_{\langle d_i\rangle} = \frac{\langle d_i\rangle}{\langle d_i\rangle_{MS}} \tag{19}$$

between the characteristic geodesic path length L, average geodesic graph distance $\langle d\rangle$ and average geodesic node distance $\langle d_i\rangle$ of a given graph and its corresponding randomized surrogate (L_{MS}, $\langle d\rangle_{MS}$ and $\langle d_i\rangle_{MS}$, respectively). In all cases and for all considered biological neural graphs, we observe λ-values very close to 1 (figure 1(d); $1.03 \leq \lambda_L \leq 1.22$, $0.98 \leq \lambda_{\langle d\rangle} \leq 1.30$ and $0.98 \leq \lambda_{\langle d_i\rangle} \leq 1.23$). In section 4.3, the various λ ratios defined in equations (17)–(19) will be used to assess the small-worldness of the investigated graphs.

4.2. Clustering

Secondly, we investigated the clustering properties of the neural graphs. The notion of "clustering" was first introduced as "transitivity" in studies of social graphs to measure the extent to which nodes are linked together in triangular relationships [22]. Specifically, three nodes A, B, and C are said to be transitive if, for existing links between nodes A and B and nodes B and C, nodes A and C are also linked. Mathematically, the graph transitivity T can be defined as the ratio between the number of closed paths of length two and the total number of paths of length two [13], i.e.

$$\begin{aligned} T &= \frac{\#\text{ closed paths of length 2}}{\#\text{ paths of length 2}} \\ &\equiv \frac{3 \times \#\text{ triangles}}{\#\text{ connected triples}}, \end{aligned} \tag{20}$$

with $0 \leq T \leq 1$. In a similar fashion, the node clustering coefficient is defined as [1, 12]

$$\begin{aligned} C_i &= \frac{2 \times \#\text{ edges in subgraph of neighbors of node } i}{a_i(a_i-1)} \\ &= \frac{1}{a_i(a_i-1)} \sum_{\substack{j=1 \\ j\neq i}}^{N_N} \sum_{\substack{h=1 \\ h\neq i,j}}^{N_N} a_{ij}a_{ih}a_{jh} \end{aligned} \tag{21}$$

($0 \leq C_i \leq 1$), and measures how likely $a_{mn} = 1$ for two neighbors m, n of node i, that is the ratio between the number of existing edges in the subgraph spanned by the neighbors of a node i and the number of possible edges present within this subgraph. The average node clustering coefficient is then given by

$$\langle C_i\rangle = \frac{1}{N_N} \sum_{i=1}^{N_N} C_i. \tag{22}$$

The clustering measures in Eqs. (20)–(22) are defined for undirected graphs only, but can be generalized to hold for digraphs as well. Here, however, the edge direction has to be taken into account. For the single triangular relationship which defines graph transitivity, a total of eight relationships can be drawn for digraphs (figure 2(a)). From these, four equivalent pairs exist, leaving four distinct definitions of global clustering coefficients for digraphs:

$$C^{(000)} = \frac{T_{\Delta}^{(000)}}{T_{<}^{(01)}}, \tag{23}$$

$$C^{(001)} = \frac{T_{\Delta}^{(001)}}{T_{<}^{(00)}}, \tag{24}$$

$$C^{(010)} = \frac{T_{\Delta}^{(010)}}{T_{<}^{(01)}}, \tag{25}$$

$$C^{(100)} = \frac{T_{\Delta}^{(100)}}{T_{<}^{(11)}}, \tag{26}$$

where $T_{\Delta}^{(abc)}$, $a, b, c \in \{0, 1\}$, denotes the number of triangles of type (abc) (figure 2(a)) and $T_{<}^{(ab)}$ the number of connected triples of type (ab) (figure 2(b)). Furthermore, the sum over all possible types of triangles and connected triples yields a total global clustering coefficient for digraphs [23]

$$C^{d} = \frac{1}{2\sum_{i=1}^{N_N}(a_i^{tot}(a_i^{tot}-1) - 2a_i^{b})} \sum_{i=1}^{N_N}\sum_{\substack{j=1\\j\neq i}}^{N_N}\sum_{\substack{h=1\\h\neq i,j}}^{N_N}(a_{ij}+a_{ji})(a_{ih}+a_{hi})(a_{jh}+a_{hj}), \tag{27}$$

where $a_i^{tot} = a_i^{in} + a_i^{out}$ and

$$a_i^{b} = \sum_{\substack{j=1\\j\neq i}}^{N_N} a_{ij}a_{ji}.$$

Note that $C^{(abc)}$ and C^{d} are set to 0 if the denominator in their respective defining equations vanishes, and that $0 \leq C^{(abc)}, C^{d} \leq 1$.

Similar to the undirected case and motivated by Eqs. (23)–(27), one can furthermore define node clustering coefficients for digraphs according to [23]

$$C_i^{(000)} = \frac{T_{\Delta i}^{(000)}}{T_{<i}^{(01)}} = \frac{1}{2(a_i^{in}a_i^{out} - a_i^{b})} \sum_{\substack{j=1\\j\neq i}}^{N_N}\sum_{\substack{h=1\\h\neq i,j}}^{N_N}(a_{ij}a_{jh}a_{hi} + a_{ih}a_{hj}a_{ji}), \tag{28}$$

$$C_i^{(001)} = \frac{T_{\Delta i}^{(001)}}{T_{<i}^{(00)}} = \frac{1}{2a_i^{in}(a_i^{in}-1)} \sum_{\substack{j=1\\j\neq i}}^{N_N}\sum_{\substack{h=1\\h\neq i,j}}^{N_N} a_{ji}a_{hi}(a_{jh}+a_{hj}), \tag{29}$$

$$C_i^{(010)} = \frac{T_{\Delta i}^{(010)}}{T_{<i}^{(01)}} = \frac{1}{2(a_i^{in}a_i^{out} - a_i^{b})} \sum_{\substack{j=1\\j\neq i}}^{N_N}\sum_{\substack{h=1\\h\neq i,j}}^{N_N}(a_{ij}a_{hj}a_{hi} + a_{ih}a_{jh}a_{ji}), \tag{30}$$

$$C_i^{(100)} = \frac{T_{\Delta i}^{(100)}}{T_{<i}^{(11)}} = \frac{1}{2a_i^{out}(a_i^{out}-1)} \sum_{\substack{j=1 \\ j\neq i}}^{N_N} \sum_{\substack{h=1 \\ h\neq i,j}}^{N_N} a_{ij}a_{ih}(a_{jh}+a_{hj}), \tag{31}$$

$$C_i^{d} \quad = \frac{1}{2(a_i^{tot}(a_i^{tot}-1)-2a_i^{b})} \sum_{\substack{j=1 \\ j\neq i}}^{N_N} \sum_{\substack{h=1 \\ h\neq i,j}}^{N_N} (a_{ij}+a_{ji})(a_{ih}+a_{hi})(a_{jh}+a_{hj}), \tag{32}$$

where $T_{\Delta i}^{(abc)}$ denotes the number of triangles of type (abc) connecting two nodes in the neighborhood of node i and node i, and $T_{<i}^{(ab)}$ the number of connected triples of type (ab) with node i being the base node and two other nodes m, n in the neighborhood of node i. Also here, $0 \leq C_i^{(abc)}, C_i^{d} \leq 1$. Finally, the average over the node clustering coefficients yields the corresponding average node clustering coefficients

$$\langle C_i^{(abc)} \rangle = \frac{1}{N_N} \sum_{i=1}^{N_N} C_i^{(abc)} \tag{33}$$

$$\langle C_i^{d} \rangle \quad = \frac{1}{N_N} \sum_{i=1}^{N_N} C_i^{d}\,. \tag{34}$$

Numerical results of the application of the various clustering coefficients to the investigated biological neural graphs are summarized in table 2 and visualized in figure 2(c)–(d). As expected, the clustering of a graph increases with its connectedness, with the graph clustering coefficients taking larger values than the corresponding average node clustering coefficients (figure 2(c)–(d), compare blue and green filled dots). Interestingly, in all cases, $C^{(000)}$ and $\langle C^{(000)} \rangle$, which correspond to the occurrence of closed geodesics of length three in digraphs, were smallest among their peers. The various clustering coefficients for the degree-preserved randomized surrogates are likewise visualized as open dots in figure 2. Also here, the average node clustering coefficients were found to be larger than the corresponding graph clustering coefficients (figure 2(d), compare open blue and green dots), but with all triangular motifs represented with equal frequency. Moreover, as for the neural graphs, a general increase of the clustering for increasing connectedness was observed in the randomized graphs. However, in comparison to the randomized graphs, the neural data exhibits a clustering which, in most cases, is less than two times larger. This finding is somewhat surprising, as previous claims of a prevalent small-world structure in biological neural graphs would demand a significantly larger clustering than in corresponding random graphs.

To quantify the dissimilarity of the clustering coefficients in the investigated neural graphs and degree-preserved randomized graphs, we defined the following ratios [6]

$$\gamma^{(abc)} = \frac{C^{(abc)}}{C_{MS}^{(abc)}}, \tag{35}$$

$$\gamma^{d} \quad = \frac{C^{d}}{C_{MS}^{d}}, \tag{36}$$

$$\bar{\gamma}^{(abc)} = \frac{\langle C_i^{(abc)} \rangle}{\langle C_i^{(abc)} \rangle_{MS}}, \tag{37}$$

clustering analysis					
	C^d	$C^{(000)}$	$C^{(001)}$	$C^{(010)}$	$C^{(100)}$
CC1	0.4384	0.4315	0.4453	0.4430	0.4339
CC2	0.5154	0.4883	0.5423	0.5361	0.4967
CE1	0.1339	0.0624	0.1429	0.1766	0.1488
CE3	0.1525	0.1093	0.1455	0.1788	0.1778
MB1	0.2199	0.1831	0.2019	0.2568	0.2413
MC1	0.3982	0.3905	0.4027	0.4075	0.3922
MC2	0.6514	0.6008	0.6056	0.7058	0.7006
MNC1	0.4838	0.4660	0.4720	0.4974	0.5006
MVC1	0.5129	0.4945	0.5131	0.5283	0.5155
MVC2	0.5129	0.4945	0.5131	0.5283	0.5155
	$\langle C_i^d \rangle$	$\langle C_i^{(000)} \rangle$	$\langle C_i^{(001)} \rangle$	$\langle C_i^{(010)} \rangle$	$\langle C_i^{(100)} \rangle$
CC1	0.5451	0.5386	0.5498	0.5490	0.5426
CC2	0.5867	0.5499	0.6118	0.6120	0.6019
CE1	0.1933	0.0851	0.2373	0.2445	0.1965
CE3	0.2424	0.1798	0.2469	0.2918	0.2591
MB1	0.3373	0.2892	0.3313	0.3811	0.3311
MC1	0.4786	0.4647	0.4784	0.4880	0.5080
MC2	0.6780	0.6125	0.6882	0.7357	0.7142
MNC1	0.5805	0.5567	0.5391	0.6065	0.6098
MVC1	0.5510	0.5276	0.5829	0.5649	0.5609
MVC2	0.5510	0.5276	0.5829	0.5649	0.5609

Table 2: Clustering analysis across neural graphs. Shown are values for the global clustering coefficients $C^{(abc)}$ and C^d, equations (23)–(27), average node clustering coefficients $\langle C^{(abc)} \rangle$ and $\langle C^d \rangle$, equations (33)–(34).

$$\bar{\gamma}^d \quad = \frac{\langle C_i^d \rangle}{\langle C_i^d \rangle_{MS}} \tag{38}$$

between the graph clustering coefficients $C^{(abc)}$, C^d and average node clustering coefficients $\langle C_i^{(abc)} \rangle$, $\langle C_i^d \rangle$ of a given digraph and its corresponding random graph ($C_{MS}^{(abc)}$, C_{MS}^d, $\langle C_i^{(abc)} \rangle_{MS}$ and $\langle C_i^d \rangle_{MS}$, respectively). In most cases, we observed γ-values smaller than 2 (figure 3). Only for the neural connectivity graphs of *C. elegans* (CE1 and CE3), the various γ take larger values up to 2.98. In accordance with the above-mentioned finding, γ-values for digraph clustering coefficients of type (000) are generally smaller, hinting at the underrepresentation of motifs describing closed geodesics of length three in biological graphs (see figure 2(d), $\gamma^{(000)}$). In the next section, γ defined in equations (35)–(38) will be used to quantify the small-worldness of the considered graphs.

4.3. Small-worldness

The notion of "small-worlds" was introduced to describe a specific class of graphs in which the typical geodesic distance (shortest path length) between two nodes is comparable to that in ER random graphs with the same number of nodes and connectedness, and in which the geodesic distance scales logarithmically with the number of nodes in the graph. Furthermore, small-worlds are endowed with a overabundance of local triangular relationships, i.e. display high clustering, when compared with their random counterparts [1]. Using ~~In~~ the notation introduced above, a graph is labeled a "small-world" if both $\lambda \geq 1$ and $\gamma \gg 1$ (e.g., see [6]).

In section 4.1, we showed that for biological neural graphs the various λ-values (equations (17)–(19)) range from 0.98 to 1.30 for the considered neural graphs (see figure 1(d)), thus suggesting a strong similarity between biological neural graphs and corresponding random graphs when considering the typical geodesic distance between nodes. Moreover, the obtained results for the average graph distance scale in accordance with the expectation for random graphs (figure 1(c), compare filled dots with solid triangles), i.e. logarithmically in the number of nodes (see equation (16) for analytical form). Thus, the first condition $\lambda \geq 1$ allowing for a characterization of biological neural graphs as small-worlds are met.

The evidence for the small-world property in the context of the various clustering coefficients is, however, much less clear. Here, we found that for the majority of the considered biological digraphs, γ takes values in the range between 1.15 and 2. Only the neural connectivity digraphs of *C. elegans* (CE1 and CE3) yield values up to 2.98 for $\gamma^{(001)}$. This finding suggests that, at least for the biological neural graphs studied here, the condition $\gamma \gg 1$ characterizing small-world graphs must be met with greater care.

In order to better quantify the level of "small-worldness", we employed a number of small-worldness indices [6, 20]. Global small-worldness indices for digraphs can be defined by

$$S^{(abc)} = \frac{\gamma^{(abc)}}{\lambda_{\langle d \rangle}} = \frac{C^{(abc)}}{C^{(abc)}_{MS}} \frac{\langle d \rangle_{MS}}{\langle d \rangle}, \tag{39}$$

$$S^{d} = \frac{\gamma^{d}}{\lambda_{\langle d \rangle}} = \frac{C^{d}}{C^{d}_{MS}} \frac{\langle d \rangle_{MS}}{\langle d \rangle}. \tag{40}$$

Similarly, two sets of local small-worldness indices are defined for digraphs as

$$\bar{S}^{(abc)} = \frac{\bar{\gamma}^{(abc)}}{\lambda_{\langle d_i \rangle}} = \frac{\langle C_i^{(abc)} \rangle}{\langle C_i^{(abc)} \rangle_{MS}} \frac{\langle d_i \rangle_{MS}}{\langle d_i \rangle}, \tag{41}$$

$$\bar{S}^{d} = \frac{\bar{\gamma}^{d}}{\lambda_{\langle d_i \rangle}} = \frac{\langle C_i^{d} \rangle}{\langle C_i^{d} \rangle_{MS}} \frac{\langle d_i \rangle_{MS}}{\langle d_i \rangle} \tag{42}$$

and

$$\tilde{S}^{(abc)} = \frac{\bar{\gamma}^{(abc)}}{\lambda_{L}} = \frac{\langle C_i^{(abc)} \rangle}{\langle C_i^{(abc)} \rangle_{MS}} \frac{L_{MS}}{L}, \tag{43}$$

$$\bar{S}^d \quad = \frac{\bar{\gamma}^d}{\lambda_L} = \frac{\langle C_i^d \rangle}{\langle C_i^d \rangle_{MS}} \frac{L_{MS}}{L}. \tag{44}$$

With these, the two conditions $\lambda \geq 1$ and $\gamma \gg 1$ defining small-world graphs are then equivalent to $\{S^\bullet, \bar{S}^\bullet, \tilde{S}^\bullet\} > 1$, where $\bullet \in \{(abc), d\}$ [6].

For each graph studied here, we calculated the various small-worldness indices defined in equations (39)–(44). In all cases, $S^\bullet$, $\bar{S}^\bullet$ and $\tilde{S}^\bullet$ take values larger than 1 (figure 4), thus placing these graphs into the class of small-worlds. However, in most of the graphs the various small-worldness indices remained far below a value of 2, with only $\bar{S}^{(001)}$, $\bar{S}^{(010)}$ for the CE3 graph reaching values up to 2.61. This suggests that the investigated neural graphs, at best, fall into the class of "borderline" small-world graphs [6].

4.4. Borderline small-worlds

In order to further explore and give meaning to the borderline small-world regime, we analytically assessed the small-worldness index of Watts-Strogatz graphs. The latter is defined by

$$S^d = \frac{C_{SW}^d}{C_{ER}^d} \frac{\langle d \rangle_{ER}}{\langle d \rangle_{SW}}, \tag{45}$$

where C_{SW}^d and C_{ED}^d denote the total clustering coefficient for small-world and Erdős-Rényi digraphs, $\langle d \rangle_{SW}$ and $\langle d \rangle_{ER}$ their average geodesic graph distance, respectively.

Recently, an analytically exact expression for the total global clustering coefficient C_{SW}^d of small-world digraphs with degree k and rewiring probability q, valid non-asymptotically for all graph sizes, was obtained [10]:

$$\begin{aligned} C_{SW}^d \; = \frac{64}{\mathcal{N}_{SW}(N_N-1)^3}\Big(& k^3(N_N-1)^3 + B_{SW}^3 A_{(3,0)} - 3kqB_{SW}^2 A_{(2,1)} \\ & + 3(kq)^2 B_{SW} A_{(1,2)} - (kq)^3 A_{(0,3)}\Big), \end{aligned} \tag{46}$$

where

$$\begin{aligned} B_{SW} &= (N_N-1)\Big(1-q+\frac{2kq}{N_N-1}\Big), \\ \mathcal{N}_{SW} &= \frac{8kN_N}{N_N-1}(2k(q^2-2q+2N_N-2)-(N_N-1)(q^2-2q+2)) \end{aligned}$$

and

$$A_{(i,j)} = \sum_{m=1}^{N_N-1} \frac{\left\{\sin\left(\frac{km}{N_N}\pi\right)\cos\left(\frac{(1+k)m}{N_N}\pi\right)\right\}^i \left\{\sin\left(\frac{(1+2k)m}{N_N}\pi\right)\right\}^j}{\left\{\sin\left(\frac{m}{N_N}\pi\right)\right\}^{-(i+j)}}. \tag{47}$$

For $q = 1$, equation (46) yields the corresponding expression for the total global clustering coefficient of ER graphs, namely

$$\begin{aligned} C_{ER}^d = \frac{64}{\mathcal{N}_{ER}(N_N-1)^3}\Big(& k^3(N_N-1)^3 + B_{ER}^3 A_{(3,0)} - 3kB_{ER}^2 A_{(2,1)} \\ & + 3k^2 B_{ER} A_{(1,2)} - k^3 A_{(0,3)}\Big) \end{aligned} \tag{48}$$

with $B_{ER} = 2k$ and

$$\mathcal{N}_{ER} = \frac{8kN_N}{N_N - 1}\Big(2k(2N_N - 3) - (N_N - 1)\Big).$$

For the average geodesic graph distance, we utilize equation (16), i.e.

$$\langle d\rangle_{ER} = \frac{\ln(N_N) - \gamma}{\ln(Co \cdot N_N)} + \frac{1}{2}, \tag{49}$$

which provides an excellent analytical approximation for the geodesic distance for Erdős-Rényi random graphs [11] (see figure 1(c), compare black dots with downward triangles). Unfortunately, no valid analytical expression or approximation for the geodesic distance has yet been proposed for Watts-Strogatz small-world graphs. The only known expression is [24]

$$\langle d\rangle_{SW} = \frac{N_N}{2k\sqrt{(kqN_N)^2 + 2kqN_N}} \tanh^{-1}\sqrt{\frac{kqN_N}{kqN_N + 2}}, \tag{50}$$

which provides reasonable approximations only in a narrow parameter regime for $q \ll 1$ and small k, and therefore will not be considered here.

To still assess the small-worldness index (45) analytically, taking the lack of an analytical expression for $\langle d\rangle_{SW}$ into account, we focused on three different approaches. We first studied the maximal value S^d can take. As $\langle d\rangle_{SW}$ is a monotonically decreasing function of q with $\langle d\rangle_{SW} = \langle d\rangle_{ER}$ for $q = 1$, we can define

$$S^d_{max} = \frac{C^d_{SW}}{C^d_{ER}}, \tag{51}$$

which provides an upper bound for S^d, i.e. $S^d \leq S^d_{max}$. We note that, due to the monotonic behavior of $\langle d\rangle_{SW}$ as function of the rewiring probability, S^d will always be smaller, especially for small q. Thus, equation (51) provides a valid upper bound of the small-worldness index for $q \ll 1$, i.e. close to symmetric ring-graph regime of the Watts-Strogatz graph model.

Secondly, we investigated the theoretical lower bound of S^d. As $\langle d\rangle_{SW}$ is maximal for $q = 0$, i.e. a symmetric ring-graph, we define

$$S^d_{min} = \frac{C^d_{SW}}{C^d_{ER}} \frac{\langle d\rangle_{ER}}{\langle d\rangle_{RG}}, \tag{52}$$

where

$$\langle d\rangle_{RG} = \frac{1}{N_N - 1}\Big(\frac{N_N - 1}{2}(a + 1) - \frac{k}{2}a^2 - \frac{k}{2}a\Big)$$

with

$$a = \Big\lfloor \frac{N_N - 1}{2k} \Big\rfloor.$$

As for a given graph $\langle d\rangle_{ER}/\langle d\rangle_{RG}$ is fixed, and C^d_{SW} is monotonically decreasing with rewiring probability q, reaching its minimum $C^d_{SW} = C^d_{ER}$ for $q = 1$, the small-worldness index S^d will always be larger, especially for larger values of q. Thus, equation (52)

provides a valid lower bound of the small-worldness index for $q \gg 0$, i.e. close to random regime of the Watts-Strogatz graph model.

Finally, we explored the behavior of the small-worldness index in a semi-analytical fashion by replacing $\langle d \rangle_{SW}$ with numerical values obtained from analyzing corresponding semi-analytical small-world models constrained by choosing N_N and $k = \lfloor \frac{1}{2} Co(N_N - 1) \rfloor$ from their neural counterparts. This allows to define

$$S^d_{sa} = \frac{C^d_{SW}}{C^d_{ER}} \frac{\langle d \rangle_{ER}}{\langle d \rangle^{num}_{SW}}, \tag{53}$$

where $\langle d \rangle^{num}_{SW}$ denotes the mean of the average geodesic graph distance obtained from numerical analysis of constructed Watts-Strogatz graphs for a given parameter set.

Figure 5(a) visualizes the borderline small-world regime (small-worldness between 1 and 3) obtained from analytical (S^d_{min}, equation (52); S^d_{max}, equation (51)) and semi-analytical (S^d_{sa}, equation (53)) considerations in a representative graph of size $N_N = 350$ for $1 \leq k \leq 50$ and $0.0001 \leq q \leq 1$. We observe that the borderline region resides, for sufficiently small value of k (i.e. sufficiently sparse graphs), only at larger values of the rewiring probability q (figure 5(a), dotted lines). This suggests, in accordance with the findings reported earlier, a dominant random component in the makeup of the corresponding graphs.

To quantify this finding, we then investigated the behavior of S^d_{min}, S^d_{max} and S^d_{sa} as function of q in the semi-analytical graph models with N_N and k chosen from their corresponding neural graphs (figure 5(b)). Specifically, we numerically assessed in each semi-analytical model the rewiring probability q_{sa} necessary to reproduce the small-worldness index S^d observed in the corresponding neural graph (figure 5(b), intersection of green dashed and solid blue line), i.e. $S^d_{sa}(q_{sa}) = S^d$. Furthermore, we estimated the rewiring probability $q_{S=max}$ at which S^d_{sa} reaches its maximum (figure 5(b), blue dotted line), in order to quantify the "optimal" small-world regime for a given graph model.

The pooled result of this analysis applied to all investigated neural graphs is presented in figure 5(c). We observe that the q_{sa} for all neural graphs is large ($q_{sa} \geq 0.32$) and can reach values of 0.68 for graphs with small connectedness. This shows that for the investigated neural graphs, under the premise that the latter are small-worlds, a significant rewiring of at least 30% and up to 70% of edges is required to construct the neural graphs with a given small-worldness index from the Watts-Strogatz model. More importantly, the rewiring probability $q_{S=max}$ at which the semi-analytical small-world model reaches its optimal small-world regime is in all investigated cases about one order of magnitude lower than q_{sa}. This shows that neural graphs, even under the premise of being small-worlds, reside far away from the optimal small-world regime of the Watts-Strogatz model, and that the latter has to be taken into account when assessing the small-worldness properties of graphs.

5. Discussion

In this work, we have conducted a comprehensive analysis of the small-world property in structural neural graphs. We started with a classical numerical quantification of the clustering coefficient and average graph geodesic distance (sections 4.1 and 4.2, respectively). From this, we then calculate in each neural graph the small-worldness index, which precisely quantifies the coexistence of high clustering and low average path length in relation to that observed in random graphs (section 4.3). We observed that, in none of the neural graphs considered here does the small-worldness index exceed a value of 3, which suggests that neural graphs fall into a "borderline" regime close to random graphs [6].

To corroborate this observation, we then employed known analytical expressions for the clustering coefficient in random and small-world graphs, and for the average geodesic distance in random graphs, and constructed semi-analytical small-world models corresponding to the investigated neural graphs (section 4.4). By studying the numerical range the global small-worldness index takes in these models as a function of the remaining free parameter, the rewiring probability, we conceived of a novel method to quantify the small-worldness property of a given graph. Taken alone, the small-worldness index introduced in [6] is meaningless, as it highly variable and strongly depends on the size and connectivity of the graph in question. Only in relation to the maximal small-worldness possible for a given graph can an objective quantification of the small-world property be made. Specifically, while the maximal small-worldness index in the semi-analytical small-world graphs is highly variable, with values ranging from 1.79 (for the macaque visual cortex, MVC1) to 9.70 (for the *C. elegans* neural connectivity graph, CE3), the actual values observed in the corresponding neural graphs are, in all cases, much lower (1.17 for MVC1, 1.68 for CE3).

In order to quantify this difference, we then needed to ask the question: which rewiring is required to obtain, in these semi-analytical small-world models, the small-worldness index observed in the neural data? We found that the observed small-worldness can only be reproduced by a rewiring that is about one order of magnitude larger than that which produces the optimal small-worldness. This finding holds for all structural neural graphs investigated here, and it means that the rewiring required to reproduce the small-worldness observed in the neural data is far larger than the rewiring of only a few edges first proposed by Watts and Strogatz as a mechanism to bring a regular ring-graph into the small-world regime. This result suggests that neural graphs reside far closer to randomness than to optimal small-worldness.

In recent years, neural connectivity diagrams of the human cerebral cortex with a higher number of nodes than the graphs considered here have been introduced [9]. There, the small-worldness index was observed to increase from 1.54 in regional connectivity to 10.64 in high-resolution connectivity. Such an increase, however, is to be expected, given the known linear scaling of small-worldness with network size [6]. Because the linear increase of small-worldness with network size will also apply to the maximal small-

worldness index, a quantification of the small-world properties must again take both the actual and maximal possible small-worldness index into account. Indeed, using the parameters of the high-resolution connectivity map from [9] in the framework of our study, we found that the maximal small-worldness index for the constructed semi-analytical model is 16.28, reached at a rewiring probability of 2% of edges, while the reported small-worldness index of 10.64 requires a rewiring of 11.5% of edges. This additional analysis thus demonstrates that the findings of this study are applicable even to the next generation of large-scale reconstructions of neural connectivity.

In conclusion, in this study we confirm the often reported finding, that the clustering of neural connectivity exceeds that of a random graph while the distance is approximately the same. We continue, however, to observe that the quantification of these measures is not enough for a complete assessment of the small-world property in a given real-world graph. Here, we propose that a meaningful assessment of the small-worldness must take the maximal small-worldness into account, as well as relate the rewiring necessary to achieve this maximal and actual small-worldness. By doing so, we have observed in this study that neural graphs, though often discussed in the context of small-world property, in fact reside far outside the optimal small-world regime.

Acknowledgments

The authors wish to thank OD Little for inspiring comments. Work supported by the CNRS and the European Community (BrainScales project, FP7-269921). LM is a PhD fellow from École des Neurosciences de Paris (ENP), and also received support from ANR (Complex-V1).

[1] Watts DJ and Strogatz SH (1998) Collective dynamics of 'small-world' networks. Nature 393: 440-442.
[2] Albert R, Hawoong J, and Barabási A (1999) Diameter of the world-wide web. Nature 401: 130-131.
[3] Amaral LAN, Scala A, Barthelemy M, and Stanley HE (2000) Classes of small-world networks. Proc Natl Acad Sci USA 97: 11149-11152.
[4] Conyon MJ and Muldoon MR (2006) The small world of corporate boards. J Business, Finance, and Accounting 33: 1321-1343.
[5] Sporns O, Zwi J (2004) The small world of the cerebral cortex. Neuroinformatics 2: 145-162.
[6] Humphries MD and Gurney K (2008) Network 'small-world-ness': A quantitative method for determining canonical network equivalence. PLoS ONE 3: e0002051.
[7] He Y, Chen ZJ, and Evans AC (2007) Small-world anatomical networks in the human brain revealed by cortical thickness from MRI. Cerebral cortex 17: 2407-2419.
[8] Li M, Chen H, Wang J, Liu F, Long Z, Wang Y, Iturria-Medina Y, Zhang J, Chunshui Y, and Chen H (2014) Handedness-and hemisphere-related differences in small-world brain networks: a diffusion tensor imaging tractography study. Brain connectivity, 4: 145-156.
[9] Hagmann P, Cammoun L, Gigandet X, Meuli R, Honey CJ, Wedeen VJ, and Sporns O (2008) Mapping the structural core of human cerebral cortex. PLoS biology 6: e159.
[10] M. Rudolph-Lilith and L. E. Muller, Algebraic approach to small-world network models. Phys. Rev. E 89, 012812 (2014).
[11] Fronczak A, Fronczak P and Hołyst JA (2004) Average path length in random networks. Phys Rev E 70: 056110.
[12] Boccaletti S, Latora V, Moreno Y, Chavez M and Hwang D-U (2006) Complex Networks: Structure and Dynamics. Phys Rep 424: 175-308.
[13] Newman M (2010) Networks: an introduction. Oxford UP.
[14] http://www.cydyns.com; http://www.newscienceportal.com/MLR
[15] Varshney LR, Chen BL, Paniaqua E, Hall DH and Chklovskii DB (2011) Structural properties of the C. elegans neuronal network. PLoS Comput Biol 7: e1001066.
[16] Maslov S and Sneppen K (2002) Specificity and stability in topology of protein networks. Science 296: 910-913.
[17] Del Genio CI, Kim H, Toroczkai Z and Bassler KE (2010) Efficient and exact sampling of simple graphs with given arbitrary degree sequence. PLoS One 5: e10012.
[18] Kim H, Del Genio CI, Bassler KE and Toroczkai Z (2012) Constructing and sampling directed graphs with given degree sequences. New J Phys 14: 023012.
[19] Newman MEJ, Strogatz SH and Watts DJ (2001) Random graphs with arbitrary degree distributions and their applications. Phys Rev E 64: 026118.
[20] Rudolph-Lilith M, Destexhe A and Muller LE (2012) Structual Vulnerability of the Nematode Worm Neural Graph. arXiv:1208.3383v1 [cond-mat.dis-nn]. Available: http://arxiv.org/abs/1208.3383v1.
[21] Rudolph-Lilith M and Muller LE (2013) Aspects of randomness in biological neural graph structures. *submitted*.
[22] Holland PW and Leinhardt S (1971) Transitivity in structural models of small groups. Comparative Group Studies 2: 107-124.
[23] Fagiolo G (2007) Clustering in complex directed networks. Phys Rev E 76: 026107
[24] M. E. J. Newman, C. Moore and D. J. Watts, Mean-field solution of the small-world network model. Phys. Rev. Lett. 84, 32013204 (2000).

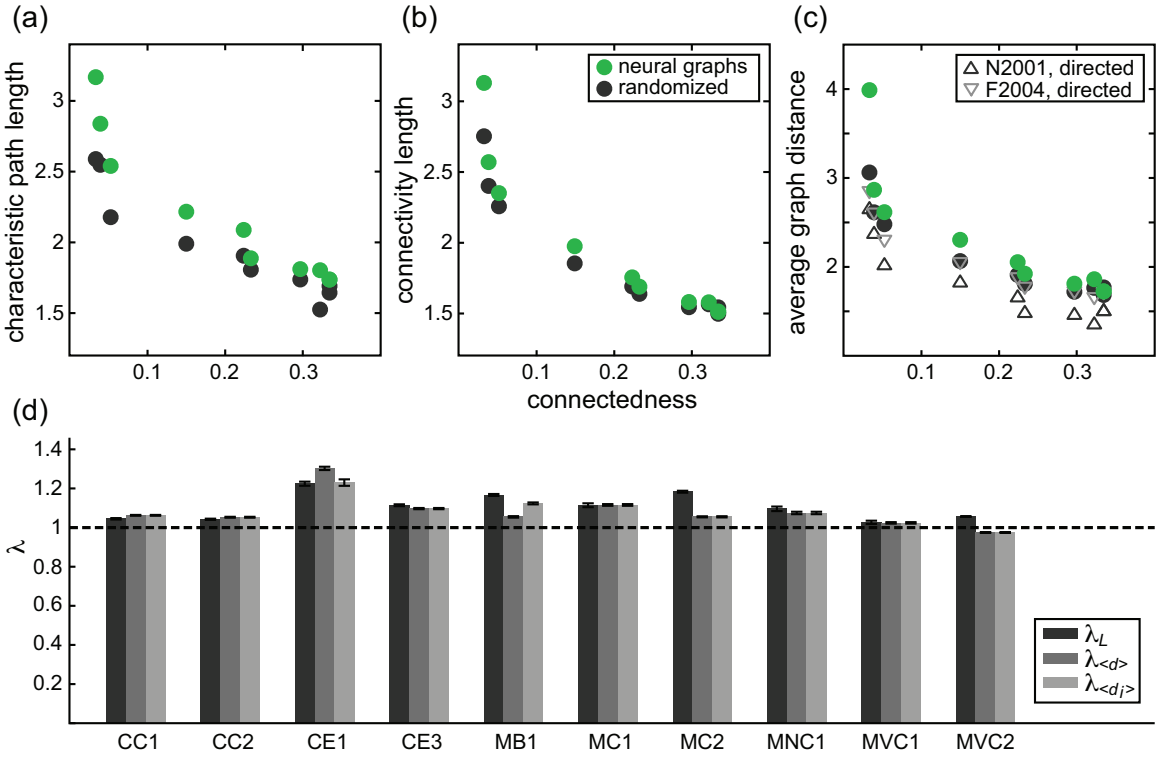

Figure 1. Geodesic distances in neural graphs. **(a-c)** Characteristic geodesic path length L (Eq. 12, **a**), geodesic connectivity length ℓ (Eq. 13, **b**) and average geodesic graph distance $\langle d \rangle$ (Eq. 10, **c**) as function of graph connectedness Co (filled dots) and their degree-preserved randomized counterparts (open dots). Theoretical prediction for the average geodesic graph distance are displayed in grey (upward triangles: Eq. 15 from [19]; downward triangles: Eq. 16 from [11]). **(d)** Ratio (mean ± SD) λ_L (Eq. 17, black), $\lambda_{\langle d \rangle}$ (Eq. 18, dark grey) and $\lambda_{\langle d_i \rangle}$ (Eq. 19, light grey) between the characteristic geodesic path length L, average geodesic graph distance $\langle d \rangle$ and average geodesic node distance $\langle d_i \rangle$ of the considered biological graphs and their corresponding EDM graphs, respectively.

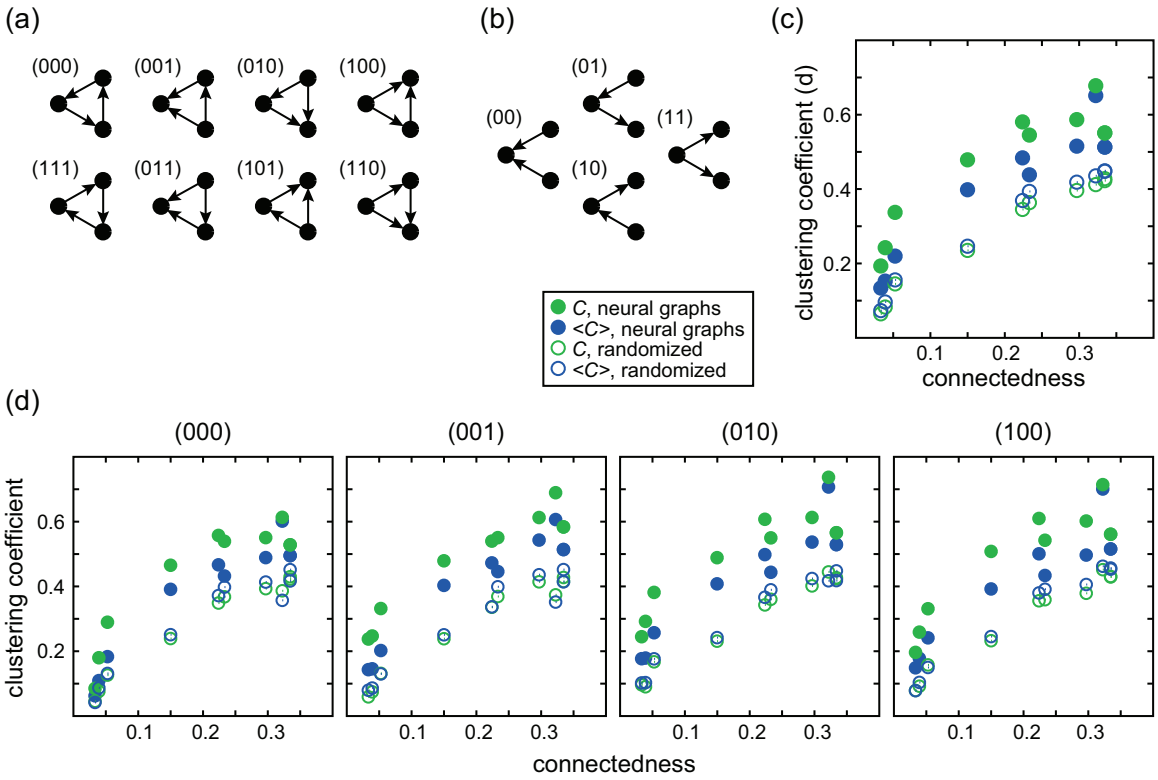

Figure 2. Clustering in neural graphs. **(a-b)** Clustering in digraphs. There are eight possible triangular relationships when taking the edge direction into account, defining triangles of type (abc) (panel **a**), and four relationships defining connected triples of type (ab) (panel **b**). From these eight triangles, four equivalent pairs exist, leaving four distinct motifs defining clustering coefficients for digraphs (Eqs. 23-27 and Eqs. 28-32). **(c)** Total graph clustering coefficient C^d (Eq. 27) and average node clustering coefficient $\langle C_i^d \rangle$ (Eq. 34) as a function of graph connectedness. Filled dots indicate the neural graphs, while open dots indicate their degree-preserved randomized counterparts. **(d)** Individual graph clustering coefficients $C^{(abc)}$ (Eqs. 23-26) and node clustering coefficients $\langle C_i^{(abc)} \rangle$ (Eq. 33), plotted as in panel **c**.

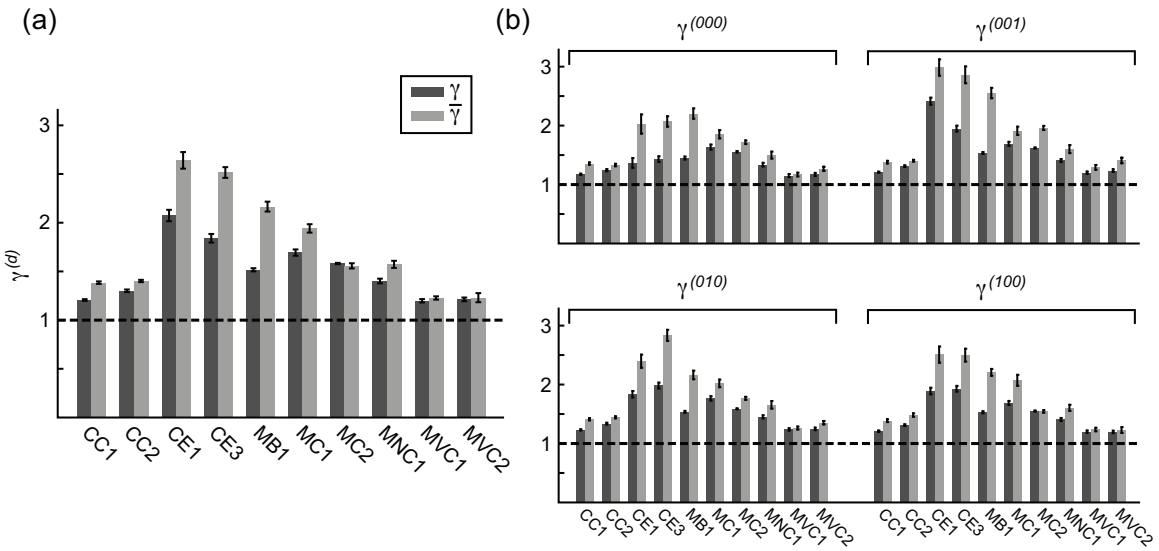

Figure 3. Clustering in neural graphs. **(a)** The ratio (mean ± SD) γ^d between the total clustering coefficient C^d in the observed and randomized graph (equations (35)–(36). $\bar{\gamma}^d$ is plotted in the same way (Eq. 38). **(b)** The ratios $\gamma^{(abc)}$ and $\bar{\gamma}^{(abc)}$ for the individual graph and node clustering coefficients are plotted as in panel **a**.

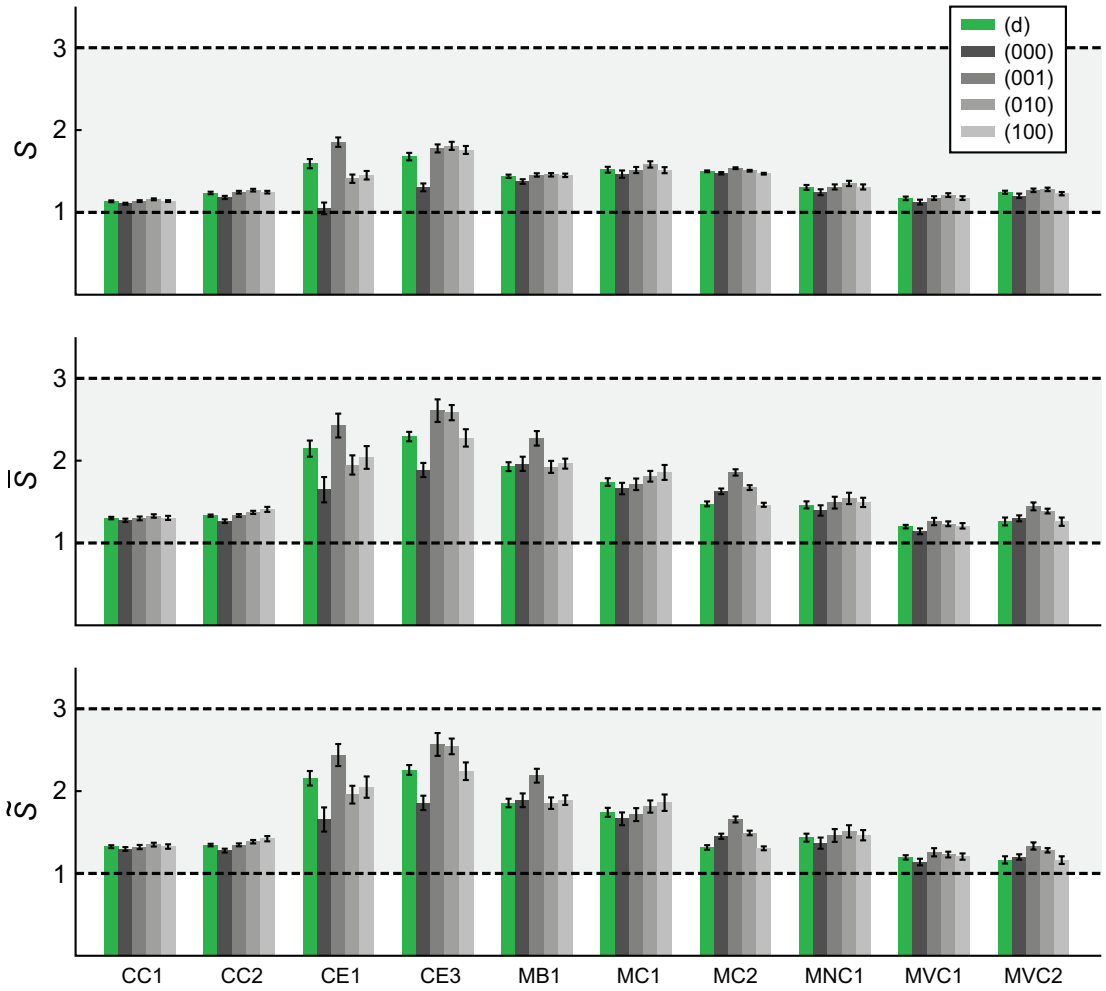

Figure 4. Small-worldness of biological neural graphs. Small-worldness indices (mean ± SD) S (Eqs. 39-44; top), $\bar{S}$ (Eqs. 41-42; middle) and $\tilde{S}$ (Eqs. 43-44; bottom) for the total graph clustering coefficient (green bar) and the individual clustering coefficients (grayscale bars). Above $\{S^{\bullet}, \bar{S}^{\bullet}, \tilde{S}^{\bullet}\} = 1$ (black dashed), a graph is considered a small-world, with $1 < \{S^{\bullet}, \bar{S}^{\bullet}, \tilde{S}^{\bullet}\} < 3$ (light grey region) defining borderline small-worlds, according to [6].

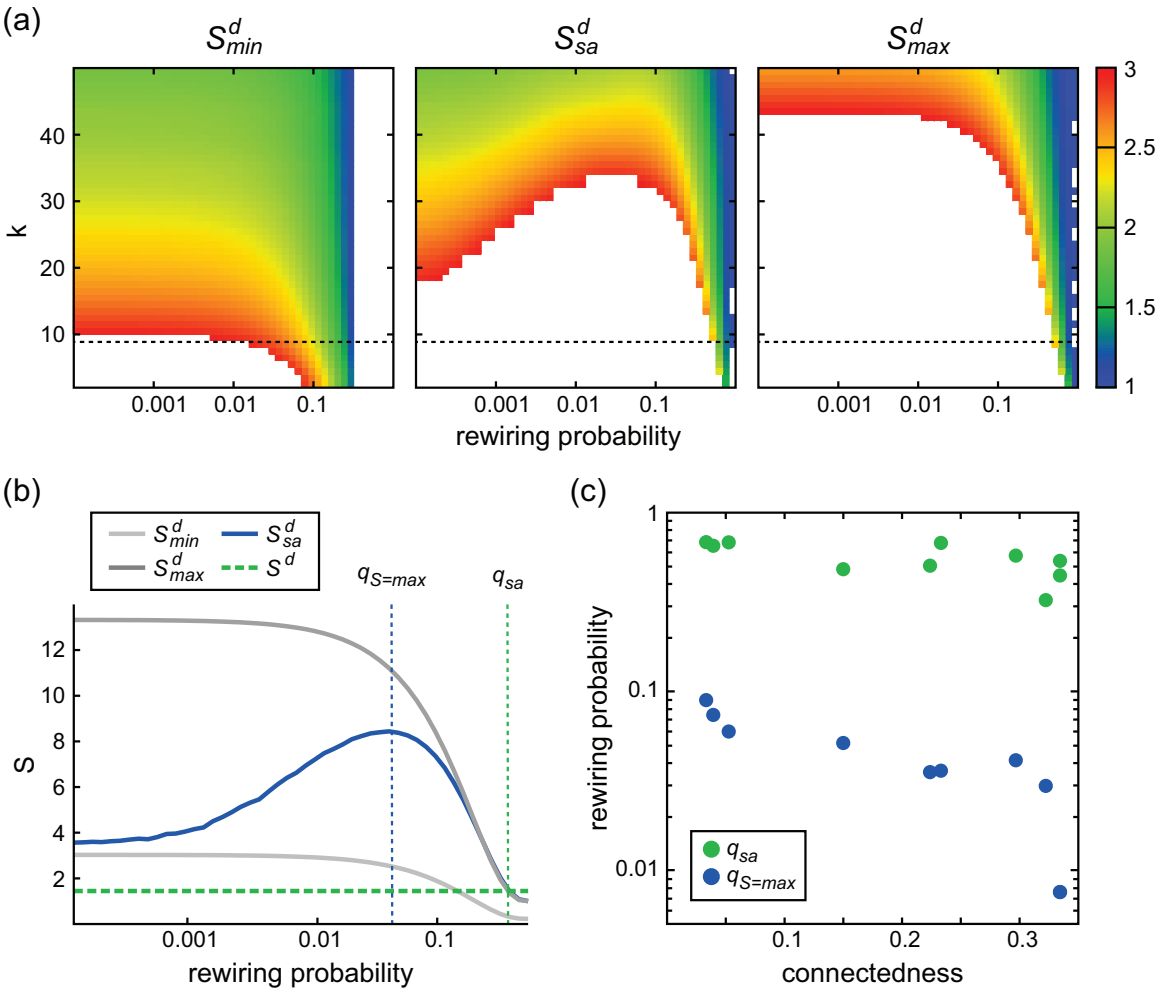

Figure 5. Analytical and semi-analytical investigation of borderline small-worlds. **(a)** Small-worldness index S^d_{min} (Eq. (52), S^d_{sa} (Eq. (53) and S^d_{max} (Eq. (51) as function of rewiring probability q and degree k of Watts-Strogatz graphs for $N_N = 350$ (corresponding to the representative MB1 neural graph). Displayed is the borderline region, defined as $1 \leq S^d_{min}, S^d_{sa}, S^d_{max} \leq 3$. The dashed line marks the degree $k = \frac{1}{2}Co(N_N - 1)$ which corresponds to the investigated MB1 graph. **(b)** Analytical small-worldness index for a small-world graph with $N_N = 350$ and $k = 9$, corresponding to the MB1 neural graph. The horizontal dashed line indicates the small-worldness index observed in the MB1 graph. The vertical dotted lines mark the q values at which the analytical curves cross the observed S^d value (cyan: S^d_{min}; red: S^d_{sa}; blue: S^d_{max}) and S^d_{sa} reaches its maximum (black). **(c)** Pooled result of the q values at which the small-worldness index observed in the investigated neural graphs is obtained (cyan: S^d_{min}; red: S^d_{sa}; blue: S^d_{max}). The black dots mark the rewiring probability in Watts-Strogatz graphs at which the maximal small-worldness is observed.

Chapter 6

Perspectives

Brief summary

In this thesis, we have traveled from topographic spiking network models, to the large-scale population activity of the monkey visual cortex, to the high-frequency oscillations in the human and monkey, and finally to the field of networks and graph theory. While the range of studies undertaken during the PhD work is large, it is my hope that a consistent method, or way of thinking about neural systems, has emerged, emphasizing mechanism in network structure and dynamics. Rather than attempt a detailed discussion of results and conclusions, in the following (and final) section, we will briefly introduce questions arising from these works, in a look toward future research.

Further questions

In **Work 1**, we observed a range of propagation speeds reported throughout the literature, varying across three orders of magnitude (from tens to thousands of millimeters per second). While the highest reported speeds are clearly due to fast conducting fibers (e.g. thalamocortical afferents, Metherate and Cruikshank, 1999), what explains the slowest propagation speeds, reported both *in vitro* (6.6 mm/s in Wu et al., 2001) and *in vivo* (50 – 70 mm/s in Xu et al., 2007; 10 – 16 mm/s in Han et al., 2008)? Here, the experiments *in vitro* may give a distinctive clue, because after pharmacological disinhibition in Wu et al. (2001), the propagation speed increases to 144 mm/s. This result can perhaps be explained simply by a shift in the excitability of single units, with the horizontal fibers (and associated synaptic delays) then setting an upper "speed limit" for the propagation. A similar explanation may furthermore be valid for the experimental results *in vivo*. We may speculate here that the observed speeds may likewise be due to an deeply anesthetized network state, which is known to affect relevant neuronal conductances (McCormick, 1992; Destexhe, 2009). This speculation is supported by our

observation that certain propagation features (specifically, the "wave reflection" effect observed in Xu et al., 2007) only occur in the most heavily anesthetized states (Muller et al., 2010). Finally, it must be noted that some differences in reported propagation speed could be due to the various quantification methods used. As discussed in **Chapter 4**, we are actively considering these questions in the context of the analysis of MEA recordings in the human and monkey.

In the context of **Work 2**, we have detected and confirmed propagating waves in VSD imaging data. A question still open, however, is to what extent the waves observed in these recordings represent supra- or sub-threshold activity. The origin of the VSD signal has been studied extensively, and is thought to come mainly from dendritic activity of superficial neurons in the cortex (Grinvald and Hildesheim, 2004; Chemla and Chavane, 2010). Thus, the key points, both for further experiments and models of cortical dynamics, are how these transient depolarizations traveling across the cortex affect local populations, and to what extent the propagation is mediated by the monosynpatic component elicited at the point of thalamocortical input. The spatial extent of spiking and subthreshold activity is known from experiments in the anesthetized cat (Bringuier et al., 1999), where the zone of depolarization evoked by thalamocortical input is observed to extent well beyond the spiking zone. The extent of these two zones, however, may change in the awake state, due to the increased synaptic bombardment of the ongoing activity (Rudolph et al., 2007; see also Ringach et al., 2002; Xing et al., 2004; Smith and Kohn, 2008; Chen et al., 2009, for values in the monkey), which typically brings neurons to a highly depolarized state characterized by transient fluctuations towards spiking threshold (Destexhe et al., 2003; Kuhn et al., 2004). This possibility should be addressed in future analysis of the multi-unit spiking activity, both in the spontaneous and stimulus-evoked case. Further, recently published experimental results (Tan et al., 2014) have provided a rare glimpse into the intracellular environment of cells in primary visual cortex of the awake, behaving monkey (with resting potentials far below spiking threshold), and provide another interesting starting point for future network models of these spatiotemporal dynamics in "low-level" sensory systems.

Finally, in the context of **Work 3**, we have demonstrated that, contrary to the often reported claims of high clustering and low path length in brain networks, neural graphs may in fact reside far outside the small-world regime. This statement has important implications for network science, where in recent years there has been a general optimism that complex networks can be described by simple models of connectivity. After observing that the small-world effect may not in fact provide an overarching description for neural connectivity, we have begun to ask what an appropriate model for neural connectivity might be. This question, and similar considerations, may form the basis for future models of neural connectivity, well-constrained by experimental data.

Appendices

Appendix A

Supporting works

Brüderle et al. (2011): *A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems*

Biol Cybern (2011) 104:263–296
DOI 10.1007/s00422-011-0435-9

ORIGINAL PAPER

Biological Cybernetics

A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems

Daniel Brüderle · Mihai A. Petrovici · Bernhard Vogginger · Matthias Ehrlich · Thomas Pfeil · Sebastian Millner · Andreas Grübl · Karsten Wendt · Eric Müller · Marc-Olivier Schwartz · Dan Husmann de Oliveira · Sebastian Jeltsch · Johannes Fieres · Moritz Schilling · Paul Müller · Oliver Breitwieser · Venelin Petkov · Lyle Muller · Andrew P. Davison · Pradeep Krishnamurthy · Jens Kremkow · Mikael Lundqvist · Eilif Muller · Johannes Partzsch · Stefan Scholze · Lukas Zühl · Christian Mayr · Alain Destexhe · Markus Diesmann · Tobias C. Potjans · Anders Lansner · René Schüffny · Johannes Schemmel · Karlheinz Meier

Received: 10 November 2010 / Accepted: 19 April 2011 / Published online: 27 May 2011

Abstract In this article, we present a methodological framework that meets novel requirements emerging from upcoming types of accelerated and highly configurable neuromorphic hardware systems. We describe in detail a device with 45 million programmable and dynamic synapses that is currently under development, and we sketch the conceptual challenges that arise from taking this platform into operation. More specifically, we aim at the establishment of this neuromorphic system as a flexible and neuroscientifically valuable modeling tool that can be used by non-hardware experts. We consider various functional aspects to be crucial for this purpose, and we introduce a consistent workflow with detailed descriptions of all involved modules that implement the suggested steps: The integration of the hardware interface into the simulator-independent model description language PyNN; a fully automated translation between the PyNN domain and appropriate hardware configurations; an executable specification of the future neuromorphic system that can be seamlessly integrated into this biology-to-hardware mapping process as a test bench for all software layers and possible hardware design modifications; an evaluation scheme that deploys models from a dedicated benchmark library, compares the results generated by virtual or prototype hardware devices with reference software simulations and analyzes the differences. The integration of these components into one hardware–software workflow provides an ecosystem for ongoing preparative studies that support the hardware design process and represents the basis for the maturity of the model-to-hardware mapping software.

D. Brüderle ( ) · M. A. Petrovici · B. Vogginger · T. Pfeil · S. Millner · A. Grübl · E. Müller · M.-O. Schwartz · D. H. de Oliveira · S. Jeltsch · J. Fieres · M. Schilling · P. Müller · O. Breitwieser · V. Petkov · J. Schemmel · K. Meier
Kirchhoff Institute for Physics, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
e-mail: bruederle@kip.uni-heidelberg.de

Present Address:
M. Schilling
Robotics Innovation Center, DFKI Bremen, Bremen, Germany

L. Muller · A. P. Davison · A. Destexhe
Unité de Neuroscience, Information et Complexité, CNRS, Gif sur Yvette, France

M. Diesmann
RIKEN Brain Science Institute and RIKEN Computational Science Research Program, Wako-shi, Japan

M. Diesmann
Bernstein Center for Computational Neuroscience, Universität Freiburg, Freiburg, Germany

M. Ehrlich · K. Wendt · J. Partzsch · S. Scholze · L. Zühl · C. Mayr · R. Schüffny
Institute of Circuits and Systems, Technische Universität Dresden, Dresden, Germany

J. Kremkow
Bernstein Center Freiburg, University of Freiburg, Freiburg, Germany

P. Krishnamurthy · M. Lundqvist · A. Lansner
Computational Biology, KTH Stockholm, Stockholm, Sweden

E. Muller
Brain Mind Institute, Ecoles Polytechniques Federales de Lausanne, Lausanne, Switzerland

T. C. Potjans
Institute of Neuroscience and Medicine (INM-6), Research Center Jülich, Jülich, Germany

T. C. Potjans
RIKEN Computational Science Research Program, Wako-shi, Japan

The functionality and flexibility of the latter is proven with a variety of experimental results.

Keywords Neuromorphic · VLSI · Hardware · Wafer scale · Software · Modeling · Computational neuroscience · PyNN

1 Introduction

1.1 Neuroscience and technology

Advances in neuroscience have often gone hand in hand with significant progress in the applied technologies, tools, and methods. While the experimental investigation of living neural tissue is indispensable for the generation of a detailed knowledge base of the brain, from which understanding of underlying principles can emerge, technological difficulties have always imposed limits to this endeavor. Until today it is not possible to study relevant observables in a sufficiently large fraction of brain tissue under realistic conditions and with a spatiotemporal resolution that is high enough to fully capture—and possibly consistently explain—the mechanisms of higher-order brain functions.

Therefore, in neuroscience, like in any other research field on dynamical systems that cannot be fully explored by experimental methods, models represent an indispensable approach to test hypotheses and theories on the real subject of interest. However, even neural modeling is significantly constrained and influenced by the set of available technologies. The spectrum of feasible experimental setups, in particular in *computational neuroscience*, directly depends on the accessible computational power. The difficulty of efficiently mapping the massive parallelism of neural computation in biological tissue to a limited number of digital *general purpose* CPUs is a crucial bottleneck in the development of large-scale computational models of neural networks, where statistics-intensive analyses or long-term observations of network dynamics can become computationally extremely expensive (see e.g., Morrison et al. 2005, 2007; Brette et al. 2006).

1.2 Neuromorphic hardware

For an alternative modeling approach, the so-called *neuromorphic engineering*, the technology-driven nature is even more obvious. In a physical, typically silicon form, neuromorphic devices mimic the structure and emulate the function of biological neural networks. This branch of neuroscience has its origins in the 1980s (Mead and Mahowald 1988; Mead 1989, 1990), and today an active community is working on analog or mixed-signal *VLSI*[1] models of neural systems (for reviews, see e.g., Renaud et al. 2007; Indiveri et al. 2009).

Dedicated implementations of said computational models are typically more power efficient compared to general purpose architectures and are well suited for e.g., embedded controllers of autonomous units like robots. Fault tolerance features observed in biological neural architectures are expected to apply to corresponding neuromorphic hardware implementations as well. This fact can offer one important way to create reliable computing components on the basis of future nano-scale hardware constituents, where current design strategies will run into serious yield problems. Moreover, the inherent parallelism of on-chip emulation of neural dynamics has the potential to overcome the aforementioned scaling limitations of pure software simulations.

Still, until today the focus of neuromorphic projects is mostly very application-specific. The majority of groups is working on neuromorphic sensors like e.g., silicon retinas and visual processing systems (Netter and Franceschini 2002; Delbrück and Liu 2004; Serrano-Gotarredona et al. 2006; Merolla and Boahen 2006; Fu et al. 2008; Gomez-Rodriguez et al. 2010) or motor control in robotics (Lewis et al. 2000). The requirement of communication with the environment is one important reason for the fact that nearly all neuromorphic devices reported so far are designed to operate in real time. However, even the projects that deal with mimicking, studying or applying neural information processing (Vogelstein et al. 2007), self-organization (Häfliger 2007; Mitra et al. 2009), or even hybrid setups coupling neuromorphic devices with living tissue (Bontorin et al. 2007) are usually focused on one type of neural architecture, one anatomical region, or one function the implemented network is supposed to fulfill.

Two main reasons for this self-limitation of neuromorphic development are the finite size of every neuromorphic device as well as the limited possibilities to change the behavior of individual cells and the network connection patterns once they have been cast into silicon. A typical approach to reduce size limitations is to scale up networks by inter-connecting multiple hardware modules (Costas-Santos et al. 2007; Berge and Häfliger 2007; Indiveri 2008; Schemmel et al. 2008). Furthermore, recent advances in neuromorphic development eventually promise to overcome the limited flexibility of hardware models by offering a sufficiently fine-grained configurability of both the neuron parameter values as well as the network connectivity (Indiveri et al. 2006, 2009; Schemmel et al. 2007, 2008, 2010; Ehrlich et al. 2007). This crucial feature allows to consider the utilization of neuromorphic systems as flexible modeling tools to approach open neuroscientific questions with new strategies (Kaplan et al. 2009; Brüderle et al. 2009, 2010; Brüderle 2009).

[1] Very large scale integration.

1.3 A novel methodological approach

The FACETS[2] research project (FACETS 2010) and its successor BrainScaleS (BrainScaleS 2010) aim at a comprehensive exploitation of the possibilities inherent to that approach. The highly interdisciplinary collaborations gather neurophysiological, theoretical, and hardware expertise in order to develop and operate a large-scale neuromorphic device that can serve as a flexible neural network emulation platform with hitherto unattained configurability and acceleration. It is planned to exploit this combination of features with experimental paradigms that are not realizable with pure software simulations, like long-term learning studies, systematic parameter explorations, and the acquisition of statistics for every tested setup.

Following this attempt, one important insight has emerged that has only rarely been addressed in the literature so far (exceptions are e.g., Dante et al. 2005; Oster et al. 2005): Any hardware device that is complex enough to serve as a useful neural modeling tool is useless without an appropriate software environment that implements a reasonable methodological framework for its operation. For any developed neuromorphic modeling platform, hard- and software have to form a functional unit. Moreover, the need for methods that have to be applied in order to make the advantages of a neuromorphic device accessible to non-hardware experts does not only refer to the actual *operation* of the device itself. Instead, already its *design process* needs to be supported and influenced by preparatory studies, e.g., with virtual versions of the future hardware–software system.

In this publication, we summarize the FACETS efforts to create a comprehensive methodological framework providing a workflow aiming to make the innovative FACETS wafer-scale hardware system a generic modeling tool that is applicable to a wide range of neuroscientific questions and accessible to the neuroscientific community.

1.4 Structure of this article

This introduction is followed by a description of the complete neuromorphic modeling infrastructure. This includes both the utilized hardware devices and the *workflow* that is in focus of this article, i.e., the framework of methods and software modules that have been developed for the design assistance, the benchmarking and the actual operation of the platform. A third section presents data and results that provide a proof of functionality for the concept as a whole. Various components of the workflow are evaluated, and the performance of benchmark model experiments performed with the complete system is studied and analyzed. The last section discusses the state of validation of the presented framework as well as its advantages and limitations considering alternative approaches. Implications and plans for future work and new perspectives arising from the presented achievements are outlined.

[2] Fast analog computing with emergent transient states.

2 The workflow components: modules and methods

The following section provides an overview over the complete infrastructure that has been developed to realize a novel neural modeling concept built around the FACETS neuromorphic hardware system. For this purpose, the neuromorphic device itself is presented in Sect. 2.1 on a level of detail that is appropriate to the method descriptions that follow. These methods are either implemented by or directly rely on an innovative software framework, which will be explained in Sect. 2.2 by means of its structure and concepts. A significant achievement for the targeted design and development of a harmonizing hardware–software unit forming the modeling platform was the collection of a set of literature-based benchmark model experiments, as summarized in Sect. 2.4.

The workflow that has been developed around these three main components is schematically depicted in Fig. 1: The library of dedicated neuroscientific benchmark models, including descriptions and measures to evaluate their correct functionality, has been established by members of the FACETS research project (FACETS 2010). For any model from within this set, a description on the basis of the simulator-independent modeling language PyNN (see Sect. 2.2.1) is available. The mentioned translation software stack performs an automated conversion of these scripts into appropriate data for the configuration and control of different hardware or software back-ends. The same stack also re-translates the resulting hardware output into the domain of its biological interpretation. During the development and optimization phase of the FACETS wafer-scale hardware system, an elaborate virtual version of the device (see Sect. 2.3) serves as a test bench for the development and tuning of all involved translation software modules.

In addition to the virtual wafer-scale device, a purely chip-based neuromorphic system (see Sect. 2.1.5) provides important information about characteristics of circuits planned to be implemented in the wafer-scale system. These ASICs thereby support the wafer-scale design process and the development of possible strategies to compensate unavoidable phenomena like transistor-level variations or noise. The outputs of all applied hardware or virtual hardware back-ends are compared with the target output descriptions included with the models in the benchmark library and with reference experiments on pure software simulators. The remaining differences are analyzed, as is exemplarily presented in Sect. 3.1.

In an ongoing optimization flow, the benchmark models are repeatedly mapped to the still evolving hardware

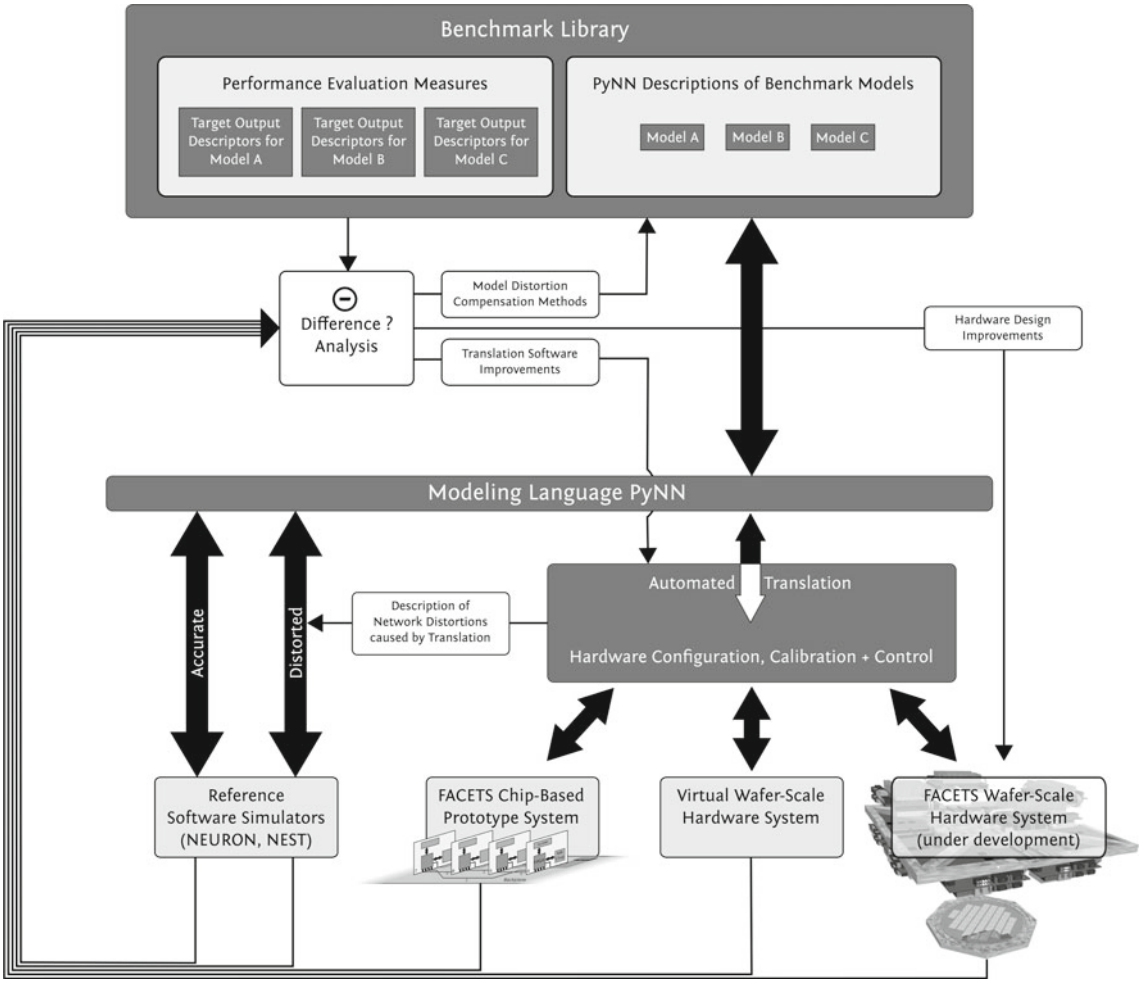

Fig. 1 Optimization workflow toward neural modeling with neuromorphic hardware devices. The main components are 1. the highly configurable FACETS neuromorphic hardware devices, 2. the software module stack that performs an automated translation of neural network experiments described with the modeling language PyNN into corresponding hardware configuration and control patterns, and 3. a benchmark library that contains a collection of neuroscientific models written in PyNN. For a detailed explanation of the complete flow and all individual steps and components see full text

substrate with the likewise continuously optimized software framework. The iteratively applied result analyses provide the fundament for improvements that close the workflow loop: The hardware design, the biology-to-hardware translation modules and optionally even the models themselves are modified such that the functional effects of remaining distortions caused by the model-to-hardware mapping process are minimized.

Hence, the first application of the presented workflow is to take novel types of hardware devices into operation. Furthermore, it can serve as a basic methodological paradigm for the actual target application of neuromorphic systems, i.e., the exploration and optimization of neural architectures by means of different optimization objectives. These include the search for computationally powerful structures or for setups that can reproduce biologically plausible dynamics.

While this section motivates and explains the workflow *as a whole* and provides descriptions of all involved components and methods, the scope of this article would be exceeded by providing detailed motivation for *all* particular choices of methods and components being part of the framework. The reasons for individual methodological or design decisions can be found in the literature referenced within the corresponding paragraphs.

2.1 The FACETS hardware system

In the following, the FACETS wafer-scale hardware system will be described with focus on conceptual and technical details that are relevant in the context of this article. More information on the hardware setup and circuitry can be found in Schemmel et al. (2008, 2010), Ehrlich et al. (2007), and Millner et al. (2010).

At the core of the FACETS wafer-scale hardware system (see Fig. 2) is an uncut wafer built from mixed-signal

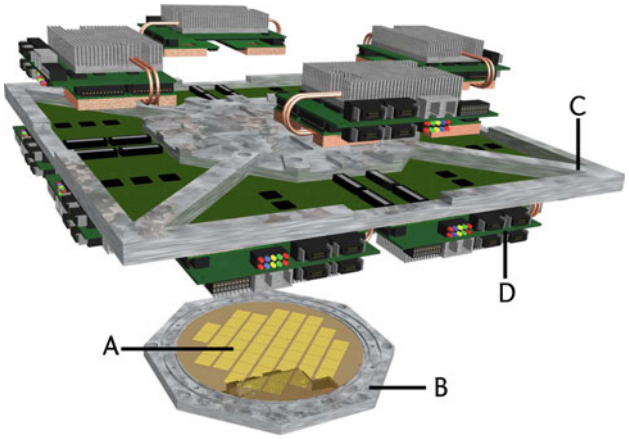

Fig. 2 The FACETS wafer-scale hardware system: Wafer (*A*) comprising HICANN building blocks and on-wafer communication infrastructure, wafer bracket (*B*), top frame (*C*) and digital inter-wafer and wafer-host communication modules (*D*)

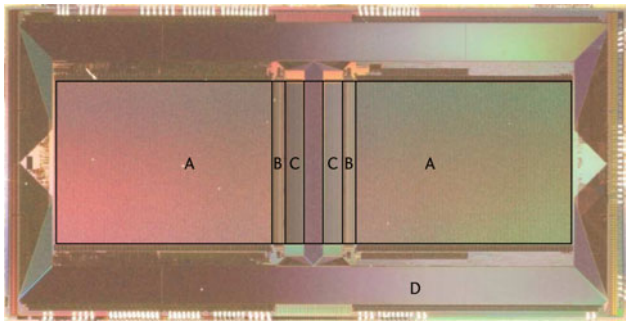

Fig. 3 A photograph of the HICANN building block with synapse arrays (*A*), neurons (*B*), floating-gate arrays (*C*), and L1 routing (*D*)

ASICs,[3] named *High Input Count Analog Neural Network* chips (*HICANNs*, Schemmel et al. 2008) that provide a highly configurable substrate which physically emulates adaptively spiking neurons and dynamic synapses. The intrinsic time constants of these VLSI model circuits are multiple orders of magnitude shorter than their biological originals. Consequently, the hardware model evolves with a speedup factor of 10^3 up to 10^5 compared to biological real time, the precise value depending on the configuration of the system.

In addition to a high-bandwidth asynchronous on-wafer event communication infrastructure, full custom digital off-wafer ASICs provide terminals for a packet-based multi-purpose communication network (Scholze et al. 2010). These so-called *Digital Network Chips (DNCs)* are backed by a flexible FPGA[4] design that handles the packet routing (Hartmann et al. 2010). The communication infrastructure is illustrated in Fig. 5. See Sect. 2.1.3 for details on the inter-chip communication scheme.

A full wafer system will comprise 384 interconnectable HICANNs, each of which implements more than 100,000 programmable dynamic synapses and up to 512 neurons, resulting in a total of approximately 45 million synapses and up to 200,000 neurons per wafer. The exact number of neurons depends on the configuration of the substrate, which allows to combine multiple neuron building blocks to increase the input count per cell.

2.1.1 Composition of the FACETS hardware system

The wafer as the main component for the FACETS wafer-scale hardware system has to be embedded into a framework that provides the electrical integration as well as the mechanical stability. The wafer has a diameter of 20 cm and will be placed into an aluminum plate which also serves as a heat sink. A multi-layer printed circuit board (PCB) is placed on top of the wafer. This PCB has to provide the fan-out of 1,500 impedance-controlled differential pairs and—in the worst case—has to deliver a total electrical power of 1,000 W to the wafer. A 14-layer fine pitch board with laser drilled microvias and a total size of 430 × 430 mm meets these requirements. The PCB will be clamped to an aluminum frame that is also used as a platform for communication devices such as the 48 DNCs and the 12 FPGA boards (see Sect. 2.1.3). Figure 2 shows a 3-D drawing of the hardware composition. All depicted electrical and mechanical components are custom-made by FACETS project members.

2.1.2 The HICANN building block

The HICANN building block shown in Fig. 3 is the neuromorphic ASIC of the FACETS wafer-scale hardware system. The inter-chip communication scheme is explained in Sect. 2.1.3.

Simplifying, the HICANN can be divided into four parts: the neuron circuits with their analog parameter storage based on floating-gate technology (Lande et al. 1996), an array of 114,688 dynamic synapses and the *Layer 1 (L1)* bus system interconnecting HICANNs on a wafer. The hardware neurons implemented by the HICANN building blocks (Millner et al. 2010) can emulate the adaptive exponential integrate-and-fire neuron model (*AdEx*, Brette and Gerstner 2005) which can produce complex firing patterns observed in biology (see e.g., Markram et al. 2004; Destexhe 2009), like spike-frequency-adaptation, bursting, regular spiking (RS), irregular spiking and transient spiking, by tuning a limited number of parameters (Naud et al. 2008). The decision to implement this particular neuron model in hardware was motivated by the large spectrum of possible and biologically relevant cell behavior realizable with a comparably compact circuitry. The latter fact is a crucial aspect when aiming at the

[3] Application-specific integrated circuit.

[4] Field programmable gate array.

integration of large numbers of neurons in one hardware system. For a neuromorphic implementation of Hodgkin-Huxley cells that consume significantly more chip area see e.g., Daouzli et al. (2008).

The model can be described by the following two differential equations for the membrane voltage V and the adaptation variable w and a reset condition specified further below:

$$-C_{\mathrm{m}}\frac{\mathrm{d}V}{\mathrm{d}t} = g_{\mathrm{l}}(V-E_{\mathrm{l}}) - g_{\mathrm{l}}\Delta_{\mathrm{t}}\mathrm{e}^{\left(\frac{V-V_{\mathrm{t}}}{\Delta_{\mathrm{t}}}\right)} + w + g_{\mathrm{e}}(t)(V-E_{\mathrm{e}}) + g_{\mathrm{i}}(t)(V-E_{\mathrm{i}}), \quad (1)$$

$$-\tau_{\mathrm{w}}\frac{\mathrm{d}w}{\mathrm{d}t} = w - a(V-E_{\mathrm{l}}). \quad (2)$$

C_{m}, g_{l}, g_{e}, and g_{i} are the membrane capacitance, the leakage conductance and the conductances for excitatory and inhibitory synaptic inputs, where g_{e} and g_{i} depend on time and on the inputs from other neurons. E_{l}, E_{i}, and E_{e} are the leakage reversal potential and the synaptic reversal potentials. The parameters V_{t} and Δ_{t} are the effective threshold potential and the threshold slope factor. The time constant of the adaptation variable is τ_{w}. The adaptation parameter a has the dimension of a conductance.

If the membrane voltage crosses a certain threshold voltage Θ, the neuron is reset:

$$V \rightarrow V_{\mathrm{reset}}, \quad (3)$$

$$w \rightarrow w + b. \quad (4)$$

The parameter b is responsible for spike-triggered adaptation.

A neuron can be constructed out of up to 64 so-called *denmem* circuits, each implementing the dynamics of the AdEx model and being connected to up to 224 synapses. This way a neuron could have synaptic inputs from up to 14,336 other cells. In addition, depressing and facilitating mechanisms of short-term synaptic dynamics (for a review, see Zucker and Regehr 2002) are implemented. A purely chip-based FACETS hardware implementation of this feature is described and applied in Bill et al. (2010).

A general limitation of neuromorphic implementations of cell models is the fact that configurable parameter values will always have limited ranges. The value ranges of all AdEx parameters configurable in the hardware implementation have been designed such that the complete set of biologically relevant firing patterns distinguished, e.g., in Naud et al. (2008) can be reproduced. If this design goal was fully achieved is currently studied with HICANN prototypes, i.e., work in progress (see also Sect. 3.4).

2.1.2.1 Hebbian Learning in the FACETS Hardware Long-term Hebbian learning in the FACETS hardware devices is implemented in every synapse as spike-timing-dependent plasticity (STDP, reviewed e.g., in Morrison et al. 2008). To ensure high flexibility in terms of mappable neuronal networks each neuron in hardware needs an appropriate number of synaptic inputs. However, due to limited die area, a trade-off between the number of synapses and the chip resources for a single synapse has to be made.

To achieve a minimal circuit size for the synapses, local correlation measurements and the local synaptic weight storage are separated from global weight *update controllers* (Schemmel et al. 2006, 2007). Causal and acausal correlations between pre- and post-synaptic spikes determine the temporal factor of the STDP rule described in Schemmel et al. (2004) and are accumulated locally until they are processed by the update controller. Synaptic weights are stored locally as digital values with a 4-bit resolution each. This resolution is again a trade-off between precision and chip resources and requires several correlated events to reach the next discrete weight value. If a sufficient amount of correlations is accumulated, the discrete weight is updated by the update controller. Since many synapses share one update controller a weight update is performed periodically with a frequency that has an upper limit determined by the circuitry (Schemmel et al. 2006). Since a reduced symmetric nearest-neighbor spike pairing scheme turned out to be one feasible approach for describing biological measurements (Burkitt et al. 2007; Morrison et al. 2008), this specific plasticity mechanism has been chosen to be implemented in hardware to further reduce the size of a synapse. Update controllers are modifying the synaptic weights by using look-up tables that are listing, for each discrete weight value, the resulting weight values in case of measured causal or acausal correlations. These look-up tables can be adapted to the weight-dependent factor of any STDP rule.

Despite its global weight update controllers, the STDP mechanism of the FACETS hardware has to be considered local to every synapse. The implementation of this particular model represents a project-wide decision on the most promising mechanism to be cast into silicon, taken in the early phase of FACETS. Recent developments in the modeling of learning and self-organization in neural networks (see e.g., Sjöström et al. 2008; Pfeiffer et al. 2010) combine such local rules with various global mechanisms like the reward-based modulation of large groups of synapses. With respect to more complex and relevant plasticity mechanisms, an extension to STDP rules with additional input parameters, e.g., membrane potentials, spike rates or global reward signals, is currently under development.

2.1.2.2 Parameter memories In contrast to most other systems, the FACETS wafer-scale hardware deploys analog floating-gate memories similar to cells developed by Lande et al. (1996) as storage devices for the analog parameters. Due to the small size of these cells, most parameters can be provided individually for a single neuron circuit. This way,

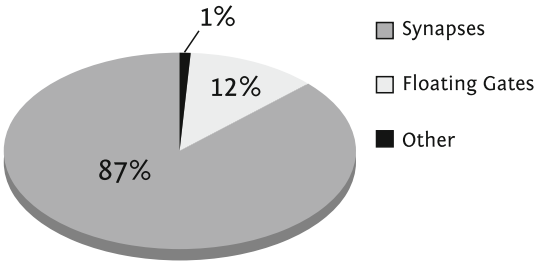

Fig. 4 Sector diagram of the parameter space to configure one HICANN chip. For a full wafer, the configuration data volume is 44 MB large

matching issues can be counterbalanced, and different types of neurons can be implemented on a single chip.

As a starting point for the parameter ranges, parameters from Brette and Gerstner (2005) and Destexhe et al (1998) have been used. The chosen ranges allow leakage time constants $\tau_{\text{mem}} = C_{\text{m}}/g_{\text{l}}$ at an acceleration factor of 10^4 between 1 and 588 ms and an adaptation time constant τ_{w} between 10 ms and 5 s in terms of biological real time. The parameters used by Pospischil et al. (2008), for example, lie easily within this range.

A substantial amount of digital memories is integrated in the chip, dominated by the synapse RAM. Each of the 114,688 synapses has 8 bit memory cells for weight and address storage. For the whole wafer, the synapse RAM alone is 38 MB large. Figure 4 shows the partitioning of the parameter memory on a HICANN building block. To compare the analog floating gates to normal digital memory, each cell has been counted as 10 bit, since this is the number of bits needed to program it.

2.1.3 Communication infrastructure

The communication infrastructure of the FACETS wafer-scale hardware is illustrated in Fig. 5. Pulse communication is generally based on the digital transmission of neural events representing action potentials, but a distinction in two network layers can be made. An asynchronous, serial protocol, named *Layer 1 (L1)* utilized by HICANNs at a wafer level provides intra-wafer action potential transmission on a high density direct interconnection grid. A second one, named *Layer 2 (L2)*, deploys the DNCs and FPGAs for synchronous, packet-based, intra- and inter-wafer communication and—compared to L1—establishes a more flexible routed network of lower density. To cope with inevitable jitter in routing delay, a time stamp is transmitted together with the address within the data packets of this network. A PC cluster that handles the mapping, configuration and control process described in Sect. 2.2 as well as the playback and recording of external stimuli to the neural network is connected to the FPGAs via multi-Gigabit Ethernet.

Activity is injected into the L1 network in the form of 6 bit serial pulse address packets by neurons that connect to the horizontal buses. Sparsely populated passive switch matrices at the intersections of horizontal and vertical buses pass the data to the vertical buses. Further sparse switch matrices connect to horizontal lines feeding synapse drivers that act as data sinks to the network. While crossing HICANN block boundaries the signals are refreshed by repeater circuits with active re-timing that are capable of driving the signals across one HICANN block. The sparseness of the switch matrices is chosen such that the repeater circuits are not overloaded while still providing maximum flexibility for implementing various neural network topologies (see Fieres et al. 2008; Schemmel et al. 2010 for more information on the underlying design decisions and analyses of the resulting limitations).

Connectivity between the HICANN blocks is established by edge connecting them in the layout. As illustrated in Fig. 5, this is only possible for eight HICANNs located within one reticle. A reticle is the largest producible unit on the wafer and no connections can be formed between reticles during standard CMOS fabrication. Wafer-scale connectivity is obtained using a post-processing method developed in the FACETS project. It offers two additional routing layers that can cover the whole wafer. By means of this technique, an inter-reticle connection pitch well below 10 μm can be achieved which facilitates the required connectivity. Furthermore, large landing pads are formed by the post-processing that connect the wafer to the system PCB via elastomeric stripe connectors (see Fig. 2 and Schemmel et al. 2010).

These stripe connectors are used to deliver all required power to the wafer. In addition, they connect high-speed communication signals between the HICANNs and the DNCs.[5] This high-speed communication interface transports configuration data as well as the above-mentioned L2 data packets. L2/L1 protocol conversion is performed inside the HICANN blocks, where L2 activity can either be injected to or read from the L1 network (see Fig. 5). The transport of the L2 packets is handled by the DNCs, which also implement a time stamp- based buffering and heap-sort algorithm (Scholze et al. 2010). Together with routing logic inside the FPGAs, the DNC–FPGA L2 network fulfills the QoS[6] demands (Philipp et al. 2009) for spiking neural networks, i.e., a constant delay at a low pulse loss rate. This is also true for inter-wafer connections routed through Ethernet switches connected to the FPGAs.

[5] For completeness it should be noted that also analog signals, e.g., selectable neuron membrane voltages, are transported through the stripe connectors.

[6] Quality of service.

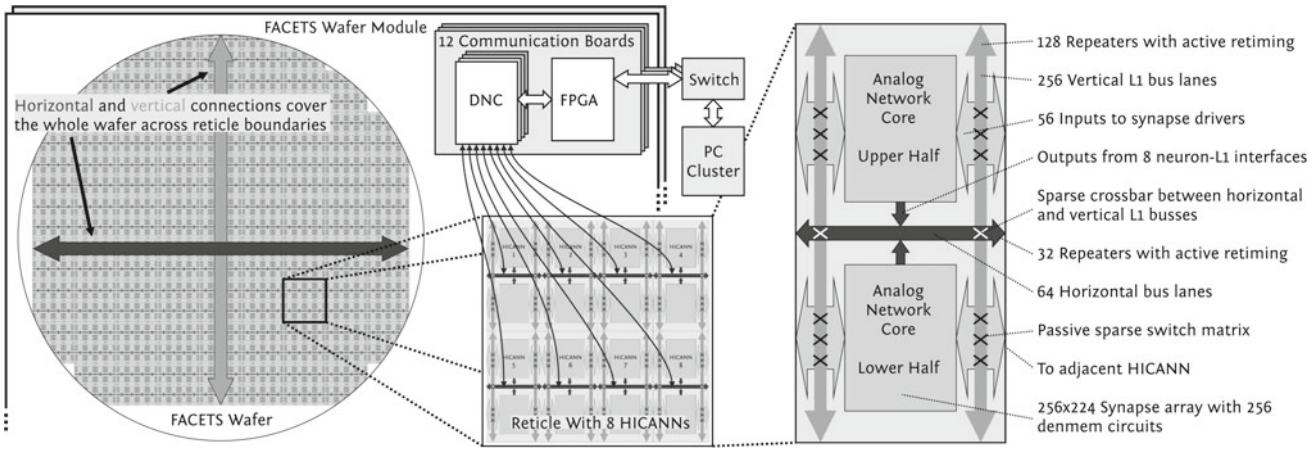

Fig. 5 Communication structure on a wafer module of the FACETS wafer-scale hardware system. Neural activity is transported horizontally (*dark gray*) and vertically (*light gray*) via asynchronous L1 buses on the HICANN building blocks. Repeater circuits at the edges of these blocks allow for a distribution of the buses over the whole wafer. Off-wafer connectivity is established by the L2 network via DNCs and FPGAs. It interfaces the L1 buses on the HICANN building blocks. Several wafer modules can be interconnected using routing functionality between the FPGAs via Ethernet switches

2.1.4 Host interface

The packet communication between wafer and host computer passes through several layers: DNCs, FPGA controller boards, and a Gigabit Ethernet layer (Norris 2003) have to be traversed. As each of the 12 FPGA controller boards (see C in Fig. 2) comprises two Gigabit ports dedicated for host communication, a total bandwidth of 24 GBit/s can be achieved. Standard networking switches concentrate these links into the required number of 10GBase-LX4 (Horak 2007) upstream ports. A standard PC cluster equipped with adequate network interface cards handles the traffic. A custom design ARQ[7]-style (Fairhurst 2002) protocol provides a reliable communication channel between the host computer and the hardware system. The FPGA controller boards act as remote terminals for these ARQ communication channels, but also provide system control functionality.

During experiments, most communication data—basically spike events—flow directly between host computer and FPGA controller boards. In contrast to this, in the initial setup stage almost all traffic—i.e., the system configuration data—is dedicated to wafer communication. In this case, the FPGA controllers act as simple transmission nodes between host computer and wafer. Both operational stages impose high demands on the communication bandwidth. The initial configuration space consumes around 50 MB (see Fig. 4). Every spike event is represented by a 15-bit time stamp and a 12-bit data field, comprising both DNC and HICANN identifiers. Thus, during an experiment approximately 1 GEvent/s can be transported to and from the host computer. At a speedup factor of 10^4, the corresponding total spike rate in the biological time domain is 100 kHz per wafer.

To meet these requirements set by the hardware scale, acceleration factor and modeling constraints, a highly scalable software implementation of the communication protocol was developed (see Sect. 2.2.10 and Schilling 2010). This multi-threaded protocol stack already provides a zero-copy API[8] to the upper software layers.

Furthermore, to support future applications, such as interfacing the FACETS hardware system to simulated environments which provide sensor output related to motor input, low round-trip times between these components are crucial. Such classes of in-the-loop experiments demand low latency communication and high bandwidth at the same time.

2.1.5 Chip-based neuromorphic system

On the development path toward the FACETS wafer-scale hardware platform, a purely chip-based neuromorphic system has been designed and built (Schemmel et al. 2006, 2007) and is in active use (Kaplan et al. 2009; Brüderle et al. 2009, 2010; Bill et al. 2010). It implements time-continuous leaky integrate-and-fire cells with conductance-based synapses and both a short-term and a long-term plasticity mechanism as described above for the wafer-scale device. Up to 16 of these ASICs, each of which provides 384 neurons and 10^5 configurable and plastic synaptic connections, can be operated individually or interconnected via a supporting backplane board. This board is connected via a single Gigabit Ethernet link to a host computer, through which multiple users

[7] Automatic repeat request.

[8] Application programming interface.

can access and use the neuromorphic devices in parallel. The possibility of remotely accessing the chips via the Internet in addition to setting up and running experiments with an available PyNN interface (see Sect. 2.2.1 and Brüderle et al. 2009) already now make this system a tool that is used for neuromorphic model exploration by users from various countries. Many circuit design strategies for the wafer-scale system are implemented for testing purposes in this chip-based device, including the STDP correlation measurements (see Sect. 2.1.2) located in every individual synapse. Basic plasticity studies supporting the design of the wafer-scale system, some of which are outlined in Sect. 2.5.3, incorporate investigations on the basis of experimental results from the chip-based devices.

2.2 Software framework

Figure 6 shows the stack of software modules that will be described in the following.

Its components seamlessly interact in performing an automated translation of arbitrary neural network experiment descriptions into appropriate data for hardware configuration and control. The same stack also automatically re-interprets the acquired hardware output into its biological interpretation. The top-level interface offered to hardware users to describe neural network setups is based on the simulator-agnostic modeling language PyNN. The concept of this approach and its advantages, especially for neuromorphic system operation, will be described in Sect. 2.2.1.

The process of mapping a PyNN description onto the configuration space of the FACETS hardware systems, including dedicated representation formats, will be described in Sects. 2.2.2–2.2.8. Sections 2.2.9 and 2.3 focus on the mapping analysis plus its testing and optimization on the basis of an elaborate virtual version of the wafer-scale hardware system. The special performance requirements for the low-level host-to-hardware communication software and the implemented corresponding solutions are outlined in Sect. 2.2.10.

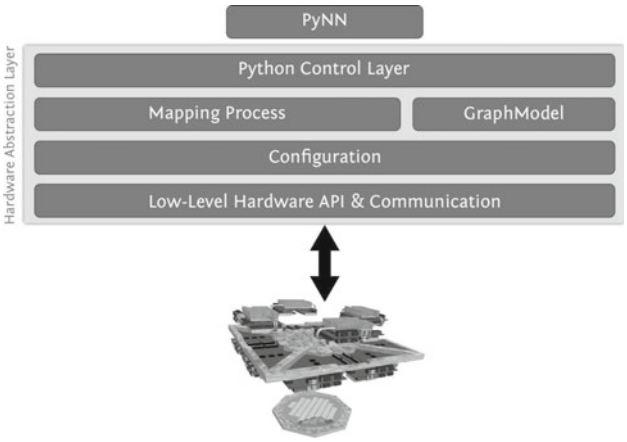

Fig. 6 Schematic of the *hardware abstraction layer*, i.e., the stack of software modules for the automated and bidirectional translation between PyNN-model descriptions and appropriate hardware configuration and control patterns. The individual modules are: a Python control layer, a mapping layer that operates on a graph-based data container (*GraphModel*), and low-level layers that deliver the generated hardware configuration patterns and control sequences via a dedicated communication protocol

2.2.1 PyNN & NeuroTools

PyNN is a simulator-independent, Python-based language designed for describing spiking neural network models (Davison et al 2008). It offers functions and classes for the setup and control of experiments, and it provides standard cell models as well as standardized dimension units. PyNN supports various software simulators like NEURON (Hines and Carnevale 2006; Hines et al. 2009), NEST (Gewaltig and Diesmann 2007; Eppler et al. 2008), Brian (Goodman and Brette 2008), and PCSIM (Pecevski et al. 2009). With PyNN, which is open source and well documented, a user can set up a neural network model, run it on any of the supported back-ends without changing the code, and directly compare the results. This provides the possibility to conveniently port experiments between different simulators, to transparently share models and results, and to verify data acquired from different back-ends (see Fig. 7).

The integration of the operating software framework for the FACETS hardware system into the PyNN concept (Brüderle et al. 2009; Brüderle 2009) is a crucial aspect of the presented neuromorphic workflow. One important motivation for this approach is to create a bridge between the communities of neuromorphic engineers and neural modelers, who have been working in rather separate projects so far. The back-end agnostic concept of PyNN, now also offering the possibility to port existing experiments between the supported software simulators and the FACETS hardware system, allows to benchmark and verify the hardware model. The API of PyNN is easy to learn, especially for scientists who have already worked with software simulators. Hence, PyNN represents an optimal way to provide non-hardware experts a convenient interface to work with the FACETS neuromorphic devices. In general, PyNN interfaces to neuromorphic systems make it possible to formulate transparent tests, benchmarks and feature requests, and therefore can influence and boost biologically oriented hardware development. They might, eventually, support the establishment of such emulation devices as useful modeling tools.

On top of PyNN, a library of analysis tools called NeuroTools (NeuroTools 2008) is being developed, which builds upon the interface and data format standards, but also exploits the possibility to incorporate third-party Python modules, e.g., for scientific computing and plotting (Oliphant 2007; Jones et al. 2001; Langtangen 2008; Hunter 2007). Thus, for

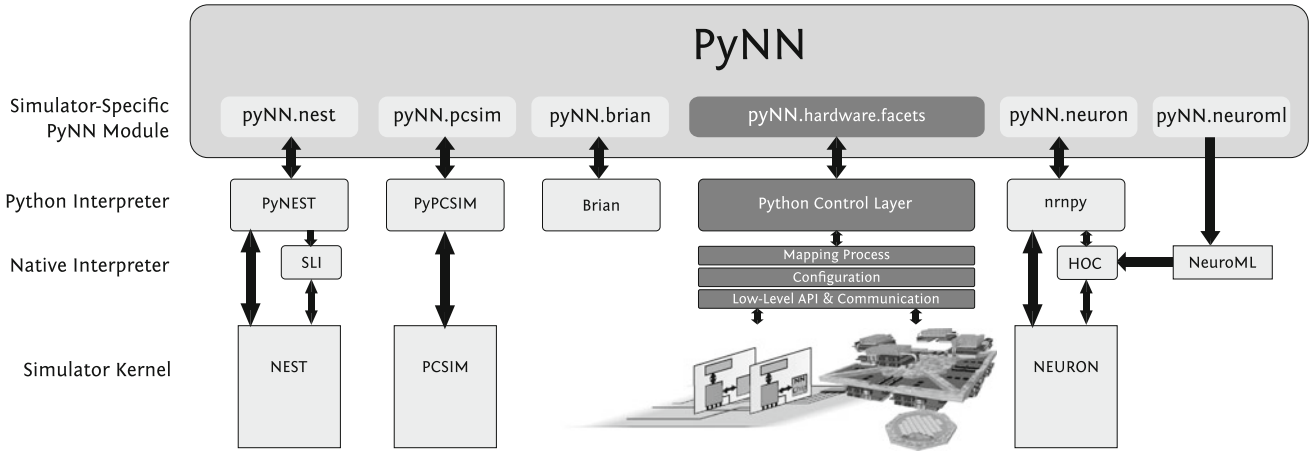

Fig. 7 Schematic of the simulator-independent modeling language PyNN. Like various established software simulators, the FACETS neuromorphic hardware systems have been integrated into the PyNN unification and standardization concept

all supported software simulators and for the FACETS neuromorphic hardware systems, all stages of neural modeling experiments—description, execution, result storage, analysis, and plotting—can be performed from within the PyNN and NeuroTools framework.

2.2.1.1 Simulations as reference for translation and calibration The hardware-specific PyNN approach incorporates quantitative bidirectional translation methods between the neuromorphic system dynamics and the biological domain, both in terms of electrical variables and the different time domains. This translation incorporates calibration routines that minimize the impact of transistor-level fixed-pattern noise on the behavior of neural and synaptic circuits. The translation and calibration scheme developed for the FACETS hardware systems directly involves reference software simulations for the biologically relevant gauging of hardware parameters, heavily exploiting the PyNN paradigm of unified setup descriptions. Section 2.2.7 provides more details on this.

2.2.2 Mapping process

The mapping process determines a valid routing network configuration and parameter value set as initial setup data for the FACETS hardware system. This takes into account topology constraints between hardware blocks such as connectivity, connection counts, priorities, and distances as well as source/target counts. Figure 8 depicts the single steps of the mapping process as described by Ehrlich et al. (2010).

The mapping is accomplished in the three main steps of *Placement*, *Routing*, and *Parameter Transformation & Calibration*, with an appropriate *Pre-* and *PostProcessing* of the configuration data. As the first three main steps are explained in more detail in the following, we will shortly summarize the functionality of the remaining parts.

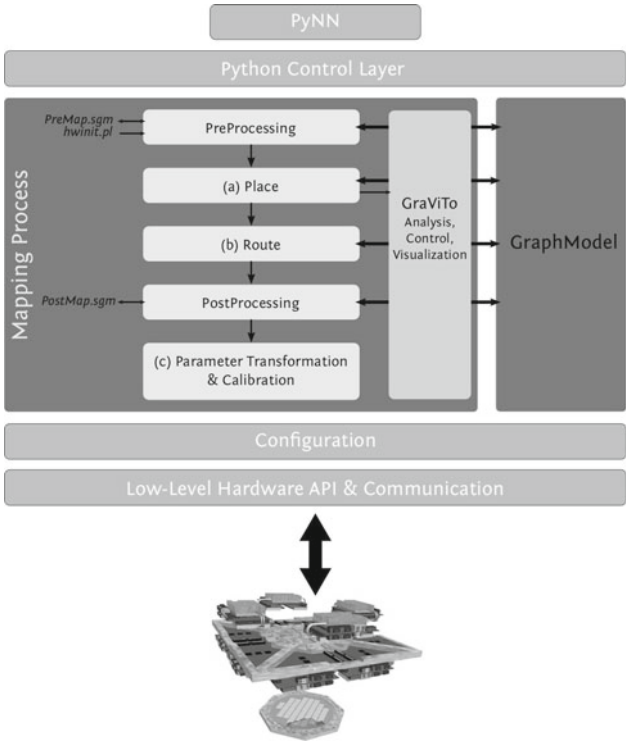

Fig. 8 Mapping PyNN neural network model descriptions onto the configuration space of the wafer-scale hardware system. The three main processing steps, all operating on one unified data container (GraphModel), are **a** the *placing* of neurons onto the available circuitry, **b** the realization of synaptic connections by appropriately configuring the available *routing* infrastructure on the device, and **c** the transformation of neuron and synapse parameters into corresponding parameter values offered by the highly configurable device. The latter step can involve calibration data to tune individual sub-circuits such that the effect of unavoidable transistor-level variations onto the mapped models is minimized

Starting with a neural architecture defined via PyNN, the first mapping step of *PreProcessing* reads in a description of the hardware (see Sect. 2.2.3), described using the novel

query language *GMPath* (see Sect. 2.2.4). It sets up an internal representation for both the hardware and the biological model in the form of a directed graph called *Graph Model* (see Sect. 2.2.4). Optionally, a so-called *PreMapping* netlist of the biological model can be streamed out into a file. Following placement and routing, the same applies for the *PostProcessing* with a *PostMapping* netlist, which includes the possibility to obtain a PyNN script that represents the (possibly distorted) network ultimately realized on the hardware back-end.

The individual steps of the process are automatically initiated and partly controllable via the PyNN module for the FACETS hardware system. Furthermore a stand-alone software named *GraViTo* is provided for the analysis of the mapping results (see Sect. 2.2.9).

2.2.3 Internal hardware description

Prior to the mapping process, we have to define the hardware in an abstract manner. For this purpose, we utilize the path language *GMPath* to set up an appropriate GraphModel (both described in Sect. 2.2.4) as a versatile internal representation.

In Fig. 9, a FACETS wafer-scale hardware setup—also applied in Ehrlich et al. (2010)—is illustrated.

As described in Sect. 2.1, the fundamental layer of the FACETS wafer-scale hardware is an array of reticles shown as light gray squares, housing the HICANN circuitry that implements neural functionality, with a second layer of DNCs above. The third and topmost layer represents a regular grid of FPGAs, colored dark gray.

2.2.4 The GraphModel container

A data model called *GraphModel* (Wendt et al. 2008) represents both the targeted biological and the configurable hardware structure within the mapping software. It can be characterized as a hierarchical hyper graph and consists of vertices (data objects) and edges (relationships between the vertices). A vertex contains a single data value. An edge can be one of the following types:

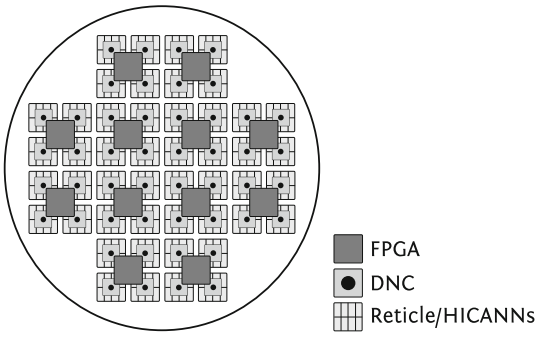

Fig. 9 Example FACETS wafer-scale hardware setup from Ehrlich et al. (2010): 12 FPGAs control 48 DNCs, which are connected to 384 HICANN ASICs

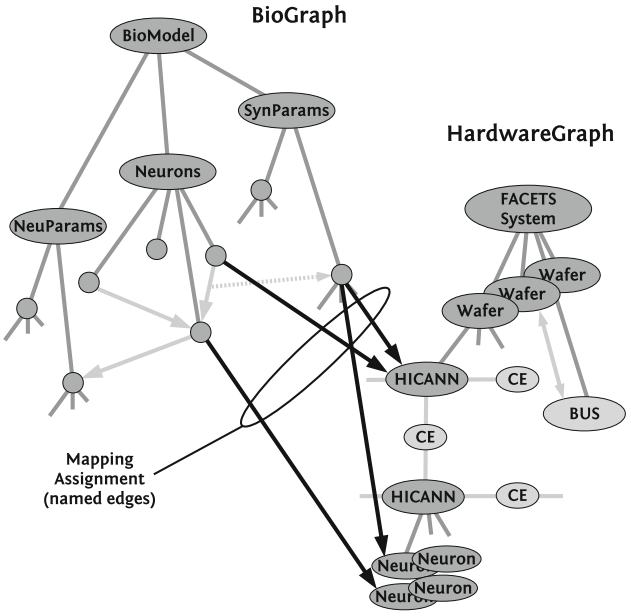

Fig. 10 A simplified example of two graph models, assigning neural elements to hardware components

hierarchical:	models a parent–child relationship, structuring the model.
named:	forms a directed and named relation between any two vertices in the model.
hyper:	assigns a vertex to a named edge, characterizing the edge in more detail.

The major advantage of this graph approach are the implementation convenience and efficiency as well as the flexibility to achieve the complex requirements from both the biological and the hardware model. Due to the structure of the graph model it can be easily (de-)serialized, providing save and restore functionality. Via the path-based query-language *GMPath* (Wendt et al. 2010) information can be dynamically retrieved from and stored to the models. The GraphModel is used to store all information during the configuration process, i.e., the models themselves, the mapping, routing and parameter transformation algorithms data and their results.

Figure 10 shows the graph model representation of a biological network (called *BioGraph*) and its hardware representation (called *HardwareGraph*), connecting elements via named edges after a placement step.

2.2.4.1 The Query Language GMPath To retrieve information from and propagate data to the graph models, the path-based query language *GMPath* was developed, providing a universal interface for placing and routing algorithms as well as for configuration, visualization, and analysis tools

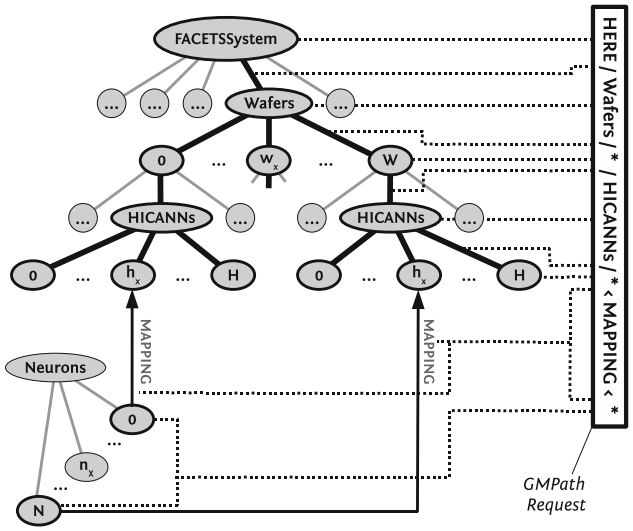

Fig. 11 An example GMPath request to retrieve all assigned neurons of the biological model

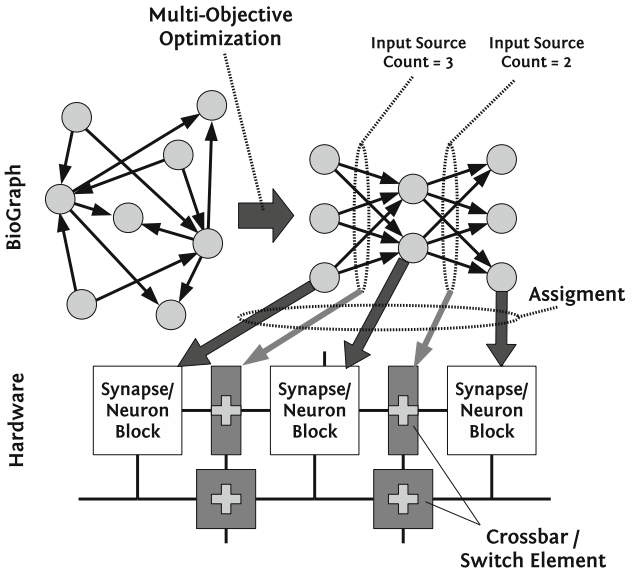

Fig. 12 An example placement, divided into an optimization and assignment step

(Wendt et al. 2010). Based on so-called navigational steps, a path request can enter the model at any point (node or edge) and addresses iteratively the logical environment by

- shifting the focus hierarchically up- or downward.
- shifting the focus back and forth along edges.
- filtering according names.
- concatenating sub-queries.

The results are lists of nodes or edges and serve the requesting software as model information input. Because of its string-based format and the ability to address nodes or edges unambiguously, the queries can be created conveniently and dynamically at runtime and can be used to extend and modify the models.

Figure 11 exemplarily shows subsequent navigational steps of an executed path request, which enters the abstract hardware model at its root, addresses all existing HICANN nodes and finally follows incoming mapping edges to their origins, the neurons of the biological model.

2.2.5 Neuron placement

The process of assigning neural elements like neurons, synapses or their parameters to distinct hardware elements is called *placement*. It can be characterized as a multi-objective optimization problem, the solution of which significantly influences the overall mapping results. Typical algorithmic approaches create clusters of cells with common properties that are mapped to the same HICANN building blocks. Possible optimization objectives are:

- minimize neural input/output variability cluster-wise.
- minimize neural connection count cluster-wise.
- comply with parameter limitations.
- comply with cluster capacities (neural capacity of hardware elements).
- minimize routing distances.

In order to achieve these objectives with user-defined weightings in acceptable computation time, a force-based optimization heuristic was developed. This algorithm balances modeled *forces* (special implementations of the optimization objectives) in an n-dimensional space until an equilibrium is reached and a final separation step assigns data objects to clusters with affine properties. Despite this problem being NP-complete, significantly improved results can be found with this algorithm in an acceptable computation time, as compared to a fast random placement.

Figure 12 illustrates a placement process, divided into an optimization step, which sorts the given biological network for optimal hardware utilization with regard to the input source variability, and an assignment step, defining the physical realization of neural elements on the hardware system.

2.2.6 Connection routing

The routing step allocates and configures the hardware resources for establishing the synaptic connections in the already placed BioGraph. Given the fixed amount of available resources it is not evident a priori whether arbitrary network topologies are always perfectly reproducible.

Synaptic connections can in principle be established via the L1 and L2 infrastructure (see Sect. 2.1.3). In the approach described here, all intra-wafer connectivity is routed exclu-

sively on L1. The L2 network is reserved for inter-wafer connections in a multi-wafer system.

The intra-wafer routing algorithms were developed in close cooperation with the wafer design (Fieres et al. 2008). Some hard-wired features of the L1 infrastructure are thus laid out to optimally suit the routing requirements. The routing itself is performed in two stages. The first stage establishes connections on a HICANN-to-HICANN level via the horizontal and vertical L1 buses, mainly by configuring the bus repeaters and sparse crossbars (see Fig. 5). In the second stage, the signals are routed from the vertical L1 bus lanes into the synapse arrays via the sparse switch matrices, the synapse drivers and the address decoders of the synapses, the latter not being shown in Fig. 5.

The algorithms were proven in various test scenarios: homogeneous randomly connected networks with up to 16,000 neurons, locally connected networks (according to Tao et al. 2004) as well as a model of a cortical column (following Binzegger et al. 2004; Kremkow et al. 2007) with 10,000 neurons. It turns out that in typical cases only a small amount of unrouted connections must be accepted, mainly due to limited resources in the second routing stage. However, it was also shown that by decreasing the density of the neuron placing the routing can be generally facilitated, at the expense of a larger portion of idle hardware synapses.

The routing algorithms proved to be also applicable for the benchmarks described later in this publication, see Sect. 2.4 and Ehrlich et al. (2010).

2.2.7 *Parameter transformation*

The parameter transformation maps parameters of given neuron and synapse models into the hardware parameter space. It is performed HICANN-wise. Biological data is first acquired from the so-called *BioGraph* of the GraphModel (see Sect. 2.2.4) and then transformed into a hardware configuration, which is stored back into the *HardwareGraph*. For an adequate biology-to-hardware translation several constraints have to be considered, such as hardware imperfections and shared or correlated parameters in the microchip.

For the membrane circuits, a two-step procedure was developed to translate the 18 biological parameters from the PyNN description to the 24 electrical parameters of the HICANN building block. The first step is to scale the biological neuron model parameters in terms of time and voltage. At this stage, the desired acceleration factor is chosen and applied to the two time constants of the neuron model. Then, the biological voltage parameters are transformed to match the voltage levels of the HICANN building block. The second step is to translate those intermediate values to appropriate hardware parameters. For this purpose, each part of the membrane circuit was characterized in transistor-level simulations, which were used to establish the translation functions between the scaled AdEx parameters and their hardware counterparts.

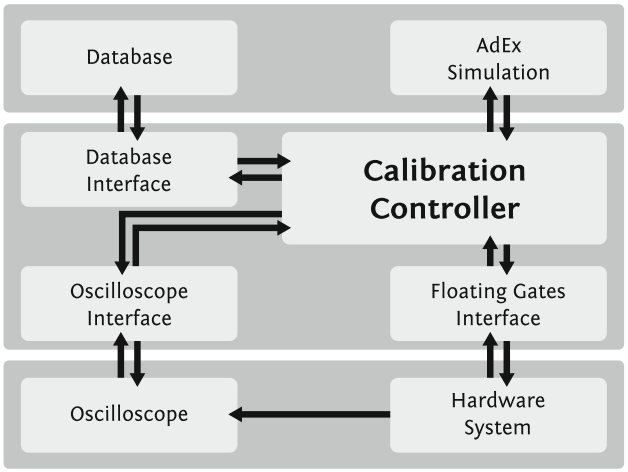

Fig. 13 Architecture of the calibration software. The main component, the calibration controller, executes the calibration algorithms and communicates with the hardware, the oscilloscope and the database via dedicated interfaces. The calibration software also incorporates an AdEx model simulator to compare software and hardware results

However, due to transistor size mismatch in the hardware, these translation functions are expected to differ from neuron to neuron. A calibration software has been developed to automatically parameterize these translation functions for each neuron. For each neuron model parameter, the software will send a spectrum of values to the HICANN building block, and measure the resulting membrane potentials of the current cell. It will then deduce the corresponding AdEx parameters from these measurements, and store the value pairs into a database. After a given number of measurement points, the algorithm will compute the relation between the hardware parameters sent to the floating gates and the AdEx parameters, and store this function into the database. Figure 13 illustrates the calibration software architecture.

Once the calibration step is done, the database can be used to automatically convert a set of biological neuron parameters to their hardware counterparts, allowing on-the-fly conversion of neuron parameters for the wafer-scale hardware system.

Concerning the synapses, there are mainly two restrictions ensuing from the chip design: 256 synapses of the same row share the maximal conductance $g_{\max}$ and the short-term plasticity mechanism, and weights are restricted to a 4-bit resolution. By averaging over all active synapses, the transformation algorithm determines $g_{\max}$ and sets the digital weights accordingly, using *stochastic rounding* to avoid systematic errors.

2.2.8 Application of the mapping flow onto the FACETS chip-based system

In order to further demonstrate the versatility of the Graph-Model-based mapping flow introduced in Sect. 2.2.4, we briefly outline the adoption of this procedure to the operation of the FACETS chip-based systems (see Sect. 2.1.5). This integration avoids code redundancy by unifying the previously independent PyNN back-ends and allows to map neural architectures onto inter-connected chips beyond single-chip boundaries (Jeltsch 2010). Due to the flexible design of the mapping framework, the translation of the PyNN description into the biological graph representation (see Sect. 2.2.2) and the placing of biological neurons onto their hardware counterparts (see Sect. 2.2.5) could be kept completely unchanged. Necessary extensions were limited to the development of a new internal hardware model that captures all features of the chip-based system as well as adapted versions of the routing and the parameter translation (described in Sects. 2.2.6 and 2.2.7, respectively) to match the different network topology and electrical parameters. Together with the low-level event distribution network established by Friedmann (2009), neural network models can now be scaled to multiple chips.

2.2.9 Mapping analysis and visualization

The application *Graph Visualization Tool—GraViTo* as described by Ehrlich et al. (2010) aids in analyzing the mapping results. GraViTo, as shown in Fig. 14, integrates several modules that display graph models in textual and graphical form and gathers statistical data. One can selectively access single nodes inside the data structure and visualize their context, dependencies and relations with other nodes in the system.

The example of the GraViTo views shows a *tree view* on the left which is utilized to browse the hierarchical structure of the graph model and examine contents and connections of individual nodes. The *3-D view* on the right provides a virtual representation of the FACETS wafer-scale hardware system for interactively browsing its architecture and configuration. It also provides a global overview over the single hardware components and the networks they form. Various statistics such as histograms for utilization of the crossbars or the synaptic connection lengths are gathered and can be displayed.

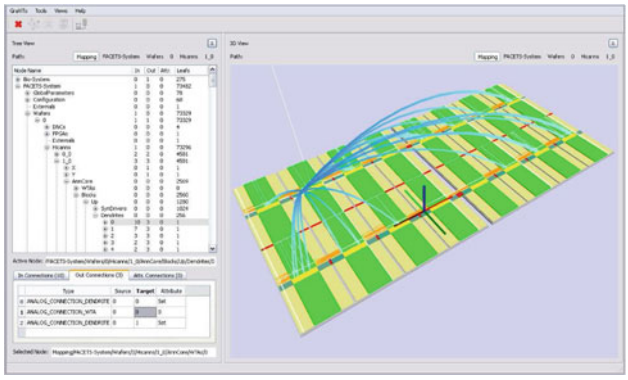

Fig. 14 Screenshot of the GraViTo application

Another option for a systematic mapping analysis arises from the previously mentioned possibility to re-translate the configured HardwareGraph contents via the mapping edges through the BioGraph into a *PostMapping* PyNN script. This script intrinsically contains all model distortions caused by the mapping process, e.g., lost synapses and discretized or clipped parameter values. Exploiting the PyNN concept, it can then be directly evaluated with a software simulator to extract possible functional consequences of the structural distortions, avoiding interferences with other effects like on-wafer communication bandwidth limitations.

2.2.10 Hardware low-level interfacing

A specialized protocol of the class of *selective ARQ*[9] *protocols* is used to provide a fast and reliable communication channel with the neuromorphic hardware device. In the OSI model[10] this corresponds to the transport layer.

Configuration and experimental data is bidirectionally transmitted via two 10 GBit/s Ethernet links per FPGA. In order to handle up to 2 GByte/s of traffic while keeping the load of the host computer as low as possible, several software techniques have been applied to the protocol implementation. Various features of existing transport protocols, notably TCP, have been implemented, including congestion avoidance, RTT[11] estimation and packet routing to keep the connection in a stable and bandwidth maximizing regime.

In matters of performance the framework is divided into three mainly independent processing *threads* (see Fig. 15, receiver thread *RX*, sender thread *TX*, and *resend* thread) to exploit the speed-up in execution of modern multiprocessor systems. Performance critical data, e.g., spike data can be placed in *shared memory* and passed to and from the hardware abstraction layers such as to avoid unnecessary copying. These shared data structures have to be protected against concurrent accesses which imposes an additional overhead in processing time. Thus, to keep the number of system calls and context switches small, access to data located in shared memory is protected by means of custom built user-space fencing and locking methods.

A purely software-stack-based test has been developed that establishes a reliable ARQ connection between two host

[9] Automatic repeat request.

[10] Open systems interconnection model.

[11] Round-trip time.

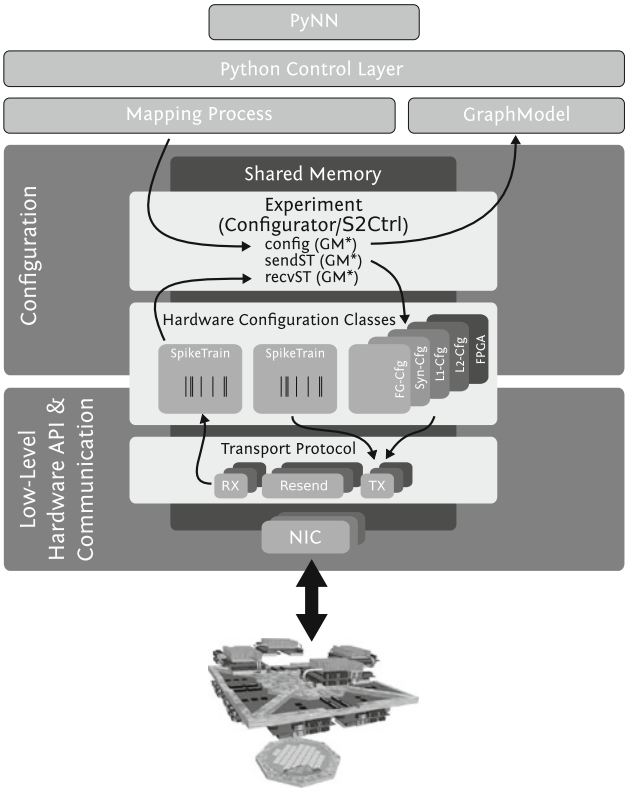

Fig. 15 Configuration and runtime control steps in the hardware abstraction layer: the *Experiment* module acquires the configuration data from the Mapping Process (see Sect. 2.2.2 and Fig. 8), generates a hardware-specific representation of this data and triggers the transfer to the hardware system

computers via 10 Gigabit Ethernet. With a hardware-specific version of this protocol, i.e., with frame sizing and protocol window size, it delivers 10 GBit/s (Schilling 2010).

2.3 Virtual hardware

An executable specification of the FACETS wafer-scale hardware system serves as a versatile tool not only during device design and optimization, but also as a test bench for all involved software layers. It is a functional model that can be used to explore the behavior and characteristics of the real wafer-scale system in its final phase of development.

2.3.1 Implementation

The so-called *virtual hardware* is a detailed simulation of the final hardware platform and has been implemented in C++/SystemC (Vogginger 2010). The virtual hardware replicates its physical counterpart in all aspects regarding functionality and configuration space. Every module of the real hardware has its functional counterpart in the virtual device, where especially the interface and communication structures accurately correspond to the physical system. It implements all analog and mixed-signal modules such as AdEx neurons and dynamic synapses (depressing and facilitating), as well as all units responsible for L1 and L2 routing. Compared to analog and RTL[12] hardware simulations, this model is tuned toward simulation speed using behavioral models of all relevant functional components. However, it is possible to replace individual modules by more sophisticated models, all the way down to simulating single wires on the chip.

The current implementation of the virtual hardware differs from the real hardware system in several aspects, most of them meeting efficiency considerations. The executable system specification is not operated from a host PC but directly from higher software layers, such that the host-to-system communication is not simulated. Furthermore, the configuration of the HICANN building block and its components is not conducted via packets received from L2, as the software implementation of the used protocol is still under development. Instead, every HICANN obtains its configuration via direct access to the GraphModel (see Sect. 2.2.4). Despite these differences the virtual hardware remains a proper replica of the FACETS wafer-scale system providing equal functionality while not suffering from hardware-specific constraints like transistor-level imperfections from the manufacturing process.

2.3.2 Analysis and verification based on virtual hardware

With its functionality and flexibility, the virtual hardware is an essential tool for the development of the software framework operating the FACETS wafer-scale hardware This includes the PyNN interface and the placement, routing and parameter transformation algorithms (see Sects. 2.2.1 and 2.2.2), which can already be tested and verified despite the real hardware not yet being available. The development of a hardware system, which shall be useful in a neural modeling context, can be strongly supported already during its design phase by determining constraints inherent to the system architecture, such as communication bottlenecks or the effect of shared and digitized parameters. Their influence can be evaluated without the interference of hardware imperfections or a missing calibration. Such studies build the basis for improvements in the hardware design or, if possible, the development of software-based corrections. The virtual hardware can be used from PyNN-like any other supported software simulator, thereby also offering an early modeler's perspective onto the capabilities of the future FACETS wafer-scale platform. Any PyNN-model, in particular the benchmark models described in Sect. 2.4, can be applied to this setup. Their output can later be analyzed and compared to reference software simulations, revealing the impact of

[12] Register transfer level.

hardware constraints onto the model behavior, e.g., the loss of certain synaptic connections during the mapping process.

2.4 Benchmark model library

We will now present a set of experiments that serve as benchmarks for the previously described mapping process. The setups are implemented in PyNN and have been contributed by FACETS project partners. They not only cover various computational aspects like memory, pattern recognition, robust information propagation in networks or dynamic switching between different functional modes, but also very different structural characteristics.

2.4.1 Layer 2/3 attractor memory model

The model used here remains faithful to the model of neocortical layers 2/3 in Lundqvist et al. (2006), and in doing so retains the modularity that is the key aspect of this architecture (Lundqvist et al. 2010). It represents a patch of cortex arranged on a hexagonal topology of N_{HC} hypercolumns each separated by 500 μm, in agreement with data from cat cerebral cortex recordings. Each hypercolumn is further subdivided into N_{MC} minicolumns, and various estimates suggest that there are about 100 minicolumns bundled into a hypercolumn (Mountcastle 1997; Buxhoeveden and Casanova 2002). For the default version of the Layer 2/3 Attractor Memory benchmark model, a total number of $N_{HC} = 9$ hypercolumns and a sub-sampling of $N_{MC} = 8$ minicolumns per hypercolumn has been used.

The arrangement of the cells in the local microcircuit together with connection probabilities is shown in Fig. 16a.

In the default variant of the model, each minicolumn consists of 30 pyramidal cells densely connected to other pyramidal cells in the same minicolumn ($P_{PYR\text{-}PYR} = 25\%$) and two *regular spiking non pyramidal* (RSNP) cells that project to $P_{RSNP\text{-}PYR} = 70\%$ of the pyramidal cells. Each hypercolumn has eight basket cells, with each pyramidal cell in a minicolumn targeting $P_{PYR\text{-}BAS} = 70\%$ of neighboring basket cells, and each basket cell targeting $P_{BAS\text{-}PYR} = 70\%$ of neighboring pyramidal cells. The extent of basket cell inhibition is limited to its home-hypercolumn (Douglas and Martin 2004). Apart from these local connections, pyramidal cells located in different hypercolumns are also connected globally ($P_{MC\text{-}MC} = 17\%$). The cartoon in Fig. 16b shows how the minicolumns in different hypercolumns, denoted by dashed lines, are connected. We developed methods to scale this architecture up or down by means of both N_{HC} and N_{MC} without losing its functionality. They are described in Sect. 2.5.1 and experimentally applied in Sect. 3.2.

Thus, a set of mutually exciting minicolumns distributed over different hypercolumns represents a stored pattern or an attractor of the network dynamics. RSNP cells in a minicolumn also receive long-range excitation. They are excited by distant pyramidal cells, given their home minicolumn is not part of the active pattern, thus inhibiting the pyramidal cells in the minicolumn. In this network, we can store as many patterns as the number of minicolumns in a hypercolumn, but by allowing for overlapping memory patterns the number of stored patterns can be increased significantly.

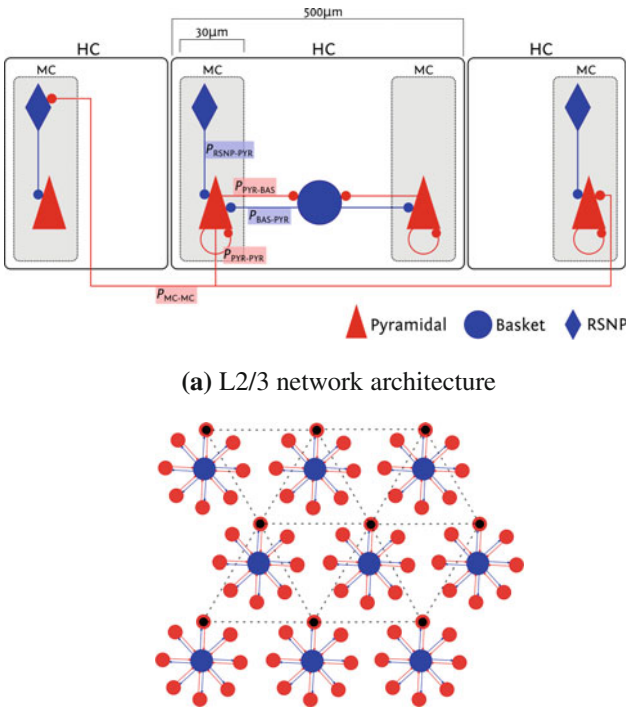

(a) L2/3 network architecture

(b) L2/3 model with 9 HC and 8 MC each

Fig. 16 Schematic detailing the network arrangement and all the excitatory and inhibitory pathways between different cell groups and their connection densities in the L2/3 Attractor Memory network model. **a** Connectivity densities of the sub-sampled network model. See the text for further description. **b** Cartoon of a network with nine hypercolumns (HC). Each hypercolumn has eight circularly arranged minicolumn (MC). The large disc at the center of each hypercolumn represents a population of basket cells. *Dashed lines* show mutually exciting minicolumns that are distributed over different hypercolumns, forming a pattern

Figure 17 shows a raster plot of the activity of the network, when all pyramidal cells are excited by individual Poisson inputs of the same rate.

Whenever an attractor becomes stronger than the others (which happens randomly), it completely suppresses their activity for a short period of time.

Pyramidal cells in an active attractor are in a so-called UP-state, where their average membrane potential is a few mV above its rest value. When plotting the trajectory of the system in potential space, with each axis representing the average membrane potential of all neurons inside an attractor, a projection along the main diagonal (the line which is

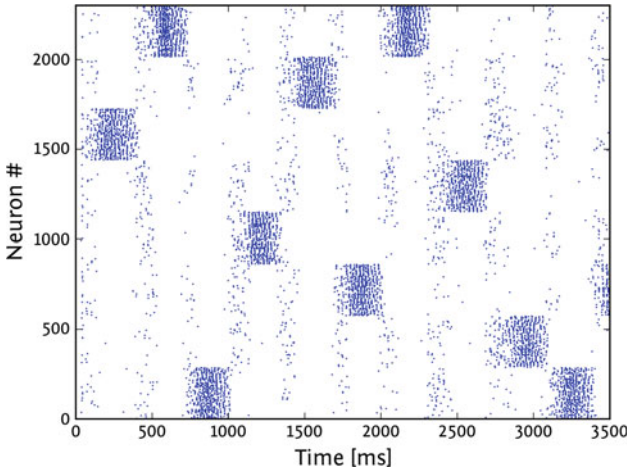

Fig. 17 Raster plot of characteristic activity of an L2/3 Attractor Memory network with nine hypercolumns and eight attractors

equidistant to all axes) will yield a typical star-like pattern (see Fig. 18).

The synaptic plasticity mechanisms are chosen such as to prevent a single attractor from becoming persistently active. Excitatory-to-excitatory synapses are modeled as depressing, which weakens the mutual activation of active pyramidal cells in time. In addition, the neurons feature an adaptation mechanism, which suppresses prolonged firing. Both mechanisms have the effect of weakening attractors over time, such that, in contrast to a classic WTA network, also weaker patterns may become active at times.

2.4.2 Synfire chains

Similar to classical *Synfire Chain models* (Diesmann et al. 1999; Aviel et al. 2003; Kumar et al. 2008, 2010), the version chosen as a mapping benchmark consists of a chain of neuron groups connected in a feedforward fashion, with a certain delay in between. This allows spiking activity to propagate along the chain in a given direction (see Fig. 19). In addition to this, the benchmark Synfire Chain model implements feedforward inhibition by subdividing each group into a *RS*, excitatory (80%) and a *fast spiking* (FS), inhibitory (20%) population (Kremkow et al. 2010a,b). Inhibitory cells are also activated by feedforward projections of excitatory cells from the previous group, but project only locally onto the excitatory population of the same group with a small delay. This allows a fine control over the duration of spiking in a single group and prevents temporal broadening of the signal as it gets passed down along the chain. In the original model of Kremkow et al. (2010b), a Synfire Chain group consists of 100 RS and 25 FS cells. Every cell, RS or FS, receives a total of 60 excitatory inputs from the previous RS population. In addition, every RS cell receives input from all 25 inhibitory neurons of the FS population within its own group.

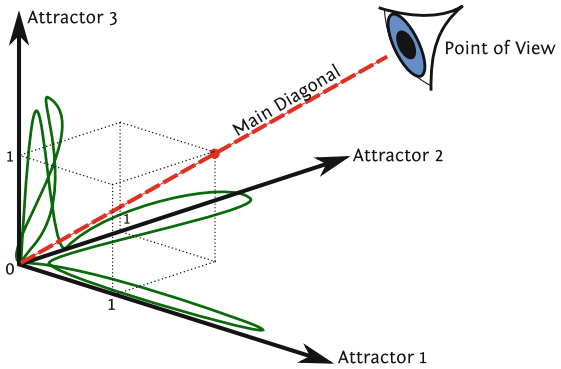

(a) Phase space trajectory projection

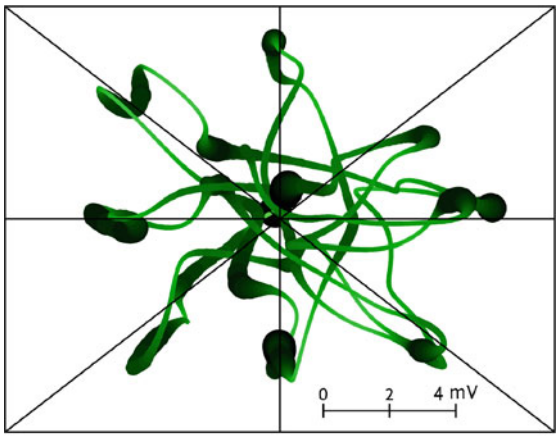

(The projection plane is perpendicular to the main diagonal.)

(b) Mean voltage trajectory

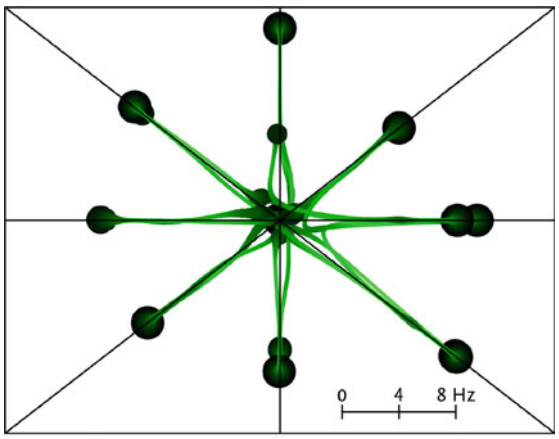

(The projection plane is perpendicular to the main diagonal.)

(c) Mean rate trajectory

Fig. 18 **a** Construction of phase space projection plots as shown e.g., in **b** and **c**: the trajectory in an n-dimensional phase space (here: $n = 3$) is projected to a hyper-plane perpendicular to the main diagonal. **b** Trajectory projection of the attractor network state evolving in eight-dimensional mean voltage and **c** mean rate phase space. Axis values represent the projected offset from a base value, which is the neuron resting potential (in mV) for the voltage traces and 0 Hz for the rate traces. The curve becomes thicker and darker as the phase space velocity decreases

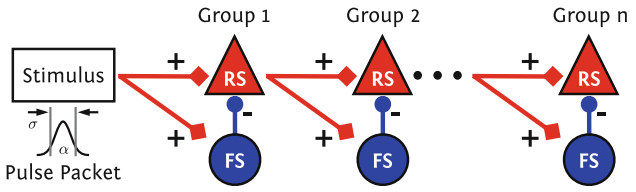

Fig. 19 Schematic of the Synfire Chain benchmark model

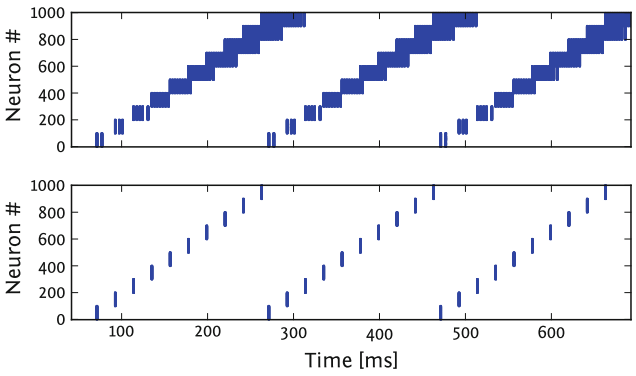

Fig. 20 Raster plot of characteristic RS activity of the Synfire Chain without (*top*) and with (*bottom*) feedforward inhibition. Note the constant spike packet width in case of the active feedforward inhibition mechanism

The inhibition is tuned such that every excitatory neuron gets to spike exactly once upon activation (see Figs. 19, 20).

Methods to scale the size of this model up or down are available and described in Sect. 2.5.1. Different architecture sizes are used to benchmark the quality of the previously described mapping process. See Sect. 3.3 for evaluation data based on scaled benchmark models.

2.4.3 Self-sustained AI states

Randomly connected networks of integrate-and-fire neurons are known to display asynchronous irregular (AI) activity states, where neurons discharge with a high level of irregularity, similar to stochastic processes, and with a low level of synchrony (Brunel 2000). These states were also found in various other network models, including those using conductance-based (Vogels and Abbott 2005) and nonlinear integrate-and-fire neuron models (Destexhe 2009). They were shown to have properties very similar to the discharge patterns observed in awake animals (El Boustani et al. 2007). As cortical neurons are characterized by nonlinear intrinsic properties (Connors and Gutnick 1990), our choice of an AI state benchmark is based on the AdEx neuron model. These nonlinear IF cells are implemented in the FACETS wafer-scale hardware (see Sect. 2.1.2) and reproduce several cell classes observed experimentally in cortex and thalamus (see Destexhe 2009).

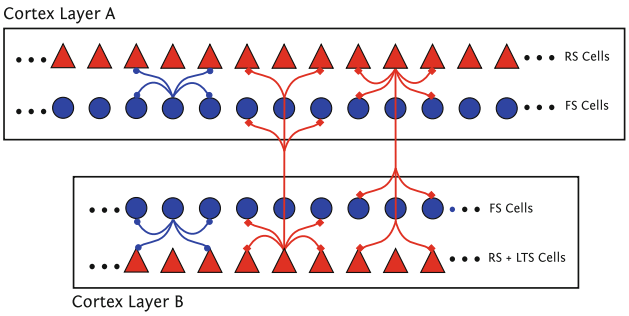

Fig. 21 Schematic of the self-sustained AI states benchmark model. It consists of two cortical layers A and B. Every layer has an excitatory and an inhibitory population, each of which contains certain sets of cell types (RS, FS, and LTS) that determine the network dynamics (see text for details). The excitatory populations project onto every other population, while the inhibitory populations only act within their layer

The particularity of the AI benchmark model is that it allows testing the influence of the various cell classes on the genesis of AI states by varying the different cellular properties. The model considers the most prominent cell classes in cerebral cortex, such as the *RS* cell, the *FS* cell, the *low-threshold spike* (LTS) cell, and the *bursting* cells of the thalamus. It was found that randomly connected networks of RS and FS cells with conductance-based synaptic interactions can sustain AI states, but only if the adaptation currents (typical of RS cells) are not too strong. With strong adaptation, the network cannot sustain AI states.

To the contrary, adding another cell class characterized by rebound responses (the LTS cell) greatly enhanced the robustness of AI states, and networks as small as about 100 neurons can self-sustain AI states with a proportion of 5% of LTS cells. Interestingly, if two of such networks (one with strong adaptation, another one with LTS cells) are reciprocally connected, the resulting two-layer network can generate alternating periods of activity and silences. This pattern is very similar to the Up- and Down-states observed in cortical activity during slow-wave sleep (Steriade 2003). Reducing the adaptation leads to sustained AI states, and is reminiscent of the transition from sleep to wakefulness, a sort of "wakening" of the network. In the context of this article, we use two variants of such networks as benchmarks: First, a network of RS, FS, and LTS cells as a minimal model of AI states. Second, a two-layer cortical network displaying Up- and Down-states. The latter is depicted in Fig. 21.

Also, this model can be scaled up and down in its size in order to benchmark the PyNN-to-hardware mapping process. In its default version, layer A consists of 1,600 excitatory RS and 400 inhibitory FS cells. Layer B contains 400 excitatory neurons, 90% of which are RS and 10% of which are LTS type, as well as 100 inhibitory FS cells. Within a single layer the connection probability is 2% for a network size of 2,000 cells. For smaller networks as for layer B the

connection probability is rescaled inversely to the network size. The inter-layer connectivity is excitatory only and has a connection probability of 1%.

2.5 Analysis based on software simulations

Compared to pure software simulators, dedicated neuromorphic hardware suffers more from limitations and imperfections, which may either directly distort the morphology of the emulated network or influence its dynamics in more subtle ways. On one hand, physical limitations such as size and number of implemented circuits or communication bandwidth impose rather inflexible constraints on parameters such as number of neurons and synapses or the amount of accessible data. On the other hand, as VLSI hardware is, inevitably, subject to manufacturing process variations, individual circuits have varying characteristics, which can only be compensated by calibration to a certain degree. As all these effects influence the dynamics of an emulated network simultaneously, it is usually very difficult to identify the connection between an individual cause and its effect. The most straightforward solution is to artificially impose individual hardware-specific distortions on software simulations, identify their impact on the network's dynamics and find, if possible, suitable compensation mechanisms.

2.5.1 Network scaling

It is very often the case that the robustness of a network scales together with its size, or, in specific cases, with the size or number of individual components. However, before analyzing the effects of distortions, it is indispensable to devise a way of scaling the (undistorted) network without influencing its dynamics. We have developed specific rules for scaling two of our three benchmark models, in order to both explore and learn how to circumvent the limitations of the hardware.

2.5.1.1 Layer 2/3 attractor emory The most obvious and natural scaling of an attractor memory network lies in changing the number of attractors, i.e., in this particular case the number of minicolumns per hypercolumn. Also, the size of the attractors can be evenly scaled by changing the number of units per attractor, i.e., the number of hypercolumns. Finally, the size of the minicolumns itself can be scaled, by varying the number of neurons per column (excitatory and inhibitory populations can be individually tuned by changing the number of pyramidal and RSNP/basket cells, respectively).

These changes would heavily affect the network dynamics, were they not accompanied by corresponding modifications in the network connectivity. The behavior of the network is likely to remain unchanged if the excitation/inhibition patterns of each neuron are kept intact. This is most easily accomplished by keeping both the excitatory and the inhibitory fan-in of each neuron constant without modifying synaptic weights. To this end, simple scaling formulae (linear with respect to size and number of the afferent populations) for the connection probabilities between populations have been derived.

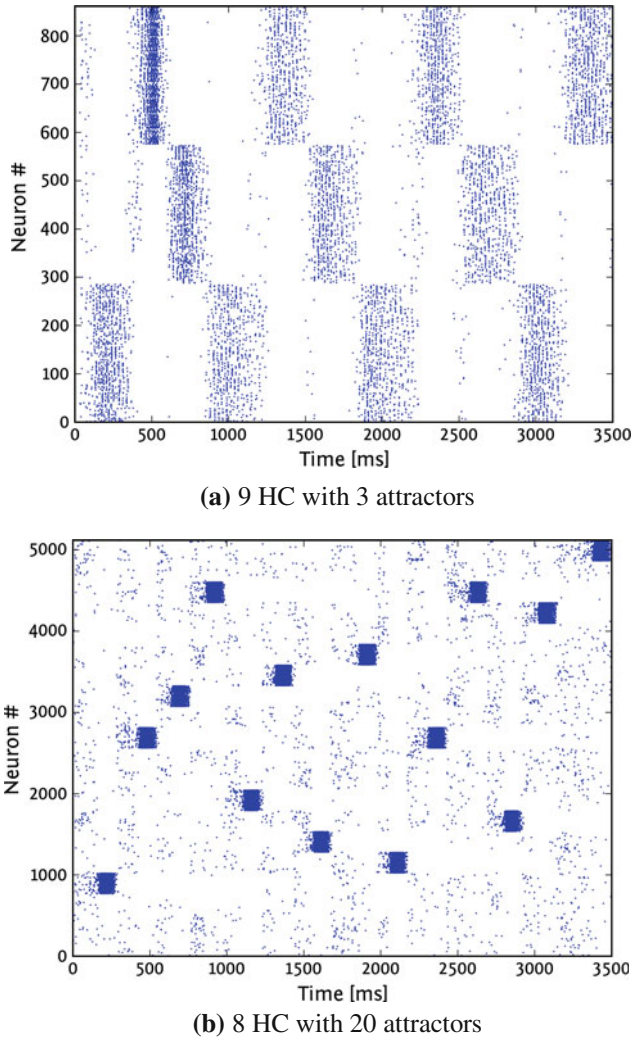

(a) 9 HC with 3 attractors

(b) 8 HC with 20 attractors

Fig. 22 L2/3 network, scaled down to nine hypercolumns (HC) with three attractors and scaled up to eight hypercolumns with 20 attractors. Note the relatively long dwell times in (**a**) compared to short dwell times in (**b**)

Figure 22 shows a scaling example where the number of attractors is varied. At first glance, it may seem that the characteristic attractor dynamics are affected, as the average attractor dwell times decrease from about 300 ms to under 200 ms. However, this is only an apparent effect, as the temporary dominance of individual attractors is a result of local fluctuations in the input. An increasing number of attractors means there is more competition among them, which in mathematical terms translates to shorter, smaller fluctuations in the input rate, therefore leading to decreasing dwell times. When only two attractors are stimulated, the

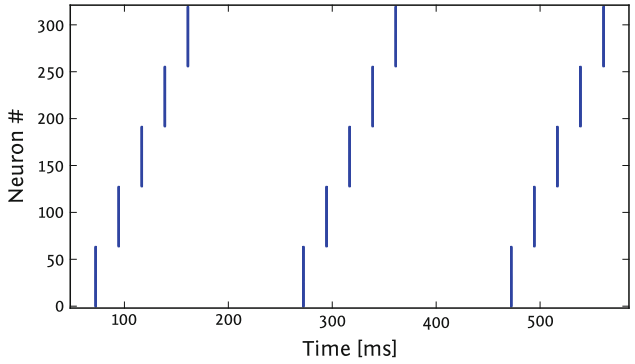

(a) 5 populations with 64 excitatory neurons each

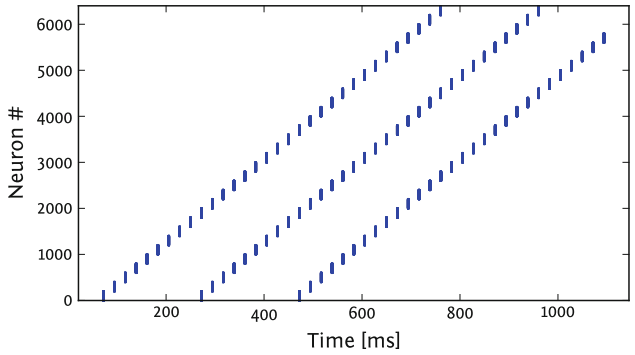

(b) 32 populations with 200 excitatory neurons each

Fig. 23 Synfire Chain scaling examples

dynamics are not influenced by the total number of attractors in the network, which supports our scaling rules.

2.5.1.2 Synfire chain with feedforward inhibition Scaling the Synfire Chain is a comparatively simple task, as there are no feedback or recurrent connections. Scaling the number of units does not require any changes in connectivity. When the number of neurons per unit is changed, the dynamics can be kept unmodified (synchronous firing within a population) if the number of inputs per neuron remains the same. Therefore, modifying a population size by a factor α simply requires that all connection probabilities are modified by a factor $1/\alpha$. Some difficulties may arise when populations become too small, as the binomial connection distribution diverges away from a symmetric Gaussian, favoring a smaller number of afferent connections and leading to activity attenuation and eventually to a break in the pulse transmission (Kumar et al. 2010). The straightforward remedy is offered by the PyNN class `FixedNumberPreConnector` which guarantees a constant but randomly distributed number of inputs. If populations become too small to accommodate the required number of connections, synaptic weights can be increased to compensate for the synaptic loss. The same can be done to cope with synapse loss resulting from the mapping process, as described in Sect. 3.1.1. Figure 23 shows a scaling example where both size and number of populations are varied.

2.5.2 Simulating hardware imperfections

For this study, we have investigated several distortion mechanisms which can be replicated in software simulations and do not require the presence of the actual hardware.

A first limitation of the hardware lies in the fact that axonal and dendritic delays cannot be programmed and the intrinsic delays caused by the hardware communication infrastructure are very small when translated to biological real time. This means that, effectively, the hardware cannot accurately reproduce the dynamics of networks which rely on delayed transmission of action potentials between two communicating neurons.

Two further distortion sources lie within the synaptic circuits of the HICANN building blocks. Since the synaptic weight resolution in a neuromorphic hardware system is limited (see Sect. 2.1.2), large differences between synaptic weights which are mapped to the same synapse driver may cause significant changes to the synapses at the lower end of the weight spectrum. Also, from our experience with the FACETS chip-based systems (Brüderle et al. 2009), we know that variations in the manufacturing process can lead to a spatial synaptic weight jitter of up to 40% of the intended value ($\sigma = 0.4\mu$, assuming a Gaussian distribution), even after calibration. This might be fatal for networks which rely on precise tuning of synaptic weights.

Because of the limited bandwidth of the communication layers (see Sect. 2.1.3), synapses may be lost during the biology-to-hardware mapping process. Ensuing distortions depend strongly on the network topology and can become arbitrarily high for sufficiently large and densely connected networks.

In addition, neuron loss can also become a key factor, not necessarily due to hardware limitations (usually, synaptic loss becomes significant long before the number of network neurons exceeds the number of neurons available on the hardware), but as an option to counteract synaptic loss by controlled reduction in the number of neurons.

Although it does not apply to the three benchmark networks we have studied, the hardware neuron model itself may become a limiting factor, when trying to map models which rely on more complex neurons. However, we consider this to be an unlikely case, as the AdEx model has been shown to successfully reproduce a wide variety of spike patterns (Brette and Gerstner 2005; Destexhe 2009; Millner et al. 2010) and has also proven very successful in emulating biological spiking activity (Naud et al. 2008). This is not the case for the FACETS chip-based neuromorphic system (see Sect. 2.1.5), which only offers simple leaky integrate-and-fire neurons. Section 3.2 describes a setup where the missing adaptation mechanism was compensated by tuning several other parameters.

2.5.3 Analysis and development of STDP in hardware

Synaptic plasticity on the highly accelerated FACETS wafer-scale hardware (for a detailed description see Sect. 2.1.2) provides a promising architecture for doing research on learning. However, so far there are no studies about the effect of low weight resolutions and limited update frequencies on the functionality of single synapses and consequently neuronal networks. In the following, two directions of study will be outlined and one detailed example will be given.

First, the question of a required minimal resolution for synaptic weights and their optimal update frequency is investigated. However, those two restrictions may be dominated by production process variations that set the upper limit for the functionality of the synapses. Production process variations cause the supposedly identical circuits for causal and acausal correlation measurement to differ due to variations in their transistors. This asymmetry limits the accuracy of detecting correlations or in other words causes a correlation offset. With respect to learning neuronal networks (e.g., Davison and Frégnac 2006), we are especially interested in the effects of hardware synapses on their ability to detect synchronous input when embedded into an appropriate architecture.

Second, the dynamics of discretized STDP are analyzed based on the assumption that the weight discretization is the most crucial restricting component influencing the dynamics of single synapses and whole networks. This analysis is carried out with respect to the equilibrium weight distribution that is obtained by evaluating an initial synaptic weight value in sequence. Within this sequence of weight evolution, the probability for causal evaluation is equal to the one for acausal evaluation. Analytical equilibrium distributions (van Rossum et al. 2000) as well as numerical equilibrium distributions of continuous weights are used as a reference.

Here, we shall discuss in detail one analysis on the effect of low resolution weights within a neuronal network. In order to isolate the functionality of a single synapse from network effects a simple network is defined (Fig. 40a). A population of pre-synaptic neurons is connected to a single post-synaptic neuron. The *Intermediate Gütig STDP model* (Gütig et al. 2003; Morrison et al. 2008) is used for the construction of the look-up table (see Sect. 2.1.2). Developing synaptic weights are compared for either correlated or uncorrelated pre-synaptic input. Correlation within the pre-synaptic population is generated by a multiple interaction process (Kuhn et al. 2003), whereas in the uncorrelated case the firing pattern of the pre-synaptic neurons are those of Poisson processes. Results for the effect of discrete weights on this network are presented in Sect. 3.2.4.

To avoid expensive changes of the chip layouts, the hardware restrictions are analyzed with preparative software simulations. Therefore, the standard STDP synapse model of the software simulator NEST (Gewaltig and Diesmann 2007) was modified by introducing a causal and acausal accumulation circuit, a digital weight value and global weight update controllers. In the following, we will call this the *hardware inspired model*. As a *reference model* another software synapse model with continuous weight values and continuous weight updates, but with a symmetric nearest-neighbor spike pairing scheme was implemented.

Further analysis with focus on the weight update frequency is in progress. In the current prototype of the HICANN building block the causal and acausal accumulation circuits will be reset commonly, if a weight update is performed. Such a common reset distorts the counterbalancing effect of the accumulation circuit receiving less correlations, because the common reset suppresses the circuit to ever elicit a weight update. Consequently the dominating accumulation circuit, in terms of eliciting weight updates, drives all synaptic weights to its corresponding boundary value. For future improvements, the performance gained by adding a second reset line to reset both accumulation circuits independently will be compared to the performance gain due to a more detailed readout of these circuits. Details about these additional studies will be presented in a publication that is in preparation.

3 Results

In the following, a summary of results is provided, all of which have been acquired by means of the workflow described in Sect. 2. The presented data demonstrate the functionality of both the PyNN-to-hardware mapping framework and the virtual wafer-scale hardware system, as the applied benchmark models are shown to exhibit the expected functionality when processed by this environment. Examples of mapping-specific distortion analyses based on reference software simulations are provided and discussed. The effect of discretized synaptic weights, as implemented in the FACETS wafer-scale hardware, is analyzed in the context of an STDP benchmark. Scalability questions regarding the graph-based data structure for the mapping process are considered on the basis of experimental data. Furthermore, we present first results of a successful application of the presented AdEx neuron circuit calibration scheme acquired with a HICANN prototype.

3.1 Benchmark results

The benchmark models and their target output descriptions described in Sect. 2.4 represent an essential tool to test and verify the workflow presented in this article on a high, functional level. This is important especially in the context of studies on neural network dynamics, where the identification of erroneous components from the analysis of spatiotemporal spike patterns can be very difficult due to a lack of insight

and intuition in the field of neural information processing principles.

3.1.1 Distortion mechanisms and compensation strategies based on software simulations

Even with a virtual version of the hardware system being available, software simulations remain a powerful analysis tool, especially since they offer access to the full range of dynamic variables of the modeled neural network, some of which may be inaccessible on the virtual hardware. In the following, we will demonstrate the effects of different distortion mechanisms via software simulations and propose several methods to either circumvent or counteract them. These methods are chosen such as to lie within the possibilities of the hardware system.

3.1.1.1 Layer 2/3 attractor memory The functionality of the L2/3 network is best determined by a combined analysis of spike and voltage data. While a visual inspection of a raster plot of all neurons usually provides a good basis for evaluation, a more thorough investigation of UP/DOWN-state statistics requires the analysis of voltages from a relatively large number of individual cells. Both the extraction of the full spike data and of multiple voltage traces are not possible on the hardware, making the use of software simulations indispensable for a proper evaluation of the effects of mapping distortions.

In order to replicate a biologically plausible pattern of axonal and dendritic delays, we have implemented a network geometry as exemplified in Fig. 24, with distance-dependent delays. When setting all delays to zero, we have observed no significant changes in the network dynamics. This is not unexpected, as this model relies more on firing rates rather than precise spike timing in order for particular attractors to become activated.

Depending on the amount of spatial synaptic weight jitter, the network shows varying levels of tolerance. For values up to 25%, the dynamics only suffer minor changes. At 50% jitter, spontaneous activation is completely suppressed, but activation via input from L4 remains functional, exhibiting the same phenomena of pattern completion and rivalry as seen in the undistorted case (see Fig. 25).

Because of its intricate connectivity, which spans both local and global scales, the layer 2/3 attractor memory network was expected to be quite sensitive to synaptic loss. Indeed, if the synapse loss is localized to certain attractors, they become either inactive (for excitatory synapse loss) or dominant (for inhibitory synapse loss). However, if synapse loss is spread homogeneously over all populations, the network becomes remarkably resilient, tolerating values as high as 40% (see Fig. 26).

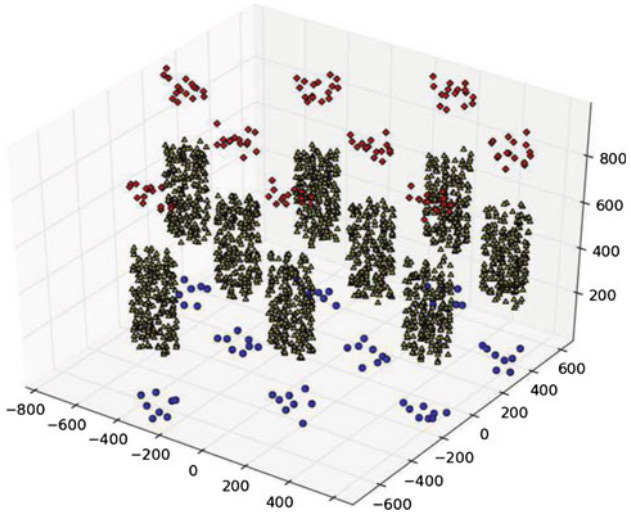

Fig. 24 Geometry of the L2/3 attractor memory model, the unit on the axes is μm

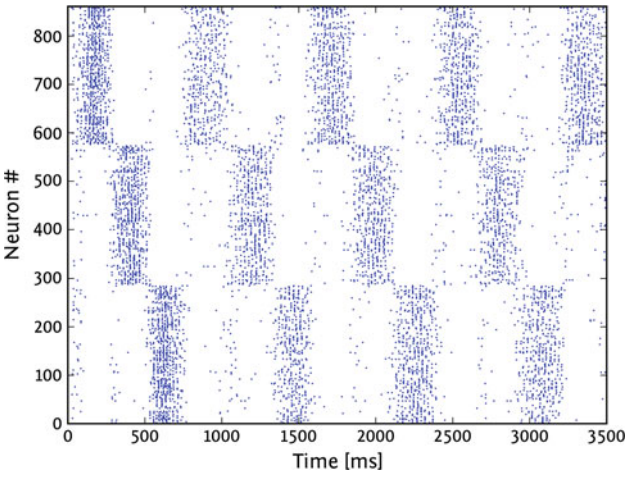

(a) Spontaneous activity at 25% spatial jitter

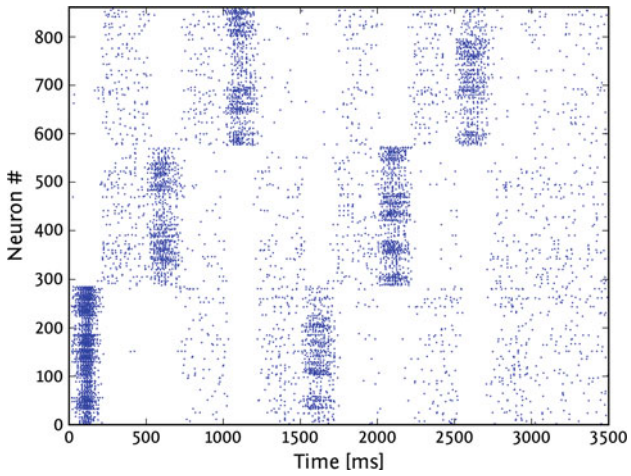

(b) L4 activation at 50% spatial jitter

Fig. 25 Effects of spatial weight jitter on a L2/3 attractor memory network with nine hypercolumns and three attractors

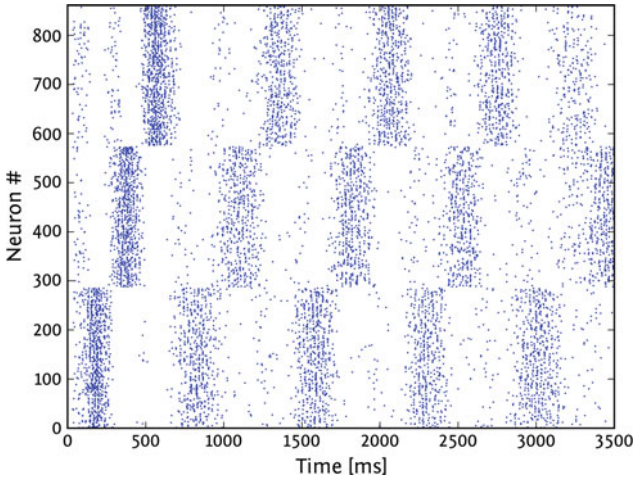

(a) Raster plot, 40% synaptic loss

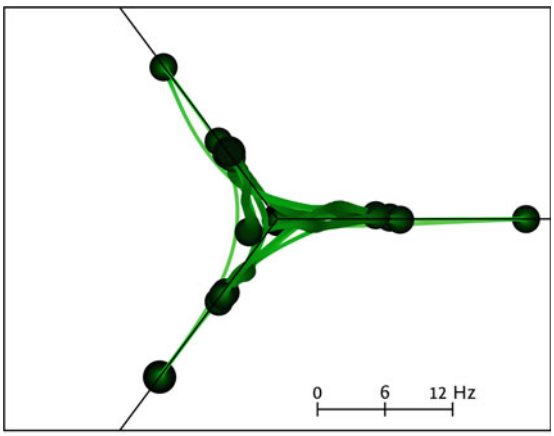

(b) Voltage star plot, 40% synaptic loss

Fig. 26 Synaptic loss tolerance of an L2/3 attractor memory network. Synaptic loss was assumed homogeneous over all populations

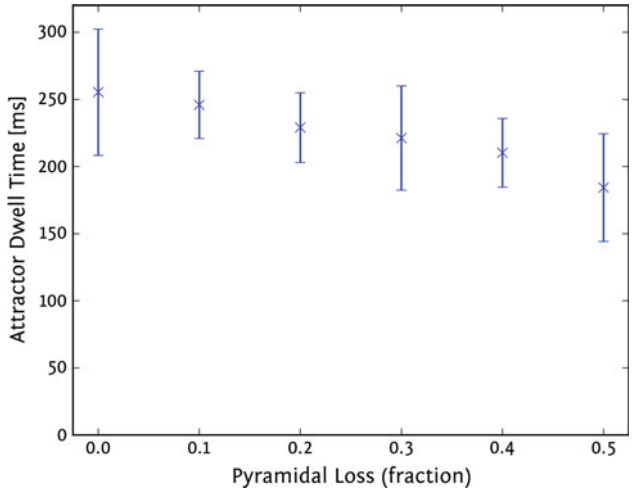

Fig. 27 Attractor dwell times versus neuron loss in the L2/3 attractor memory network

In contrast to synaptic loss, the loss of pyramidal neurons (which make up about 87% of the network) has only little effect on the network dynamics, even up to values as high as 50%, regardless of the number of minicolumns or hypercolumns present (see Fig. 27). It is, for example, possible to have a functioning network with as little as 12 pyramidal cells per minicolumn. This circumstance has major consequences for synapse loss compensation.

When synapse loss increases beyond a certain limit, intra-attractor excitation and inter-attractor inhibition become too weak for attractor dynamics to emerge. The total number of synapses scales linearly with the total number of neurons (when network scaling conserves the afferent fan-in, as described in Sect. 2.5.1), so reducing the neuron count represents a straightforward way of circumventing synapse loss. This can be achieved by reducing the number of attractors (which may, however, not always be desirable) or by reducing the number of neurons per attractor by decreasing either the total number of hypercolumns or the neuron count per minicolumn.

Elimination of pyramidal neurons (without re-scaling the fan-in) is a much more efficient method in terms of synapse number reduction, as the total synapse count has an approximately quadratic dependence on the number of neurons per minicolumn. Since, for this particular network model, attractor dynamics are largely insensitive to pyramidal cell elimination, as described above, this becomes a method of choice when dealing with harsh bandwidth limitations.

Especially in cases where synapse loss is relatively small and inhomogeneous, afferent synaptic input can be restored by increasing the corresponding synaptic weights (see Fig. 28). While it is always possible to hereby establish the required average firing rates of individual populations, this compensation mechanism needs to be used cautiously, as it can influence spike train statistics and neuron input–output curves.

3.1.1.2 Synfire Chain with feedforward inhibition The Synfire Chain model presented in Sect. 2.4.2 relies heavily on delayed transmission of action potentials between inhibitory and excitatory populations. Eliminating these delays causes afferent EPSPs and IPSPs to overlap, possibly leading to suppression of efferent spikes (Kremkow et al. 2010a). This makes a direct mapping of the model to the FACETS wafer-scale hardware impossible. However, as the hardware offers the possibility of tuning synaptic time constants of individual neurons, it is possible to compensate for missing delays by adjusting the rising flank of EPSPs for the inhibitory neurons. This can be achieved by increasing the corresponding synaptic time constants and decreasing the corresponding synaptic weights simultaneously (see Fig. 29).

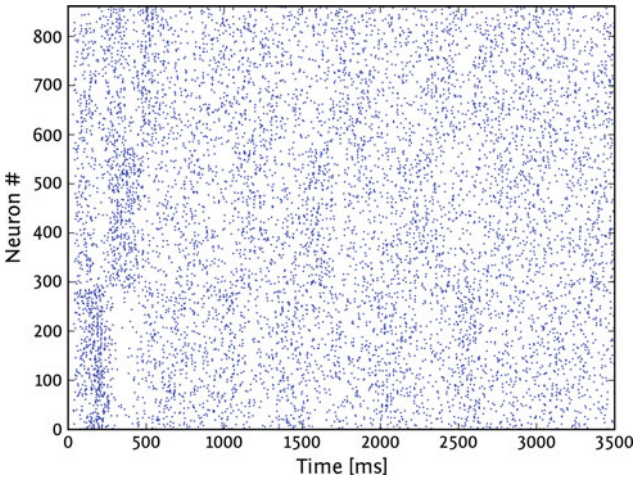

(a) 60% synapse loss, uncompensated

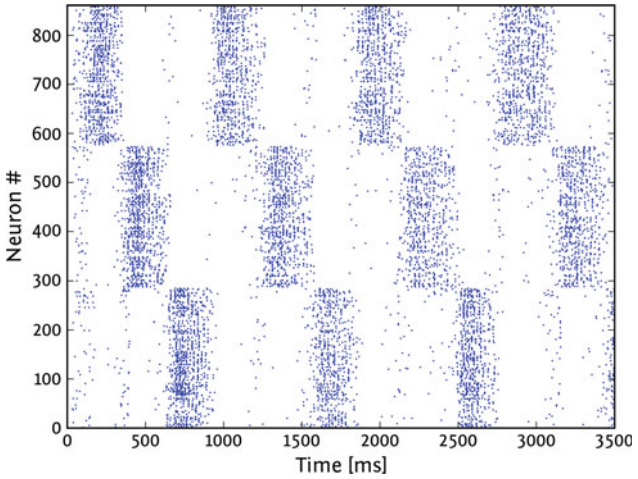

(b) 60% synapse loss, compensation by modified weights

Fig. 28 **a** High synaptic loss destroys attractor dynamics. Several methods for compensating or counteracting this effect are presented in the text above. **b** Shows the results of compensation by modified synaptic weights

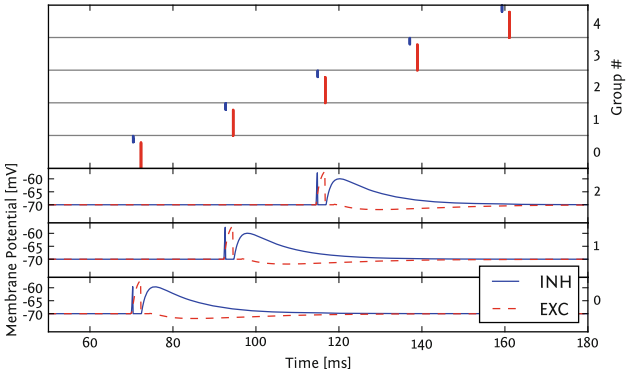

(a) Synfire Chain with delayed spike propagation

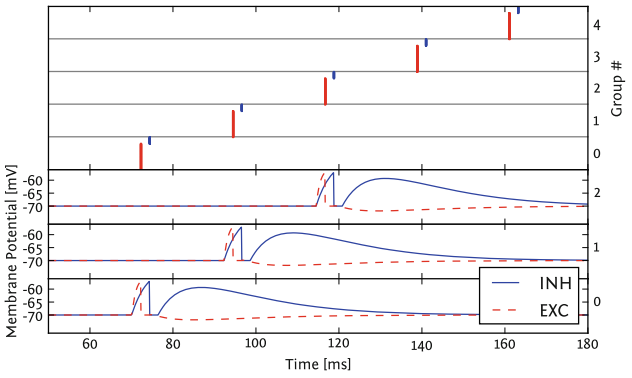

(b) Synfire Chain with modified synapses, no delays

Fig. 29 **a** Delayed spike propagation is essential in the original Synfire Chain model. **b** By modifying synaptic parameters (see text for details), effective afferent spike times can be reproduced without propagation delays

Spatial synaptic weight jitter may effectively cancel itself out for large numbers of synapses, but can lead to skewed afferent input, especially in smaller networks. Depending on the amount of spatial jitter (variance of the underlying Gaussian, see Sect. 2.5.2), this might lead to individual excitatory neurons not firing, as a consequence of a too low average afferent weight. This causes a chain reaction which leads to an increasing number of silent neurons for every subsequent population in the Synfire Chain, ultimately causing the activity to die out (see Fig. 30a).

Synapse loss has qualitatively the same effect, only manifests itself much stronger, as it is not symmetrically distributed around zero. Even relatively low values of around 2% completely suppress the propagation of the signal after only few iterations (see Fig. 30b). Both distortion mechanisms can be effectively compensated by increasing excitatory synaptic weights (see Fig. 31). Since all excitatory neurons within a population spike only once, simultaneously, modification of synaptic weights does not affect spike train statistics.

The other obvious way of compensating synapse loss is by decreasing the overall network size, as described for the L2/3 attractor memory network. This can be achieved by decreasing either the number of populations or their size. Since the mapping algorithm tends to cluster neurons from the same population together on neighboring HICANNs, reducing population sizes is more efficient for reducing the required communication bandwidth.

3.1.2 Analysis based on virtual hardware

The benchmark experiments were simulated with the virtual hardware, thereby verifying the functionality of the whole software workflow and the general capability of the system to serve as a neural modeling platform.

Before applying the benchmark experiments (see Sect. 2.4), we determined the maximum reachable bandwidth of the L2 links (FPGA-to-DNC and DNC-to-FPGA) with the aid of the virtual hardware. We have found that—despite the Poisson distribution of spiking activity—the achieved

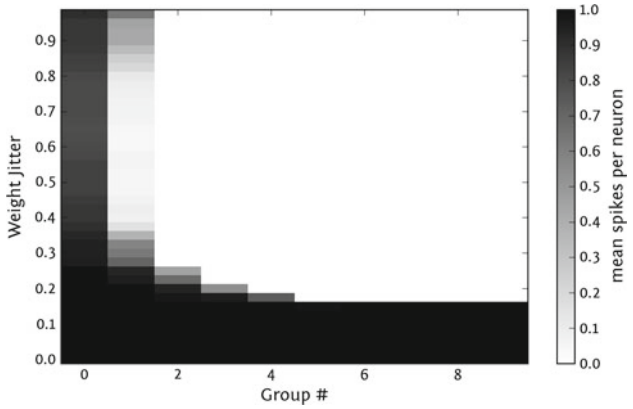

(a) Impact of spatial weight jitter on signal propagation

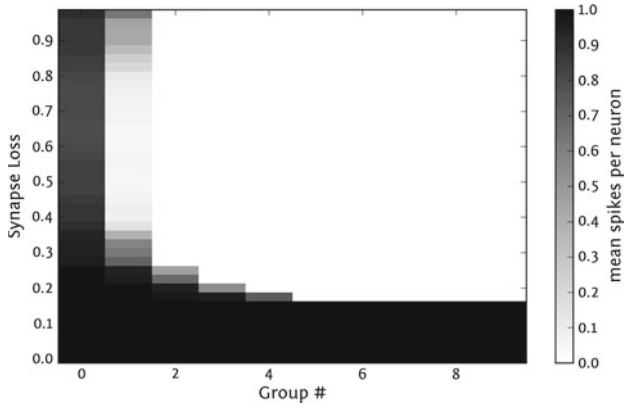

(b) Impact of synapse loss on signal propagation

Fig. 30 **a** Sufficiently high spatial weight jitter causes a breakdown of signal propagation along the Synfire Chain. **b** Synapse loss is even more critical, completely attenuating the signal after only few iterations

bandwidth corresponds to the one theoretically expected from the data link speed and pulse packet sizes.

The gathered results were used to enhance the routing of L2 pulse events to the wafer (see Sect. 2.2.6), which distributes external spike sources over all available L2 links, such that the bandwidth provided by a given hardware setup is fully exploited. The application of this is crucial, when it comes to the realization of network models with either high input rates as the Layer 2/3 attractor memory model, or highly correlated input as for the Synfire Chain, where hundreds of spikes need to be delivered within a very small time window. Having these limitations in mind, one can choose the size of the hardware setup properly before actually mapping neural experiments onto the FACETS wafer-scale hardware, in a way that all requirements are considered in terms of spatial and temporal issues (i.e., neuron/synapse and bandwidth resources).

3.1.2.1 Synfire Chain The Synfire Chain model with feed-forward inhibition was successfully run on the virtual FACETS wafer-scale hardware. The stable propagation of pulse volleys from one group to the next is plotted in Fig. 32. In

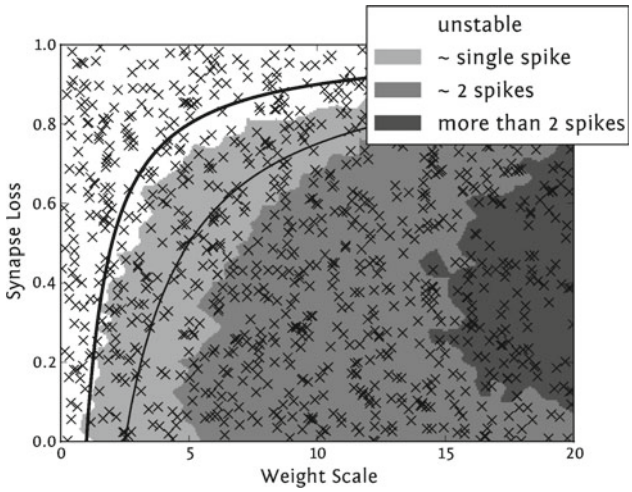

Fig. 31 Synapse loss blocks the propagation of the signal along the Synfire chain (*white zone*). A synapse loss probability of p can be compensated very efficiently by scaling the weights by a factor of $\alpha/(1-p)$ (*thick black line*: $\alpha = 1$, *thick black line*: $\alpha = 2.5$). Note that the *thin black line* stays within the *light gray* area that denotes a stable propagation with one spike per neuron. Increasing the synaptic weights even more, effectively overcompensating the synapse loss, results in multiple spikes being fired by the excitatory neurons (*dark gray* and *black zone*). The total number of spikes per burst is limited by the refractory period and the time until the first inhibitory spike, meaning that the signal does not broaden indefinitely

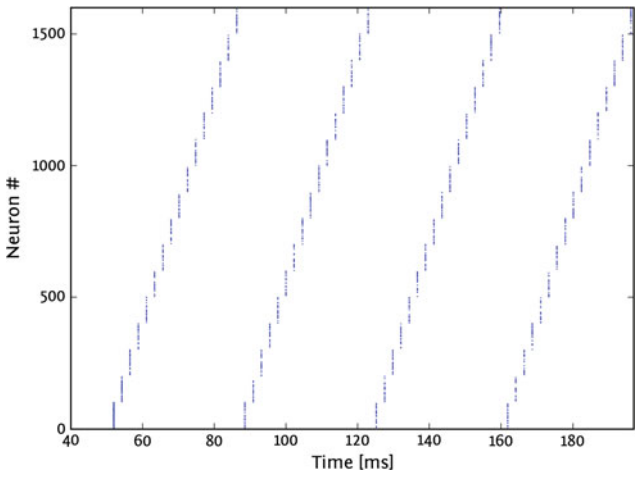

Fig. 32 Synfire Chain with 16 groups connected to a loop, simulated on the virtual hardware

this case, the network consisted of 16 groups with 100 excitatory and 25 inhibitory each, the groups were connected to form a loop such that the activity would be sustained indefinitely. However, this model proved to be very sensitive to distortions: If more than two neurons of a group do not fire, the Synfire Chain will die out immediately, because the local inhibition comes up too early and prevents the excitatory cells from spiking. This happens also due to the lack of synaptic delays in the current implementation of the hardware, as only L1 connections are used for the routing of neural events

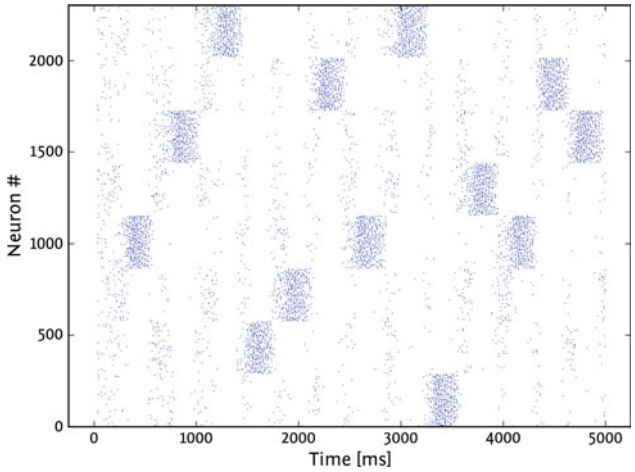

Fig. 33 Raster plot of the layer 2/3 attractor memory model simulated on the virtual hardware: firing activity is shown only for pyramidal cells, default size with nine hypercolumns and eight attractors

within the network. For upscaled versions of this model and a restricted hardware size, where the mapping process yields a synapse loss larger than 5%, the functionality could not be sustained according to the high sensitivity of the used parameters, i.e., the pulse volley only reached the second group. Nevertheless, we were able to regain the benchmark functionality by compensating the synapse loss through either strengthened weights or downscaled neuron populations (see Sect. 3.1.1).

3.1.2.2 Layer 2/3 attractor memory model The presented software framework also performed very well when mapping such an elaborated neural architecture like the layer 2/3 attractor memory model onto the (virtual) FACETS wafer-scale hardware: Fig. 33 shows the spike output of the default model with 2,376 neurons simulated on a virtual wafer snippet containing 8×2 reticles. This successful replication of the benchmark's dynamics not only underscores the correct operation of the placing and routing algorithms (see Sects. 2.2.5 and 2.2.6), but also indicates that the transformation from biological to hardware models (see Sect. 2.2.7) works properly and does not distort the model's behavior, concretely in this example a variety of different short-term plasticity settings could be transferred to shared hardware configurations.

3.1.2.3 Self-sustained AI states The cortical network with self-sustaining AI states was also successfully realized on the virtual hardware. The single-layer cortical model was implemented for different sizes and parameter sets, where the model functionality was preserved without distortions. The two-layer cortical network exhibiting Up- and Down-states was also realized at the default size with 2,500 cells and varying adaptation parameters, see Fig. 34 for an exemplary raster plot together with a reference software simulation with NEURON.

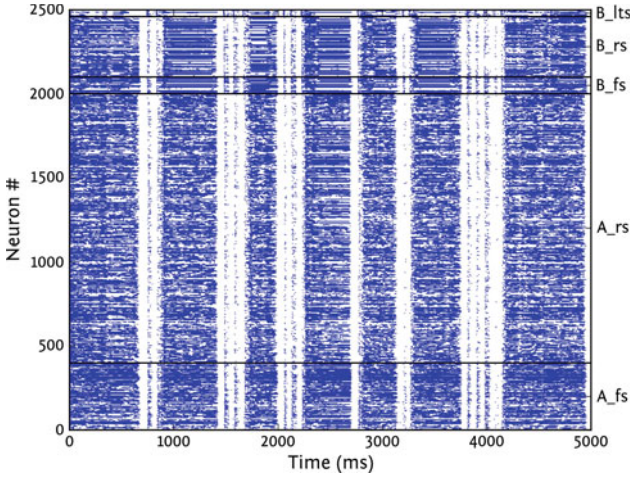

(a) NEURON

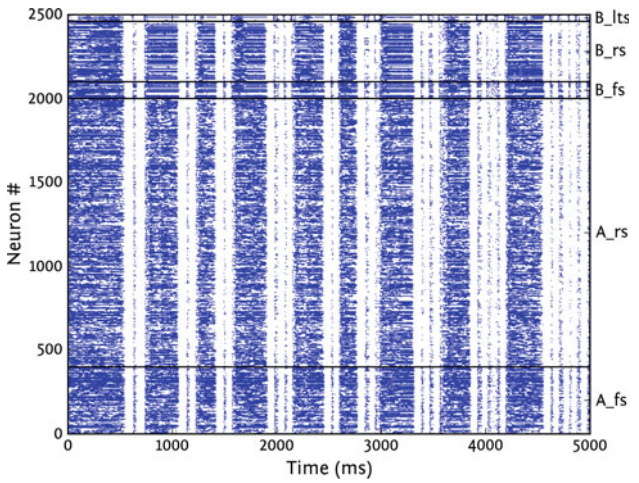

(b) Virtual Hardware

Fig. 34 Raster plot of the two-layer cortical network exhibiting Up- and Down- states simulated with NEURON (**a**) and on the virtual hardware (**b**). *Horizontal lines* depict the limits between RS, FS, and LTS neurons of layers A and B. The first cortical layer consists of 2,000 cells, the second of 500 cells. 10% of all cells are initially stimulated to induce asynchronous irregular firing in the whole network. The first layer is per se not self-sustaining, i.e., the activity dies out after a while, the second smaller layer is able to sustain its activity due to a large number of LTS cells. The sparse connectivity between the two layers assures that the activity in the first is reactivated by excitatory input from the second layer

3.2 Cross-platform implementation of a benchmark model

As a demonstration of the versatility of the methodological framework discussed in the previous sections, this section will present the implementation of one of our benchmark models on three different back-ends: the software simulator NEST, the FACETS chip-based system and the virtual wafer-scale hardware. For this purpose, we have chosen the L2/3 attractor memory model, due to its challenging connectivity patterns and the interesting high-level functionality it

displays. Because of the limited size of the chip back-end, the original model needed to undergo some profound changes, which will be detailed in the following sections.

3.2.1 FACETS chip-based neuromorphic system

One ASIC in the current version of the FACETS chip-based system as described in Sect. 2.1.5 offers 192 fully interconnectable leaky integrate-and-fire neurons with static synapses. Since the original model requires 2,376 adapting neurons interconnected through plastic synapses, we had to heavily modify the network configuration in order to keep its functionality. Reducing the total number of neurons from 2,376 to 192 was done following the scaling rules described in Sect. 2.5.1. In this context, the observation that pyramidal cells can be lost without significantly affecting the dynamics of the network became extremely useful. In order to provide relatively long dwell times, we have chosen a setup with only three attractors and four hypercolumns (i.e., four minicolumns per attractor). The number of basket cells per hypercolumn was reduced from the original 8 to 6, while the number of pyramidal cells per hypercolumn was reduced from the original 30 to 12. The number of RSNP cells per minicolumn remained constant at 2. Thus, this setup implements the original model architecture with exactly 192 neurons. See Fig. 35 for a schematic of the resulting architecture.

Due to a lack of neural adaptation and synaptic plasticity (which are both crucial in the original model, as they limit the pyramidal cell UP-state duration), we needed to adapt the neuron parameters (leak conductance, reset, rest, threshold, and reversal potentials) and synapse characteristics (weight, decay time constant) in such a way as to retain as much as possible of the original dynamics, on average. One additional constraint which needed to be taken into account was the limited range of synaptic weights available on the neuromorphic chip. We were able to compensate this, to some extent, by modifying the connection densities among the neuron populations.

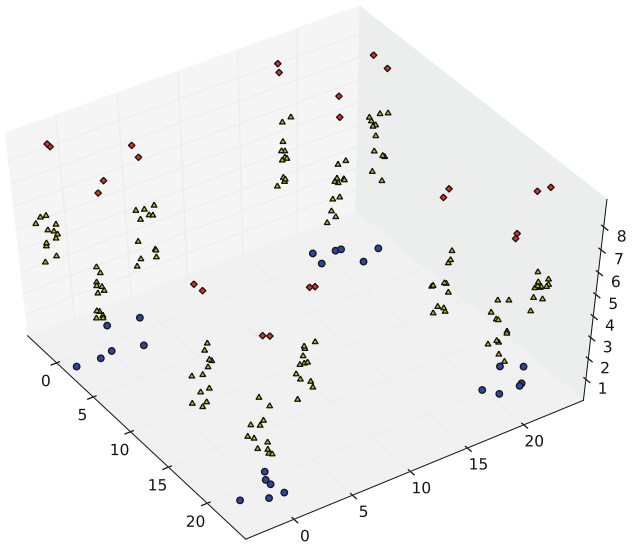

Fig. 35 Geometry of the scaled-down L2/3 attractor memory network model. Note the greatly reduced number of hypercolumns (HC) and of pyramidal cells per minicolumn (MC) as compared to Fig. 24

One important consequence is that because the network is unable to adapt, its dynamics change significantly. If one would only remove adaptation and plasticity, without changing other parameters, the first attractor which becomes activated would last indefinitely. Therefore, the removal of these two mechanisms needs to be accompanied by a reduction of intra- and inter-columnar excitation. This in turn causes the network to become much more input driven, which manifests itself in an extreme sensitivity of attractor dwell times toward the momentary input activity. Dwell times become more erratic and even small changes in the average input rate cause attractors to become either dominant or virtually inactive.

Also, due to the limited input bandwidth of the ASIC (for the chosen architecture: 64 channels at about 80 Hz), some degree of input correlation was inevitable, as each of the 144 pyramidal cells requires a Poisson stimulation of 300 Hz. In order to maintain attractor stability, we have chosen to have no overlapping inputs for different attractors (and thus zero correlation, for the Poisson input we have used). This, on the other hand, leads to an increased input correlation among pyramidal cells belonging to the same attractor, which, in absence of adaptation, tends to prolong attractor dwell times.

Figure 36 shows L2/3 attractor memory benchmark results acquired with the FACETS chip-based system: (a) Attractors have been excited by Poisson input with an overall rate of 750, 1, and 550 Hz, respectively. Note the relatively long dwell times, which are mostly due to high correlations among pyramidal inputs within an attractor. The discrepancy in the input activity needed to ensure a balanced activation of attractors 1 and 3 is due to hardware manufacturing fluctuations, which appear to be very complex in nature, often interacting with each other and being highly dependent on the ongoing activity on the chip. (b) Attractors have been excited by Poisson input with an overall rate of 750, 700, and 650 Hz, respectively. Note the large fluctuations in attractor dwell time due to lack of adaptation which leads to strongly input-driven dynamics. Also note that, in contrast to Fig. 36, when attractor 2 becomes active, the input activity of attractor 3 required an increase of 100 Hz in order to achieve balanced activation. The most likely cause is capacitive cross-talk between the analog circuits, which varies depending on the throughput rate. (c) Attractors have been excited by Poisson input with an overall rate of 700, 800, and 600 Hz, respectively. Note that only a slight increase of the attractor 2 input rate,

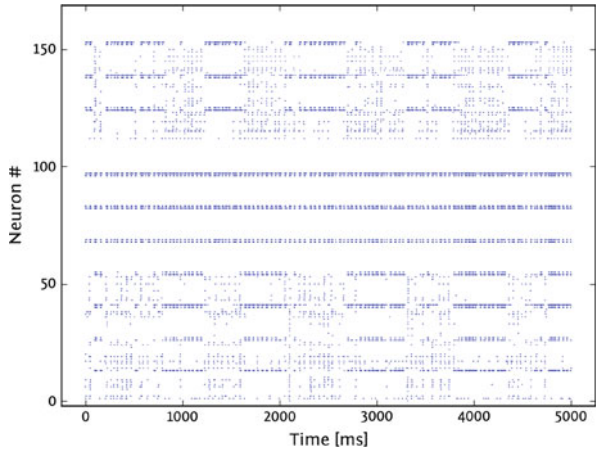

(a) Poisson input with an overall rate of 750, 1 and 550 Hz

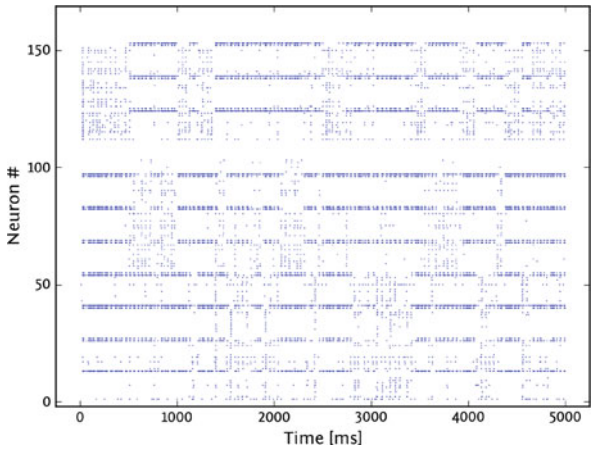

(b) Poisson input with an overall rate of 750, 700 and 650 Hz

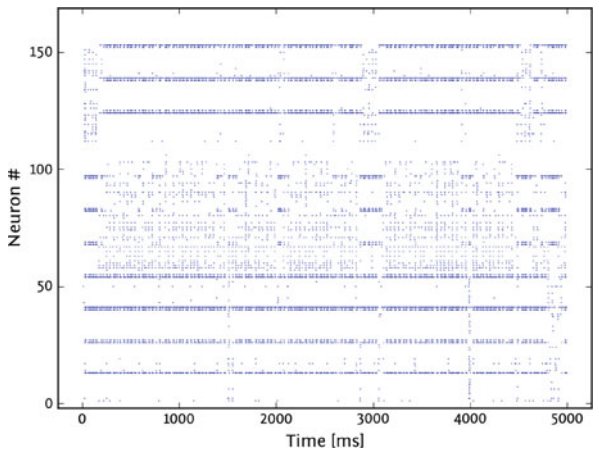

(c) Poisson input with an overall rate of 700, 800 and 600 Hz

Fig. 36 Raster plots of scaled-down L2/3 attractor memory networks on the FACETS chip-based neuromorphic system. For explanation see text

with respect to the other attractors, results in almost complete dominance of attractor 2. Again, this is due to the lack of adaptive mechanisms.

3.2.2 Virtual hardware

The scaled-down version of the L2/3 attractor memory network model was successfully implemented on the *virtual hardware* (see Sect. 2.3). No changes had to be applied to the model in order to realize it on the FACETS wafer-scale virtual hardware, as the HICANN building block implements AdEx-type neurons, which include the dynamics of the leaky integrate-and-fire neurons from the chip-based system. The scaled-down model passed through the whole mapping process described in Sects. 2.2.2, 2.2.5, 2.2.6, and 2.2.7 and was finally mapped and simulated on a snippet of 2×2 reticles of a wafer.

The results of the virtual hardware simulation of the scaled-down L2/3 attractor memory network can be seen in Fig. 37, where the individual attractors were stimulated with different rates. Depending on the specific stimulation, the network reliably exhibits the same behavior as reference software simulations, which are described in the following section.

Figure 38 shows the 3-D visualization of the network model and its mapping onto the wafer with the GraViTo software (see Sect. 2.2.9).

3.2.3 NEST simulator

The same PyNN script as in the previous sections was eventually used with the software simulator NEST, as a means of comparison to an *ideal* simulation environment. By providing identical parameter settings, one can hereby gain a good perspective for gauging the effects of the hardware-inherent fluctuations.

The results are practically identical to the ones from the virtual hardware, perhaps not surprisingly, as it is not subject to hardware-specific manufacturing process fluctuations (see Figs. 37, 39).

Also, due to its small size, does the network not pose any challenge to the mapping algorithm, making the hardware realization a perfect replica of its software counterpart. Still, the successful emulation offers a convincing proof of the efficacy of our mapping work flow.

The chip-based neuromorphic device, on the other hand, is subject to the full range of hardware-specific distortions. Nevertheless, the resulting network dynamics agree well with the NEST results, requiring only small adjustments in the input activity. These results are expected to greatly improve on the wafer-scale hardware, thanks to the superior architecture of the HICANN units. Also, a much more complex neuron model and the availability of both short- and long-term synaptic plasticity mechanisms will make the wafer-scale hardware much more capable of emulating biologically accurate network models.

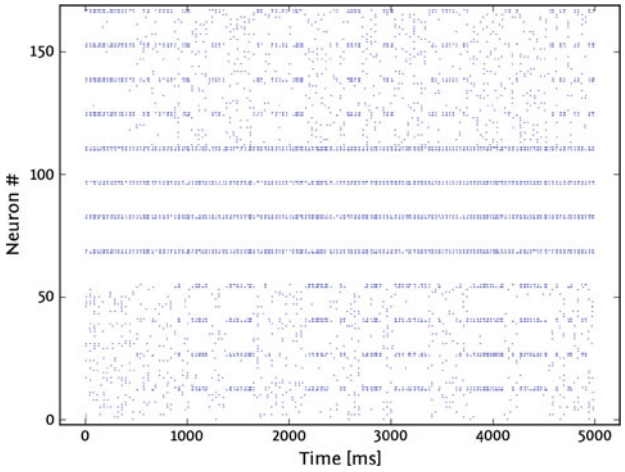

(a) Poisson input with overall rate of 550, 1 and 550 Hz

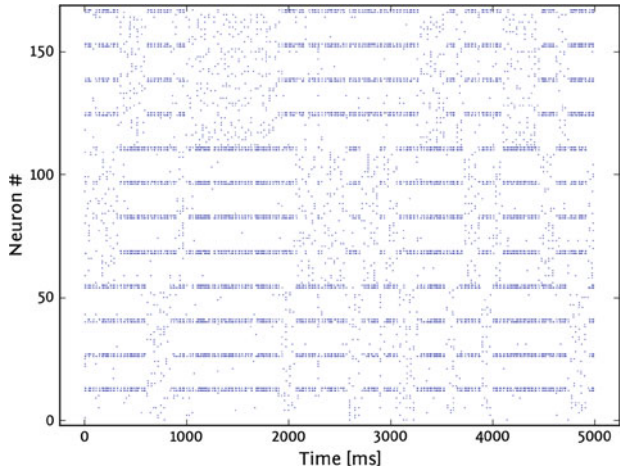

(b) Poisson input with overall rate of 700 Hz

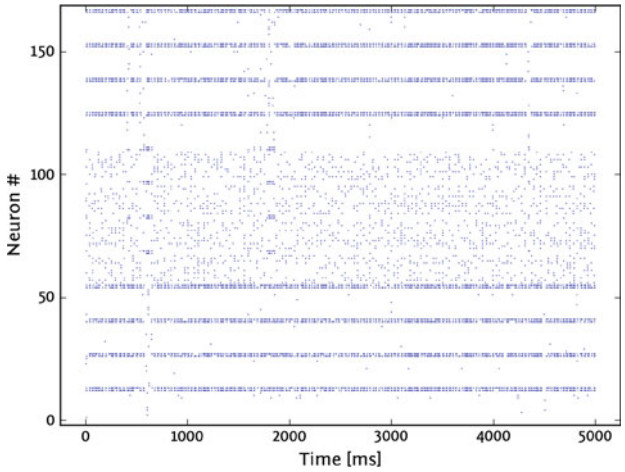

(c) Poisson input with overall rate of 700, 800 and 700 Hz

Fig. 37 Scaled-down L2/3 attractor memory network simulated with the *virtual hardware*. Attractors have been excited by Poisson input with different overall rates

As a conclusion, we note that the software results are in very good agreement with the ones generated by our hardware back-ends, thus supporting our work flow concept and

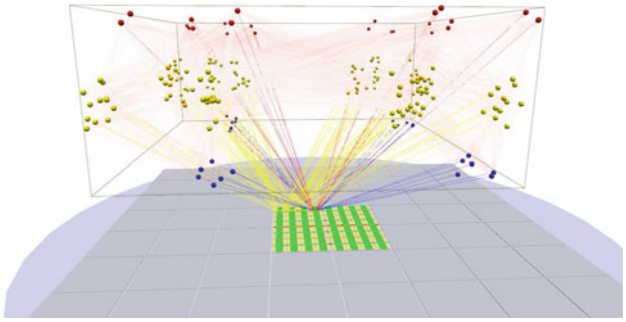

Fig. 38 3-D view of the scaled-down L2/3 attractor memory network model (see Fig. 35) and its mapping to the wafer generated by the GraVi-To software (Sect. 2.2.9)

solidifying the position of our neuromorphic hardware as a universal modeling tool. The particularly appealing feature, especially from a neural modeling perspective, is the seamless transition from software simulation to hardware emulation, which, from the perspective of the PyNN user, is accomplished by modifying a single line of code.

3.2.4 Analysis and development of STDP in hardware

The effects of discrete synaptic weights on networks is analyzed by means of a simple network (Fig. 40a). Ten pre-synaptic neurons are connected to one post-synaptic neuron using both the reference and the hardware inspired STDP synapses as described in Sect. 2.5.3. In order to analyze effects of discrete weights isolated from other hardware-specific constraints the weight update frequency is set equal to the time resolution $h = 0.1$ ms of the software simulator. The spike rates of the pre-synaptic neurons are adapted in such way that the post-synaptic neuron is firing at about 10 Hz. In case of correlated pre-synaptic neurons their correlation coefficient is $c = 0.05$ (Kuhn et al. 2003). Varying the input spike rates or the correlation coefficient does not change the conceptual outcome. As the currently implemented hardware synapses have a weight resolution of 4 bits, this resolution is used to test the performance of the hardware.

Figure 40 shows the mean weight traces for runs with correlated pre-synaptic neurons as well as for separate runs with uncorrelated pre-synaptic neurons. In case of hardware inspired STDP synapses (Fig. 40c) the standard deviations of the mean weight traces are much larger than those of the reference STDP synapses. These increased deviations are due to the large weight steps between adjacent discrete weights. Applying a t-test to the synaptic weight distribution after 1,000 s shows that the hardware inspired STDP synapses can nevertheless distinguish between uncorrelated and correlated input. For hardware inspired STDP synapses the probability that the synaptic weights of both populations are separated is $p = 0.02$, compared to $p = 2 \cdot 10^{-8}$ for the reference

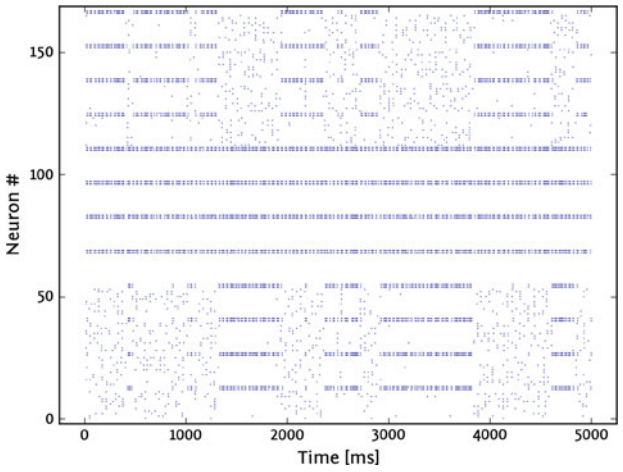

(a) Poisson input with overall rate of 550, 1 and 550 Hz

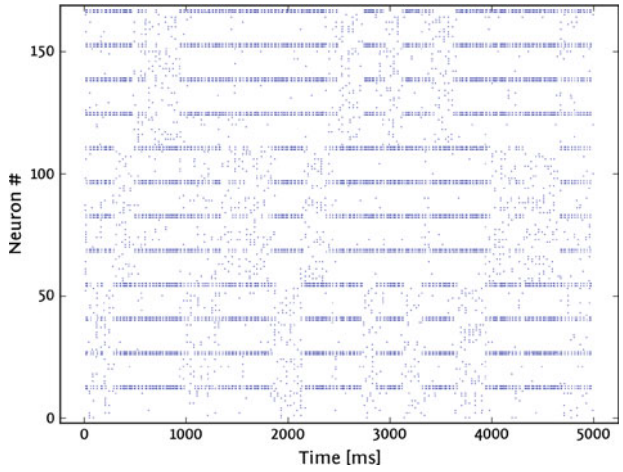

(b) Poisson input with overall rate of 700 Hz

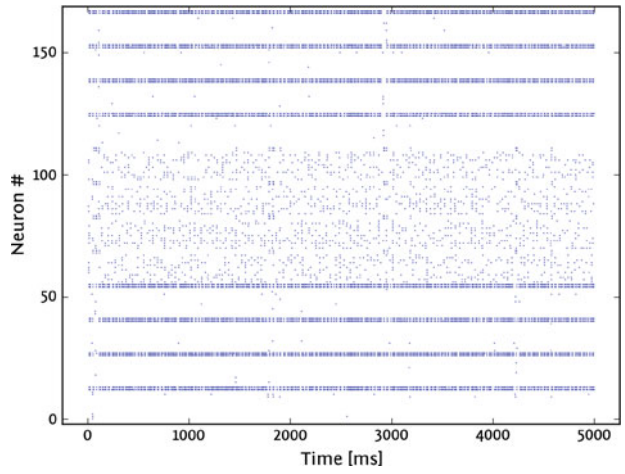

(c) Poisson input with overall rate of 700, 800 and 700 Hz

Fig. 39 Scaled-down L2/3 attractor memory network simulated with the NEST software. Attractors have been excited by Poisson input with different overall rates

STDP synapses. This ability of distinction determines the ability of detecting synchronous input, which is fundamental for most STDP applications. For correlation coefficients as low as $c = 0.05$ a resolution of 4 bits is still sufficient to detect synchronous pre-synaptic firing.

3.3 Software performance

The usability of any hardware modeling platform strongly depends on the time needed for configuration and reprogramming, thus the benchmarks introduced in Sect. 2.4 also serve as tests for *scalability* in terms of time and space.

Figure 41 shows that the space consumption for the *BioModel* data grows almost linearly depending on the number of neurons and the synaptic density ρ_{Syn}. Thus, for the given benchmarks, the model sizes for networks with a neuron count $N_{BIO} \leq 10^5$ and an approximate average $\rho_{Syn} \leq 10\%$ stay within an acceptable limit of 10 GB. Furthermore, the placement algorithms, in spite of the cubical problem, grow below $O(n^2)$ and as such fulfill the requirement of a reasonable runtime for complex mapping problems (Ehrlich et al. 2010).

3.4 HICANN prototype calibration

In order to provide an example of its functionality, the calibration framework described in Sect. 2.2.7 has been used to reproduce a biologically relevant tonic spiking neuron on the HICANN prototype. First, a reference simulation of a tonic spiking neuron using the AdEx model was created. For this simulation, the adaptation and exponential terms were disabled. The simulated neuron showed a firing rate of 53.4 Hz, which due to the speedup factor of the HICANN system corresponds to 534 kHz in the hardware domain.

The calibration was performed on a hardware neuron, and the calibration data was stored in the database. Then, the biological parameters from the reference simulation were sent to the database, which provided the necessary hardware parameters in return. The floating gates of the corresponding neuron on the HICANN prototype were then programmed with these values. The results are shown in Fig. 42. After calibration, the hardware neuron showed a firing rate of 536 kHz, which is very close to the reference simulation.

4 Discussion

Within the FACETS research project, a novel type of neuromorphic hardware system has been developed. The device combines massive acceleration, large network sizes and a high configurability with the possible advantages inherent to analog designs, such as power efficiency and a time-continuous operation. Following this strategy, neuromorphic engineering has the potential to step out of its niche and provide new and relevant input to neuroscience, e.g., toward the understanding of cortical dynamics. Still, as we noticed during the

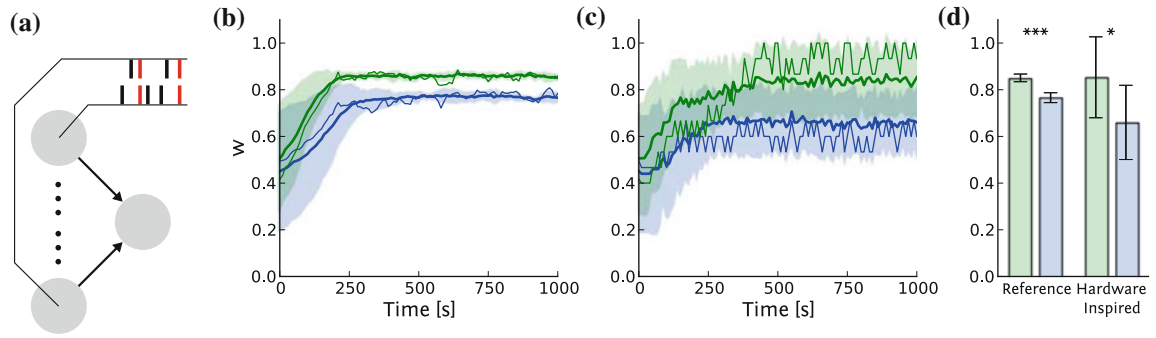

Fig. 40 **a** STDP evaluation network layout with variable correlation strength c. *Gray circles* represent neurons, *arrows* synapses. The spike trains of neurons contain random (*black*) and correlated (*red*) spikes. **b** Mean (*thick line*) and standard deviation (*shaded area*) of reference STDP synapses with correlated (*green*) and uncorrelated (*blue*) pre-synaptic neurons. *Thin lines* are single example traces. **c** Like **b**, but with hardware inspired STDP synapses. **d** Weight distributions after 1000 s for reference ($p < 0.001$) and hardware inspired ($p < 0.05$) STDP synapses

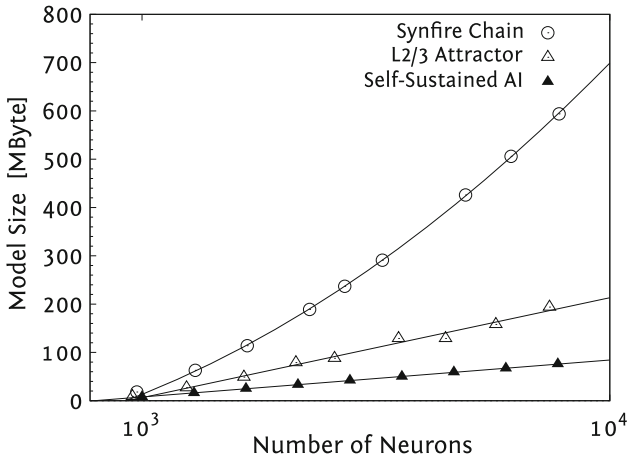

Fig. 41 Mapping process scaling in terms of *BioModel* sizes (Ehrlich et al. 2010)

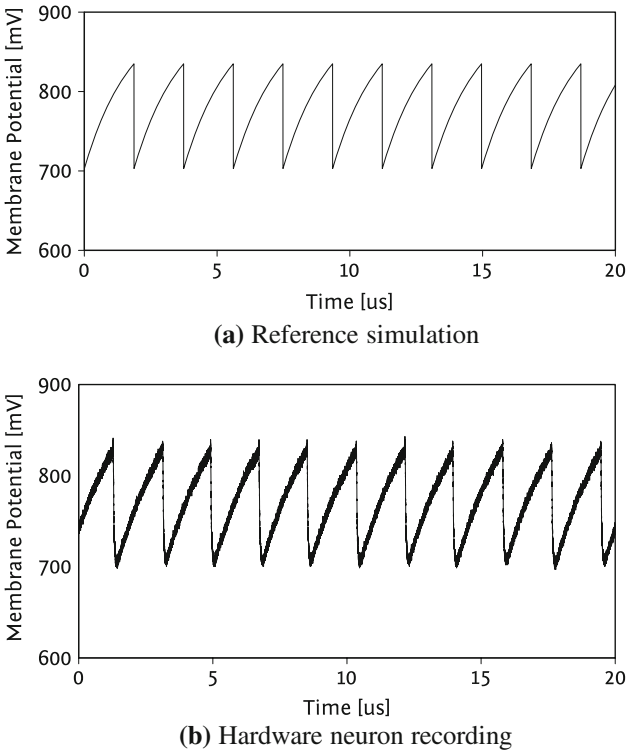

Fig. 42 Comparison between the reference simulation (**a**) and the membrane recording of the hardware neuron after calibration (**b**)

development of our specific system and during first experiments with prototypes, the originally available techniques and tools were clearly insufficient for exploiting the potential of such devices in a modeling context. It is our experience that the quality of interfaces that make hardware flexibility actually usable is as essential as the electronic substrate itself.

The presented work approaches this challenge by introducing a methodological framework that establishes a balance between the configuration complexity and potential of a novel hardware system on the one hand and the usability and spectrum of possible applications for modelers without a hardware background on the other. This neuromorphic modeling workflow has been depicted both in its conceptual whole and by means of detailed component descriptions. It represents one major outcome of the inter-disciplinary collaboration among FACETS partners, thereby integrating expertise and progress in the fields of physiologically well-founded cortex modeling, hardware engineering and community-driven software development. The multitude of the described components and their structured interaction reflects the comprehensiveness we are aiming at.

We showed experimental data that provide a proof of mature functionality of the implemented stack of model-to-hardware translation tools. The experimental results of mapping distortion studies on the basis of our virtual wafer-scale hardware system and reference software simulations represent examples of ongoing analysis work that continuously improves our software layer stack, the hardware design

and our neuromorphic modeling experience. A dedicated follow-up publication focusing on these analysis efforts is in preparation. In particular, this work soon to be published will focus on computational aspects and address many questions that remain open at this point, especially concerning the computational and functional limitations that are imposed to the network models by the presented concepts.

The architectures in the presented benchmark collection already now cover a wide spectrum of computationally interesting aspects and relevant connectivity structures. However, although the workflow presented in this article enables us to successfully realize these benchmarks with the FACETS hardware devices, further validation of the introduced concepts will be required. Important features like the synaptic plasticity mechanisms and large regions of the technically available hardware configuration space have not yet been systematically explored with our workflow in network contexts. And, as soon as alternative neuromorphic devices with a comparable degree of configurability and size (but possibly different solutions and components) will be available from other groups, all applicable aspects of the described neuromorphic workflow will have to be tested also with these platforms.

In addition to such necessary investigations, we plan to extend the set of models that we use to benchmark and tune our workflow. So far the realization of a large variety of biologically relevant structures has been the primary goal of iteratively applying the depicted optimization process. A second focus will be put on computationally powerful architectures in general, independent of their biological plausibility. Building on this work, the presented methodological framework with the neuromorphic hardware system at its core will eventually be used to approach open neuroscientific questions.

Acknowledgements The research leading to these results has received funding by the Sixth Framework Programme of the European Community (EC) under grant agreement no. 15879 (FACETS). Marc-Olivier Schwartz is supported by a doctoral fellowship in the Seventh Framework Programme of the EC under grant agreement no. 237955 (FACETS-ITN). Lyle Muller is supported by a doctoral fellowship from the École des Neurosciences de Paris (ENP, Paris School of Neuroscience).

References

Aviel Y, Mehring C, Abeles M, Horn D (2003) On embedding synfire chains in a balanced network. Neural Comput 15(6):1321–1340

Berge HKO, Häfliger P (2007) High-speed serial AER on FPGA. In: Proceedings of the 2007 IEEE international symposium on circuits and systems (ISCAS), pp 857–860

Bill J, Schuch K, Brüderle D, Schemmel J, Maass W, Meier K (2010) Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity. Front Comp Neurosci 4(129)

Binzegger T, Douglas RJ, Martin KAC (2004) A quantitative map of the circuit of cat primary visual cortex. J Neurosci 24(39):8441–8453

Bontorin G, Renaud S, Garenne A, Alvado L, Le Masson G, Tomas J (2007) A real-time closed-loop setup for hybrid neural networks. In: Proceedings of the 29th annual international conference of the IEEE engineering in medicine and biology society (EMBS2007)

BrainScaleS (2010) Project website. http://www.brainscales.eu

Brette R, Gerstner W (2005) Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. J Neurophysiol 94:3637–3642

Brette R, Rudolph M, Carnevale T, Hines M, Beeman D, Bower JM, Diesmann M, Morrison A, Goodman PH, Harris FCJr, Zirpe M, Natschlager T, Pecevski D, Ermentrout B, Djurfeldt M, Lansner A, Rochel O, Vieville T, Muller E, Davison AP, El Boustani S, Destexhe A (2006) Simulation of networks of spiking neurons: a review of tools and strategies. J Comput Neurosci 23(3):349–398

Brüderle D (2009) Neuroscientific modeling with a mixed-signal VLSI hardware system. PhD thesis, Ruprecht-Karls-Universität, Heidelberg

Brüderle D, Müller E, Davison A, Muller E, Schemmel J, Meier K (2009) Establishing a novel modeling tool: a python-based interface for a neuromorphic hardware system. Front Neuroinform 3(17)

Brüderle D, Bill J, Kaplan B, Kremkow J, Meier K, Müller E, Schemmel J (2010) Simulator-like exploration of cortical network architectures with a mixed-signal VLSI system. In: Proceedings of the 2010 IEEE international symposium on circuits and systems (ISCAS), pp 2784–2787

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J Comput Neurosci 8(3):183–208

Burkitt A, Gilson M, Hemmen J (2007) Spike-timing-dependent plasticity for neurons with recurrent connections. Biol Cybern 96(5):533–546

Buxhoeveden D, Casanova M (2002) The minicolumn and evolution of the brain. Brain Behav Evol 60:125–151

Connors B, Gutnick M (1990) Intrinsic firing patterns of diverse neocortical neurons. Trends Neurosci 13:99–104

Costas-Santos J, Serrano-Gotarredona T, Serrano-Gotarredona R, Linares-Barranco B (2007) A spatial contrast retina with on-chip calibration for neuromorphic spike-based AER vision systems. IEEE Trans Circuits Syst 54(7):1444–1458

Dante V, Del Giudice P, Whatley A (2005) Hardware and software for interfacing to address-event based neuromorphic systems. Neuromorp Eng 2(1):5–6

Daouzli A, Saighi S, Buhry L, Bornat Y, Renaud S (2008) Weights convergence and spikes correlation in an adaptive neural network implemented on vlsi. In: Proceedings of the international conference on bio-inspired systems and signal processing (BIOSIGNALS), pp 286–291

Davison AP, Frégnac Y (2006) Learning crossmodal spatial transformations through spike-timing-dependent plasticity. J Neurosci 26(21):5604–5615

Davison AP, Brüderle D, Eppler JM, Kremkow J, Muller E, Pecevski D, Perrinet L, Yger P (2008) PyNN: a common interface for neuronal network simulators. Front Neuroinform 2:(11)

Delbrück T, Liu SC (2004) A silicon early visual system as a model animal. Vis Res 44(17):2083–2089

Destexhe A (2009) Self-sustained asynchronous irregular states and Up/Down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. J Comput Neurosci 3:493–506

Destexhe A, Contreras D, Steriade M (1998) Mechanisms underlying the synchronizing action of corticothalamic feedback through inhibition of thalamic relay cells. J Neurophysiol 79:999–1016

Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. Nature 402: 529–533

Douglas RJ, Martin KAC (2004) Neuronal circuits of the neocortex. Annu Rev Neurosci 27:419–451

Ehrlich M, Mayr C, Eisenreich H, Henker S, Srowig A, Grübl A, Schemmel J, Schüffny R (2007) Wafer-scale VLSI implementations of pulse coupled neural networks. In: Proceedings of the international conference on sensors, circuits and instrumentation systems (SSD-07)

Ehrlich M, Wendt K, Zühl L, Schüffny R, Brüderle D, Müller E, Vogginger B (2010) A software framework for mapping neural networks to a wafer-scale neuromorphic hardware system. In: Proceedings of ANNIIP 2010, pp 43–52

El Boustani S, Pospischil M, Rudolph-Lilith M, Destexhe A (2007) Activated cortical states: experiments, analyses and models. J Physiol (Paris) 101:99–109

Eppler JM, Helias M, Muller E, Diesmann M, Gewaltig MO (2008) PyNEST: a convenient interface to the NEST simulator. Front Neuroinform 2:12

FACETS (2010) Fast analog computing with emergent transient states—project website. http://www.facets-project.org

Fairhurst G (2002) RFC 3366: advice to link designers on link automatic repeat request (ARQ). http://www.rfc-editor.org/rfc/rfc3366.txt

Fieres J, Schemmel J, Meier K (2008) Realizing biological spiking network models in a configurable wafer-scale hardware system. In: Proceedings of the 2008 international joint conference on neural networks (IJCNN)

Friedmann S (2009) Extending a hardware neural network beyond chip boundaries. Diploma thesis (English), Ruprecht-Karls-Universität, Heidelberg, HD-KIP-09-41, http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=1938

Fu Z, Culurciello E, Lichtsteiner P, Delbrück T (2008) Fall detection using an address-event temporal contrast vision sensor. In: Proceedings of the 2008 IEEE international symposium on circuits and systems (ISCAS), pp 424–427

Gewaltig MO, Diesmann M (2007) NEST (neural simulation tool). Scholarpedia 2(4):1430

Gomez-Rodriguez F, Miro-Amarante L, Diaz-del Rio F, Linares-Barranco A, Jimenez G (2010) Real time multiple objects tracking based on a bio-inspired processing cascade architecture. In: Proceedings of 2010 IEEE international symposium on circuits and systems (ISCAS), pp 1399–1402

Goodman D, Brette R (2008) Brian: a simulator for spiking neural networks in Python. Front Neuroinform 2(5)

Gütig R, Aharonov R, Rotter S, Sompolinsky H (2003) Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. J Neurosci 23(9):3697–3714

Häfliger P (2007) Adaptive WTA with an analog VLSI neuromorphic learning chip. IEEE Trans Neural Netw 18(2):551–572

Hartmann S, Schiefer S, Scholze S, Partzsch J, Mayr C, Henker S, Schüffny R (2010) Highly integrated packet-based AER communication infrastructure with 3Gevent/s throughput. In: IEEE international conference on electronics, circuits and systems, ICECS, Dec 2010, pp 952–955

Hines ML, Carnevale NT (2006) The NEURON book. Cambridge University Press, Cambridge

Hines ML, Davison AP, Muller E (2009) NEURON and Python. Front Neuroinform 3(1)

Horak R (2007) Telecommunications and data communications handbook. Wiley-Interscience, New York

Hunter JD (2007) Matplotlib: a 2D graphics environment. IEEE Comput Sci Eng 9(3):90–95

Indiveri G (2008) Neuromorphic VLSI models of selective attention: from single chip vision sensors to multi-chip systems. Sensors 8(9):5352–5375

Indiveri G, Chicca E, Douglas R (2006) A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity. IEEE Trans Neural Netw 17(1):211–221

Indiveri G, Chicca E, Douglas R (2009) Artificial cognitive systems: from VLSI networks of spiking neurons to neuromorphic cognition. Cogn Comput 1(2):119–127

Jeltsch S (2010) Computing with transient states on a neuromorphic multi-chip environment. Diploma thesis (English), Ruprecht-Karls-Universität, Heidelberg, HD-KIP 10-54, http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=2095

Jones E, Oliphant T, Peterson P (2001) SciPy: open source scientific tools for Python. http://www.scipy.org/

Kaplan B, Brüderle D, Schemmel J, Meier K (2009) High-conductance states on a neuromorphic hardware system. In: Proceedings of the 2009 international joint conference on neural networks (IJCNN)

Kremkow J, Kumar A, Rotter S, Aertsen A (2007) Emergence of population synchrony in a layered network of the cat visual cortex. Neurocomputing 70:2069–2073

Kremkow J, Aertsen A, Kumar A (2010a) Gating of signal propagation in spiking neural networks by balanced and correlated excitation and inhibition. J Neurosci 30(47):15760–15768

Kremkow J, Perrinet L, Masson G, Aertsen A (2010b) Functional consequences of correlated excitatory and inhibitory conductances. J Comput Neurosci 28(3):579–594

Kuhn A, Aertsen A, Rotter S (2003) Higher-order statistics of input ensembles and the response of simple model neurons. Neural Comput 15(1):67–101

Kumar A, Rotter S, Aertsen A (2008) Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. J Neurosci 28(20):5268–5280

Kumar A, Rotter S, Aertsen A (2010) Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. Nat Rev Neurosci 11(9):615–627

Lande T, Ranjbar H, Ismail M, Berg Y (1996) An analog floating-gate memory in a standard digital technology. In: Proceedings of fifth international conference on microelectronics for neural networks, pp 271–276

Langtangen HP (2008) Python scripting for computational science, 3rd edn. Springer, Berlin

Lewis MA, Etienne-Cummings R, Cohen AH, Hartmann M (2000) Toward biomorphic control using custom aVLSI chips. In: Proceedings of the international conference on robotics and automation. IEEE Press

Lundqvist M, Rehn M, Djurfeldt M, Lansner A (2006) Attractor dynamics in a modular network of neocortex. Netw Comput Neural Syst 17(3):253–276

Lundqvist M, Compte A, Lansner A (2010) Bistable, irregular firing and population oscillations in a modular attractor memory network. PLoS Comput Biol 6(6)

Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G, Wu C (2004) Interneurons of the neocortical inhibitory system. Nat Rev Neurosci 5(10):793–807

Mead CA (1989) Analog VLSI and neural systems. Addison Wesley, Reading

Mead CA (1990) Neuromorphic electronic systems. Proc IEEE 78:1629–1636

Mead CA, Mahowald MA (1988) A silicon model of early visual processing. Neural Netw 1(1):91–97

Merolla PA, Boahen K (2006) Dynamic computation in a recurrent network of heterogeneous silicon neurons. In: Proceedings of the 2006 IEEE international symposium on circuits and systems (ISCAS)

Millner S, Grübl A, Schemmel J, Meier K, Schwartz M-O (2010) A VLSI implementation of the adaptive exponential integrate-and-fire neuron model. In: Advances in neural information processing systems (NIPS), vol 23, pp 1642–1650

Mitra S, Fusi S, Indiveri G (2009) Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI. IEEE Trans Biomed Circuits Syst–3(1):32–42

Morrison A, Mehring C, Geisel T, Aertsen A, Diesmann M (2005) Advancing the boundaries of high connectivity network simulation with distributed computing. Neural Comput 17(8):1776–1801

Morrison A, Aertsen A, Diesmann M (2007) Spike-timing-dependent plasticity in balanced random networks. Neural Comput 19(6):1437–1467

Morrison A, Diesmann M, Gerstner W (2008) Phenomenological models of synaptic plasticity based on spike timing. Biol Cybern 98(6):459–478

Mountcastle VB (1997) The columnar organization of the neocortex. Brain 120(4):701–722

Naud R, Marcille N, Clopath C, Gerstner W (2008) Firing patterns in the adaptive exponential integrate-and-fire model. Biol Cybern 99(4):335–347

Netter T, Franceschini N (2002) A robotic aircraft that follows terrain using a neuromorphic eye. In: Conf. intelligent robots and system, pp 129–134

NeuroTools (2008) Website. http://neuralensemble.org/trac/NeuroTools

Norris M (2003) Gigabit ethernet technology and applications. Artech House, Boston

Oliphant TE (2007) Python for scientific computing. IEEE Comput Sci Eng 9(3):10–20

Oster M, Whatley AM, Liu SC, Douglas RJ (2005) A hardware/software framework for real-time spiking systems. In: Proceedings of the 2005 international conference on artificial neural networks (ICANN)

Pecevski DA, Natschläger T, Schuch KN (2009) PCSIM: a parallel simulation environment for neural circuits fully integrated with Python. Front Neuroinform 3:(11)

Pfeiffer M, Nessler B, Douglas RJ, Maass W (2010) Reward-modulated hebbian learning of decision making. Neural Comput 22(6): 1399–1444

Philipp S, Schemmel J, Meier K (2009) A QoS network architecture to interconnect large-scale VLSI neural networks. In: Proceedings of the 2009 international joint conference on neural networks (IJCNN), pp 2525–2532

Pospischil M, Toledo-Rodriguez M, Monier C, Piwkowska Z, Bal T, Frégnac Y, Markram H, Destexhe A (2008) Minimal hodgkin-huxley type models for different classes of cortical and thalamic neurons. Biol Cybern 99(4):427–441

Renaud S, Tomas J, Bornat Y, Daouzli A, Saighi S (2007) Neuromimetic ICs with analog cores: an alternative for simulating spiking neural networks. In: Proceedings of the 2007 IEEE symposium on circuits and systems (ISCAS)

Schemmel J, Meier K, Muller E (2004) A new VLSI model of neural microcircuits including spike time dependent plasticity. In: Proceedings of the 2004 international joint conference on neural networks (IJCNN), IEEE Press, pp 1711–1716

Schemmel J, Grübl A, Meier K, Muller E (2006) Implementing synaptic plasticity in a VLSI spiking neural network model. In: Proceedings of the 2006 international joint conference on neural networks (IJCNN), IEEE Press

Schemmel J, Brüderle D, Meier K, Ostendorf B (2007) Modeling synaptic plasticity within networks of highly accelerated I&F neurons. In: Proceedings of the 2007 IEEE international symposium on circuits and systems (ISCAS), IEEE Press, pp 3367–3370

Schemmel J, Fieres J, Meier K (2008) Wafer-scale integration of analog neural networks. In: Proceedings of the 2008 international joint conference on neural networks (IJCNN)

Schemmel J, Brüderle D, Grübl A, Hock M, Meier K, Millner S (2010) A wafer-scale neuromorphic hardware system for large-scale neural modeling. In: Proceedings of the 2010 IEEE international symposium on circuits and systems (ISCAS), pp 1947–1950

Schilling M (2010) A highly efficient transport layer for the connection of neuromorphic hardware systems. Diploma thesis, Ruprecht-Karls-Universität, Heidelberg, HD-KIP-10-09, http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=2000

Scholze S, Henker S, Partzsch J, Mayr C, Schüffny R (2010) Optimized queue based communication in VLSI using a weakly ordered binary heap. In: Proceedings of the 2010 international conference on mixed design of integrated circuits and systems (MIXDES)

Serrano-Gotarredona R, Oster M, Lichtsteiner P, Linares-Barranco A, Paz-Vicente R, Gómez-Rodríguez F, Riis HK, Delbrück T, Liu SC, Zahnd S, Whatley AM, Douglas RJ, Häfliger P, Jimenez-Moreno G, Civit A, Serrano-Gotarredona T, Acosta-Jiménez A, Linares-Barranco B (2006) AER building blocks for multi-layer multi-chip neuromorphic vision systems. In: Weiss Y, Schölkopf B, Platt J (eds) Advances in neural information processing systems, vol 18. MIT Press, Cambridge, pp 1217–1224

Sjöström PJ, Rancz EA, Roth A, Häusser M (2008) Dendritic excitability and synaptic plasticity. Physiol Rev 88(2):769–840

Steriade M (2003) Neuronal substrates of sleep and epilepsy. Cambridge University Press, Cambridge

Tao L, Shelley M, McLaughlin D, Shapley R (2004) An egalitarian network model for the emergence of simple and complex cells in visual cortex. PNAS 101:366–371

van Rossum MCW, Bi GQ, Turrigiano G (2000) Stable hebbian learning from spike timing-dependent plasticity. J Neurosci 20: 8812–8821

Vogels TP, Abbott LF (2005) Signal propagation and logic gating in networks of integrate-and-fire neurons. J Neurosci 25(46): 10786–10795

Vogelstein RJ, Mallik U, Vogelstein JT, Cauwenberghs G (2007) Dynamically reconfigurable silicon array of spiking neuron with conductance-based synapses. IEEE Trans Neural Netw 18:253–265

Vogginger B (2010) Testing the operation workflow of a neuromorphic hardware system with a functionally accurate model. Diploma thesis, Ruprecht-Karls-Universität, Heidelberg, HD-KIP-10-12, http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=2003

Wendt K, Ehrlich M, Schüffny R (2008) A graph theoretical approach for a multistep mapping software for the FACETS project. In: Proceedings of the 2008 WSEAS international conference on computer engineering and applications (CEA), pp 189–194

Wendt K, Ehrlich M, Schüffny R (2010) GMPath—a path language for navigation, information query and modification of data graphs. In: Proceedings of ANNIIP 2010, pp 31–42

Zucker RS, Regehr WG (2002) Short-term synaptic plasticity. Annu Rev Physiol 64:355–405

Petrovici et al. (2014): *Bridging the gap between software simulation and emulation on neuromorphic hardware*

Noname manuscript No.
(will be inserted by the editor)

Bridging the gap between software simulation and emulation on neuromorphic hardware

An investigation of causes, effects and compensation of network-level anomalies in a mixed-signal waferscale neuromorphic modeling platform

Mihai A. Petrovici · Bernhard Vogginger · Paul Müller · Oliver Breitwieser · Mikael Lundqvist · Lyle Muller · Matthias Ehrlich · Alain Destexhe · Anders Lansner · René Schüffny · Johannes Schemmel · Karlheinz Meier

Received: date / Accepted: date

Abstract Advancing the size and complexity of neural network models leads to an ever increasing demand for computational resources for their simulation. Neuromorphic devices offer a number of advantages over conventional computing architectures, such as high emulation speed or ow power consumption, but this usually comes at the price of reduced configurability and precision. In this article, we investigate the consequences of several such effects that are common to neuromorphic devices, more specifically limited hardware resources, limited parameter configurability and parameter variations due to fixed pattern noise. Our final aim is to provide an array of methods for coping with such inevitable distortions. As a testbench, we use an executable system specification (ESS) of the BrainScaleS neuromorphic system, which has been designed as a universal emulation back-end for neuroscientific modeling. We address the most essential limitations of this device in detail and study their effects on three prototypical benchmark network models within a well-defined, systematic workflow. For each network model, we start by defining quantifiable functionality measures by which we then assess the effects of typical hardware-induced distortions, both in idealized software simulations and on the ESS. For those effects that cause intolerable deviations from the original network dynamics, we suggest generic compensation mechanisms and demonstrate their effectiveness. Both the suggested workflow and the investigated compensation mechanisms are largely back-end independent and provide a generic methodological environment for configurable neuromorphic emulation devices.

Keywords Neuromorphic Hardware · Configurability · Fixed-Pattern Noise · Synapse Loss · Compensation of Distortions

M. A. Petrovici · P. Müller · O. Breitwieser · J. Schemmel · K. Meier
Kirchhoff Institute for Physics
Ruprecht-Karls-Universität Heidelberg, Germany
Tel.: +49 6221 549847
E-mail: mpedro@kip.uni-heidelberg.de

M. Ehrlich · R. Schüffny · B. Vogginger
Institute of Circuits and Systems, Technische Universität Dresden, Germany

M. Lundqvist · A. Lansner
Computational Biology, KTH Stockholm, Sweden

A. Destexhe · L. Muller
Unité de Neuroscience, Information et Complexité, CNRS, Gif sur Yvette, France

1 Introduction

1.1 Modeling and computational neuroscience

The limited availability of detailed biological data has always posed a major challenge to the advance of neuroscientific understanding. The formulation of theories about information processing in the brain has therefore been predominantly model-driven, with much freedom of choice in model architecture and parameters. As more powerful mathematical and computational tools became available, increasingly detailed and complex cortical models have been proposed. However, because of the manifest nonlinearity and sheer complexity of interactions that take place in the nervous system, analytically treatable ensemble-based models can only partly cover the vast range of activity patterns and behavioral phenomena that are characteristic for biological nervous systems (Laing and Lord, 2009). The high level of model complexity often required for computational proficiency and biological plausibility has led to a rapid

development of the field of computational neuroscience, which focuses on the simulation of network models as a powerful complement to the search for analytic solutions (Brette et al, 2007).

The feasibility of the computational approach has been facilitated by the development of the hardware devices used to run neural network simulations. The brisk pace at which available processing speed has been increasing over the past few decades, as allegorized by Moore's Law, as well as the advancement of computer architectures in general, closely correlate to the size and complexity of simulated models. Today, network models with tens of thousands of neurons are routinely simulated on desktop machines, with supercomputers allowing several orders of magnitude more (Djurfeldt et al, 2008; Helias et al, 2012). However, as many authors have pointed out (see e.g. Morrison et al (2005), Brette et al (2007)), the inherently massively parallel structure of biological neural networks becomes progressively difficult to map to conventional architectures based on digital general-purpose CPUs, as network size and complexity increase.

Conventional simulation becomes especially restrictive when considering long time scales, such as are required for modeling long-term network dynamics or when performing statistics-intensive experiments. Additionally, power consumption can quickly become prohibitive at these scales (Bergman et al, 2008; Hasler and Marr, 2013).

1.2 Neuromorphic Hardware

The above issues can, however, be eluded by reconsidering the fundamental design principles of conventional computer systems. The core idea of the so-called neuromorphic approach is to implement features (such as connectivity) or components (neurons, synapses) of neural networks directly *in silico*: instead of calculating the dynamics of neural networks, neuromorphic devices contain physical representations of the networks themselves, behaving, by design, according to the same dynamic laws. An immediate advantage of this approach is its inherent parallelism (emulated network components evolve in parallel, without needing to wait for clock signals or synchronization), which is particularly advantageous in terms of scalability. First proposed by Mead in the 1980s (Mead and Mahowald, 1988; Mead, 1989, 1990), the neuromorphic approach has since delivered a multitude of successful applications (Renaud et al, 2007; Indiveri et al, 2009, 2011).

By far the largest number of neuromorphic systems developed thus far are highly application-specific, such as visual processing systems (Serrano-Gotarredona et al, 2006; Merolla and Boahen, 2006; Netter and Franceschini, 2002; Delbrück and Liu, 2004) or robotic motor control devices (Lewis et al, 2000). Several groups have focused on more biological aspects, such as the neuromorphic implementation of biologically-inspired self-organization and learning (Häfliger, 2007; Mitra et al, 2009), detailed replication of Hodgkin-Huxley neurons (Zou et al, 2006) or hybrid systems interfacing analog neural networks with living neural tissue (Bontorin et al, 2007).

These devices, however, being rather specialized, can not match the flexibility of traditional software simulations. Adding configurability comes at a high price in terms of hardware resources, due to various hardware-specific limitations, such as physical size and essentially two-dimensional structure. So far there have only been few attempts at realizing highly configurable hardware emulators (Indiveri et al, 2006; Vogelstein et al, 2007; Rocke et al, 2008; Schemmel et al, 2010; Furber et al, 2012). This approach alone, however, does not completely resolve the computational bottleneck of software simulators, as scaling neuromorphic neural networks up in size becomes non-trivial when considering bandwidth limitations between multiple interconnected hardware devices (Costas-Santos et al, 2007; Berge and Häfliger, 2007; Indiveri, 2008; Fieres et al, 2008; Serrano-Gotarredona et al, 2009).

1.3 The BrainScaleS hardware system

A very efficient way of interconnecting multiple VLSI[1] modules is offered by so-called wafer-scale integration. This implies the realization of both the modules in question and their communication infrastructure on the same silicon wafer, the latter being done in a separate, post-processing step. The BrainScaleS wafer-scale hardware (Schemmel et al, 2010) uses this process to achieve a high communication bandwidth between individual neuromorphic cores on a wafer, thereby allowing a highly flexible connection topology of the emulated network. Together with the large available parameter space for neurons and synapses, this creates a neuromorphic architecture that is comparable in flexibility with standard simulation software. At the same time, it provides a powerful alternative to software simulators by avoiding the abovementioned computational bottleneck, in particular owing to the fact that the emulation duration does not scale with the size of the emulated network, since individual netowrk components operate, inherently, in parallel. An additional benefit which is inherent to this specific VLSI implementation is the

[1] Very Large Scale Integration

high acceleration with respect to biological real-time, which is facilitated by the high on-wafer bandwidth. This allows investigating the evolution of network dynamics over long periods of time which would otherwise be strongly prohibitive for software simulations.

1.4 Hardware-Induced Distortions: A Systematic Investigation

Along with the many advantages it offers, the neuromorphic approach also comes with limitations of its own. These have various causes that lie both in the hardware itself and the control software. The neural network emulated by the hardware device can therefore differ significantly from the original model, be it in terms of neuron dynamics, connectivity between populations or individual neuron or synapse parameters.

Due to the complexity of state-of-the-art neuromorphic platforms and their control software, as well as the vast landscape of emulable neural network models, a thorough and systematic approach is essential for providing reliable information about causes and effects of hardware-induced distortions in model dynamics and for ultimately designing effective compensation methods. In this article, we design and perform such a systematic analysis and compensation for several hardware-specific distortion mechanisms.

First and foremost, we identify and quantify the most important sources of model distortions. We then proceed to investigate their effect on network functionality.

In order to cover a wide range of possible network dynamics, we have chosen three very different cortical network models to serve as benchmarks. In particular, these models implement several prototypical cortical paradigms of computation, relying on winner-take-all structures (attractor networks), precise spike timing correlations (synfire chains) or balanced activity (self-sustained asynchronous irregular states). For every emulated model, we define a set of functionality criteria, based on specific aspects of the network dynamics. Most importantly, these criteria need to be precisely quantifiable, in order to facilitate an accurate comparison between software simulations and hardware emulations. These values are then measured, in software simulations, for various relevant realizations (i.e. for different network sizes, numbers of functional units etc.) of the original, undistorted network.

Because multiple distortive effects occur simultaneously in hardware emulations, it is often difficult, if not impossible, to understand the relationship between causes and effects of modifications in the network dynamics. Therefore, we investigate the effects of individual distortion mechanisms by implementing them, separately, in software simulations. As before, we perform these analyses over a wide range of network realizations, since - as we will show later - these may strongly influence the examined distortive effects.

After having established the relationship between structural network distortions caused by hardware effects and their consequences for network dynamics, we demonstrate various compensation techniques in order to restore the original network behavior.

In the final stage, for each of the studied models, we simulate an implementation on the hardware backend by running an appropriately configured executable system specification, which includes the full panoply of hardware-induced distortions. Using the proposed compensation techniques, we then attempt to deal with all these effects simultaneously.The results from these experiments are then compared to results from software simulations, thus allowing a comprehensive assertion of the effectivity of our proposed compensation techniques, as well as of the capabilities and limitations of the neuromorphic emulation device.

1.5 Article Structure

In Sec. 2, we describe our testbench neuromorphic modeling platform with its most relevant components, as well as the essential layers of the operation workflow. We continue by explaining the causes of various distortion effects that are expected to be common for similar mixed-signal neuromorphic devices. In the same section, we also introduce the executable system specification of the hardware, which we later use for experimental investigations. Sec. 3 contains the description of the three benchmark models, including characteristic aspects of their dynamics which we later use as quality controls. For each model, we discuss the effects of individual hardware-induced distortions as observed in software simulations, propose various compensation strategies and investigate their efficacy in maintaining the functionality of the benchmark models. Subsequently, we apply these methods to large-scale neuromorphic emulations and examine the results. Following the discussion in **??**, the Appendices contain detailed descriptions of hardware and model parameters, as well as of the analysis tools and algorithms we have used.

2 Neuromorphic testbench and investigated distortions

In this section we introduce the BrainScaleS neuromorphic wafer-scale hardware system and its executable system specification, henceforth called the ESS, as the testbench for our studies. The system's hardware and software components are only described on an abstract level, while highlighting the mechanisms responsible for distortions of the emulated networks. Finally, we identify the three most relevant distortion types as being synapse loss, synaptic weight noise and non-configurable axonal delays.

2.1 The BrainScaleS wafer-scale hardware

Fig. 1 shows a photograph of the BrainScaleS wafer-scale hardware system: the 8 inch silicon wafer contains 196 608 neurons and 44 million plastic synapses implemented in mixed-signal VLSI circuitry. Due to the high integration of the circuits, the capacitances and thus the intrinsic time constants are small, so that neural dynamics take place approximately 10 000 faster than biological real time. The principal building block of the wafer is the so-called HICANN[2] chip Schemmel et al (2010, 2008). During chip fabrication one is limited to a maximum area, a reticle[3], thus usually such a wafer is cut into individual chips after production. For the BrainScaleS system, however, the wafer is left intact, and additional wiring is applied onto the wafer's surface in a post-processing step. This process establishes connections betwen all 384 HICANN blocks that allow a very high bandwidth for on-wafer pulse-event communication (Schemmel et al, 2008). The neuromorphic wafer is accompanied by a stack of digital communication modules for the connection of the wafer to the host PC and to other wafers (Sec. 2.1.2).

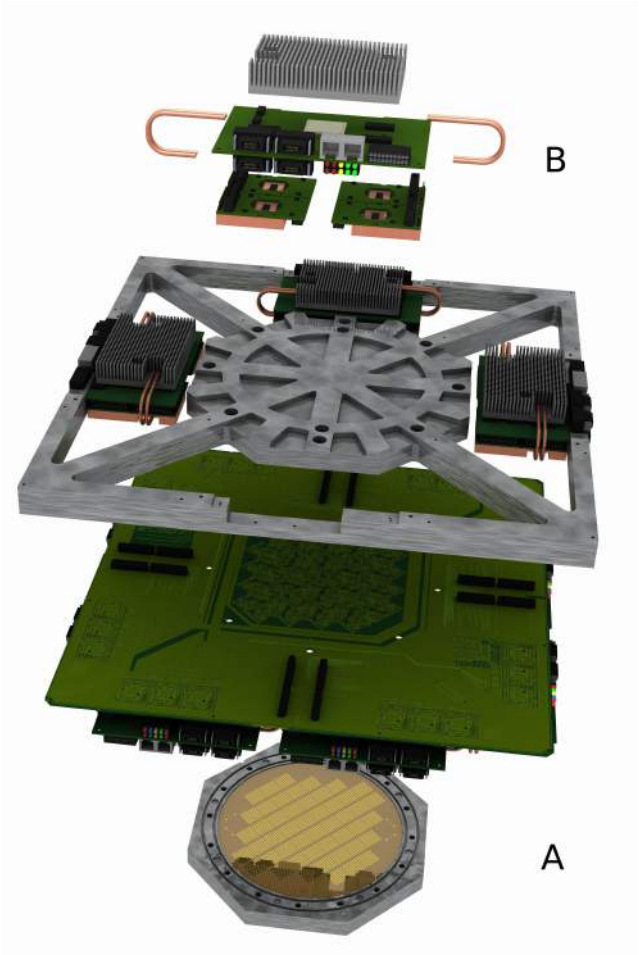

Fig. 1: The BrainScaleS wafer-scale hardware system: (**A**) Wafer comprising HICANN building blocks and on-wafer communication infrastructure covered by an aluminium plate, (**B**) digital inter-wafer and wafer-host communication modules. Also visible: mechanical and electrical support.

2.1.1 HICANN building block

On the HICANN chip (lower left of Fig. 2), one can recognize two symmetric blocks which hold the analog core modules: Most of the area is occupied by the synapse array with 224 rows and 256 columns. All synapses in a column are connected to one of the 256 neuron circuits located at the center of the chip. For each two adjacent synapse rows, there is one *synapse driver* that forms the input for pre-synaptic pulses to the synapse array. Synapse drivers are evenly distributed to the left and right side of one synapse array (56 per side). A grid of horizontal and vertical buses enables the routing of spikes from neuron circuits to synapse drivers.

Up to 64 neuron circuits can be interconnected to form neurons with up to 14336 synapses. The neurons emulate the dynamics of the Adaptive-Exponential Integrate-and-Fire model (AdEx, Brette and Gerstner (2005)) in analog circuitry, defined by equations for the membrane voltage V, the adaption current w and a reset condition

[2] High Input Count Analog Neural Network
[3] a reticle is the maximum area that can be simultaneously exposed during photolitography

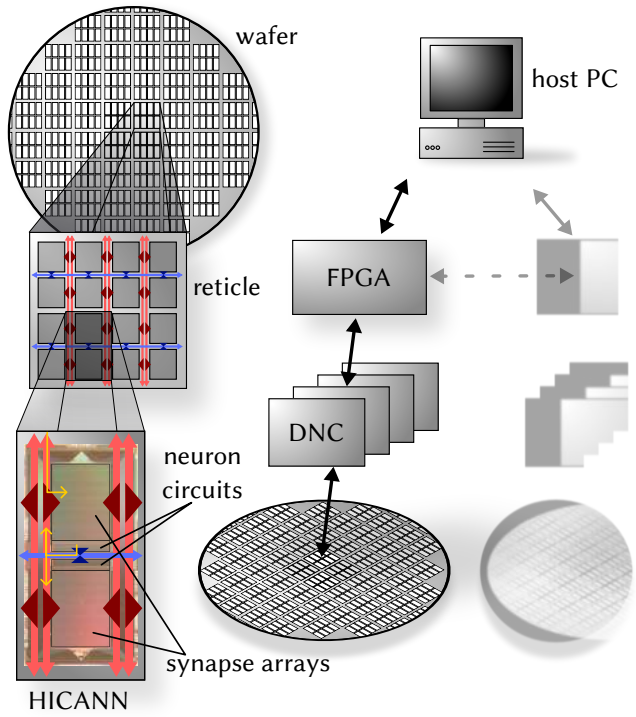

Fig. 2: Architecture of the BrainScaleS wafer-scale hardware system. Left: The HICANN building block has two symmetric halves with synapse arrays and neuron circuits. Neural activity is transported horizontally (blue) and vertically (red) via asynchronous buses that span over the entire wafer. Exemplary spike paths are shown in yellow on the HICANN: The incoming spike packet is routed to the synapse drivers. In the event that a neuron spikes, it emits a spike packet back into the routing network. Right: Off-wafer connectivity is established by a hierarchical packed-based network via DNCs and FPGAs. It interfaces the on-wafer routing buses on the HICANN building blocks. Several wafer modules can be interconnected using routing functionality between the FPGAs.

that applies when a spike is triggered:

$$C_\mathrm{m}\frac{dV}{dt} = -g_\mathrm{L}(V-E_\mathrm{L}) + g_\mathrm{L}\Delta_\mathrm{T}\exp\left(\frac{V-E_\mathrm{T}}{\Delta_\mathrm{T}}\right) - w + I^\mathrm{syn} \tag{1}$$

$$\tau_w\frac{dw}{dt} = a(V-E_\mathrm{L}) - w \tag{2}$$

$$\text{if } V \geq E^\mathrm{spike}: \quad \begin{cases} V \to E^\mathrm{r} \\ w \to w + b \end{cases}, \tag{3}$$

where C_m, g_L and E_L denote the membrane capacitance, leak conductance and leak potential, respectively, E_T and Δ_T represent the spike initiation threshold and the threshold slope factor and τ_w and a represent the adaptation time constant and coupling parameter. When V reaches a certain threshold value E^spike, a spike is emitted and the membrane potential is reset to E^r. At the same time, the adaptation variable is increased by a fixed amount b, thereby allowing for spike-frequency adaptation. An absolute refractory mechanism is supported by clamping V to its reset value for the refractory time τ_refrac. The generated spikes are transmitted digitally to synapse drivers (analog multiplier), synapses (digital multiplier) and finally other neurons, where postsynaptic conductance courses are generated and summed up linearly, resulting in the synaptic current I^syn:

$$I^\mathrm{syn} = \sum_{\text{synapses } i} g_i(E_i^\mathrm{rev} - V) \tag{4}$$

$$\tau^\mathrm{syn}\frac{dg_i}{dt} = -g_i + w_i^\mathrm{syn}\sum_{\text{spikes } s}\delta(t-t_s) \quad . \tag{5}$$

Here, g_i represents the synaptic conductance and E_i^rev the synaptic reversal potential of the i-th synapse, τ^syn the time constant of the exponential decay and w^syn the synaptic weight. In the hardware implementation (Millner et al, 2010), each neuron features two of such synaptic input circuits, which are typically used for excitatory and inhibitory input. Nearly all parameters of the neuron model and the synaptic input circuits are individually adjustable by means of analog storage banks based on floating gate technology (Lande et al, 1996). In the hardware neuron, both the circuit for the adaption mechanism and the exponential term circuit can be effectively disconnected from the membrane capacitance, such that a simple Leaky Integrate-and-Fire (LIF) model can also be emulated. The hardware membrane capacitance is fixed to one of two possible values. As the parameters controlling the temporal dynamics of the neuron such as g_L and the time constants are configurable within a wide range, the hardware is able to run at a variable speedup factor ($10^3 - 10^5$) compared to biological real time. In particular, the translation of the membrane capacitance between the hardware and the biological domain can be chosen freely due to the independent configurability of both membrane and synaptic conductances, thereby effectively allowing the emulation of point neurons of arbitrary size - within the limits imposed by the hardware parameter ranges.

In contrast to neurons, where each parameter is fully configurable within the specified ranges, the *synaptic weights* are adjustable by a combination of analog and digital memories. The synaptic weight w^syn is proportional to a row-wise adjustable analog parameter g_max and to a 4-bit digital weight specific to each synapse. The g_max of two adjacent rows can be configured to be a fixed multiple of each other. This way, two synapses of adjacent rows can be combined to offer a weight resolution of 8 bits, at the cost of halving the number of synapses for this synapse driver.

Long-term learning is incorporated in every synapse through spike-timing-dependent plasticity (STDP, Bi and Poo, 1998). The implemented STDP mechanism follows a pairwise update rule with programmable update functions (Morrison et al, 2008). As STDP is not contained in the models investigated in this article (Sec. 3), we refer to Brüderle et al (2011) and Schemmel et al (2006, 2007) for details on the hardware implementation and to Pfeil et al (2012) for an applicability study of these circuits.

In contrast to the long-term learning, the implemented *short-term plasticity* mechanism (STP) decays over several hundreds of milliseconds. It is motivated by the phenomenological model by Markram et al (1998) and depends only on the pre-synaptic activity, therefore being implemented in the synapse driver. For every incoming spike, a synapse only has access to a portion U of the recovered partition R of its total synaptic weight $w_{\text{max}}^{\text{syn}}$, which then instantly decreases by a factor $1-U$ and recovers slowly along an exponential with the time constant τ_{rec}, thus emulating synaptic depression. Facilitation is implemented by replacing the fixed U with a running variable u, which increases with every incoming spike by an amount $U(1-u)$ and then decays exponentially back to U with the time constant τ_{facil}:

$$w_{n+1}^{\text{syn}} = w_{\text{max}}^{\text{syn}} R_{n+1} u_{n+1} \tag{6}$$

$$R_{n+1} = 1 - [1 - R_n(1-u_n)] \exp\left(-\frac{\Delta t}{\tau_{\text{rec}}}\right) \tag{7}$$

$$u_{n+1} = U + u_n(1-U) \exp\left(-\frac{\Delta t}{\tau_{\text{facil}}}\right) \tag{8}$$

with Δt being the time interval between the nth and $(n+1)$st afferent spike. In contrast to the original Tsodyks-Markram (TSO) mechanism, the hardware implementation does not allow simultaneous depression and facilitation. See Sec. A.1 for details about the hardware implementation and the translation of the original model to the hardware STP.

All of the neuron and synapse parameters mentioned above are affected by fixed-pattern noise due to transistor-level mismatch in the manufacturing process. While for most parameters this distortion can be reduced to a minimum by proper calibration, it is not possible to the same extent for synaptic weights due to their limited configurability. Therefore, synaptic weight noise is likely to be the dominant distortive effect of fixed-pattern noise.

For technical details about the HICANN chip and its components we refer to Schemmel et al (2010, 2008).

2.1.2 Communication infrastructure

The infrastructure for pulse communication in the wafer-scale system is supplied by a two-layer approach: While the on-wafer network routes pulses between neurons on the same wafer, the off-wafer network connects the wafer to the outside world, i.e. to the host PC or to other wafers.

The backbone of the *on-wafer communication* consists of a grid of horizontal and vertical buses enabling the transport of action potentials by a mixture of time division and space division multiplexing. Each HICANN building block contains 64 horizontal buses at its center and 128 vertical buses located on each side of the synapse blocks, as can be seen on the left hand side of Fig. 2. A bus can carry the spikes of up to 64 source neurons by transmitting a serial 6-bit signal encoding the currently sending neuron (with an ID from 0 to 63). Neurons inject their pulses into a horizontal bus; by enabling a static switch of a sparse crossbar between horizontal and vertical buses, the injected serial signal can be made available to a vertical bus next to the synapse array. Another sparse switch matrix allows to feed the signals from the vertical buses into the synapse array, more precisely into the synapse drivers which represent the data sinks of the routing network. The bus lanes do not end at the HICANN border but run over the whole wafer by edge-connecting the HICANN building blocks. We note that, due to electrotechnical reasons, the switches could not be implemented as full matrices, thus their sparseness was chosen as a compromise still providing maximum flexibility for implementing various neural network topologies (Fieres et al, 2008; Schemmel et al, 2010). Both the sparseness of the switches and the limited number of horizontal and vertical buses represent a possible restriction for the connectivity of network models. If an emulated network requires a connectivity that exceeds the on-wafer bus capacity, some synapses will be impossible to map to the wafer and will therefore be lost.

Pulse propagation delays in the routing network are small, distance-dependent and not configurable: the time between spike detection and the onset of a post-synaptic potential (PSP) has been measured as 120 ns for a recurrent connection on a HICANN. The additional time needed to transmit a pulse across the entire wafer is typically less than 100 ns (Schemmel et al, 2008), hence the overall delay sums up to 1.2 - 2.2 ms in the biological time domain, assuming a speedup factor of 10^4. Also, in case of synchronous bursting of the neurons feeding one bus, some pulses are delayed with respect to others, as they are processed successively. In rare cases some pulses may be completely discarded, e.g., when the to-

tal rate of all 64 neurons feeding one bus exceeds 10 kHz for longer than 6.4 ms (in biological real-time).

A hierarchical packet-based network provides the infrastructure for *off- and inter-wafer communication*. All HICANNs on the wafer are connected to the surrounding system and to other wafers via 12 pulse communication subgroups (PCS). Each PCS consists of one FPGA[4] and 4 ASICs[5] that were designed for high-bandwidth pulse-event communication (so-called Digital Network Chips or DNCs). Being the only communication link to/from the wafer, the off-wafer network also transports the configuration and control information for all the circuits on the wafer. As depicted in Fig. 2, the network is hierarchically organized: one FPGA is connected to four DNCs, each of which is connected to 8 HICANNs of a reticle. Each FPGA is also connected to the host PC and potentially to up to 4 other FPGAs. When used for pulse-event communication, an FPGA-DNC-HICANN connection supports a throughput of 40 Mevents/s (Scholze et al, 2011b) with a timing precision of 4 ns. In the biological time domain, this corresponds to monitoring the spikes of all 512 neurons on a HICANN firing with a mean rate of 8 Hz each with a resolution of 0.04 ms. The same bandwidth is available simultaneously in the opposite direction, allowing a flexible network stimulation with user-defined spiketrains. For each FPGA-DNC-HICANN connection there are 512 pulse addresses that have to be subdivided into blocks of 64 used for either stimulation or recording. For all technical details about the PCS, the FPGA design and the DNC, we refer to Scholze et al (2011a), Hartmann et al (2010) and Scholze et al (2010).

Although the off-wafer communication interface allows the interconnection of multiple wafers, we restrict our studies here to the use of a single wafer.

2.2 Software framework

The utilized software stack (Brüderle et al, 2011) allows the user to define a network description and maps it to a hardware configuration.

The network definition is accomplished by using PyNN (Davison et al, 2008), a simulator-independent API[6] to describe spiking neural network models. It can interface to several simulation platforms such as NEURON (Hines et al, 2009) or NEST (Eppler et al, 2008) as well as to neuromorphic hardware platforms (Brüderle et al, 2009; Galluppi et al, 2010).

[4] Field Programmable Gate Array
[5] Application Specific Integrated Circuit
[6] Application Programming Interface

The mapping process (Ehrlich et al, 2010; Brüderle et al, 2011) translates the PyNN description of the neural network structure, as well as its neuron and synapse models and parameters, in several steps into a neuromorphic device configuration. This translation is constrained by the architecture of the device and its available resources.

The first step of the mapping process is to allocate static structural neural network elements to particular neuromorphic components during the so-called *placement*. Subsequently, a *routing* step is executed for establishing connections in between the placed components. During the final *parameter transformation* step, all parameters of the network components (neurons and synapses) are translated into hardware parameters. This includes the application of previously obtained *calibration* data to reduce mismatches between ideal neuromorphic circuitry behavior and real analogue signal hardware behavior.

The objective of the mapping process is to find a configuration of the hardware that best reproduces the neural network experiment specified in PyNN. The most relevant constraints are sketched in the following:

Each hardware neuron circuit has a limited number of 224 incoming synapses. By interconnecting several neuron circuits one can form "larger" neurons with more incoming synapses (Sec. 2.1.1), with the trade-off that the overall number of neurons is reduced. Still, each hardware synapse can not be used to implement a connection from an arbitrary neuron but only from a subset of neurons, namely the 64 source neurons whose pulses arrive at the corresponding synapse driver. For networks larger than 10 000 neurons it is the *limited number of inputs* to one HICANN that becomes even more restricting, as there are only 224 synapse drivers, yielding a maximum of 14366 different source neurons for all neurons that are placed to the same HICANN. Hence, one objective of the mapping process is to reduce this number of source neurons per HICANN, thus increasing the number of realized synapses on the hardware. In general, this criterion is met when neurons with common pre-synaptic partners are placed onto the same HICANN and neurons with common targets inject their pulses into the same on-wafer routing bus.

All of the above, as well as the limited number of on-wafer routing resources (Sec. 2.1.2) make the mapping optimization an NP-hard problem. The used placement and routing algorithms, which improve upon the ones described in Brüderle et al (2011) and Fieres et al (2008) but are far from being optimal, can minimize the effect of these constraints only to a certain degree. Thus, depending on the network model size, its connectivity, and the choice of the mapping algorithms, *synapses are*

lost during the mapping process; in other words, some synapses of a network defined in PyNN will be inexistent in the corresponding network emulated on the hardware. For an estimation of the amount of synapse loss, we first scaled all three benchmark models to sizes between 1000 and 100 000 neurons and mapped them onto the hardware using a simple, not optimized placement strategy. The results strongly depend on the size and the connectivity structure of the emulated network. In order to allow a comprehensive discussion within this study, we then used various placement strategies, sometimes optimizing the mapping by hand to minimize the synapse loss, or purposely using a wasteful allocation of resources to generate synapse loss.

2.3 Executable system specification (ESS)

The ESS is a detailed simulation of the hardware platform (Ehrlich et al, 2007; Brüderle et al, 2011) that replicates the topology and dynamics of the communication infrastructure as well as the analog synaptic and neuronal components.

The simulation encompasses a numerical solution of the equations that govern the hardware neuron and synapse dynamics (Eq. 1 to 5 and 19 to 21) and a detailed reproduction of the digital communication infrastructure at the level of individual spike transmission in logical hardware modules. The ESS is a *specification* of the hardware in the sense that its configuration space faithfully maps the possible interconnection topologies, parameter limits, parameter discretization and shared parameters. Being executable, the ESS also covers dynamic constraints, such as the consecutive processing of spikes which can lead to spike time jitter or spike loss. Variations in the analog circuits due to production variations are not simulated at transistor level but are rather artificially imposed on synaptic weights. All of this allows to simultaneously capture the complex dynamic behavior of the hardware and comply with local bandwidth limitations, while allowing relatively quick simulations due to the high level of abstraction. Simulations on the ESS can be controlled using PyNN (Sec. 2.2), similarly to any other PyNN-compatible back-end. Both for the real hardware and for the ESS, the mapping process translates a PyNN network into a device configuration, which is then used as an input for the respective back-end. One particular advantage of the ESS is that it allows access to state variables which can otherwise not be read out from the real hardware, such as the logging of lost or jittered time events.

2.4 Investigated distortions

Reviewing the hardware and software components of the BrainScaleS wafer-scale system (Sec. 2.1 and 2.2) leaves us with a number of mechanisms that can affect or impede the emulation of neural network models:

- neuron and synapse models are cast into silicon and can not be altered after chip production
- limited ranges for neuron and synapse parameters
- discretized and shared parameters
- limited number of neurons and synapses
- restricted connectivity
- synapse loss due to non-optimal algorithms for NP-hard mapping
- fixed-pattern noise due to transistor level mismatch
- non-configurable pulse delays and jitter
- limited bandwidth for stimulation and recording of spikes

It is clear that, for all of the above distortion mechanisms, it is possible to find a corner case where network dynamics are influenced strongly. However, a few of these effects stand out: on one hand, they are of such fundamental nature to mixed-signal VLSI that they are likely to play some role in any similar neuromorphic device; on the other hand, they are expected to influence any kind of emulated network to some extent. We have therefore directed our focus towards these particular effects, which we summarize in the following. In order to allow general assessments, we investigate various magnitudes of these effects, also beyond the values we expect for our particular hardware implementation.

Neuron and synapse models While some network architectures employ relatively simple neuron and synapse models for analytical and/or numerical tractability, others rely on more complex components in order to remain more faithful to their biological archetypes. Such models may not allow a straightforward translation to those available on the hardware, requiring a certain amount of fitting. In our particular case, we search for parameters to Eq. 1 to 5 and 19 to 21 that best reproduce reproduce low-level dynamics (e.g. membrane potential traces for simple stimulus patterns) and then tweak these as to optimally reproduce high-level network behaviors. Additionally, further constraints are imposed by the parameter ranges permitted by the hardware (Tab. 5).

Synapse loss Above a certain network size or density, the mapping process may not be able to find enough hardware resources to realize every single synapse. In a first stage, we model synapse loss as homogeneous, i.e.,

each synapse is deleted with a fixed probability between 0 and 50 %. To ease the analysis of distortions, we make an exception for synapses that mediate external input, since, in principle, they can be prioritized in the mapping process such that the probability of losing them practically vanishes. Ultimately however, the compensation methods designed for homogeneous synapse loss are validated against a concrete mapping scenario.

Non-configurable axonal delays Axonal delays on the wafer are not configurable and depend predominantly on the processing speed of digital spikes within one HICANN, but also on the physical distance of the neurons on the wafer. In all simulations, we assume a constant delay of 1.5 ms for all synaptic connections in the network, which represents an average of the expected delays when running the hardware with a speedup of 10^4 with respect to real time.

Synaptic weight noise Calibration methods exist that allow a strong reduction of process variations affecting most neuron and synapse parameters. For individual synaptic weights, however, such a calibration is not applicable due to high costs and the limited 4-bit resolution. Therefore, synaptic weights are strongly affected by fixed pattern noise. In both software and ESS simulations, we model this effect by drawing synaptic weights from a Gaussian centered on the target synaptic weight with a variance-to-mean-ratio between 0 and 50 %. Occasionally, this leads to excitatory synapses becoming inhibitory and vice versa, which can not happen on the hardware. Such weights are clipped to zero[7]. For ESS simulations we assume a synaptic weight noise of 20 %, as test measurements on the hardware indicate that the fixed-pattern noise can not be reduced to below this level.

3 Effects and compensation of distortions

In the following, we analyze the effects of hardware-induced distortions on a set of neuronal network models and propose adequate compensation mechanisms for restoring the original network dynamics. The aim of these studies is twofold. On one hand, we propose a generic workflow which can be applied for different neural network models regardless of the neuromorphic substrate, assuming it posesses a certain degree of configurability (Fig. 3). On the other hand, we seek to characterize the universality of the BrainScaleS neuromorphic

[7] This clipping effectively leads to an increase of the mean of the distribution. For 50 % noise the mean is increased by 0.425 %, so that this effect can be neglected.

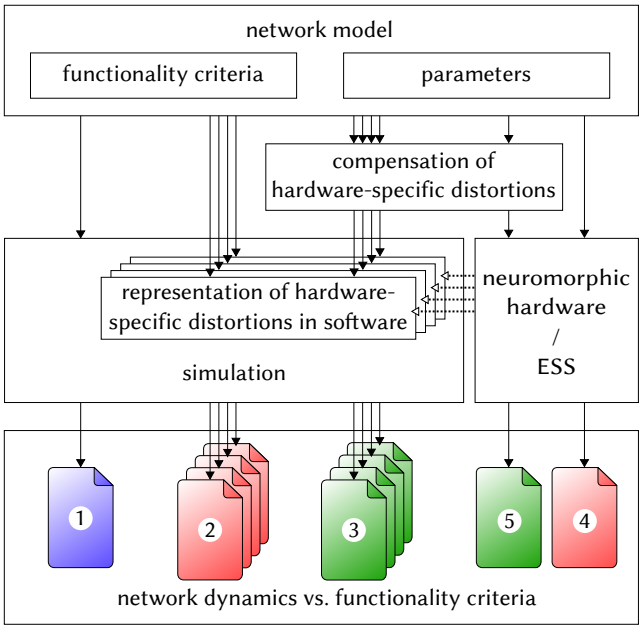

Fig. 3: Schematic of the workflow we have used for studying and compensating hardware-induced distortions of network dynamics. (1) A given network model is defined by providing suitable parameters (for its connectivity and components) and well-defined functionality criteria. (2) The distortions that are expected to occur natively on the hardware back-end are implemented and studied individually in software simulations. (3) Compensation methods are designed and tested, with the aim of recovering the original network dynamics as determined by the functionality criteria. (4) The network model is run on the hardware (here: the ESS) without any compensation to evaluate the full effect of the combined distortion mechanisms. (5) The compensation methods are combined and applied to the hardware (here: the ESS) simulation in order to restore the original network dynamics.

device by assessing its capability of emulating very different large-scale network models with minimal, if any, impairment to their functionality.

In order to allow a comprehensive overview, the set of benchmark experiments is required to cover a broad range of possible network architectures, parameters and function modi. To this end, we have chosen three very different network models, each of which highlights crucial aspects of the biology-to-hardware mapping procedure and poses unique challenges for the hardware implementation. In order to facilitate the comparison between simulations of the original model and their hardware implementation, all experimental setups were implemented in PyNN, running the same set of instructions on either simulation back-end.

For each of our benchmark models we define a set of specific well-quantifiable functionality measures. Assuming a broad range of hardware-induced distortions to various network parameters, their effects on these

measures are investigated in software simulations[8] and various changes to the model structure are proposed in order to recover the original functionality. The feasibility of these compensation methods is then studied for the BrainScaleS neuromorphic platform with the help of the ESS described in Sec. 2.3.

3.1 Cortical layer 2/3 attractor memory

As our first benchmark, we have chosen an attractor network model of the cerebral cortex which exihbits characterisic and well-quantifiable dynamics, both at the single-cell level (membrane voltage UP and DOWN states) and for entire populations (gamma band oscillations, pattern completion, attentional blink). For this model, the mapping to the hardware was particularly challenging, due to the complex neuron and synapse models required by the original architecture on the one hand, as well as its dense connectivity on the other. In particular, we observed that the shape of synaptic conductances strongly affects the duration of the attractor states. As expected for a model with relatively large populations as functional units, it exhibits a pronounced robustness to synaptic weight noise. Homogeneous synapse loss, on the other hand, has a direct impact on single-cell dynamics, resulting in significant deviations from the expected high-level functionality, such as the attenuation of attentional blink. As a compensation for synapse loss, we suggest two methods: increasing the weights of the remaining synapses in order to maintain the total average synaptic conductance and reducing the size of certain populations and thereby decreasing the total number of required synapses. After mapping to the hardware substrate, synapse loss is not homogeneous, due to the different connectivity patterns of the three neuron types required by the model. However, we were able to apply a population-wise version of the suggested compensation methods and demonstrate their effectiveness in recovering the previously defined target functionality measures.

3.1.1 Architecture

As described in Lundqvist et al (2006) and Lundqvist et al (2010), this model (henceforth called L2/3 model) implements a columnar architecture (Mountcastle, 1997; Buxhoeveden and Casanova, 2002). The connectivity is compliant with data from cat cerebral cortex connectivity (Thomson et al, 2002). The key aspect of the model is its modularity, which manifests itself on two levels. On a large scale, the simulated cortical patch is represented by a number N_{HC} of hypercolumns (HCs) with a diameter of 500 µm each, arranged on a hexagonal topology. On a smaller scale, each HC is further subdivided into a number N_{MC} of minicolumns (MCs), with various estimates suggesting a number of about 100 MCs bundled into a single HC (Mountcastle, 1997; Buxhoeveden and Casanova, 2002). However, such MCs should first and foremost be seen as functional units, and could, in biology, also be a group of distributed, but highly interconnected cells (Song et al, 2005; Kampa et al, 2006; Perin et al, 2011). In the model, each MC consists, in turn, of 30 pyramidal (PYR), 2 regular spiking non-pyramidal (RSNP) and 1 basket (BAS) cells (Peters and Sethares, 1997; Markram et al, 2004). Within each MC, PYR neurons are mutually interconnected, with 25% connectivity, such that they will tend to be co-active and code for similar input.

The functional units of the network, the MCs, are connected in global, distributed patterns containing a set of MCs in the network (Fig. 4). These global connections, carried out by PYR cells, are thought to arise from learning, and form dynamical attractors in the network that can be re-activated by partial or corrupted stimulation (Lundqvist et al, 2006). This notion that pre-wired functional units with stereotyped connectivity connect on a more global scale during learning has recently gained experimental support (Perin et al, 2011). Here the attractors, or patterns, were formed before the onset of simulation and contained one MC from each HC. We have only considered the case of orthogonal patterns, which implies that no two attractors share any number of MCs. Due to the mutual excitation within an attractor, the network is able to perform pattern completion, which means that whenever a subset of MCs in an attractor is activated, the activity tends to spread throughout the entire attractor.

Pattern rivalry results from competition between attractors mediated by short and long-range connections via inhibitory interneurons. Each HC can be viewed as a soft winner-take-all (WTA) module which normalizes activity among its constituent MCs (Lundqvist et al, 2010). This is achieved by the inhibitory BAS cells, which receive input from the PYR cells from the 8 closest MCs and project back onto the PYR cells in all the MCs within the home HC. Apart from providing long-range connections to PYR cells within the same pattern, the PYR cells within an MC project onto RSNP cells in all the MCs which do not belong to the same pattern and do not lie within the same HC. The inhibitory RSNP cells, in turn, project onto the PYR cells in their respective MC. The effect of this connectivity is a disynaptic inhibition between competing patterns.

[8] All simulations were performed with NEST (Diesmann and Gewaltig, 2002) or Neuron (Hines and Carnevale, 2003).

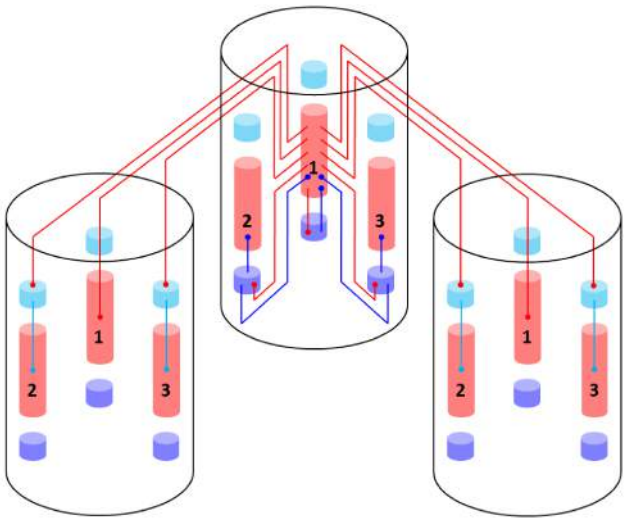

Fig. 4: Pseudo-3D schematic of the layer 2/3 model architecture. Excitatory (PYR) cell populations are represented as red cylinders, inhibitory populations as blue ones (BAS: dark, RSNP: light). A minicolumn (MC) consists of three vertically aligned populations: one PYR, one BAS and one RSNP. Multiple MCs are grouped into hypercolumns (HCs, transparent cylinders). MCs with the same ID (one per HC) form a so-called pattern or attractor. When active, all PYR cells belonging to an attractor excite each other via short-range (within an MC) and long-range (between MCs) connections. The inhibition of PYR cells belonging to other attractors occurs via inhibitory interneurons: locally (within an HC) through BAS cells and globally (between HCs) through RSNP cells. Only a subset of connections are shown, namely those which are mainly used during active periods of attractor 1.

Fig. 4 shows a schematic of the default architecture, emphasizing the connectivity pattern described above. Since a network with the size of the cat striate cortex would require the simulation of some 20 million neurons and 100 billion synapses, we have used a subsampled model as a default architecture. It consists of $N_{\mathrm{HC}} = 9$ HCs, each containing $N_{\mathrm{MC}} = 9$ MCs, yielding a total of 2673 neurons. Due to its modular structure, this default model can easily be scaled up or down in size with preserved dynamics, as described in the Appendix (Sec. B.4).

When a pattern receives enough excitation, its PYR cells enter a state reminiscent of a so-called local UP-state (Cossart et al, 2003), which is characterized by a high average membrane potential, several mV above its rest value, and elevated firing rates. Pattern rivalry leads to states where only one attractor may be active (with all its PYR cells in an UP-state) at any given time. Inter-PYR synapses feature an STP mechanism which weakens the mutual activation of PYR cells over time and prevents a single attractor from becoming persistently active. Additionally, PYR neurons exhibit spike-frequency adaptation, which also suppresses prolonged firing. These mechanisms impose a finite lifetime on the attractors such that after their termination more weakly stimulated or less excitable attractors can become active, in contrast to what happens in classical WTA networks.

The inputs to the layer 2/3 PYR cells arrive from the cortical layer 4, which is represented by 5 cells per MC. The layer 4 cells project onto the layer 2/3 PYR cells and can be selectively activated by external Poisson spike trains. Additionally, the network receives unspecific input representing activity in various connected cortical areas outside the modeled patch. This input is modeled as diffuse noise and generates a background activity of several Hz. Depending on the level of this diffuse noise, attractors could either become spontaneously active (high noise) or require stimulation (low noise) to activate. In the first case the network operates in a bistable regime where all miniciolumns are active at a very low rate as default (Djurfeldt et al, 2008; Lundqvist et al, 2010).

More details on the model architecture, as well as neuron and synapse parameters, can be found in the Appendix (Sec. B.1).

3.1.2 Functionality criteria

Fig. 5 shows some characteristic dynamics of the L2/3 model, which have also been chosen as functionality criteria and are described below.

The core functionality of the original model is easily identifiable by its distinctive display of spontaneously activating attractors in e.g. raster plots (**A**) or voltage star plots (**D**, for an explanation of star plots see Sec. B.8). However, in particular for large network sizes, spontaneous attractors become increasingly sparse. Additionally, many further indicators of functionality can be found, such as the average membrane potential or the gamma oscillations observed in UP states. Finally, when receiving L4 stimulation in addition to the background noise, the original model displays important features such as pattern completion and attentional blink, which need to be reproducible on the hardware as well. Consequently, we consider several measures of functionality throughout our analyses.

When an attractor becomes active, it remains that way for a characteristic dwell time τ_{on}. The dwell time depends strongly on the neuron and synapse parameters (as will be discussed in the following sections) and only weakly on the network size (**C**, **F**), since the scaling rules ensure a constant average fan-in for each neuron type. Conversely, this makes τ_{on} sensitive to hardware-induced variations in the average synaptic input. The detection of active attractors is performed automati-

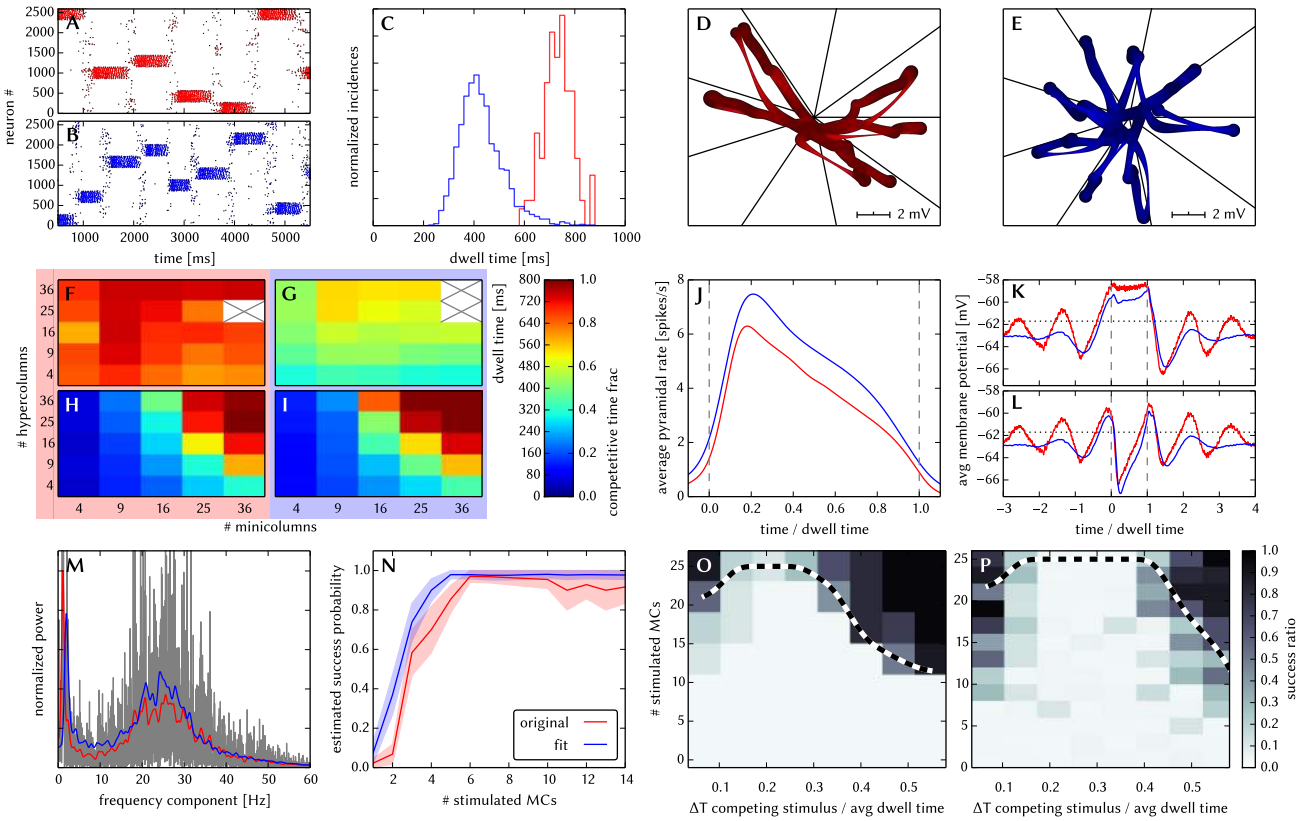

Fig. 5: Comparison between original and adapted L2/3 network models. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. Measurements from the original model are depicted (or highlighted) in red, while those from the adapted model are depicted (or highlighted) in blue. (**A**, **B**): Raster plots of spiking activity. Attractors activate spontaneously only due to diffuse background noise. Only PYR cells are shown. The MCs are ordered such that those belonging to the same attractor (and *not* those within the same HC) are grouped together. **C**: Attractor dwell time distributions. The shorter average dwell times in the adapted model are caused by sharper PSPs which miss the long NMDA time constants. (**D**, **E**): Star plots of average PYR cell voltages from a sample of 5 PYR cells per MC. Details on this representation of multidimensional data can be found in Sec. B.8. (**F**, **G**): Average dwell time for various network sizes. (**H**, **I**): Fraction of time spent in competitive states (i.e. no active attractors) for various network sizes. While dwell times remain relatively constant, competition times increase with network size, suppressing spontaneous attractors in large networks. (**J**): Average firing rate of PYR within an active period of their parent attractor. (**K**): Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**L**): Average voltage of PYR cells before, during and after an attractor they do not belong to is active (DOWN state). For subplots **J**, **K** and **L**, the abscissa has been subdivided into multiples of the average attractor dwell time in the respective simulations. The oscillations of the average voltages occur due to spike-triggered adaptation: after an active period, PYR cells need to recover before being able to support an active period of their home attractor, during which time they are inhibited by other active attractors. The more pronounced attenuation of the oscillations in the adapted model happens due to a higher relative variability of dwell times (compare subplot **C**). In subplots **K** and **L** the dotted line indicates the leak potential E_L of the PYR cells. (**M**): Smoothed power spectrum of PYR firing rate averaged over all MCs. The grey curve in the background represents the unsmoothed spectral density for the original model. Attractor switches ($\approx$ 2 Hz) and gamma oscillations ($\approx$ 25 Hz) can be clearly observed. (**N**): Pattern completion in a 25HC×25MC network. Estimated probability of an attractor to fully activate (success ratio) as a function of the number of stimulated constituent MCs, measured over 25 trials per abscissa value. (**O**, **P**): Attentional blink in a 25HC×25MC network. Two attractors are stimulated (the second one only partially, i.e. a certain number of constituent MCs) with a temporal lag of ΔT in between. Activation probability of the second attractor and $p = 0.5$ iso-probability contours, measured over 14 trials per (ΔT, #MCs) pair. A detailed description of the data and methods used for all figures concerning the L2/3 model can be found in Sec. B.1 to B.7.

cally using the spike data (for a description of the algorithm, see Sec. B.5).

We describe the periods between active attractors as competition phases and the time spent therein as the total competition time. The competition time varies strongly depending on the network size (**H**). One can observe that the competition time is a monotonically increasing function of both N_{HC} and N_{MC}. For an increasing number of HCs, i.e., a larger number of neurons in every pattern, the probability of a spontaneous activation of a sufficiently large number of PYR cells decreases. For an increasing number of MCs per HC, there is a larger number of competing patterns, leading to a reduced probability of any single pattern becoming dominant.

When an attractor becomes active, the average spike rate of its constituent PYR cells rises sharply and then decays slowly until the attractor becomes inactive again (**J**). Two independent mechanisms are the cause of this decay: neuron adaptation and synaptic depression. The characteristic time course of the spike rate depends only weakly on the size of the network.

As described in Sec. 3.1.1, PYR cells within active attractors enter a so-called local UP state, with an increased average membrane potential and an elevated firing rate (**K**). While inactive or inhibited by other active attractors, PYR cells are in a DOWN state, with low average membrane potential and almost no spiking at all (**L**). In addition to these characteristic states, the average PYR membrane potential exhibits oscillations with a period close to τ_{on}. These occur because the activation probability of individual attractors is an oscillatory function of time as well. In the immediate temporal vicinity of an active period (i.e., assuming an activation at $t = 0$, during $[-\tau_{\mathrm{on}}, 0) \cup [\tau_{\mathrm{on}}, 2\tau_{\mathrm{on}})$) the same attractor must have been inactive, since PYR populations belonging to an activated attractor need time to recuperate from synaptic depression and spike-triggered adaptation before being able to activate again.

An essential emerging feature of this model are oscillations of the instantaneous PYR spike rate in the gamma band within active attractors (**M**). The frequency of these oscillations are independent of size and rather depend on excitation levels in the network (Lundqvist et al, 2010). Although the gamma oscillations might suggest periodic spiking, it is important to note that individual PYR cells spike irregularly ($\langle \mathrm{CV_{ISI}} \rangle = 1.36 \pm 0.36$ within active attractors).

Apart from these statistical measures, two behavioral properties are essential for defining the functionality of the network: the pattern completion and attentional blink mentioned above. The pattern completion ability of the network can be described as the successful activation probability of individual patterns as a function of the number of stimulated MCs (**N**). Similarly, the attentional blink phenomenon can also be quantified by the successful activation rate of an attractor as a function of the number of stimulated MCs if it is preceded by the activation of some other attractor with a time lag of ΔT (**O**). For small ΔT, the second attractor is completely "blinked out", i.e., it can not be activated regardless of the number of stimulated MCs. To facilitate the comparison between different realizations of the network with respect to attentional blink, we consider the 50% iso-line, which represents the locus of the input variable pair which leads to an attractor activation ratio of 50%.

A detailed description of the data and methods used for these figures can be found in the Appendix (Sec. B.1 to Sec. B.7).

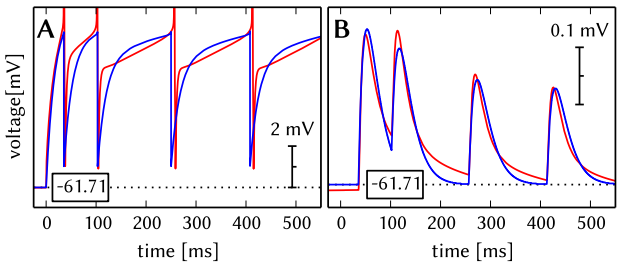

Fig. 6: Comparison of original neuron (multi-compartment HH) and synapse (NMDA+AMPA) dynamics (red) to the fitted dynamics of hardware-compatible neuron (point AdEx) and synapse (single decay time constant) models (blue). (**A**) Membrane potential of PYR cells under spike-inducing current stimulation. While the precise membrane potential time course of the original neuron model can not be reproduced by a single-compartment AdEx neuron, spike timing and especially firing rates can be recovered. (**B**) PSPs generated by PYR→PYR synapses between MCs where the spikes from **A** were used as input. As a replacement for the multiple synaptic time constants in the original model, we have chosen an intermediate value for τ^{syn}, which constitutes the main reason for the difference in PSP shapes. Additionally, the combination of STP and saturation in the original model had to be replaced by STP alone.

3.1.3 Neuron and synapse model translation

A particular feature of this benchmark model is the complexity of both neuron and synapse models used in its original version. Therefore, the first required type of compensation concerns the parameter fitting for the models implemented on the hardware. Some exemplary results of this parameter fit can be seen in Fig. 6. More details can be found in the Appendix (Sec. B.2).

Neurons There are several differences between the Hodgkin-Huxley-type neurons of the original model and the

AdEx neurons on the waferscale hardware. The most evident difference concerns the modeling of spikes. However, in the subthreshold regime, the membrane voltage is determined predominantly by the synaptic input current (leaky integrator dynamics) and is therefore independent of the spiking mechanism, rendering the voltage time courses for the two models very similar under the assumption of identical synaptic currents. When the membrane potential becomes high enough to trigger a spike, the typical membrane potential time course of a Hodgkin-Huxley neuron can be, in principle, well approximated by the exponential term in the AdEx equation (Brette and Gerstner, 2005). However, when fitting for spike timing rather than membrane potential, we found that spike times were best reproduced when eliminating the exponential term, i.e. setting $\Delta_T = 0$.

Adaptation is an essential feature of both the PYR and the RSNP cells in the original model, where it is caused by an accumulation of calcium entering the cell and causing the activation of K_{Ca} channels. Since the voltage-gated Ca channels only open at very high membrane voltages, which only occur during an action potential, this mechanism is functionally similar to spike-triggered adaptation, which in the AdEx model is given by the spike-triggered quantal increase b of the adaptation variable w. For RSNP cells, in order to reproduce the correct equilibrium spike frequency, setting the adaptation parameter a to a nonzero value was also required.

One further difference resides in the original neurons being modeled as having several compartments, whereas the hardware only implements point neurons. The passive neuron properties (membrane capacitances and leak conductances) were therefore determined by fitting the membrane potential time course under stimulation by a step current which was not strong enough to elicit spikes.

Synapses We have performed an initial estimation of synaptic weights and time constants by fitting the membrane potential time course of the corresponding neurons in a subthreshold regime. However, two important differences remain between the synapses in the original model and the ones available on our hardware.

In the original model, PYR-PYR and PYR-RSNP synapses contain two types of neurotransmitters: Kainate/AMPA and NMDA (see Tab. 8). Due to the vastly different time constants for neurotransmitter removal at the postsynaptic site (6 ms and 150 ms, respectively), the PSPs have a characteristic shape, with a pronounced peak and a long tail. This shape can not be precisely reproduced with only a single decay time constant.

One further difference lies in the saturating nature of the postsynaptic receptor pools after a single presynaptic spike. In principle, this behavior could be emulated by the TSO plasticity mechanism by setting $U = 1$ and $\tau_{rec} = \tau^{syn}$. However, this would conflict with the TSO parameters required for modeling short-term depression of PYR synapses and would also require parameters outside the available hardware ranges.

For these reasons, we have further modified synaptic weights and time constants by performing a behavioral fit, i.e., by optimizing these parameters towards reproducing the correct firing rates of the three neuron types in two scenarios - first without and then subsequently with inhibitory synapses. Because the original model was characterized by relatively long and stable attractors, we further optimized the excitatory synapse time constants towards this behavior.

Post-fit model behavior Fig. 5 shows the results of the translation of the original model to hardware-compatible dynamics and parameter ranges. Overall, one can observe a very good qualitative agreeement of characteristic dynamics with the original model. In the following, we discuss this in more detail and explain the sources of quantitative deviations.

When subject to diffuse background noise only, the default size network clearly exhibits its characteristic spontanous attractors (**B**). Star plots exhibit the same overall traits, with well-defined attractors, characterized by state space trajectories situated close to the axes and low trajectory velocities within attractors (**E**). Attractor dwell times remain relatively stable for different network sizes, while the competition times increase along with the network size (**G** and **I**). The average value of dwell times, however, lies significantly lower than in the original (**C**). The reason for this lies mainly in the shape of EPSPs: the long EPSP tails enabled by the large NMDA time constants in the original model caused a higher average membrane potential, thereby prolonging the activity of PYR cells.

Within attractors, active and inactive PYR cells enter well-defined local UP and DOWN states, respectively (**K** and **L**). Before and after active attractors, the dampened oscillations described in Sec. 3.1.2 can be observed. In the adapted model, attenuation is stronger due to a higher coefficient of variation of the dwell times ($\frac{\sigma}{\mu} = 0.20$ as compared to 0.08 in the original model).

Average PYR firing rates within active attractors have very similar time courses (**J**), with a small difference in amplitude, which can be attributed to the difference in EPSP shapes discussed earlier. Both low-frequency switches between attractors (< 3 Hz, equivalent to the incidence rate) and high-frequency gamma

oscillations arising from synchronous PYR firing (with a peak around 25 Hz) can be clearly seen in a power spectrum of the PYR firing rate (**M**).

Pattern completion occurs similarly early, with a steep rise and nearly 100% success rate starting at 25% of stimulated MCs per attractor (**N**). Attentional blink follows the same qualitative pattern (**P**, **Q**), although with a slightly more pronounced dominance of the first activated attractor in the case of the adapted network, which happens due to the slightly higher firing rates discussed above.

Having established the quality of the model fit and in order to facilitate a meaningful comparison, all following studies concerning hardware-induced distortions and compensation thereof use data from the adapted model as reference.

3.1.4 Synapse loss

The effects of homogeneous synapse loss and the results of the attempted compensation are depicted in Fig. 7. More detailed plots can be found in the Appendix (Fig. 25).

Effects With increasing synapse loss, the functionality of the network gradually deteriorates. Attractors become shorter or disappear entirely, with longer periods of competition in between (**D**, **K**, **O**).

While average excitatory conductances are only affected linearly by synaptic loss, inhibitory conductances feel a compound effect of synapse loss, as it affects both afferent and efferent connections of inhibitory interneurons. Therefore, synapse loss has a stronger effect on inhibition, leading to a net increase in the average PYR membrane potential (**R**, **S**). Additionally, since all connections become weaker, the variance of the membrane potential becomes smaller, as observed in the corresponding star plots as well (**E**). The weaker connections also decrease the self-excitation of active attractors while decreasing the inhibition of inactive ones, thereby leading to shorter attractor dwell times (**P**). Somewhat surprisingly, the maximum average PYR firing rate in active attractors remains almost unchanged when subjected to synapse loss. However, the temporal evolution of the PYR firing rate changes significantly (**Q**).

The pattern completion ability of the network suffers particularly in the region of weak stimuli, due to weaker internal excitation of individual attractors. The probability of triggering a partially stimulated pattern can drop by more than 50% (**T**). Due to the decreased stability of individual attractors discussed above, rival attractors are easier to excite, thereby significantly suppressing the attentional blink phenomenon (**U**).

Compensation As a first-order approximation, we can consider the population average of the neuron conductance as the determining factor in the model dynamics. For synapses with exponential conductance courses, the average conductance generated by the ith synapse is proportional to both synaptic weight w_{ij} and afferent firing rate ν_j. Because conductances sum up linearly, the total conductance that a neuron from population i receives from some other population j is, on average (see Eq. 30)

$$\langle g \rangle = N_j p_{ij} \langle w_{ij} \rangle \langle \nu_j \rangle \tau^{\text{syn}} \quad , \tag{9}$$

where N_j represents the size of the presynaptic population and p_{ij} represents the probability of a neuron from the presynaptic population to project onto a neuron from the postsynaptic population. Since homogeneous synapse loss is equivalent to a decrease in p_{ij}, we can compensate for synapse loss that occurs with probability p_{loss} by increasing the weights of the remaining synapses by a factor $1/(1 - p_{\text{loss}})$. Fig. 7 shows the results of this compensation strategy for $p_{\text{loss}} = 0.5$. In all aspects, a clear improvement can be observed. The remaining deviations can be mainly attributed to two effects. First of all, preserving the average conductance by compensating homogeneous synapse loss with increasing synaptic weights leads to an increase in the variance of the membrane potential (Eq. 29). Secondly, finite population sizes coupled with random elimination of synapses lead to locally inhomogeneous synapse loss and further increase the variability of neuronal activity.

Instead of compensating for synapse loss after its occurrence, it is also possible to circumvent it altogether after having estimated the expected synapse loss in a preliminary mapping run. For the L2/3 model, this can be done without altering the number of functional units (i.e., the number of HCs and MCs) by changing the size of the PYR cell populations. For this approach, however, the standard scaling rules (Sec. B.4) need to be modified. These rules are designed to keep the average number of inputs per neuron constant and would increase the total number of PYR-incident synapses by the same factor by which the PYR population is scaled. This would inevitably lead to an increased number of shared inputs per PYR cell, with the immediate consequence of increased firing synchrony. Instead, when reducing the PYR population size, we compensate for the reduced number of presynaptic partners by increasing relevant synaptic weights instead of connection probabilities. This modified downscaling leads to a net reduction of the total number of synapses in the network, thereby potentially reducing synaptic loss between all populations. Fig. 7 shows the effects of scaling down the PYR population size until the total remaining num-

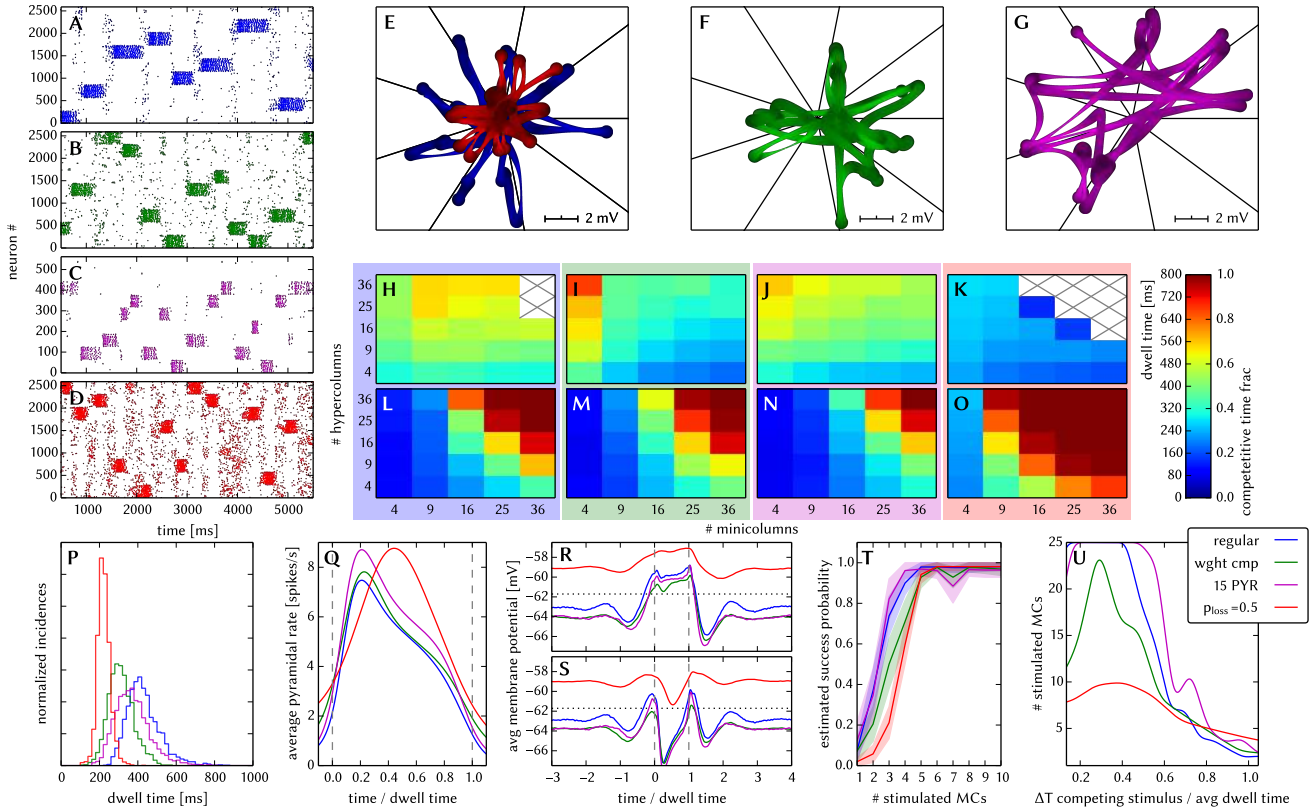

Fig. 7: Compensation of homogeneous synaptic loss in the L2/3 model. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. Here, we use the following color code: blue for the original model, red for the distorted case (50% synapse loss), green for the compensation via increased synaptic weights and purple for the compensation by scaling down the size of the PYR populations. (**A** - **D**) Raster plots of spiking activity. The MCs are ordered such that those belonging to the same attractor (and *not* those within the same HC) are grouped together. Synapse loss weakens the interactions within and among MCs, causing shorter dwell times and longer competition times. Both compensation methods successfuly counter these effects. These phenomena can also be observed in subplots **H**-**P**. (**E** - **G**) Star plots of average PYR voltages from a sample of 5 PYR cells per MC. Synapse loss leads to a less pronounced difference between the average PYR membrane potential within and outside of active attractors. After compensation, the differences between UP and DOWN states become more pronounced again. These phenomena can also be observed in subplots **R** and **S**. (**H** - **K**) Average dwell time for various network sizes. (**L** - **O**) Fraction of time spent in competitive states (i.e. no active attractors) for various network sizes. (**P**) Distributions of dwell times. (**Q**) Average firing rate of PYR cells within an active period of their parent attractor. (**R**) Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**S**) Average voltage of PYR cells before, during and after an attractor they do not belong to is active (DOWN state). For subplots **Q**, **R** and **S**, the abscissa has been subdivided into multiples of the average attractor dwell time in the respective simulations. In subplots **R** and **S** the dotted line indicates the leak potential E_L of the PYR cells. (**T**) Pattern completion in a 25HC×25MC network. Estimated activation probability from 25 trials per abscissa value. Synapse loss shifts the curve to the right, i.e., more MCs need to be stimulated to achieve the same probability of activating their parent attractor. Both compensation methods restore the original behavior to a large extent. (**U**) Attentional blink in a 25HC×25MC network: $p = 0.5$ iso-probability contours, measured over 14 trials per (ΔT, #MCs) pair. Synapse loss suppresses attentional blink, as inhibition from active attractors becomes to weak to prevent the activation of other stimulated attractors. Compensation by increasing the weight of the remaining synapses alleviates this effect, but scaling down the PYR population sizes directly reduces the percentage of lost synapses and is therefore more effective in restoring attentional blink.

ber of synapses is equal to the realized number of synapses in the distorted case (50% of the total number of synapses in the undistorted network). More detailed plots of the effects of PYR population downscaling can be found in Fig. 26. The two presented compensation methods can also be combined to further improve the final result, as we show in Sec. 3.1.7.

3.1.5 Synaptic weight noise

One would not expect the synaptic weight noise to affect the L2/3 model strongly, as it should average out over a large number of connections between the constitutent populations. It turns out that the surprisingly strong impact of synaptic weight noise is purely due

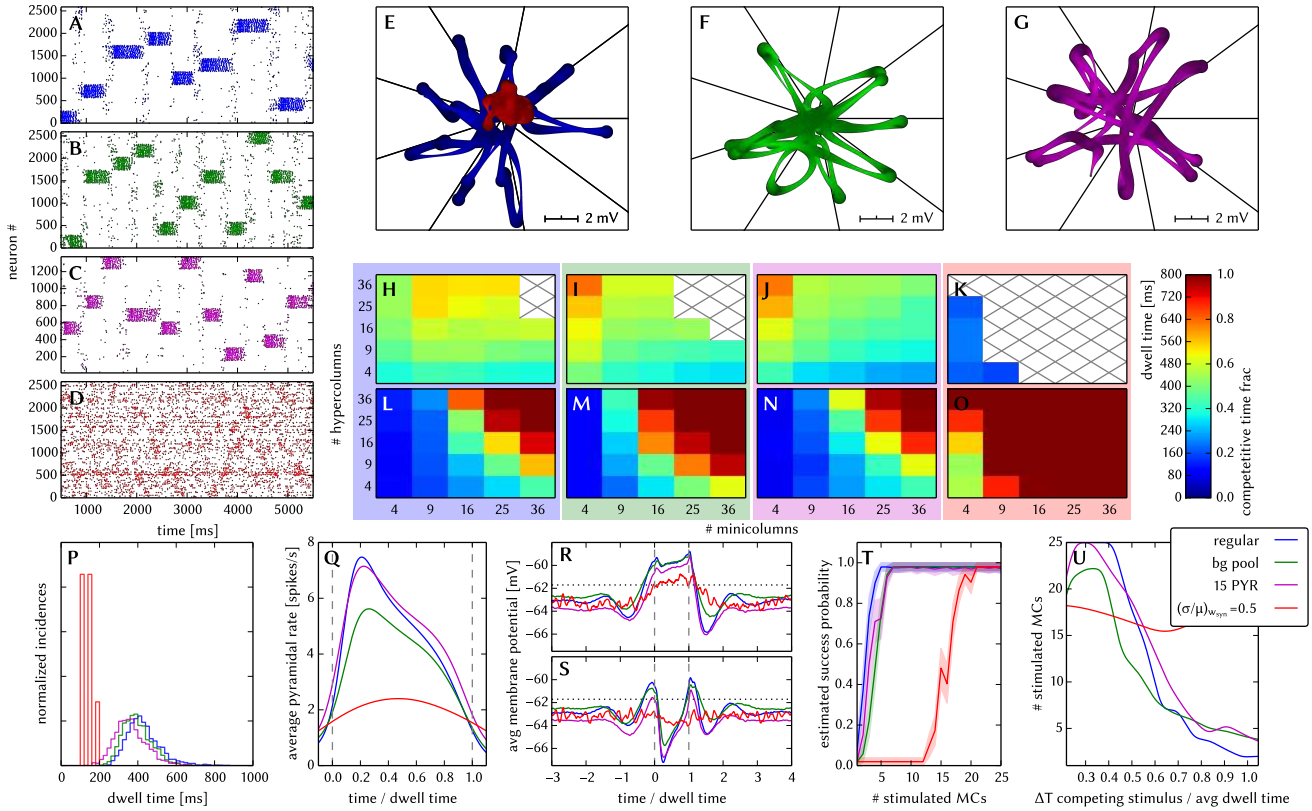

Fig. 8: Compensation of synaptic weight noise in the L2/3 model. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. Here, we use the following color code: blue for the original model, red for the distorted case (50% synaptic weight noise), green for the compensation via multiple background sources per PYR cell and purple for the same compensation method, but with scaled down PYR populations. Altogether, we note that the observed effects happen almost exclusively due to each PYR cell receiving background input via a single synapse. When compensated via the inclusion of multiple background sources, the network exhibits remarkable robustness towards synaptic weight noise. (**A** - **D**) Raster plots of spiking activity. The MCs are ordered such that those belonging to the same attractor (and *not* those within the same HC) are grouped together. When each PYR cell has a single background source, high levels of synaptic weight noise cause some PYR cells to become completely silent, while others spike disproportionately often. This can completely disrupt the stability of attractors, resulting in largely random spiking, with long competition times between the occasional appearance of weak, unstable attractors. The inclusion of multiple background sources per PYR cell efficiently counters these effects. This compensation strategy works just as well for downscaled PYR populations. The phenomena described above can also be observed in subplots **H**-**P**. (**E** - **G**) Star plots of average PYR voltages from a sample of 5 PYR cells per MC. The disrupted attractor behavior and erratic PYR spiking result in weak fluctuations of average PYR voltages with essentially no clear UP or DOWN states. After compensation, the differences between UP and DOWN states become more pronounced again. These phenomena can also be observed in subplots **R** and **S**. (**H** - **K**) Average dwell time for various network sizes. (**L** - **O**) Fraction of time spent in competitive states (i.e. no active attractors) for various network sizes. (**P**) Distributions of dwell times. (**Q**) Average firing rate of PYR cells within an active period of their parent attractor. (**R**) Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**S**) Average voltage of PYR cells before, during and after an attractor they do not belong to is active (DOWN state). For subplots **Q**, **R** and **S**, the abscissa has been subdivided into multiples of the average attractor dwell time in the respective simulations. In subplots **R** and **S** the dotted line indicates the leak potential E_L of the PYR cells. (**T**) Pattern completion in a 25HC×25MC network. Estimated activation probability from 25 trials per abscissa value. Due to erratically firing PYR cells in the distorted network, much stronger stimulation is needed to guarantee the appearance of an attractor. Compensation restores the original behavior to a large extent. (**U**) attentional blink in a 25HC×25MC network: $p = 0.5$ iso-probability contours, measured over 14 trials per (ΔT, #MCs) pair. Due to the highly unstable attractors in the distorted network, attentional blink is completely suppressed. Compensation restores blink, but not to its original strength, due to the synaptic weight noise within the network itself.

to the implementation of background stimulus in this model and can therefore be easily countered.

Effects The relative deviation of the total synaptic conductance scales with $\langle g \rangle / \mathrm{Var}\,[g] \sim 1/\sqrt{\nu_{\mathrm{input}}} \sim 1/\sqrt{N}$ (see Eq. 29), where ν_{input} is the total input frequency and N the number of presynaptic neurons. Therefore, interactions between large populations are not expected to be strongly affected by synaptic weight noise.

The only connections where an effect is expected are the RSNP→PYR connections, because the presynaptic RSNP population consists of only 2 neurons per MC. However, long-range inhibition also acts by means of a second-order mechanism, in which an active MC activates its counterpart in some other HC, which then in turn inhibits all other MCs in its home HC via BAS cells. This mechanism masks much of how synaptic weight noise affects RSNP→PYR connections.

Nevertheless, synaptic weight noise appears to have a strong effect on network dynamics (Fig. 8, red curves). The reason for that lies in the way the network is stimulated. In the original model, each PYR cell receives input from a single Poisson source. This is of course a computational simplification and represents diffuse noise arriving from many neurons within other cortical areas. However, having only a single noise source connected by a single synapse to the target neuron makes the network highly sensitive to synaptic weight noise.

Compensation The compensation for this effect was done by increasing the number of independent noise sources per neuron, thereby reducing the statistically expected relative noise conductance variations per PYR cell. The only limitation lies in the total number of available external spike sources and the bandwidth supplied by the off-wafer communication network (Sec. 2.1.2). Once this limit is reached, the number of noise inputs per PYR cell can still be increased even further if PYR cells are allowed to share noise sources. Given a total number of available Poisson sources N and a noise population size of n sources per PYR cell, the average pairwise overlap between two such populations is n^2/N. Therefore, as long as the average overlap remains small enough, the overlap-induced spike correlations will not affect the network dynamics.

In our example (Fig. 8, green curves), we have chosen $n = 100$, while the total number of Poisson sources is set at $N = 5000$. Note how this relatively simple compensation method efficiently restores most functionality criteria. The most significant remaining differences can be seen in pattern completion and attentional blink (**T**, **U**) and appear mainly due to the affected RSNP→PYR connections.

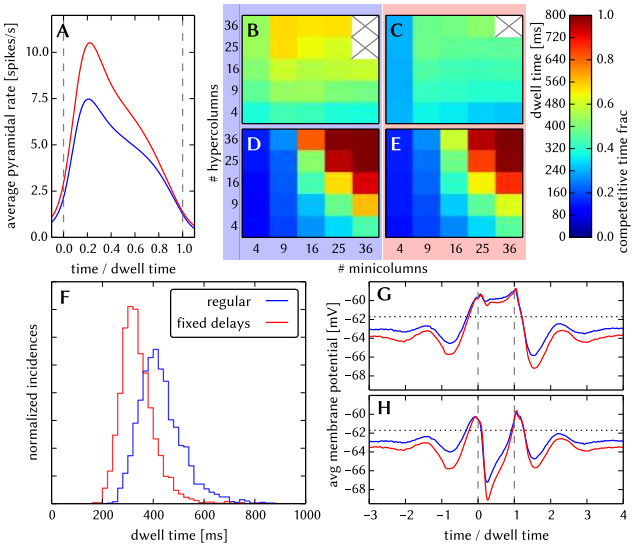

Fig. 9: Effects of fixed axonal delays on the L2/3 model. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. Data from the regular and distorted models is depicted (or highlighted) in blue, and red, respectively. (**A**) Average firing rate of PYR cells within an active period of their parent attractor. (**B**, **C**) Average dwell time for various network sizes. (**D**, **E**) Fraction of time spent in competitive states (i.e. no active attractors) for various network sizes. (**F**) Distributions of dwell times. (**G**) Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**H**) Average voltage of PYR cells before, during and after an attractor they do not belong to is active (DOWN state). For subplots **A**, **G** and **H**, the abscissa has been subdivided into multiples of the average attractor dwell time in the respective simulations. In subplots **G** and **H** the dotted line indicates the leak potential E_{L} of the PYR cells.

In addition to the investigation of synaptic weight noise on the default model, we repeated the same experiments for the model with reduced PYR population sizes (Fig. 8, purple curves), which we have previously suggested as a compensation method for synaptic weight noise (Sec. 3.1.4). The fact that PYR population reduction does not affect the network functionality in the case of (compensated) synaptic weight noise is an early indicator for the compatibility of the suggested compensation methods when all distortion mechanisms are present (Sec. 3.1.7).

3.1.6 Non-configurable axonal delays

In the original model, axonal delays between neurons are proportional to the distance between their home HCs. At an axonal spike propagation velocity of 0.2 m/ms, the default (9HC×9MC) network implements axonal delays distributed between 0.5 and 8 ms. While PYR cells within an MC tend to spike synchronously in gamma waves, the distribution of axonal delays reduces synchronicity between spike volleys of different MCs.

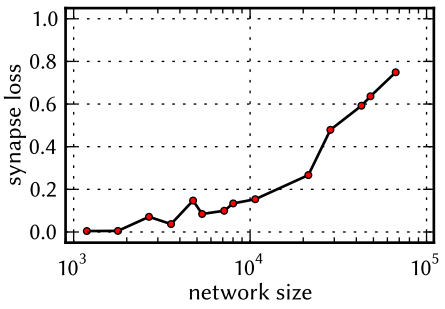

Fig. 10: Synapse Losses after mapping the L2/3 model.

Fixed delays, on the other hand, promote synchronicity, thereby inducing subtle changes to the network dynamics (Fig. 9). The synchronous arrival of excitatory spike volleys causes PYR cells in active attractors to spike more often (**A**). Their higher firing rate in turn causes shorter attractor dwell times, due to their spike frequency adaptation mechanism (**B**, **C**, **F**). During an active attractor, the elevated firing rate of its constituent PYR cells causes a higher firing rate of the inhibitory interneurons belonging to all other attractors. This, in turn, leads to a lower membrane potential for PYR cells during inactive periods of their parent attractor (**G**, **H**). As these effects are not fundamentally disruptive and also difficult to counter without significantly changing other functional characteristics of the network, we chose not to design a compensation strategy for this distortion mechanism in the L2/3 network.

3.1.7 Full simulation of combined distortions

In a final step, we emulate the L2/3 model on the ESS (Sec. 2.3), and compensate simultaneously for all of the effects discussed above. We first investigate how much synapse loss to expect for different network sizes, and then realize the network at two different scales in order to investigate all of the chosen functionality criteria. The default network (9HC×9MC) is used to analyze spontaneous attractors, while a large-scale model (25HC×25MC) serves as the test substrate for pattern completion and pattern rivalry.

Synapse loss The synapse loss after mapping the L2/3 model onto the BrainScaleS hardware is shown in Fig. 10 for different sizes, using the scaling rules defined in Sec. B.4. Synapse loss starts to occur already at small sizes and increases rapidly above network sizes of 20 000 neurons. The jumps can be attributed to the different ratios between number of HCs and number of MCs per HC (Tab. 16).

Table 1: Projection-wise synapse loss in % of the L2/3 model after the mapping process for the default (9HC×9MC) and large-scale (25HC×25MC) network. See text for the respective differences between the distorted (dist) and compensated (dist) networks.

	9HC×9MC		25HC×25MC	
projection	dist	comp	dist	comp
PYR → PYR (local)	21.1	21.0	0.9	0.3
PYR → PYR (global)	20.8	21.2	8.0	0.4
PYR → RSNP	22.6	21.9	37.0	28.8
PYR → BAS	8.2	7.6	15.0	0.2
BAS → PYR	23.3	39.4	0.5	0.2
RSNP → PYR	22.7	39.9	0.0	3.9
L4 → PYR	44.1	45.4	15.5	2.3
background → PYR	32.3	31.3	17.3	1.3
total	22.2	25.2	17.9	9.8

Small-scale model The default model (9HC×9MC) can, in principle, be mapped onto the hardware without any synapse loss if the full wafer is available for use. Nevertheless, in some scenarios, a full wafer might not be available, due to faulty components or part of its area being used for emulating other parts of a larger parent network. We simulate this scenario by limiting the usable wafer area to 4 reticles (out of a total of 48 on the full wafer). With the reduced available hardware size, the available pulse bandwidth of the off-wafer communication network decreases as well, such that diffusive background noise can not be modeled with one individual Poisson source per neuron. Hence, each pyramidal neuron receives input from 9 out of 2430 background sources. The total synapse loss for the given network setup amounts to 22.2 % and affects different projection types with varying strength (Tab. 1). Also external synapses are lost, since, in contrast to the synapse loss study (Sec. 3.1.4), they have not been prioritized in the mapping process in this case. Additionally, we applied 20 % synaptic weight noise and simulated the network with a speedup factor of 12 000. The behavior on the ESS is shown in Fig. 11. The distorted network shows no spontaneous attractors (**C**), which can be mainly attributed to the loss of over 32 % of the background synpases. To recover the original network behavior, we first increased the number of background neurons per cell from 9 to 50 to compensate for synaptic weight noise, and also scaled the weights by $1/(1 - p_{\text{loss}})$ for each projection type with extracted synapse loss values p_{loss} (Tab. 1), following the synapse loss compensation method described in Sec. 3.1.4. These measures effectively restored the attractor characteristics of the network (Fig. 11). The attractor dwell times remained a bit smaller than for the regular network (**G**), which can be ascribed to the non-configurable delays (Sec. 3.1.6).

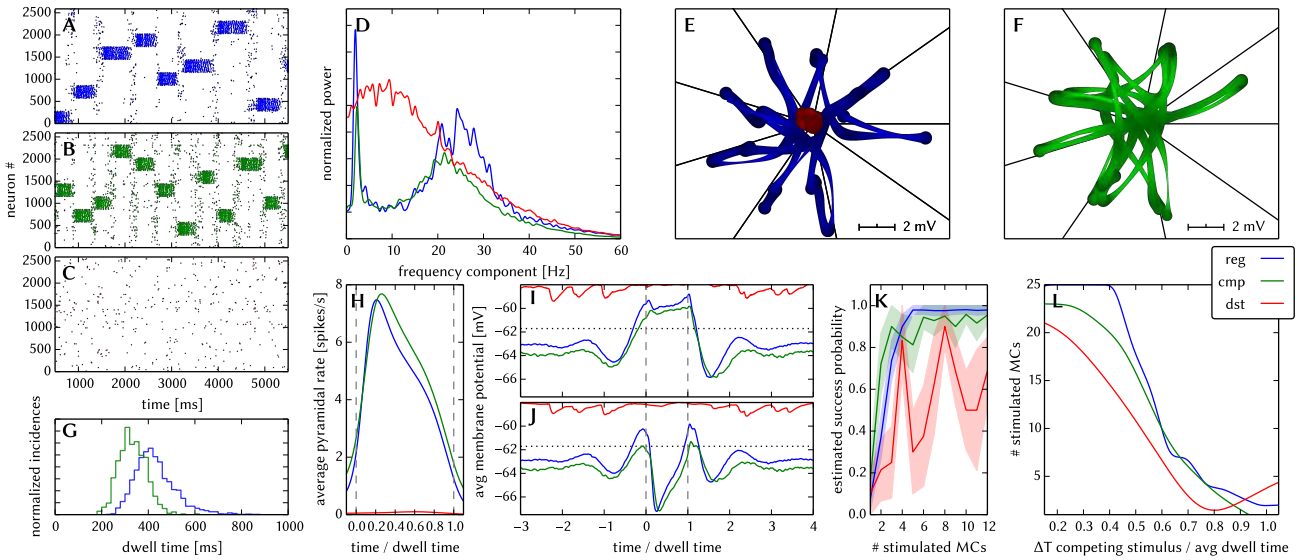

Fig. 11: ESS emulation of the L2/3 model. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. Here, we use the following color code: blue for the original model, red for the distorted case on the ESS (with 20 % synaptic weight noise and ≈ 20 % synapse loss), and green for the compensated case on the ESS. (**A** - **C**) Raster plots of spiking activity. The MCs are ordered such that those belonging to the same attractor (and *not* those within the same HC) are grouped together. A synapse loss of 32 % on the background synpases (see Tab. 1) is the main reason for which no spontaneous attractors are evoked. For this reason, there are no red curves in **G**, **H**, **I** and **J**. Applying a weight compensation and increasing the number of background sources from 9 to 50 effectively restores the original behavior. (**D**) Power spectrum of global activity. Since no spontaneous attractors are evokes, neither attractor switching (∼ 3 Hz) nor gamma oscillations (∼ 25 Hz) can be observed. The spectrum of the distorted network complies with the asynchronous irregular firing observed in **C**. Compensation restores both of the characteristic peaks in the spectrum. (**E** and **F**) Star plots of average PYR voltages from a sample of 5 PYR cells per MC. The disrupted attractor behavior results in a weak fluctuations of average PYR voltages with essentially no clear UP or DOWN states. After compensation, the differences between UP and DOWN states become more pronounced again. (**G**) Distributions of dwell times. The disrupted network effectively shows no spontaneous attractors. As expected from the software simulations, the dwell times remain, on average, slightly shorter after compensation. (**H**) Average firing rate of PYR cells within an active period of their parent attractor. The higher firing rates after compensation are caused by the fixed, short delays, which promote synchronous firing and therefore stronger mutual excitation among PYR cells. (**I**) Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**J**) Average voltage of PYR cells before, during and after an attractor they do not belong to is active (DOWN state). For subplots **H**, **I** and **J**, the abscissa has been subdivided into multiples of the average attractor dwell time in the respective simulations. In subplots **I** and **J** the dotted line indicates the leak potential E_L of the PYR cells. (**K**) Pattern completion in a 25HC×25MC network. Estimated activation probability from 25 trials per abscissa value. (**L**) Attentional blink in a 25HC×25MC network: $p = 0.5$ iso-probability contours, measured over 14 trials per (ΔT, #MCs) pair. Since the distorted network showed no spontaneous attractors (**C**), we used the average dwell time from the pattern completion experiment (**K**) for normalization.

Large-scale model The ability of the network to perform pattern completion and exhibit pattern rivalry was tested on the ESS for the large-scale model with 25 HCs and 25 MCs per HC. From the start, we use a background pool with 5000 Poisson sources and 100 sources per neuron to model the diffusive background noise, as used for the compensation of synaptic weight noise (Sec. 3.1.5). As with the small-scale network, the synapse loss of 17.9 % shows significant heterogeneity (Tab. 1), and affects mainly projections from PYR to inhibitory cells, but also connections from the background and L4 stimulus. In contrast to the idealized case in Sec. 3.1.4, where each synapse is deleted with a given probability, the synapse loss here happens for entire projections at the same time, i.e. all synapses between two populations are either realized completely or not at all. We note that the realization of all PYR-RSNP synapses is a priori impossible, as each RSNP cell has $24 \times 24 \times 30 = 17280$ potential pre-synaptic neurons (cf. scaling rules in Sec. B.4), which is more than the maximum possible number of pre-synaptic neurons per HICANN (14336, see Sec. 2.1.1). The simulation results with 20 % synaptic weight noise for pattern completion and pattern rivalry are shown in Fig. 11 **K** and **L** (red curves). In both cases the network functionality is clearly impaired. In particular, the ability of an active pattern to suppress other patterns is noticeably detoriated, which can be traced back to the loss of 37 % of PYR-RNSP connections.

In order to restore the functionality of the network we used a two-fold approach: First, we attempted to reduce the binary loss of PYR-RSNP projections by reducing the number of PYR cells per MC from 30 to 20, which decreases the total number of neurons in the network, as well as the number of potential presynaptic neurons per RNSP cell. The synapse loss was thereby reduced to 28.8 % for PYR-RSNP projections and was eliminated almost completely for all other projections (Tab. 1). Secondly, we compensated for the remaining synapse loss by scaling the synaptic weights as described in Sec. 3.1.4.

After application of these compensation mechanisms, we were able to effectively restore the original functionality of the network. Both pattern completion and attentional blink can be clearly observed. The small remaining deviations from the default model can be attributed to the inhomogeneity of the synapse loss and the fixed delays on the wafer.

3.2 Synfire chain with feed-forward inhibition

Our second benchmark network is a model of a series of consecutive neuron groups with feed-forward inhibition, called *synfire chain* from here on (Kremkow et al, 2010b). This network acts as a selective filter to a synchronous spike packet that is applied to the first neuron group of the chain. The behavior of the network is quantified by the dependence of the filter properties on the strength and temporal width of the initial pulse. Our simulations show that synapse loss can be compensated in a straightforward manner. Further, the major impact of weight noise on the network functionality stems from weight variations in background synapses, which can be countered by modification of synaptic and neuronal parameters. The effect of fixed axonal delays on the filtering properties of the network can be countered only to a limited extent by modifying synaptic time constants and the strength of local inhibition. Simulations using the ESS show that the developed compensation methods are applicable simultaneously. Furthermore, they highlight some further sources of potential failure of pulse propagation that originate from bandwidth limitations in the off-wafer communication infrastructure.

3.2.1 Architecture

Feed-forward networks with a convergent-divergent connection scheme provide an ideal substrate for the investigation of activity transport. Insights have been gained regarding the influence of network characteristics on its response to different types of stimulus (Aertsen et al, 1996; Diesmann et al, 1999; Vogels and Abbott, 2005). Similar networks were also considered as computational entities rather than purely as a medium for information transport (Abeles et al, 2004; Schrader et al, 2010; Kremkow et al, 2010a). The behavior of this particular network has been shown to depend on the connection density between consecutive groups, on the balance of excitation and inhibition as well as on the presence and magnitude of axonal delays in Kremkow et al (2010b). This makes it sensitive to hardware-specific effects such as an incomplete mapping of synaptic connectivity, the variation of synaptic weights, bandwidth limitations which cause loss of individual spike events and limited availability of adjustable axonal delays and jitter in the spike timing that may be introduced by different hardware components.

The feed-forward network comprises a series of successive neuron groups, each group containing one excitatory and one inhibitory population. The excitatory population consists of 100 regular-spiking (RS), the inhibitory of 25 fast-spiking (FS) cells. Both cell types are modeled as LIF neurons with exponentially shaped synaptic conductance without adaptation, as described in Sec. 2.1.1. Both RS and FS neurons are parameterized using identical values (Tab. 17).

Each excitatory population projects to both populations of the consecutive group while the inhibitory population projects to the excitatory population in its local group (Fig. 12 **A**). There are no recurrent connections within the RS or FS populations. In the original publication (Kremkow et al, 2010b), each neuron was stimulated independently by a Gaussian noise current. Because the hardware system does not offer current stimulus for all neurons, all neurons in the network received stimulus from independent Poisson spike sources. Since for both Gaussian current stimulus and the diffusion limit of Poisson stimulus (high input rates, low synaptic weights) the membrane potential is stationary Gaussian, with an autocorrelation dominated by the membrane time constant, the only remaining differences are due to the finite, but small, synaptic time constants. The rate and synaptic weight of the background stimulus were adjusted to obtain similar values for the mean and variance of the membrane potential, resulting in a firing rate of 2 kHz with a synaptic weight of 1 nS.

The initial synchronous stimulus pulse is emitted by a population of spike sources, which has the same size and connection properties as a single RS population within the network. A temporally localized *pulse packet* was used as a stimulus, whereby each of the 100 spike sources emitted a_0 spikes that were sampled from a Gaussian distribution with a common mean time and a given standard deviation σ_0. The variables (σ_i, a_i) are

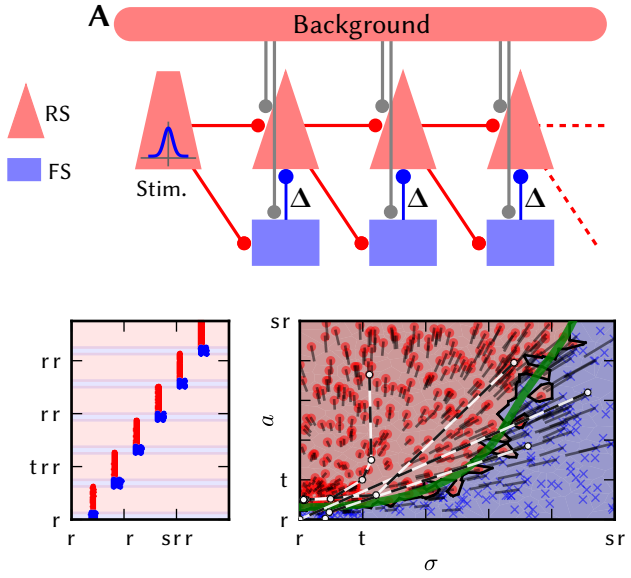

Fig. 12: Synfire chain network. (**A**) Connectivity of the synfire chain with feed-forward inhibition (Kremkow et al, 2010b). Excitatory projections are shown in red, inhibitory in blue. In the default realization the network consists of six consecutive groups. The local FS → RS projection has an adjustable delay Δ, which affects the network dynamics. The intergroup delay is set to 20 ms for visualization purposes following the previous work; this has no influence on the filter properties because the delay of both intergroup projections is equal. The background stimulus is realized using random Gaussian current (original) and Poisson background spikes (adapted version for the hardware). The parameters for neurons and connections are given in Tab. 17 and Tab. 18. (**B**) Exemplary raster plot of the network behavior. The first group receives a pulse packet with $a = 1$ and $\sigma_0 = 1\,\text{ms}$, which propagates as a synchronous spike volley along the chain. (**C**) Characterization of the network behavior in the (σ, a) state space. Each marker represents the initial stimulus parameters (σ_0, a_0). The stimulus parameters were selected randomly from the region ($a_0 < 10$, $\sigma_0 \leq 10\,\text{ms}$). The region with ($a_0 < 2$, $\sigma_0 \leq 2\,\text{ms}$) was simulated more frequently to increase the resolution near the convergence points of the propagation. The marker color is linearly scaled with the activity in the last group, a_6, being blue for $a_6 = 0$ and red for $a_6 = 1$ and is set to red for $a_6 > 1$. To improve visibility, the background is colored according to the color of the nearest marker, red for $a_6 \geq 0.5$ and blue otherwise. Experiments where the RS group did not fire are marked as ×. The gray lines originating from each marker denote the direction towards the pulse volley parameters (σ_1, a_1). The green line shows a fit to the separatrix between zero and nonzero activity at the last group of the synfire chain (see Sec. 3.2.2 for details). This approximation is used to compare the behavior of different modifications of the original network. The dashed black and white lines show four exemplary trajectories through the (σ, a) state space.

later used to describe the characteristics of the activity in the ith group of the chain, referring to the temporal pulse width and number of spike pulses per neuron, respectively.

3.2.2 Functionality criteria

The functionality of the feed-forward network is assessed by examining the propagation of a synchronous pulse after the stimulus is applied to the first group in the chain (Fig. 12 **B**). The propagation is quantified by applying initial stimuli of varying strength $a_0 \in [0, 10]$ and temporal spread $\sigma_0 \in [0\,\text{ms}, 10\,\text{ms}]$. For each synfire group $i \in \{1, ..., 6\}$, the activation is determined by setting a_i to the number of emitted spikes divided by the number of neurons in the RS population. σ_i is the standard deviation of the spike pulse times. Typically, the resulting "trajectory" in the (σ, a) space (Fig. 12) is attracted to one of two fixed points: either near $(0\,\text{ms}, 1)$, i.e., the pulse packet propagates as a synchronous spike volley, and $(0\,\text{ms}, 0)$, i.e., the propagation dies out (e.g., Fig. 13 **A**).

The network behavior is characterized by the separating line between successful and extinguished propagation in the state space (σ, a) of the initial stimulus; this line will be called *separatrix* from here on. The differentiation between extinguished and successful propagation is defined by $a_6 \geq 0.5$ resp. $a_6 < 0.5$ in the last (6th) group. This is justified because in the undistorted case, a is clustered around the values 0 and 1 (Fig. 27). Due to the statistic nature of the connectivity, background stimulus and pulse packet, the macroscopic parameters σ and a do not fully determine the behavior of the system. This means that in the reference simulation, there is a small region around the separatrix where the probability of a stable propagating pulse is neither close to zero nor to one. Thus, in addition to the location of the separatrix (Sec. C.3.2), the width of this region is taken as a functionality criterion.

The background stimulus is adjusted such that the spontaneous firing rate in the network is below 0.1 Hz, in accordance with Kremkow et al (2010b). In cases in which distortions induce a much stronger firing, the spike trains are filtered before the analysis by removing spikes which appear not to be within a spike volley (Sec. C.3.4).

3.2.3 Synapse loss

Homogeneous synapse loss affects the strength of both excitatory and inhibitory projections equally. Additionally, the number of incoming spikes seen by a single neuron varies as synapses are removed probabilistically, in

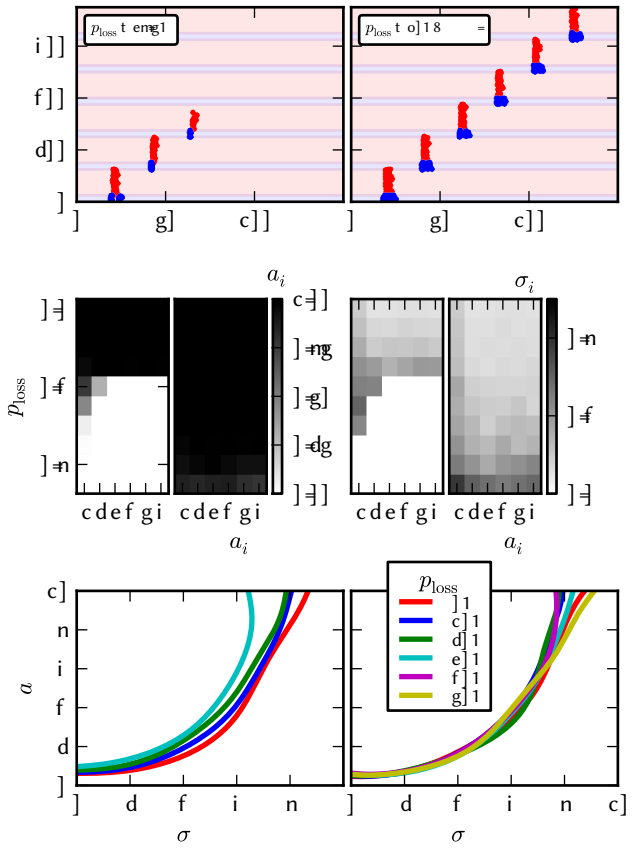

Fig. 13: Effect and compensation of synapse loss for the synfire chain network. (**A**) Synfire network with 37.5% synapse loss applied to all connections within the network. External connections (synchronous stimulus and background) are not affected. (**B**) Raster plot with active compensation at 90% synapse loss. (**C**) Activation a_i in each group i with varying values of synapse loss. (**D**) a_i as in **C** but with active compensation. (**E**) Pulse width σ_i in each group i with varying values of synapse loss. (**F**) Pulse width as in **E** but with active compensation. (**G**) Comparison of approximated separatrix locations for synapse loss values from 0% to 50%. The lines for 40% and 50% are missing because no stable region exists. (**H**) Approximated separatrix locations with active compensation.

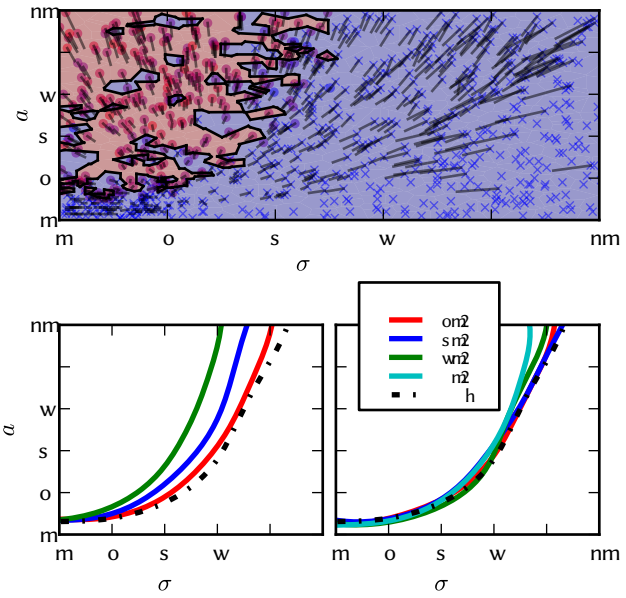

Fig. 14: Effect of synaptic weight noise on the synfire chain model. The spike data for all three plots was filtered to remove spontaneous spikes in individual neurons, which stem from weight increase in some background synapses due to weight noise. (The filter parameters were $T = 10\,\text{ms}, N = 25$ cf. Sec. C.3.4) (**A**) State space at 80% weight noise. The set of inputs that evokes activity in the last group is patchy as a consequence of the distortion. In the compensated case the separation is sharp again, as shown in Fig. 29. (**B**) Approximate separatrix locations for smaller values of weight noise. (**C**) Approximate separatrix locations for the compensated case.

contrast to the undistorted model with a fixed number of incoming connections for each neuron type (Tab. 18). Synapse loss was applied to all internal connections as well as to the connection from the synchronized stimulus population to the first group in the network; background stimulus was not affected (cf. Sec. 2.4). Fig. 13 **A** shows a single experiment with synapse loss of 37.5 %, contrasting with the undistorted case (Fig. 12 **A**). Above a certain value of synapse loss, the signal fails to propagate to the last group. As shown in Fig. 13 **C** and **E** for one stimulus parameter set, successful propagation stops at a synapse loss value between 30% and 40%. The pulse width increases with rising synapse loss due to the increasing variation of synaptic conductance for individual neurons (**E**). The effect is reversed by increasing all synaptic weights in the network by a factor of $1/(1 - p_{\text{loss}})$, with p_{loss} being the probability of synapse loss. This compensation strategy can effectively counter synapse loss of up to 90 % (**B**, **D**) and the pulse width increase is shifted to larger values of synapse loss (**F**). The distortion has only a minor effect on the a-value of the separatrix in the depicted region (**G**). However, the location of the separatrix at $\sigma_0 = 0$ rises with synapse loss until it reaches the fixed point at approx. $(0.1\,\text{ms}, 1)$, at which point a bifurcation occurs and the the attractor region for $(0.1\,\text{ms}, 1)$ disappears (as described in Diesmann et al (2001) for the case of varying weights). In the compensated case, the separatrix locations are identical with the undistorted case within the measurement precision.

3.2.4 Synaptic weight noise

The effect of synaptic weight noise is shown in Fig. 14. Similarly to the effect of synapse loss, the region of sta-

ble propagation shrinks (**B**); additionally, the border between the regions of stable and extinguished propagations becomes less sharp (**A**). This is caused by two effects: Varying strength of the background stimulus, and varying strength of the synaptic connections within the network. The first effect is significant because the background stimulus to each neuron is provided through a single synapse. Thus, the resting potential of each neuron is shifted, significantly changing its excitability and, in some cases, inducing spontaneous activity. One possibility of countering this effect is to utilize several synapses for background stimulus thereby averaging out the effect of individual strong or weak synapses, as has been done in the case of the L2/3 model in Sec. 3.1.5. Here, a different method was employed: The resting potential E_L was raised while simultaneously lowering the synaptic weight from the background stimulus. The parameters were chosen in such a way that the mean and variance of the distribution of membrane voltages in each neuron population was kept at the value of the undistorted network:

$$\langle V \rangle \approx w_0 \cdot \langle K \rangle + E_\mathrm{L} \quad \text{and} \tag{10}$$

$$\mathrm{Var}\,[V] \approx w_0^2 \mathrm{Var}\,[K] + \mathrm{Var}\,[w]\left(\langle K \rangle^2 + \mathrm{Var}\,[K]\right) \ , \tag{11}$$

where $K(t) = \sum_{\mathrm{spk}\ j} \kappa(t - t_j)$ represents the effect of the background stimulus, κ being the PSP kernel, and $\mathrm{Var}\,[w] = w_0^2\sigma^2$ appears due to fixed-pattern weight noise. In the distorted case, the width of this distribution is a combined effect of the random background stimulus and the fixed-pattern weight variation, while in the original case it originates from the stochasticity of the stimulus only. In the undistorted case, $\mathrm{Var}\,[w]$ is 0, and only the first term contributes to $\mathrm{Var}\,[V]$. With increasing σ^2, the contribution of the second term to $\mathrm{Var}\,[V]$ increases, which is compensated by changing w_0 accordingly, keeping $\mathrm{Var}\,[V]$ at the original level. This, in turn, changes $\langle V \rangle$, which is compensated by a change of E_L.

The effect of synaptic weight noise within the network itself is less significant compared to its impact on the noise stimulus. Fig. 14 **C** shows that removing the effect of background stimulus noise alone is sufficient to counteract synaptic noise values of up to 50%.

3.2.5 Non-configurable axonal delays

Fig. 15 **A** shows the effect of varying axonal delays between the inhibitory and excitatory population of a single synfire group. As was shown in Kremkow et al (2010b), the delay can be employed to control the position of the separatrix between stable and unstable propagation. Because the axonal delay is not configurable for on-wafer connections, a different method is

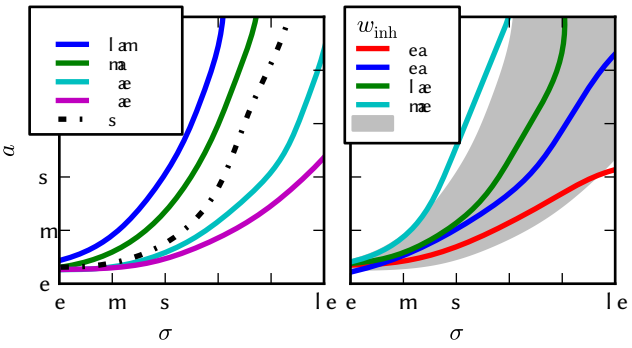

Fig. 15: Delays in the synfire chain model. **(A)** Reproduction of Kremkow et al (2010b), fig. 4c. The location of the separatrix is modified by changing the axonal delay of local inhibition. For a value of 0.4 ms, no stable region is present. **(B)** The location of the separatrix is modified by varying weights for synapses taking part in local inhibition. The axonal delay of local inhibition was fixed at 1.5 ms and the inhibitory time constant was increased by a factor of 3. The gray region shows the range of the separatrix location for delay values from 1.2 ms to 7 ms (the range in plot **A**) as reference.

required to regain the ability to control the separatrix. While Sec. 3.2.3 and Sec. 3.2.4 show that synaptic weight noise and synapse loss can influence the location of the separatrix, a method is required that is independent of those distortion mechanisms. Diesmann (2002) shows that several parameters, including group size and noise level, can modify the separatrix location, albeit for a model without feed-forward inhibition. Here, we investigate to which extent parameter modification can reproduce the effect of variable delays. For very short delays (in this case, 0.1 ms, not shown), stable propagation does not occur, because the onset of local inhibition is nearly synchronous with the onset of external excitation. This effect was countered by increasing the synaptic time constant and simultaneously decreasing the synaptic weight for local inhibition, thus extending the duration of inhibition that acts on the RS population. The inhibitory synaptic time constant was increased by a factor of 3 while simultaneously reducing the synaptic weight of the inhibitory projection. Fig. 15 **B** shows the result of the compensation for 1.5 ms local inhibition delay. For both values of axonal delay, the location of the separatrix can be controlled by changing the weight of inhibition. However, its shape differs from the delay-induced case because of the modified delay mechanism of inhibition. Reduction of the weight beyond a certain point is not possible, as balanced inhibition is required for network functionality (Kremkow et al, 2010b). It is important to note that this kind of compensation is specific to the state space region which

is examined, and that it can not be extended to arbitrarily large delays.

3.2.6 Full simulation of combined distortions

At last, we simulate the synfire chain with the ESS and compensate simultaneously for all the distortions addressed above. Before running ESS simulations, we have verified the compatibility of the proposed compensation strategies for different distortion mechanisms in software simulations dealing with the simultaneous incidence of synaptic weight noise, synapse loss and non-configurable axonal delays (Sec. C.3.1). We proceed with a quantification of synapse loss after mapping the synfire chain for different network sizes to the hardware. For the ESS simulations we limit the model to very few hardware resources to artificially generate synapse loss, such that all of the above distortion mechanisms are present. Additional hardware simulations investigating the influence of spike loss and jitter on the network functionality are provided in Sec. C.3.5.

Synapse loss We mapped the synfire chain at different network sizes onto the BrainScaleS wafer-scale hardware in order to quantify the synapse loss. The results are shown in Fig. 16 **A**. For this purpose we developed network scaling rules that depend on the number and the size of the synfire groups (Sec. C.2). Due to its modular structure and feed-forward connectivity scheme, there is no synapse loss for networks with up to 10 000 neurons. However, for network sizes above 30 000 neurons, the ratio of lost synapses increases abruptly. With increasing network size more neurons have to be mapped onto one HICANN thereby reducing the number of hardware synapses per neuron. Moreover, as the group size grows with the network size (cf. Tab. 19), also the number of pre-synaptic neurons for all neurons mapped onto one HICANN increases, so that the maximum number of inputs to a HICANN, i.e. the synapse drivers, becomes a limiting constraint. The combination of both factors unavoidably leads to synapse loss.

Distorted and compensated simulation For the ESS simulation, we applied the following modifications to the benchmark network: originally, each cell in the network receives Poisson background stimulus from an individual source with 2000 Hz. Because the off-wafer pulse routing network does not support such high bandwidths (cf. Sec. 2.1.2), we reduce the total number of background sources from 750 to 192 and let each neuron receive input from 8 sources, while decreasing the Poisson rate by a factor of 8, using the same mechanism as for the compensation of fixed pattern noise in the L2/3 model (cf. Sec. 3.1.5). For the same reason, the network was emulated with a speedup factor of 5000 compared to biological real-time, whereby the effective bandwidth for stimulation and recording is doubled with respect to the normal operation with a speedup of 10 000. As seen before, no synapse loss occurs for small networks. However, as discussed for the L2/3 model in Sec. 3.1.7, one can consider situations where only a small part of the wafer is available for experiments, or where some neurons or synaptic elements are defective or missing a calibration. Therefore, in order to generate synapse loss in the feed-forward network, we limited the network to only 8 out of 48 reticles of the wafer and furthermore declare half of the synapse drivers as not available. This resulted in a total synapse loss of 27.4 %. As with the L2/3 model, the synapse loss was not homogeneous but depended strongly on the projection type (Tab. 2).

Table 2: Projection-wise synapse loss of the synfire chain with feed-forward inhibition after the mapping process.

projection	synapse loss [%]
Pulse Packet $\rightarrow$ RS_0	21.3
Pulse Packet $\rightarrow$ FS_0	12.7
$RS_n \rightarrow RS_{n+1}$	32.4
$RS_n \rightarrow FS_{n+1}$	32.0
$FS_n \rightarrow RS_n$	20.8
Poisson background $\rightarrow$ ALL	0
total	27.4

We simulated the synfire chain with default neuron and synapse parameters on the ESS with 20 % synaptic weight noise and the above synapse loss. The (σ, a) state space (Fig. 16 **B**) shows no stable point of propagation. This can be mainly attributed to the small and non-configurable axonal delays which are in the range of 0.6 ms to 1.2 ms for the chosen speedup factor of 5000.

In order to recover the original behavior, we applied the previously developed compensation methods described in Sec. 3.2.3 to 3.2.5. Synapse loss was compensated separately for each projection type using Tab. 2. We were able to compensate for all distortions while still maintaining control over the position of the separatrix (Fig. 16 **C**).

However, we encountered some abnormalities as can be seen in Fig. 16 **D** showing the (σ, a) state space for one of the separatrices: For $\sigma \approx 3$ ms and $a > 7$ one can recognize a purple region indicating that not all RS cells of the last group spiked. In reality, spikes occurred for all RS cells in the simulated hardware network, but not all spikes were recorded because they were lost in the off-wafer communication network (Sec. 2.1.2). For very small σ_0 an additional effect can appear: input bandwidth limitations can result in very dense pulse volleys

not being propagated through the synfire chain, as can be seen e.g. for the blue point with $\sigma_0 = 0.0361\,\text{ms}$ and $a_0 = 7.6$ in the upper left of **D**. In that particular case the large majority of input spikes were lost in the off-wafer communication network so that they did not even reach the first synfire group. We remark that this effect only appeared for σ_0 smaller than 0.1 ms.

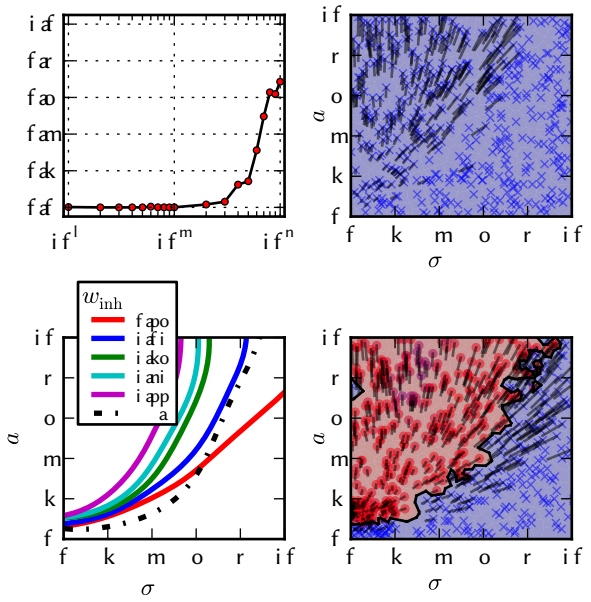

Fig. 16: Distorted and compensated simulations of the feed-forward synfire chain on the ESS: (**A**) Synapse loss after mapping the model with different numbers of neurons onto the BrainScaleS System. (**B**) (σ,a) state space on the ESS with default parameters, 20 % weight noise, and 27.4 % synapse loss. (**C**) After compensation of all distortions, different separatrices are possible by setting different values of the inhibitory weight. (**D**) Compensated state space belonging to the blue separatrix in **C**. w refers to the synaptic weight of local inhibition.

3.3 Self-sustained asynchronous irregular activity

Our third benchmark is a cortically inspired network with random, distance-dependent connectivity which displays self-sustained asynchronous and irregular firing (short: "AI network"). We define functionality measures on several levels of abstraction, starting from single network observables such as the network firing rate, the correlation coefficient and the coefficient of variation, the properties of the power spectrum of the network activity, up to global behavior such as the dependence of network dynamics on the internal synaptic weights g_{inh} and g_{exc}. We test two compensation strategies based

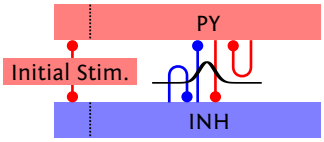

Fig. 17: Schematic of the connectivity of the random cortical network. Excitatory PY and inhibitory INH neurons are connected randomly with a spatial, Gaussian connection probability profile. The connection properties are given in Sec. D. A small part of the network is stimulated in the beginning of the experiment.

on a mean field approach and on iterative modification of individual neuron parameters. While the first method offers a way to control the mean firing rate in the presence of synapse loss, the second is applicable to the synapse loss and weight noise distortions simultaneously, in contrast to the other presented compensation methods. The distortion of axonal delays does not significantly affect the network functionality because the intrinsic hardware delay is approximately equal to the delay utilized in the model. A scaling method for the network size is introduced and the effectivity of the second compensation method was demonstrated using the ESS on a large network with mapping-induced synapse loss and imposed fixed-pattern synapse noise.

3.3.1 Architecture

Self-sustained states in spiking neural networks are known to be exquisitely sensitive to the correlation dynamics generated by recurrent activity (Kumar et al, 2008; El Boustani and Destexhe, 2009). Because of this sensitivity, a model of self-sustained activity within the asynchronous-irregular self-sustained activity regime can provide a strong comparison between hardware and software platforms, by requiring the hardware network to reproduce the low firing, weakly correlated, and highly irregular dynamics of this state. Notably, it is often observed that this activity regime provides a good match to the dynamics observed experimentally in the awake, activated cortex (Destexhe and Pare, 1999; Brunel, 2000; Destexhe et al, 2003). Additionally, one can note that the self-sustained activity regime provides an interesting test of the BrainScaleS hardware system, as in this state, the model network is not driven by external Poisson input, but has dynamics dominated by interally generated noise (Destexhe and Contreras, 2006), beyond the initial brief Poisson stimulation to a small percentage of the network.

Here, we used a reduced model based on that published in (Destexhe, 2009). Neurons in the network followed the AdEx equations 1 to 3, modeling regular spiking pyramidal cells (PY) with spike frequency adap-

tation (Connors and Gutnick, 1990) and fast spiking inhibitory cells (INH) with relatively little spike frequency adaptation. Instead of explicitly modeling the thalamocortical or corticocortical networks, as in the previous work, we have chosen to modify the model, simplifying it to a single two-dimensional toroidal sheet and adding local connections and conduction delays. The addition of local connectivity follows the experimental observation that horizontal connections in neocortex project, for the most part, to their immediate surroundings (Hellwig, 2000), while the choice of linear conduction delays reflects electrophysiological estimates of conduction velocity in these unmyelinated horizontal fibers, in the range of 0.1 to 0.5 $\mathrm{m\,s^{-1}}$ (Hirsch and Gilbert, 1991; Murakoshi et al, 1993; Bringuier et al, 1999; González-Burgos et al, 2000; Telfeian and Connors, 2003). Propagation delays are known to add richness to the spatiotemporal dynamics of neural network models (Roxin et al, 2005), and in this case are observed to expand the region in the 2D space spanned by the excitatory and inhibitory conductances that supports self-sustained activity, albeit only slightly. The spatiotemporal additions to the network model of (Destexhe, 2009) incorporate structure that allows making general comparisons between the population dynamics in the model and in multichannel recordings *in vivo*, and variants of this network model have been used to explain observed differences in the spatial structure of activity across different brain states (Muller and Destexhe, 2011).

3.3.2 Functionality criteria

The global functionality criterion for this network consists of the ability to sustain activity in an asynchronous and irregular activity regime. The activity is considered self-sustained upon persistence to the end of the chosen simulation period. The activity characteristics are quantified for the pyramidal cells using the mean and variance of the firing rates, the irregularity of individual spike trains ($\mathrm{CV_{ISI}}$, Eq. 31), the synchrony via the correlation coefficient (CC, Eq. 32) and the power spectrum of the excitatory activity. The implementation details are given in Sec. D.2.

These criteria were evaluated for a range of excitatory and inhibitory synaptic weights g_{exc} and g_{inh} for the default network consisting of 3920 neurons. Fig. 18 (**A**) shows the region in the (g_{exc}, g_{inh}) parameter space that allows self-sustained activity, which is achieved at pyramidal firing rates above 8 Hz (**G**).

The coefficient of variation of the firing rates across neurons ($\mathrm{CV_{rate}}$) is small ($\mathrm{CV_{rate}} < 0.2$, see the zero-distortion data in Fig. 19 **B**), as all neurons have identical numbers of afferent synapses with identical weights in each network realization. In addition to the parameter space plots in the top row of Fig. 18, we plot the other criteria against the mean firing rate in the bottom row and recognize the latter as the principal property of each state that mostly determines all other criteria.

The activity is irregular ($\mathrm{CV_{ISI}} > 1$) across all states (**C**) and is mainly determined by the network firing rate: the $\mathrm{CV_{ISI}}$ first increases with the firing rate, then saturates and decreases for rates higher than 50 Hz (**H**). Over the entire parameter space, the spike trains of the pyramidal cells are only weakly correlated, with a CC between 0.01 and 0.03.

The average CC increases with the firing rate, which can be attributed to local areas in which neurons synchronize over short time periods. At last, we look at the power spectrum of the global pyramidal activity, exemplarily for the (9 nS, 90 nS) state in (**F**). As a comparison for further studies we follow (Brunel, 2000) and use the position of the non-zero peak in the power-spectrum, which is shown for each (g_{exc}, g_{inh}) point (**E**) and as a function of the firing rate (**J**): The position of the power spectrum peak frequency (Sec. D.2) increases linearly with the mean firing rate.

3.3.3 Non-configurable axonal delays

For the analysis of the effects of non-configurable delays we repeated the (g_{exc}, g_{inh}) sweep with all axonal delays set to 1.5 ms, cf. Sec. 2.4. The distortion did not affect any of the functionality criteria, as each neuron still received synaptic input comparable to the undistorted case. One might expect the delay distortion to have an influence on the power spectrum of global activity as we switched from distance-dependent delays to a globally constant delay of 1.5 ms as it changes the temporal correlation of the effect of a neuron on all of its efferents. In fact, the power spectra did not change significantly, which can be explained as follows: In the undistorted case, the *average* of all distance-dependent delays in the network amounts to 1.55 ms (cf. Fig. 31), which is close to the constant distortion value of 1.5 ms we use to model the non-configurable delays on the hardware. In this particular case, the hardware delay matches the average delay in the network such that no distortion is introduced. Accordingly, parameter space sweeps on the ESS yielded the same results.

In Sec. D.4.2 we provide further simulations on the influence of the distribution of delays on the behavior of the network, showing that the effect of the distance-dependent delays is small and that it is mostly the average delay which matters. In our case, this delay exactly corresponds to the average delays on the wafer when

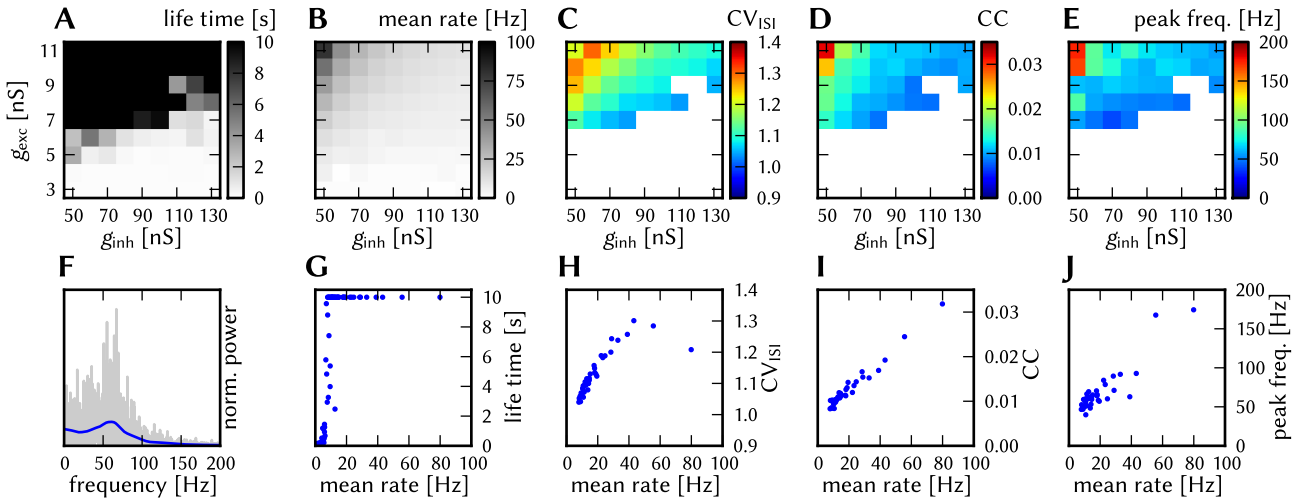

Fig. 18: Behavior of the undistorted AI network. On the top: survival time (**A**), mean firing rate (**B**), coefficient of variance $\mathrm{CV_{ISI}}$ (**C**), coefficient of correlation CC (**D**) and position of peak in power spectrum of global activity (**E**) in the parameter space for g_{exc} and g_{inh} for the default network with 3920 neurons without any distortions. (**F**) Power spectrum of the global pyramidal activity for the state ($g_{\mathrm{exc}} = 9\mathrm{nS}$, $g_{\mathrm{inh}} = 90\mathrm{nS}$). The population activity was binned with a time of 1ms, the raw spectrum is shown in gray, the blue curve shows a Gauss-filtered ($\sigma = 5\,\mathrm{Hz}$) version for better visualization. The position of the peak in the filtered version was used for (**E**). In (**G - J**) the dependence of single criteria on the mean firing rate is shown: survival time (**G**), $\mathrm{CV_{ISI}}$ (**H**), CC (**I**), position of peak in power spectrum (**J**). In the last three plots only surviving states of the (g_{exc}, g_{inh}) space were considered.

running at a speedup of 10 000 compared to biological real-time, such that there is no need for a compensation here.

We note that for variants of this benchmark, where the average network delay is higher or lower than 1.5 ms, there exists a simple but effective compensation strategy by just modifying the speedup of the emulation on the hardware, such that the average network delay is directly mapped onto the hardware delay. We can assume a modified experiment where the average delay amounts to 3 ms. By choosing a speedup of 20 000, this delay can be directly mapped to the 150 ns average delay on the hardware. Such a change of emulation speed is not arbitrary, as one has to make sure that the neural dynamics can still be emulated at the chosen speed (cf. supported parameter ranges in Tab. 5). This must hold also for potential other networks that are interconnected to the AI network. Also, this method would entail that the available bandwidth for external stimulation is reduced by a factor of 2, which is no issue for this self-sustaining kind of network but must be considered for other network models that require a large proportion of the off-wafer bandwidth. We would not expect such a confinement for the internal pulse communication between neurons, as the maximum bandwidth of the on-wafer routing network was designed to support speedup factors up to 10^5.

3.3.4 Synaptic weight noise

To analyze the influence of the fixed-pattern noise in the hardware (cf. Sec. 2.4) we applied synaptic weight noise in a range from 10 % to 50 %. The most significant results are shown in Fig. 19: The region of self-sustained states in the (g_{exc}, g_{inh}) space was increased with synaptic weight noise, cf. the circles in (**C**) marking states that survived with 50 % synaptic weight noise but not in the undistorted case. The firing rate increases with distortion (**A**): the change is the stronger the lower g_{exc} and diminishes for states with an already high firing rate in the undistorted case (**C**). Synaptic weight noise leads to an increase of the variation of firing rates ($\mathrm{CV_{rate}}$), with the change being stronger for high population firing rates (**B**). The $\mathrm{CV_{ISI}}$ as a function of firing rates remains unchanged for low rates, but decreases for higher firing rates in proportion to the distortion level (**E**). The weight noise introduces randomness into the network, thereby reducing synchrony: The pairwise correlation between neurons decreases linearly with amount of weight distortion (**F**). The power spectrum of the global activity is not affected by the weight noise.

3.3.5 Synapse loss

Synapse loss has a similar influence on the network behavior as synaptic weight noise: Fig. 20 shows the

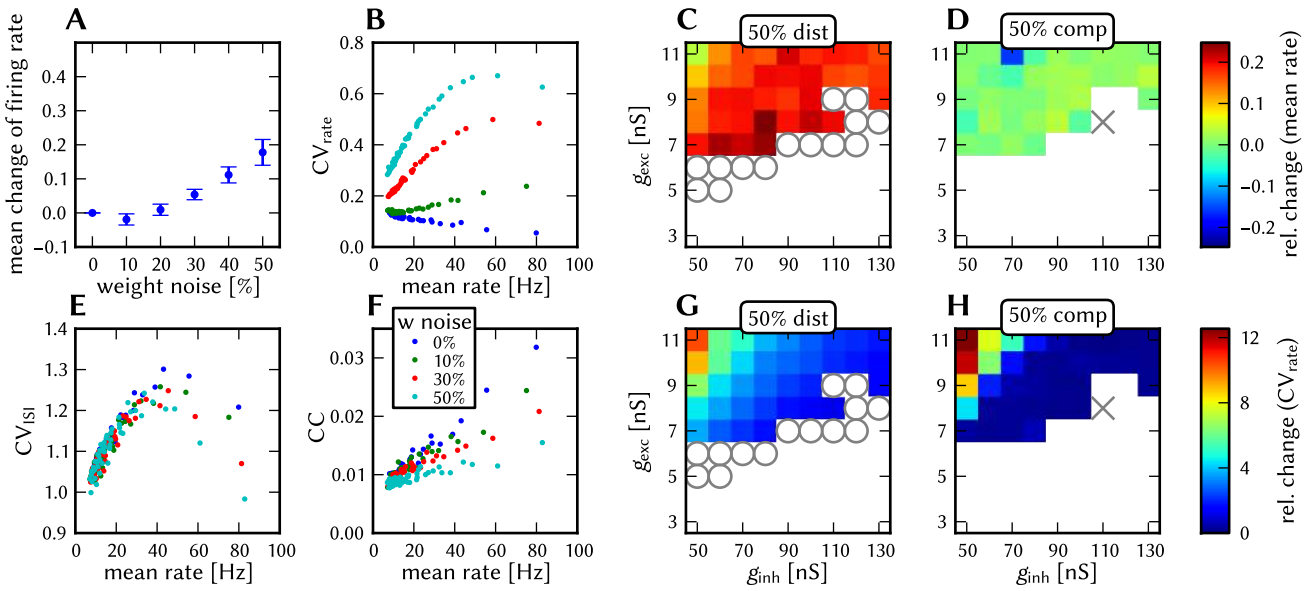

Fig. 19: Effect and compensation of synapse weight noise in the AI network: (**A**) Relative change of the firing rate with respect to the undistorted network averaged over all sustained states for varying synapse weight noise. (**B**) $\mathrm{CV}_{\mathrm{rate}}$ as a function of mean rate for every survived state for varying synapse weight noise. (**C** and **D**) Relative change of the firing rate with respect to the undistorted for each state for 50 % synapse weight noise(**C**) and compensated (**D**). (**E**) $\mathrm{CV}_{\mathrm{ISI}}$ as a function of mean rate for varying synapse weight noise. (**F**) CC as a function of mean rate for varying synapse weight noise. (**G** and **H**) Relative change of $\mathrm{CV}_{\mathrm{rate}}$ with respect to the undistorted for each state for 50 % synapse weight noise(**G**) and compensated (**H**). In (**C** and **D**) and (**G** and **H**): A cross marks a state that was sustained in the undistorted but not sustained in the compared case. A circle marks a state that was not sustained in the original but sustained in the compared case.

results of the g_{exc}-g_{inh} sweeps for synapse loss values between 10 % and 50% (cf. Sec. 2.4). The region of sustained states increases with synapse loss but not as strongly as for weight noise (**C**). The firing rate increases with synapse loss (**A**) Compared to the change caused by synaptic weight noise, however, the effect is much stronger for synapse loss. The same holds for the variance of the firing rates across the pyramidal neurons, which again increases with synapse loss, as can be seen in (**B**). Note that the $\mathrm{CV}_{\mathrm{rate}}$ first increases with the mean rate, then reaches a maximum and finally saddles for high rates. We remark that for high synapse loss, some neurons did not fire at all. Both the irregularity and the correlation of firing decrease with increasing synapse loss, leaving the network still in an asynchronous irregular state (**E** and **F**). Synapse loss shows no effect on the power spectrum of global pyramidal activity.

3.3.6 Compensation strategies

The distortions analyzed in the previous section leave two major criteria that need to be recovered: The population firing rate and the variation of firing rates across the population. We consider the other effects (change of CC, $\mathrm{CV}_{\mathrm{ISI}}$, peak frequency in power spectrum) as minor because they are mainly determined by the mean rate and discard them in the following.

One apparent approach for recovering the original firing rate is to change the strengths of the synaptic weights g_{exc} and g_{inh}. Considering the conducted $(g_{\mathrm{exc}}, g_{\mathrm{inh}})$ parameter space sweeps, we could simply select the distorted state that best matches the criteria of the undistorted reference. However, this method requires to scan g_{exc} and g_{inh} over a wide range to finally get to the desired result. Preferably, one wants to have a compensation method that can be applied to a single experiment and works without huge parameter sweeps.

Mean field compensation for rate change The mean firing rate in the network rises with an increasing synapse loss value. This effect can be understood using a mean-field approach (see, e.g., Kumar et al, 2008) in which the response rate of a single neuron's firing rate is assumed to be a function of the mean network firing rate.

$$\nu_i = f\left(\nu_{\mathrm{in,exc}}, \nu_{\mathrm{in,inh}}\right) \tag{12}$$

With this ansatz, which is similar to the approach in (Brunel, 2000) where the afferent neurons are replaced by independent Poisson processes with equal instantaneous rate in a sparse random network, the mean firing

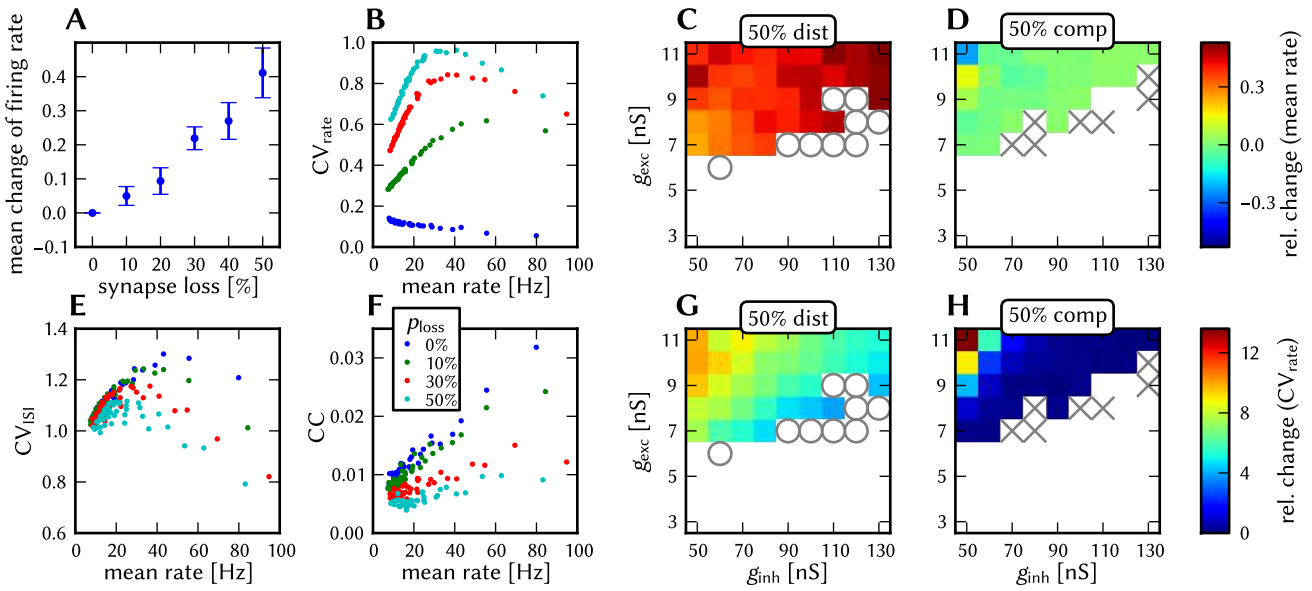

Fig. 20: Effect and compensation of synapse loss in the AI network: (**A**) Relative change of the firing rate with respect to the undistorted network averaged over all sustained states for varying synapse loss. (**B**) $\mathrm{CV}_{\mathrm{rate}}$ as a function of mean rate for every survived state for varying synapse loss. (**C**, **D**) Relative change of the firing rate with respect to the undistorted case for each state for 50 % synapse loss (**C**) and compensated (**D**). (**E**) $\mathrm{CV}_{\mathrm{ISI}}$ as a function of mean rate for varying synapse loss. (**F**) CC as a function of mean rate for varying synapse loss. (**G**, **H**) Relative change of $\mathrm{CV}_{\mathrm{rate}}$ with respect to the undistorted case for each state for 50 % synapse loss (**G**) and compensated (**H**). In **C**, **D**, **G** and **H**: A cross marks a state that was sustained in the undistorted but not sustained in the compared case. A circle marks a state, that was not sustained in the original but sustained in the compared case.

rate in a self-sustained state can be calculated as a stable, self-consistent solution of the gain function being equal to the firing rate of a single neuron:

$$\hat{\nu}(p_{\mathrm{loss}}) = f\left(N_{\mathrm{exc}}(1-p_{\mathrm{loss}})\hat{\nu}, N_{\mathrm{inh}}(1-p_{\mathrm{loss}})\hat{\nu}\right) \tag{13}$$

Here, N_{exc} and N_{inh} are the number of pre-synaptic connections of a given neuron, and p_{loss} is the modeled synapse loss value. Fig. 21 **A** shows the gain function (right-hand side of Eq. 13) of PY and INH neurons for $p_{\mathrm{loss}} = 0$ yielding the stable solution $\tilde{\nu}(0) \approx 14\,\mathrm{Hz}$ as the intersection of the $y = x$ diagonal and the gain function. Analogously, the solution for $p_{\mathrm{loss}} = 0.5$ can be determined as the intersection with the $y = 2x$ line (considering $\nu_{\mathrm{in}}(p_{\mathrm{loss}}) = p_{\mathrm{loss}} \cdot \nu_{\mathrm{in}}$). The result justifies the assumption of the mean firing rate of inhibitory and excitatory neurons being equal for $p_{\mathrm{loss}} < 0.5$.

The parameter change that is necessary to restore the original mean firing rate can be calculated using the following relationship for the time scaling of the solution of a differential equation:

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}, t) \tag{14}$$

$$\mathbf{y}(t) := \mathbf{x}(\alpha t) \tag{15}$$

$$\dot{\mathbf{y}}(t) = \alpha\dot{\mathbf{x}}(\alpha t) = \alpha\mathbf{F}(\mathbf{y}(t), \alpha t) \tag{16}$$

$$\Rightarrow \tilde{\mathbf{F}}(\mathbf{x}, t) := \alpha\mathbf{F}(\mathbf{x}, \alpha t) \tag{17}$$

Assuming that $\mathbf{x}$ is the state of the dynamic variables within a network, $\mathbf{y}$ describes a network which follows the same time dependence with the dynamics scaled by the factor α in time. As the given random cortical network shows self-sustained behavior, the transition from $\mathbf{F}$ to $\tilde{\mathbf{F}}$ requires only the modification of internal network parameters, because there is no external input (which would have to also be modified otherwise). In particular, the transition encompasses scaling τ_{m}, τ_x^{syn}, τ_{refrac}, τ_{w} and the synaptic delays by α, while leaving the conductance jump after each presynaptic PSP unchanged. α is calculated from the measured gain function (cf. Fig. 21) via

$$\alpha = \frac{\tilde{\nu}(p_{\mathrm{loss}})}{\tilde{\nu}(0)} \tag{18}$$

The resulting firing rate with and without compensation is shown in Fig. 21 **C**. The results also show that the variance of the firing rates across neurons grows with rising synapse loss due to the increasing difference in connectivity within the networks. An extension of the mean-field-based compensation to this kind of inhomogeneous connectivity would be impractical, as it requires knowledge of the actual network realization (which is available only after the mapping step) and the measurement of Fig. 21 **A** for all occuring counts of

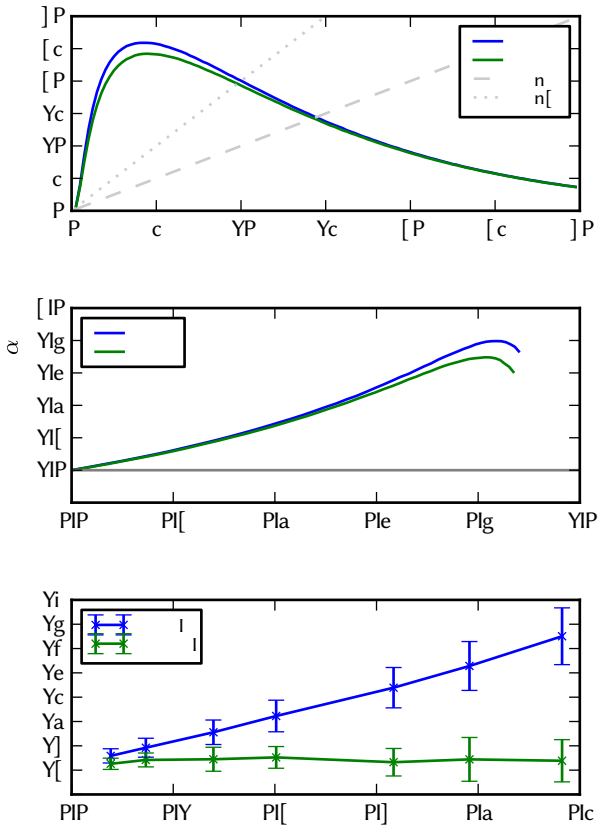

Fig. 21: Mean-field-based compensation method for the AI network. (**A**) Mean firing rate of a single PY and INH neuron given a poisson stimulus by the external network with a given rate. (**B**) Compensation factor α calculated from the data in **A**. (**C**) Compensation applied to the self-sustained network (with parameters $g_{\text{inh}} = 90\,\text{nS}$, $g_{\text{exc}} = 9\,\text{nS}$). The error bars denote the standard deviation of mean firing rates across all neurons. "orig." marks the original network without compensation, in "comp." the neuron parameters were modified according to the compensation factor. The scaling of internal delays had only minimal effect on the firing rate (not shown)

presynaptic inhibitory and excitatory neurons. Thus, a different method is considered in Sec. 3.3.6.

In conclusion, this method can be applied when the actual synapse loss value and the mean response function of a single neuron is known. It only depends on the single neuron response properties; the amount of synapse loss has to be known a priori, but not the complete network dynamics. The method depends on the ability to modify synaptic delays according to the scaling rule. However, for the given network, this scaling has only a minimal effect on the mean firing rate.

Iterative compensation The iterative compensation method aims at reducing two distortion effects: the change of the mean firing rate of the pyramidal neurons and its variance across neurons, which are both apparent for synapse loss and synaptic weight noise. It relies on the controlability of the hardware neuron parameters allowing to fine tune the AdEx parameters for every individual neuron (Sec. 2.1.1). The iterative compensation functions as follows: We start with the results of the reference and the distorted network. From the reference simulation we extract the target mean rate ν^{tgt} of the neurons in a population. For each neuron in the distorted network, we compare its actual firing rate against ν^{tgt}, and modify the excitability of the neuron in proportion to the difference between target and measured firing rate. The distorted network with modified neuron parameters is then simulated and the output is compared again to the reference network. This iterative compensation step is repeated until the characteristics of the last step approximately match those of the reference simulation. In our simulations, we modified the spike initiation threshold E_{T}, with its change $\Delta E_{\text{T}} = c_{\text{comp}}(\nu^{\text{tgt}} - \nu^{\text{act}})$ being proportional to the difference between the actual and the target rate. We found that, when choosing the compensation factor c_{comp} appropriately, 10 iterations are sufficient to restore the mean and variance of the firing rates in the undistorted network. While the compensated mean rate exactly corresponds to ν^{tgt}, the compensated CV_{rate} is higher than in the reference network, but reliably below the 1.2-fold of the reference value. The iterative compensation applied in the following is described in detail in Sec. D.3. We remark that the proposed iterative compensation requires a controllable, deterministic mapping, which guarantees that in each iteration the neurons and synapses are always mapped onto the same hardware elements. Furthermore, the complete compensation process needs to be repeated for each network instance. In fact, we perform a calibration of the apparent distortions (fixed pattern noise and synapse loss) similar to Pfeil et al (2013) in order to reduce their effects. Hence, whenever we change the random seed that is used to generate the probabilistic connectivity between the neurons, the iterative compensation needs to be run anew. Thus, a reference from a non-distorted simulation or, e.g., from theory is needed. However, once obtained, the result of the compensation can be used for long-running simulations or as part of a larger compound network.

3.3.7 Results of iterative compensation

Synaptic weight noise In order to verify the iterative compensation strategy we applied it to the distorted parameter space with 50 % synaptic weight noise. The results are shown in Fig. 19, which displays the rela-

tive difference of the mean and variance of the firing rates with respect to the reference simulation in **D** and **H**. The region of sustained activity in the (g_{exc}, g_{inh}) parameter space of the compensated network matches the one of the reference simulation very well. The mean and variance of firing rates could be successfully recovered for most of the states; with the exception of states with a mean rate higher than 25 Hz, where both criteria still differ notably from the reference after 10 iterations (upper left regions in the parameter spaces). We expect that the performance of the iterative compensation for those states could be further improved by tuning the compensation factor c_{comp} (Sec. D.3) for high firing rates. The other criteria such as $\mathrm{CV}_{\mathrm{ISI}}$ and peak frequency could be fully recovered, following the assumption made earlier, that those criteria mainly depend on the firing rate. However, the coefficient of pairwise cross-correlation (CC) of the compensated networks is lower than in the reference simulation, i.e., the randomness introduced by the fixed-pattern noise is still effective.

Synapse loss The results of the application of the iterative compensation strategy to the (g_{exc}, g_{inh}) parameter space with 50 % synapse loss are shown in Fig. 20 (**D** and **H**), displaying the relative difference of mean and variance of firing rates. The compensation was not as effective as for synaptic weight noise: Some states with a low base firing rate were unstable (marked with a cross), i.e. the network did not survive until the end of simulation. As before, the mean and variance of firing rates can be successfully restored for low and medium base firing rates. Again, for high firing rates, the iterative compensation only performed moderately (upper left regions in the parameter spaces **D** and **H**). The other criteria show the same behavior as in the weight noise compensation, i.e. the peak frequency and $\mathrm{CV}_{\mathrm{ISI}}$ are in good match with the reference while the pairwise correlation (CC) decreased due to the randomness introduced by the synapse loss. We repeated the iterative compensation for the parameter space with 30 % synapse loss: The results (not shown) are comparable to the 50 % case, but exhibit fewer unstable states, i.e., there were more combinations of g_{exc} and g_{inh} whose compensated network survived.

Conclusion We conclude that the iterative compensation of distorted networks works for both synapse loss and fixed-pattern noise. The compensation also works when both are present at the same time, see Sec. D.4.3 for details. While there seems to be no limit for fixed-pattern noise, compensation of synapse-loss induced distortions is only possible up to a certain degree, as the network tends to become less stable with fewer synapses involved.

3.3.8 Full simulation of combined distortions

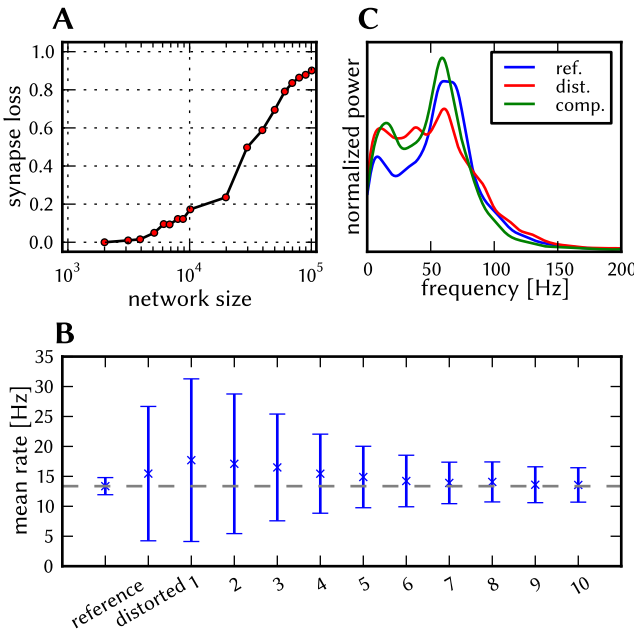

Fig. 22: AI network on the ESS: (**A**) Synapse loss after mapping the network with different sizes onto the BrainScaleS system (**B**) Iterative compensation of the large-scale network with 22 445 neurons on the ESS: evolution of mean and standard deviation of firing rates for 10 iterations (**C**) Gauss-filtered power spectrum of global activity of the pyramidal neurons in the large-scale network. Reference spectrum shown in blue (simulated with Nest), distorted and compensated spectra in red resp. green, both simulated with the ESS.

In a last step the iterative compensation method designed for the AI network was tested in ESS simulations. Like for the other two models we forced distortions to test the developed compensation strategies. Therefore, we scaled up the network such that a significant fraction of synapses is lost during the mapping process. This large-scale network was then emulated on the ESS and compared to the undistorted reference simulation with NEST. Afterwards, we applied the compensation strategy developed in the previous section to restore the original behavior of the AI network.

Synapse loss Mapping such homogeneous networks that lack any modularity represents the worst-case scenario for the mapping process, as they have little room for optimization. In Fig. 22 **A** the relative synapse loss is plotted for various network sizes using the scaling method described in Sec. D.1.2. One can see that already for

low numbers of neurons some synapse loss occurs, although there are sufficient hardware synapses and synapse drivers: due to the sparseness of the on-wafer routing switches some routing buses don't find a free switch to connect to its respective target HICANNs, such that synapses are lost. A kink in the graph of the synapse loss can be seen at around 20 000 neurons, where at least 64 neurons are mapped onto one HICANN (cf. Tab. 6). In such a network with random connectivity it is merely possible to find 64 neurons whose pool of pre-synaptic neurons is smaller than 14 336, which is the maximum number of pre-synaptic neurons per HICANN, such that synapse loss must occur. Recall that there is a maximum of 14 336 pre-synaptic neurons for all neurons mapped onto one HICANN. As the connectivity in the AI network is probabilistic, the chance to find groups of 64 neurons whose pool of pre-synaptic neurons is smaller than 14 336 is close to zero.

Large-scale network In order to produce a demanding scenario, we scaled the model to a size of 22 445 neurons (Sec. D.1.2). The size was chosen such that the network almost occupies an entire wafer, while mapping up to 64 neurons onto one HICANN. This large-scale network has a total of approximately 5.6 million synapses. The statistics of the reference simulation can be found in Tab. 4 and are in accordance with the scaling behavior investigated in the Appendix, Sec. D.4.1.

Distorted network In the above scenario, 28.1 % of synapses were lost during the mapping process (for projection-wise numbers see Tab. 3). We remark that the synapse loss at this size is higher than during the synapse loss sweep in Fig. 22 **A**, as we used a sequence of mapping algorithms that guarantees a balance between synapse loss of excitatory and inhibitory connections. Still, there were slightly more inhibitory connections lost than excitatory ones (Tab. 3). Additionally, we applied a fixed-pattern noise of 20 % to the synaptic weights in the ESS simulation. The result of the latter can be found in Tab. 4: the network still survived until the end of the simulation, but the firing rate and its variance increased compared to the reference simulation, which complies with the prediction of the distortion analysis.

Compensated network We then used the iterative compensation method from Sec. 3.3.6 to compensate the abovementioned distortions and repeated the ESS simulation with the modified network. The evolution of the firing rates over 10 iterations is shown in Fig. 22 **B**: One can clearly see how, step by step, the firing rate approaches the target rate and that at the same time the variance of firing rates decreases. The statistics of the final iteration are listed in Tab. 4: It was possible to fully recover the target mean rate. The variation of firing across neurons (CV_{rate}) was significantly reduced from 0.726 to 0.212 but was still twice as large as in the reference network. The other functionality criteria match the reference simulation very well (Tab. 4), as does the power spectrum of global activity in Fig. 22 **C**.

Table 3: Projection-wise synapse loss of the large-scale AI network after the mapping process. PY: excitatory pyramidal neurons. INH: fast spiking inhibitory cells, STIM: external Poisson sources for initial stimulation

projection	synapse loss [%]
PY → PY	26.9
PY → INH	28.1
INH → PY	31.1
INH → INH	33.4
STIM → PY	77.5
STIM → INH	89.4
total	28.1

Table 4: Statistics of the large-scale AI network: reference (ref.) simulated with NEST, distorted (dist.) and compensated (comp.) with the ESS

criteria	ref.	dist.	comp.
Rate [Hz]	13.4	15.5	13.6
CV_{rate}	0.107	0.726	0.212
CV_{ISI}	1.12	1.11	1.09
CC	0.00103	0.0011	0.00166
Peak Frequency[Hz]	60.3	60.7	59.0

4 Conclusions

In this study, we have presented a systematic comparison between neural network simulations carried out with ideal software models and a specific implementation of a neuromorphic computing system. The results for the neuromorphic system were obtained with a detailed simulation of the hardware architecture. The core concept is, essentially, a functionalist one: neural networks are defined in terms of functional measures on muliple scales, from individual neuron behavior up to network dynamics. The various neuron and synapse parameters are then tuned to achieve the target performance in terms of these measures.

The comparison was based on three cortically inspired benchmark networks: a layer 2/3 columnar architecture, a model of a synfire chain with feed-forward inhibition and a random network with self-sustained, irregular firing activity. We have chosen these specific network architectures for two reasons. First of all, they

implement very different, but widely acknowledged computational paradigms and activity regimes found in neocortex: winner-take-all modules, spike-correlation-based computation, self-sustained activity and asynchronous irregular firing. Secondly, due to their diverse properties and structure, they pose an array of challenges for their hardware emulation, being affected differently by the studied hardware-specific distortive effects.

All three networks were exposed to the same set of hardware constraints and a detailed comparison with the ideal software model was carried out. The agreement was quantified by looking at several chosen microscopic and and macroscopic observables on both the cell and network level, which we dubbed "functionality criteria". These criteria were chosen individually for each network and were aimed at covering all of the relevant aspects discussed in the original studies of the chosen models.

Several hardware constraint categories have been studied: the dynamics of the embedded neuron and synapse models, limited parameter ranges, synapse loss due to limited hardware resources, synaptic weight deviations due to fixed-pattern noise and the lack of configurable axonal delays. The final three effects were studied in most detail, as they are expected to affect essentially every hardware-emulated model. The investigated distortive effects were studied both individually, as well as combined, similarly to the way they would occur on a real hardware substrate. As expected, above certain magnitudes of the hardware-induced distortions, substantial deviations of the functionality criteria were observed.

For each of the three network models and for each type of distortion mechanism, several compensation strategies were discussed, with the goal of tuning the hardware implementation towards maximum agreement with the ideal software model. With the proposed compensation strategies, we have shown that it is possible to considerably reduce, and in some cases even eliminate the effects of the hardware-induced distortions. We therefore regard this study as an exemplary workflow and a toolbox for neuromorphic modelers, from which they can pick the most suitable strategy and eventually tune it towards their particular needs.

In addition to the investigated mechanisms, several other distortive effects are routinely observed on neuromorphic hardware. A (certainly not exhaustive) list might include mismatch of neuron and synapse parameters, shared parameter values (i.e., not individually configurable for each neuron or synapse) or limited parameter programming resolution. These mechanisms are highly back-end-specific and therefore difficult to generalize. However, although they are likely to pose individual challenges by themselves, some of their ultimate effects on the target network functionality can be alleviated with the compensation strategies proposed here.

Our proposed strategies aim at neuromorphic implementations that compete in terms of network functionality with conventional computers but offer major potential advantages in terms of power comsumption, simulation speed and fault tolerance of the used hardware components. If implemented successfully, such neuromorphic systems would serve as fast and efficient simulation engines for computational neuroscience. Their potential advantages would then more than make up for the overhead imposed by the requirement of compensation.

From this point of view, hard-ware-induced distortions are considered a nuisance, as they hinder precise and reproducible computation. In an alternative approach, one might consider the performance of the system itself at some computational task as the "fitness function" to be maximized. In this context, some particular architecture of an embedded model, together with an associated target behavior, would then become less relevant. Instead, one would design the network structure specifically for the neuromorphic substrate or include training algorithms that are suitable for such an inherently imperfect back-end. The use of particular, "ideal" software models as benchmarks might then given up altogether in favor of a more hardware-oriented, stand-alone approach. Here, too, the proposed compensation strategies can be actively embedded in the design of the models or their training algorithms.

The hardware architecture used for our studies is, indeed, suited for both approaches. It will be an important aspect of future research with neuromorphic systems to develop procedures that tolerate or even actively embrace the temporal and spatial imperfections inherent to all electronic circuits. These questions need to be addressed by both model and hardware developers, in a common effort to determine which architectural aspects are important for the studied computational problems, both from a biological and a machine learning perspective.

Acknowledgements We would like to thank Eric Müller for his invaluable help with the software infrastructure and Jens Kremkow for his support with the synfire chain model. This research was supported by EU grant #269921 (BrainScaleS) and the Manfred Stärk Foundation.

Appendix A Neuromorphic hardware

A.1 Short-term plasticity

As mentioned in Sec. 2.1.1, the hardware short-term plasticity mechanism is an implementation of the phenomenological model by Markram et al (1998). We first describe the hardware STP model and then provide the translation between the original and the hardware model.

Model description Unlike the theoretical model (Markram et al, 1998), which allows the occurrence of both depression and facilitation at the same time, the hardware implementation does not allow their simultaneous activation. The ongoing pre-synaptic activity is tracked with a time-varying active partition I with $0 \leq I \leq 1$, which decays exponentially to zero with time constant τ_{stdf}. Following a pre-synaptic spike, I is increased by a fixed fraction $U_{\text{SE}}(1 - I)$, resulting in the following dynamics for the active partition:

$$I_{n+1} = [I_n + U_{\text{SE}}(1 - I_n)] \exp\left(-\frac{\Delta t}{\tau_{\text{stdf}}}\right) \quad , \tag{19}$$

with Δt being the time interval between the nth and $(n+1)$st afferent spike.

This active partition can be used to model depressing or facilitating synapses as follows:

$$w_{\text{STP}}^{\text{depression}} = 1 - \lambda \cdot I \tag{20}$$

$$w_{\text{STP}}^{\text{facilitation}} = 1 + \lambda \cdot (I - \beta) \quad . \tag{21}$$

Here, w_{STP}^{x} corresponds to a multiplicative factor to the static synaptic weight, with λ and β being configurable variables, and x denotes the mode being either depression or facilitation.

According to Eq. 8 the n-th effective synaptic weight is then given by

$$w_n^{\text{syn}} = w_{\text{static}} w_{\text{STP}}^{x} \quad . \tag{22}$$

Due to a technical limitation, the change of synaptic weights by STP can not be larger than the static weight, such that $0 \leq w_{\text{STP}}^{x} \leq 2$. We refer to Schemmel et al (2008) for details of the hardware implementation of STP and to Bill et al (2010) for neural network experiments on neuromorphic hardware using this STP model.

Table 5: Parameter ranges of the BrainScaleS wafer-scale hardware. All ranges correspond to a membrane capacitance of $C_{\text{m}} = 0.2\,\text{nF}$ and a hardware speedup of 10^4 compared to real time. It is possible to choose an arbitrary value for C_{m}, but then the ranges of parameters a, b and of the synaptic weights are multiplied by $\frac{C_{\text{m}}}{0.2\,\text{nF}}$.

Description	Name	Min	Max	Unit	Comment
Neuron (Adaptive Exponential Integrate&Fire)					
Absolute refractory period	τ_{refrac}	0.16	10.0	ms	
Spike detection potential	E^{spike}	-125.0	45.0	mV	
Reset potential	E^{r}	-125.0	45.0	mV	
Leakage reversal potential	E_{L}	-125.0	45.0	mV	
Membrane time constant	τ_{m}	9	105	ms	
Adaptation coupling param	a	0	10.0	nS	adaptation can be fully disabled
Spike triggered adapt. param	b	0	86	pA	
Adaptation time constant	τ_w	20.0	780.0	ms	
Threshold slope factor	Δ_{T}	0.4	3.0	mV	exponential spike generation can be
Spike initiation threshold	E_{T}	-125.0	45.0	mV	fully disabled
Excitatory reversal potential	$E^{\text{rev,e}}$	-125.0	45.0	mV	
Inhibitory reversal potential	$E^{\text{rev,i}}$	-125.0	45.0	mV	
Exc. synaptic time constant	$\tau^{\text{syn,e}}$	1.0	100.0	ms	
Inh. synaptic time constant	$\tau^{\text{syn,i}}$	1.0	100.0	ms	
Synapses					
Weight	w^{syn}	0	0.300	µS	4-bit resolution
Axonal delay (on-wafer)	delay	1.2	2.2	ms	not configurable
Short Term Plasticity					
Utilization of synaptic efficacy	U	0.11	0.47		possible values: $[\frac{1}{9}, \frac{3}{11}, \frac{5}{13}, \frac{7}{15}]$
Recovery time constant	τ_{rec}	40.0	900.0	ms	One of the two time constants has to
Facilitation time constant	τ_{facil}	35.0	200.0	ms	be set to 0.0. Available range depends on U (maximum range given).
Stimulus					
External spike sources	ν	0.0	4000	Hz	cf. Scholze et al (2011b)

Transformation from original model The original model by Markram et al (1998) (Eq. 8) can be translated to the hardware model (Eq. 19 to 21) when one of the two time constants (τ_{rec} or τ_{facil}) is equal to zero.

For depression only ($\tau_{\text{facil}} = 0$), the nth synaptic weight is given by (cf. Eq. 8):

$$w_n^{\text{syn}} = w_{\text{max}}^{\text{syn}} R_n U \quad . \tag{23}$$

The time course of R can be exactly represented by $(1 - I)$ if the scaling factor λ of the short-term plasticity mechanism is set to 1. Additionally, the static synaptic weight w_{static} has to be adapted such that the applied synaptic weights are equal, giving us the following transformation: $\tau_{\text{stdf}} = \tau_{\text{rec}}$, $U_{\text{SE}} = U$, $\lambda = 1$ and $w_{\text{static}} = w_{\text{max}}^{\text{syn}} U$.

For facilition only ($\tau_{\text{rec}} = 0$), the recovered partition remains fully available all the time ($R = 1 = \text{const}$) and only the utilization varies with time. Thus the nth synaptic weight is given by:

$$w_n^{\text{syn}} = w_{\text{max}}^{\text{syn}} u_n \quad . \tag{24}$$

The time course of u now has to be emulated by the right-hand side of Eq. 21; more precisely, we use I to represent the course of $u-U$. Additionally we set $U_{\text{SE}} = U$ and $\tau_{\text{stdf}} = \tau_{\text{facil}}$, and level the limits for the synaptic weights. In the original model, u is always between U and 1, while for the hardware model the STP factor is limited to values between 0 and 2 due to technical reasons. By setting $\lambda = 1$ and considering that I is always within 0 and 1, the supplied range for $w_{\text{STP}}^{\text{facilitation}}$ is $[1 - \beta, 2 - \beta]$. In order to match the range of applied weights of both models, we need to solve the following system of equations:

$$\begin{aligned} (1 - \beta) \cdot w_{\text{static}} &= U \cdot w_{\text{max}}^{\text{syn}} \\ (2 - \beta) \cdot w_{\text{static}} &= 1 \cdot w_{\text{max}}^{\text{syn}} \quad . \end{aligned}$$

Solving for w_{static} and β yields

$$\begin{aligned} w_{\text{static}} &= (1 - U) \cdot w_{\text{max}}^{\text{syn}} \\ \beta &= \frac{1 - 2U}{1 - U} \quad . \end{aligned}$$

A.2 Parameter ranges

Here, we provide a full list of available parameter ranges for the BSS waferscale platform in Tab. 5. As mentionend in Sec. 2.1.1, one has the choice between two different capacitances in the hardware neuron. The parameter ranges specified in Tab. 5 correspond to the big capacitance (2.6 pF). When using the small capacitance (0.4 pF) some parameter ranges change: the limits of τ_{m} are multiplied by $\frac{0.4}{2.6}$, the ranges for a, b, and the synaptic weight are divided by $\frac{0.4}{2.6}$.

In Tab. 6 we show how the tradeoff between total neuron number and maximum fan-in per neuron is realized on this device.

Table 6: List of typical usage scenarios of the wafer-scale hardware system. One can either opt for many neurons with few synapses or for fewer neurons but a higher connection density.

Nr of Neurons	Synapses/ Neuron	DenMems/ Neuron	Neurons/ HICANN
196 608	224	1	512
98 304	448	2	256
49 152	896	4	128
24 576	1792	8	64
12 288	3584	16	32
6144	7168	32	16
3072	14 336	64	8

Appendix B Cortical layer 2/3 attractor memory

B.1 Original model parameters

In Tab. 7, we summarize the parameters and characteristics of the original model, as found in Lundqvist et al (2006). These have served as the basis for the model fit, for which the parameters can be found in the next subsection.

B.2 Fitted Hardware-Compatible Parameters

Tab. 11, Tab. 12 and Tab. 13 contain all parameters required for the fits described in Sec. 3.1.3.

The diffuse background stimulus was generated by Poisson spike sources at a total rate of 300 Hz per PYR cell.

Apart from random noise, the PYR cells further receive input from other PYR cells in cortical layer 4. The input intensity was calculated from the number of cells in layer 4 likely to project onto layer 2/3, which was estimated to be around 30 with a rate of approximately 10 Hz and a connection density of 25 % (Lundqvist et al, 2006).

Therefore, a Poisson process with 75 Hz was used for each PYR cell input. Since we used static synapses for the Poisson input, the synaptic weights for source-PYR connections were chosen as 30% of PYR-PYR connections within the MCs. This was verified for compliance the original model from (Lundqvist et al, 2006), which uses 7 to 8 sources per stimulated PYR cell with a rate of 10 Hz each and depressing synapses. For each stimulus event in the pattern completion and rivalry experiments (described below), layer 4 cells were set to fire for 60 ms. In each stimulated MC, 6 PYR cells were targeted from layer 4.

Table 10: Original network structure: connection probabilities

within an MC	
PYR → PYR	0.25
RSNP → PYR	0.70
between MCs inside the same HC	
PYR → BAS	0.70
BAS → PYR	0.70
between MCs in different HCs	
PYR → PYR	0.30
PYR → RSNP	0.17

Table 13: Stimulus parameters for the L2/3 model.

Background	
# of sources per PYR	1
rate	300 Hz
weight	0.000 224 µS
Shared background pool	
# of sources per PYR	100 out of 5000 total
rate	3 Hz
weight	0.000 224 µS
L4	
# of sources per MC	5
$p_{\mathrm{L4}\to\mathrm{PYR}}$	0.75
weight	0.001 237 5 µS (30% local PYR→PYR)

B.3 Delays

Each connection within the same MC was set to have constant synaptic delay of 0.5 ms. Additionally, axonal delays for connections between different MCs were realized by taking into account their spatial distance at an average axonal propagation speed of 200 µm/ms. Both the HCs in the whole network as well as the MCs within a single HC are laid out on a hexagonal grid with a edge length of 500 µm (HC↔HC) / 60 µm (MC↔MC). In the default network (9HC×9MC) this leads to delays between 0.5 ms and 8 ms.

B.4 Scaling

Due to the modularity of this network model, several straightforward possibilities exist for increasing or decreasing its size without affecting its basic functionality. One can vary the total number of neurons simply by modifying the number of cells per MC. One can also vary the number of MCs per attractor by varying the total number of HCs. And finally, one can change the number of attractors by changing the number of MCs per HC accordingly.

All such changes need to be accompanied by corresponding modifications in connectivity in order to preserve the network dynamics. This has been done by keeping the average input current per neuron within an active attractor constant, which is equivalent to conserving the fan-in for every neuron from every one of its afferent populations and leads to the scaling rules shown in Tab. 15. In order to facilitate a comparison with the original results from Lundqvist et al (2006) and Lundqvist et al (2010), we have only considered homogeneous changes, meaning that all modules (MCs, HCs) were equal in size and symmetrically connected.

Table 7: Original neuron parameters

Parameter	PYR	RSNP	BAS	Unit
g_{ext}	0.082	0.15	0.15	$\mu S/mm^2$
E_{leak}	-75	0.15	-75	mV
E_{Na}	50	50	50	mV
E_{Ca}	150	150	150	mV
E_K	-80	-80	-80	mV
$E_{Ca(NMDA)}$	20	20	20	mV
g_L	0.74	0.44	0.44	$\mu F/mm^2$
C_m	0.01	0.01	0.01	$\mu F/mm^2$
Soma diameter ± stdev	21 ± 2.1	7 ± 0.7	7 ± 0.7	μm
g_{Na} initial segment	2500	2500	2500	$\mu S/mm^2$
g_K initial segment	83	5010	5010	$\mu S/mm^2$
g_{Na} soma	150	150	150	$\mu S/mm^2$
g_K soma	250	1000	1000	$\mu S/mm^2$
g_{NMDA}	75.0	75.0	-	$\mu S/mm^2$
Ca_V influx rate	1.00	1.00	1.00	$mV^{-1}ms^{-1}mm^{-2}$
Ca_{NMDA} influx rate	2.96	0.0106	-	$s^{-1}mV^{-1}\mu S^{-1}$
Ca_V decay rate	6.3	4	-	s^{-1}
Ca_{NMDA} decay rate	1	1	-	s^{-1}
g_K (Ca_V)	29.4	105	0.368	nS
g_K (Ca_{NMDA})	40	40	-	nS
# compartments	6	3	3	
Dendritic area (relative soma)	4	4	4	
Initial segment area (relative soma)	0.1	0.1	0.1	

Table 8: Original synapse parameters

Pre → Post	Type	Duration [s]	τ_{raise} [s]	τ_{decay} [s]	E^{rev} [mV]	U	τ_{rec} [s]	[mV]
PYR → PYR	Kainate/AMPA	0.0	0.0	0.006	0	0.25	0.575	-
PYR → PYR	NMDA	0.02	0.005	0.150	0	0.25	0.575	0.020
PYR → BAS	Kainate/AMPA	0.0	0.0	0.006	0	-	-	-
PYR → RSNP	Kainate/AMPA	0.0	0.0	0.006	0	-	-	-
PYR → RSNP	NMDA	0.02	0.005	0.150	0	-	-	0.020
BAS → PYR	GABA	0.0	0.0	0.006	-85	-	-	-
RSNP → PYR	GABA	0.0	0.0	0.006	-85	-	-	-

Table 9: Original network structure: number of neurons per functional unit

	HCs	MCs	PYR	BAS	RSNP	total neurons
per MC	-	-	30	1	2	33
per HC	-	8	240	8	16	264
network total	9	72	2160	72	144	2376

The connections to the BAS cells required special treatment for two reasons. Firstly, during an active state, they receive input from a single MC, but are excited by all MCs in a HC during the competition period between active attractors. Only one aspect can be preserved when scaling and we have considered the dynamics during UP states as most important, leading to a "PYR → BAS" scaling rule independent of N_{MC}. Secondly, because PYR cells in MCs only project to the nearest 8 BAS cells, there are always precisely 8 active BAS cells per HC within an active attractor, which yields a simple "BAS → PYR" scaling rule. When decreasing the number of attractors however, the number of existing BAS cells per HC also decreases, making an appropriate connection density scaling necessary. This is the reason for the two different "BAS → PYR" scaling rules found in Tab. 15.

Tab. 16 shows the combinations of N_{HC} and N_{MC} used for the quantification of synapse loss after mapping the L2/3 model onto the hardware in Fig. 10. In these mapping sweeps the diffusive background noise was modeled, as for the large-scale network ported to the ESS (Sec. 3.1.7), with a background pool of 5000 Poisson sources and every PYR cell receiving input from 100 of the sources.

Table 11: Fitted neuron parameters for the L2/3 model.

Parameter	PYR	RSNP	BAS	Unit	Comment
C_{m}	0.179	0.0072	0.00688	nF	from the fits in Fig. 23 **A-C**
$E^{\mathrm{rev,e}}$	0.0	0.0	0.0	mV	difference to original model compensated by synaptic weights
$E^{\mathrm{rev,i}}$	-80.0	-	-	mV	difference to original model compensated by synaptic weights
τ_{m}	16.89	15.32	15.64	ms	from the fits in Fig. 23 **A-C**
τ_{refrac}	0.16	0.16	0.16	ms	minimum available in hardware at the used speedup
$\tau^{\mathrm{syn,e}}$	17.5	66.6	6.0	ms	see paragraph "Synapses"
$\tau^{\mathrm{syn,i}}$	6.0	-	-	ms	see paragraph "Synapses"
V_{reset}	-60.7	-72.5	-72.5	mV	from the fits in Fig. 23 **D-F**
E_{L}	-61.71	-57.52	-56.0	mV	from the fits in Fig. 23 **D-F**
a	0.0	0.28	0.0	nS	see fig from the fit in Fig. 23 **B**
b	0.0132	0.00103	0.0	nA	from the fits in Fig. 23 **D**, **E**
ΔT	0.0	0.0	0.0	mV	from the fits in Fig. 23 **D-F**
τ_w	196.0	250.0	0.0	ms	from the fits in Fig. 23 **D**, **E**
E^{spike}	-53.0	-51.0	-52.5	mV	from the fits in Fig. 23 **D-F**
V_T	-	-	-	mV	not used since $\Delta T = 0$

Table 12: Fitted synapse parameters for the L2/3 model.

Pre-Post	type	weight [μS]	τ^{syn} [ms]	U	τ_{rec} [ms]	τ_{facil} [ms]
PYR-PYR (local)	exc	0.004125	17.5	0.27	575.	0.
PYR-PYR (global)	exc	0.000615	17.5	0.27	575.	0.
PYR-BAS	exc	0.000092	6.0	-	-	-
PYR-RSNP	exc	0.000024	66.6	-	-	-
BAS-PYR	inh	0.0061	6.0	-	-	-
RSNP-PYR	inh	0.0032	6.0	-	-	-
background-PYR	exc	0.000224	17.5	-	-	-

Table 14: Average firing rates (in Hz) of the different cell types of the L2/3 model with only certain synapses active.

setup no.	active synapses	ν_{PYR}	ν_{RSNP}	ν_{BAS}
1	background-PYR, PYR-BAS, PYR-RSNP	0.738 ± 0.096	57.946 ± 6.993	4.655 ± 1.081
2	same as 1 + BAS-PYR	0.174 ± 0.021	13.430 ± 1.910	1.119 ± 0.441
3	same as 1 + RSNP-PYR	0.257 ± 0.037	20.375 ± 2.536	1.783 ± 0.954
4	same as 2 + 3 + PYR-PYR (local)	0.200 ± 0.030	14.679 ± 2.261	1.258 ± 0.544
5	same as 2 + 3 + PYR-PYR (global)	0.204 ± 0.078	14.954 ± 5.680	1.337 ± 0.625

Table 15: Scaling rules for the connection densities of the L2/3 model. N_x represents the number of units of type x (the original values are found in Tab. 9). p represents the original connection probability as found in table Tab. 10. Whenever a scaled probability $\tilde{p}$ exceeded 1, it was clipped to 1, but the weights of the corresponding synapses were also increased by $w^{\tilde{\mathrm{syn}}} = w^{\mathrm{syn}} \cdot \tilde{p}$.

Connection	Scaled conn. prob. $\tilde{p}$
PYR → PYR (same MC)	$29/(N_{\mathrm{PYR}} - 1) \cdot p$
PYR → PYR (different MC)	$30/N_{\mathrm{PYR}} \cdot 8/(N_{\mathrm{HC}} - 1) \cdot p$
PYR → RSNP	$30/N_{\mathrm{PYR}} \cdot 8/(N_{\mathrm{HC}} - 1) \cdot p$
PYR → BAS	$30/N_{\mathrm{PYR}} \cdot p$
RSNP → PYR	$2/N_{\mathrm{RSNP}} \cdot p$
BAS → PYR (Enlarging)	$1/N_{\mathrm{BAS}} \cdot p$
BAS → PYR (Shrinking)	$1/N_{\mathrm{BAS}} \cdot 8/N_{\mathrm{MC}} \cdot p$

B.5 UP-state detection

One crucial element of the analysis is the detection of UP-states from which various other properties such as dwell times, competition times as well as average spike-rates in UP- and DOWN-states are determined. The method of choice for detecting UP-states is based on the fact that the mean spike rate of an attractor during an UP-state is much higher than the spike rate in all remaining patterns in their corresponding DOWN-states, whereas – in times of competition – two or more attractors have elevated but rather similar spike rates. A measure which quantifies this relationship is the standard deviation σ of all mean spike rates per attractor *at a given time* t. The attractor with index i is then said to be in an UP-state at time t if the following relation

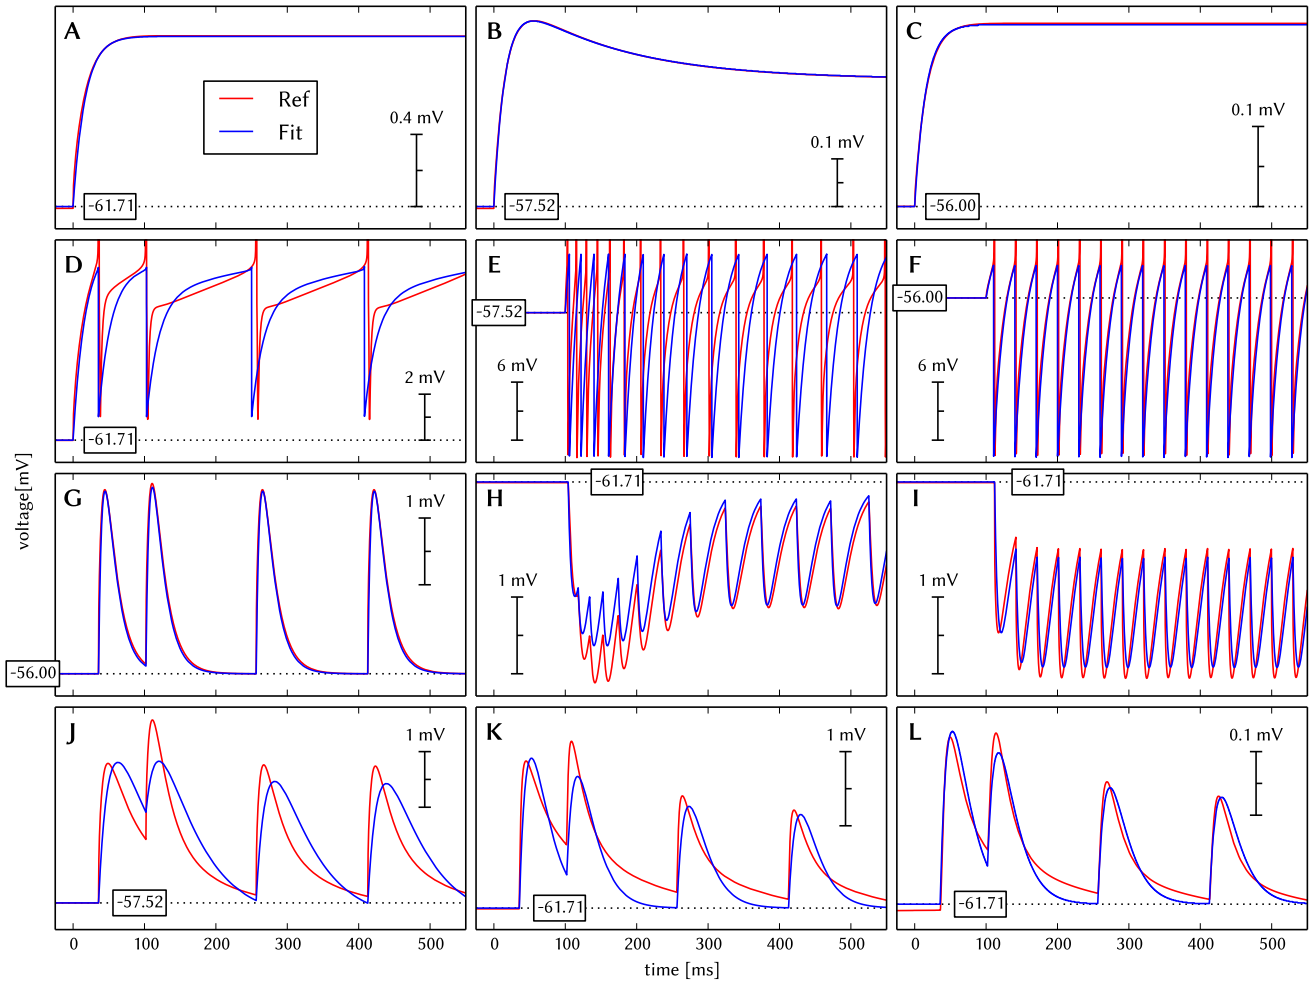

Fig. 23: Comparison of original neuron and synapse dynamics to the fitted dynamics of hardware-compatible models. (**A - C**) Membrane potential of the three different cell types (PYR, RSNP, and BAS, respectively) under subthreshold current stimulation. These were used to determine the rest voltage E_{L}, total equivalent membrane capacitance C_{m}, membrane time constant τ_{m} and the adaptation coupling parameter a. (**D - F**) Membrane potential of the three different cell types (PYR, RSNP, and BAS, respectively) under spike-inducing current stimulation. While the precise membrane potential time course of the original neuron model can not be reproduced by a single-compartment AdEx neuron, it was possible to reproduce the spike timing and especially average firing rates quite accurately. A small deviation of spiking frequency can be observed for RSNP cells during the first 50 ms - in the original model, they adapt slower than their AdEx counterparts. From these fits, the values for the absolute refractory period τ_{refrac}, reset voltage V_{reset}, threshold voltage V_T, slope factor ΔT, spike-triggered adaptation b and adaptation time constant τ_w were extracted. (**G - L**) PSP fit results for all synapse types of the L2/3 model (PYR→BAS, RSNP→PYR, BAS→PYR, PYR→RSNP, PYR→PYR within a MC, and PYR→PYR between MCs, respectively). The output spikes from **D - F** have been used as input. These fits were used to determine synaptic weights w^{syn}, time constants τ^{syn} and the TSO parameters U and τ_{refrac}. Because the hardware synapses only support a single conductance decay time constant, as opposed to the two different time constants in the original model for AMPA/kainate and NMDA, we have chosen an intermediate value for τ^{syn}, which constitutes the main reason for the difference in PSP shapes. A second reason lies in the saturating nature of synaptic conductances in the original model, which can not be emulated on the hardware without affecting the required TSO parameters (see Sec. 3.1.3).

holds:

$$r_i(t) > c \cdot \sigma(t) > \max_{r \in \{1,\dots,N_{\mathrm{MC}}\} \setminus i} r_k(t) \quad , \tag{25}$$

where $r_i(t)$ is the rate of attractor i at time t and c is a numerical constant which is set to 1.

This method of detection has several advantages: it is based exclusively on spike trains (and not voltages or conductances, which are more difficult to read out and require much more storage space), it has a clear notion of there being at most one UP-state at any given time and it is completely local (in time), meaning that a very large value somewhere on the time axis cannot bias the detection at other times.

Table 16: Scaling table for the L2/3 model used for the synapse loss estimation in Fig. 10

N_{HC}	N_{MC}	total neurons
18	2	1188
9	6	1782
27	3	2673
18	6	3564
36	4	4752
9	18	5346
18	12	7128
27	9	8019
18	18	10 692
18	36	21 384
36	24	28 512
36	36	42 768
27	54	48 114
45	45	66 825

In small networks with randomly spiking neurons, it might happen by chance that all but one of the spike rates lie below the (approximately) constant standard deviation. These falsely detected UP-states are very short and can thus easily be filtered out by requiring a minimal duration for UP states, which we chose at $100\,ms$. This value was chosen after investigating dwell time histograms, as it distinguishes reliably between random fluctuations and actual active attractors.

B.6 Pattern Completion

Pattern completion is a basic property of associative-memory networks. By only stimulating a subset of PYR cells pwithin a pattern, the complete pattern is recalled. The activity first spreads within the stimulated MCs, turning them dominant in their corresponding HCs. After that, the activity spreads further to other HCs – while the already dominating MCs stabilize each other through mutual stimulation – activating the whole pattern while suppressing all others. All PYR cells in the corresponding attractor hence enter an UP-state.

To verify the pattern completion ability of the network, a series of simulations was performed. In order to reduce the occurrence of spontaneously activating attractors – which would interfere with the activation of the stimulated attractor – competition was investigated in larger networks of size 25HC×25MC, as they exhibit almost no spontaneous attractors (the competition time fractions are much higher, see Fig. 5 **H**).

For each network, all of the 25 patterns were stimulated one by one in random order. The time between consecutive stimuli was chosen to be 1000 ms to ensure minimal influence between patterns. The number of stimulated MCs (one per HC) was varied over the course of multiple simulations.

After simulation, each network was analyzed for successfully activated patterns. An activation attempt was said to be successful if the stimulated pattern was measured as active within 200 ms after the stimulus onset. If another pattern was active up to 75 ms or if the stimulated pattern had already been active between 20 − 500 ms prior to the stimulus onset, the attempt was deemed invalid and ignored during the calculation of success ratios. This was done to take into account the fact that a pattern is more difficult to activate when another one is already active or while it is still recovering from a prior activation. From all valid attempts the success probability (assuming a binomial distribution of successful trials) was estimated using the Wilson interval

$$\tilde{p} = \frac{1}{1+\frac{z^2}{n}}\left[\hat{p} + \frac{z^2}{2n} \pm z\sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z^2}{4n^2}}\right] \tag{26}$$

where $\hat{p}$ represents the success ratio, n the number of valid attempts and $z = 1$ the desired quantile.

For most experiments (regular, synaptic weight noise and homogeneous synaptic loss) the number of invalid activations was always below 5 (out of 25). The only exception was the PYR population size scaling: starting at 15 PYR cells, the validity rate roughly halves for every reduction in size (by 5 PYR cells per step) due to the increased occurrence of spontaneous attractors. For simulations carried out on the ESS, only 10 patterns out of 25 were stimulated. Out of these 10 attempts, only 5 were valid, on average.

B.7 Pattern Rivalry / Attentional Blink

Another important feature of the L2/3 model is its ability to reproduce the attentional blink phenomenon, i.e., the inability of one pattern, stimulated by layer 4 input, to terminate another already active pattern and become active itself. This phenomenon was investigated through a series of different networks of same size as in Sec. B.6 (25HC×25MC). For each network, 24 out of 25 patterns were randomly assigned to 12 pairs. Afterwards, pattern rivalry was tested on all of these pairs in intervals of 1000 ms.

Let the two patterns in each pair be denoted A and B. In order to guarantee an immediate activation of pattern A, 6 out of 25 HCs were stimulated (as then all completion attempts are successful, see Fig. 5 **N**). Then, after a certain delay ΔT, pattern B was stimulated with a varying amount of HCs. Both the number of stimulated HCs as well as the delay ΔT were varied for each network.

The same way as in Sec. B.6, each network was then analyzed as to whether pattern B was successfully activated or not. If the competing pattern B was activated within 200 ms after the stimulus onset and stayed active for at least 100 ms, the attempt was counted as successful, otherwise it was deemed unsuccessful. As before, attempts during which spontaneously activated patterns intervened were ignored. From all successful and unsuccessful attempts, the success probability was then estimated the same way as in pattern completion, using Eq. 26.

The validity ratios for pattern rivalry are not significantly different from those discussed in Sec. B.6. Most experiments (regular, synaptic weight noise and homogeneous synaptic loss) have 10 to 12 valid attempts (out of 12). As before, for the PYR population size scaling experiments, the number of valid attempts dropped progressively (8.2 ± 1.7, 4.8 ± 2.1 and 2.2 ± 1.5 valid attempts for 15, 10 and 5 PYR per MCs respectively). Simulations carried out on the ESS had an average of 4 (distorted case) and 6 (compensated case) valid attempts (out of 10).

Different network configurations have been compared in terms of *attentional blink* by estimating the 0.5 iso-probability contour in the following way. For every delay ΔT, the transition point from below to above 0.5 probability for successful activation of the second pattern was estimated by linearly interpolating between the two nearest data points with a success ratio of above and below 0.5, respectively. In case there were several such transition points only the one with the highest stimulus was considered. If no transition point could be identified, the transition was fixed at at either 25 or 0 stimulated MCs, depending on whether all success ratios were above or below 0.5. When there were no valid attempts for a certain delay/stimulus pair, its success probability estimate was replaced by the median of all valid activation attempts for that particular time delay ΔT (this only occurred sporadically in ESS and PYR population size scaling with less than 15 PYR cells per MC). After identifying the transition point for every time delay ΔT, intermediate values were interpolated linearly. Finally, the interpolated values were Gauss-filtered ($\mu = 0.25 \times$ step size for ΔT in the dataset) to better approximate the true 0.5 iso-probability contour.

B.8 Star plots

While the spiking activity of many cells can be visualized quite well in raster plots, illustrating the temporal evolution of their membrane potentials is less straightforward. Here, we have chosen to use so-called star plots

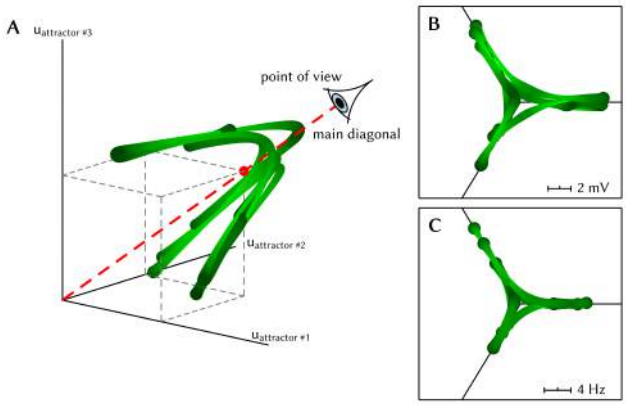

Fig. 24: Visualization of the star plot as a projection in the case of a three-dimensional state space. (**A**) Illustration of the view point with average membrane voltage data plotted in three-dimensional Cartesian coordinates. The data was taken from a (9HC×3MC)-network and covers a 2.5 s period of network activity. (**B**) Resulting star plot from regular view point. (**C**) Star plot of the corresponding average attractor rate data.

for visualizing both average voltages and average firing rates of entire cell populations.

In a system evolving in an abstract space with 3 dimensions, a star plot represents the orthogonal projection of the state space trajectory along the main diagonal of the corresponding Cartesian coordinate system onto a plane perpendicular to it. For n dimensions, points $\mathbf{x}$ in the star plot are no longer projections of the states $\mathbf{z}$, but are rather calculated as

$$\mathbf{x} = \sum_{i=1}^{n} z_i \left(\cos \frac{2\pi i}{n}, \sin \frac{2\pi i}{n} \right) \tag{27}$$

A visualization for $n = 3$ is illustrated in Fig. 24.

In case of the L2/3 network, the number of dimensions is given by the number of attractors, with each axis describing some particular feature of the corresponding attractor (such as the average voltage or spike rate of the constituent PYR cells).

In addition to the position in state space, the state space velocity is also encoded in a star plot by both the thickness and the color of the trajectory. Especially in the case of the L2/3 network, this can be very useful in visualizing e.g. attractor stability or competition times. Here, both line thickness and lightness were chosen proportional to $(\text{const} + e^{-|d\mathbf{x}|/dt})$, with $\mathbf{x}$ being the position in state space.

Fig. 25 **B** and **C** show two characteristic examples of star plots used for visualizing the dynamics of the L2/3 network.

B.9 Average synaptic conductance due to Poisson stimulation

For a single Poisson source with rate ν_i connected to the neuron by a synapse with weight w_i and time constant τ^{syn}, the conductance course can be viewed as a sum of independent random variables, each of them representing the conductance change caused by a single spike. In the limit of large ν_i, the central limit theorem guarantees the convergence of the conductance distribution to a Gaussian, with moments given by

$$\begin{aligned}
\langle g_i^{\text{syn}} \rangle &= \sum_{\text{spk } s} \langle w_i \Theta(t - t_s) \exp\left(-\frac{t - t_s}{\tau^{\text{syn}}}\right) \rangle \\
&= \lim_{T \to \infty} \frac{\langle N \rangle}{T} w_i \int_0^T \exp\left(-\frac{t}{\tau^{\text{syn}}}\right) dt \\
&= w_i \nu_i \tau^{\text{syn}} \quad . \qquad (28) \\
\text{Var}\left[g_i^{\text{syn}}\right] &= \sum_{\text{spk } s} \text{Var}\left[w_i \Theta(t - t_s) \exp\left(-\frac{t - t_s}{\tau^{\text{syn}}}\right)\right] \\
&= \lim_{T \to \infty} \langle N \rangle \left\{ \langle \left[w_i \Theta(t) \exp\left(-\frac{t}{\tau^{\text{syn}}}\right)\right]^2 \rangle \right. \\
&\left. + \langle \left[w_i \Theta(t) \exp\left(-\frac{t}{\tau^{\text{syn}}}\right)\right] \rangle^2 \right\} \\
&= \lim_{T \to \infty} \nu_i T \left\{ \frac{1}{T} w_i^2 \int_0^T \exp\left(-2\frac{t}{\tau^{\text{syn}}}\right) dt \right. \\
&\left. - \frac{1}{T^2} \left[\int_0^T \exp\left(-\frac{t}{\tau^{\text{syn}}}\right) dt\right] \right\} \\
&= \frac{w_i^2 \nu_i \tau^{\text{syn}}}{2} \quad . \qquad (29)
\end{aligned}$$

Since conductances sum up linearly, N Poisson sources lead to an average conductance of

$$\begin{aligned}
\langle g^{\text{syn}} \rangle &= \langle \sum_{i=1}^{N} g_i^{\text{syn}} \rangle \\
&= N \langle w \rangle \langle \nu \rangle \tau^{\text{syn}} \quad . \qquad (30)
\end{aligned}$$

B.10 Detailed simulations of synapse loss and PYR population reduction

Fig. 25 and 26 show the effects of various levels of synapse loss and PYR population reduction, respectively.

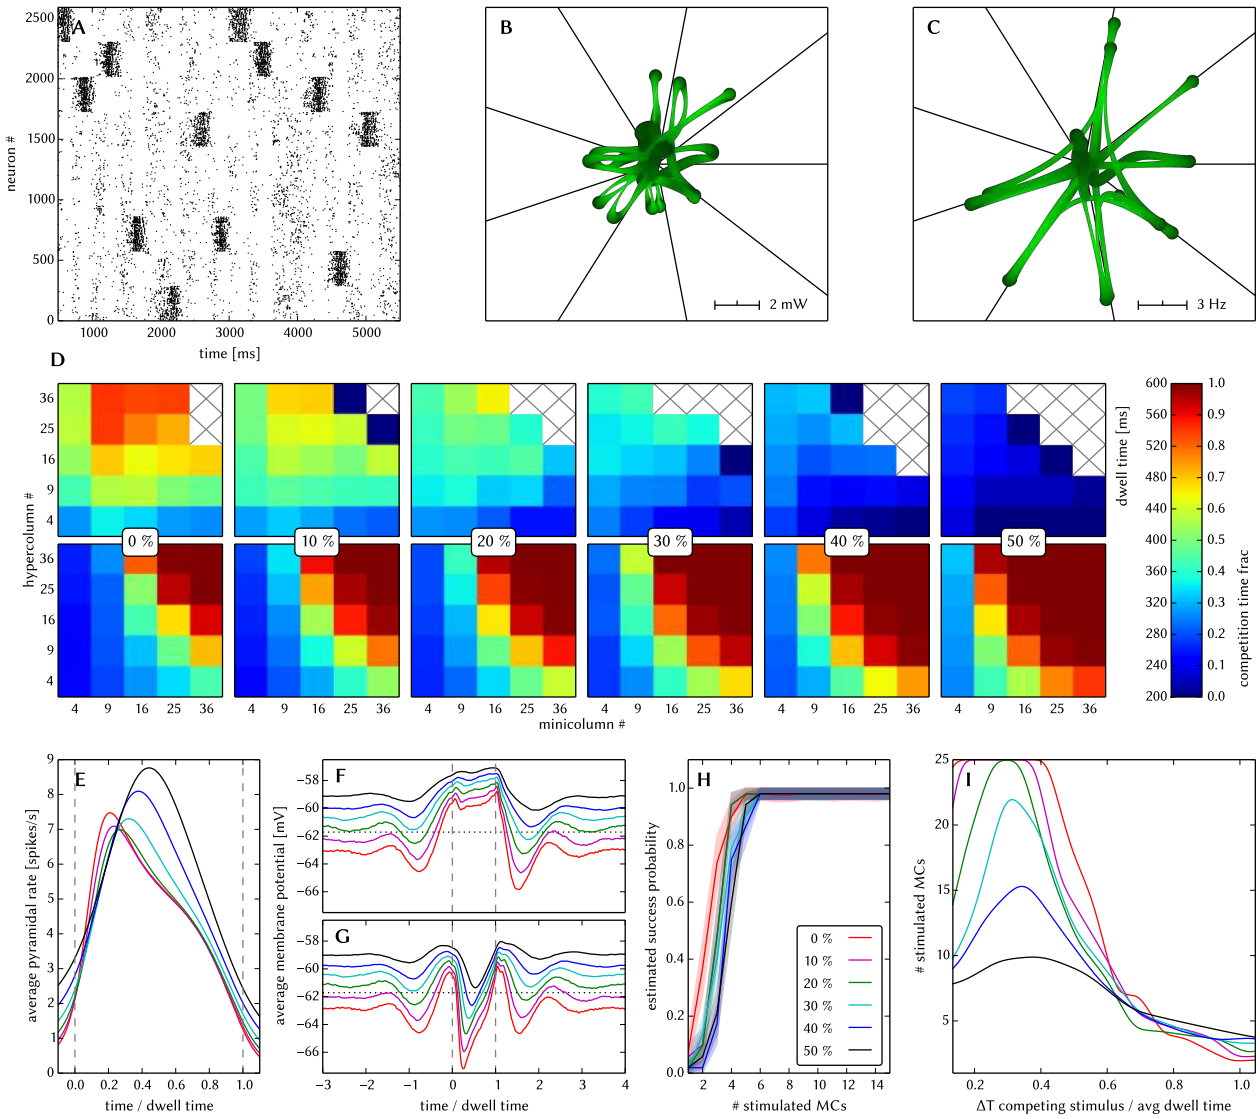

Fig. 25: Effects of homogeneous synapse loss on the L2/3 model. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. The topmost 3 figures exemplify the dynamics of the network at 50% synapse loss, all other figures show the effects of various degrees of synapse loss (0-50%). (**A**) Raster plot of spiking activity. Only PYR cells are shown. The MCs are ordered such that those belonging to the same attractor (and *not* those within the same HC) are grouped together. (**B**) Star plot of average PYR cell voltages from a sample of 5 PYR cells per MC. (**C**) Star plot of average PYR cell firing rates. (**D**) Average dwell times and relative competition times for various network sizes. (**E**) Average firing rate of PYR cells during an UP state. (**F**) Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**G**) Average voltage of PYR cells before, during and after an attractor they do not belong to is active. For the previous three plots, the abscissa has been subdivided into multiples of the attractor dwell time. In subplots **F** and **G** the dotted line indicates the leak potential E_{L} of the PYR cells. (**H**) Pattern completion in a 25HC×25MC network. (**I**) Attentional blink in a 25HC×25MC network: $p = 0.5$ iso-probability contours.

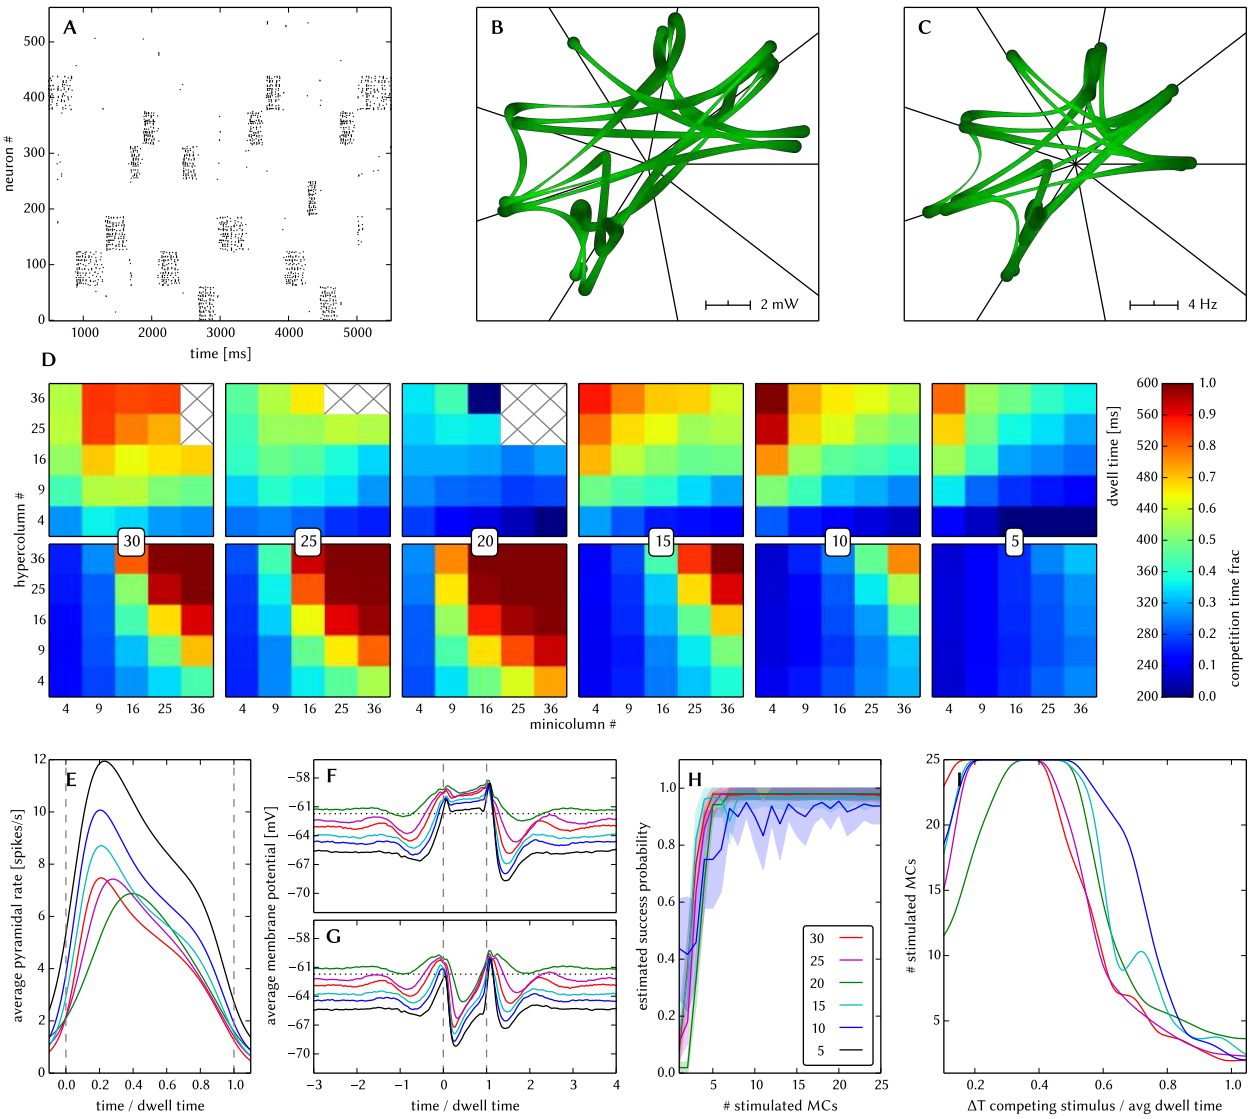

Fig. 26: Effects of PYR population size scaling on the L2/3 model. Unless explicitly stated otherwise, the default network model (9HC×9MC) was used. The topmost 3 figures exemplify the dynamics of the network at 50% of its original PYR population size, all other figures show the effects of various degrees of PYR population reduction (0-50%). (**A**) Raster plot of spiking activity. Only PYR cells are shown. The MCs are ordered such that those belonging to the same attractor (and *not* those within the same HC) are grouped together. (**B**) Star plot of average PYR cell voltages from a sample of 5 PYR cells per MC. (**C**) Star plot of average PYR cell firing rates. (**D**) Average dwell times and relative competition times for various network sizes. (**E**) Average firing rate of PYR cells during an UP state. (**F**) Average voltage of PYR cells before, during and after their parent attractor is active (UP state). (**G**) Average voltage of PYR cells before, during and after an attractor they do not belong to is active. For the previous three plots, the abscissa has been subdivided into multiples of the attractor dwell time. In subplots **F** and **G** the dotted line indicates the leak potential E_L of the PYR cells. (**H**) Pattern completion in a 25HC×25MC network. (**I**) Attentional blink in a 25HC×25MC network: $p = 0.5$ iso-probability contours. In **H** and **I**, the dataset for 5 PYR cells per MC was omitted because of its extremely low validity rate.

Appendix C Synfire chain with feed-forward inhibition

C.1 Model parameters

The neuron and connectivity parameters are given in Tab. 17 and Tab. 18.

Table 17: Neuron parameters used in the synfire chain benchmark model

Parameter	Value	Unit
C_{m}	0.29	nF
τ_{refrac}	2	ms
E^{spike}	-57	mV
E^{r}	-70	mV
E_{L}	-70	mV
τ_{m}	10	ms
$E^{\mathrm{rev,e}}$	0	mV
$E^{\mathrm{rev,i}}$	-75	mV
$\tau^{\mathrm{syn,e}}$	1.5	ms
$\tau^{\mathrm{syn,i}}$	10	ms

Table 18: Projection properties for the feed-forward synfire chain

Projection	weight µS	incoming synapses	delay ms
$\mathrm{RS}_n \rightarrow \mathrm{RS}_{n+1}$	0.001	60	20
$\mathrm{RS}_n \rightarrow \mathrm{FS}_{n+1}$	0.0035	60	20
$\mathrm{FS}_n \rightarrow \mathrm{RS}_n$	0.002	25	4

C.2 Network scaling

In the default setup studied in this article, the synfire chain consists of 6 groups of 125 neurons (100 excitatory and 25 inhibitory). In order to quantify the amount of synapse loss after mapping the network to the BrainScaleS wafer-scale hardware for different network sizes, we define the following network scaling rules. When increasing the network size, we vary both the number of synfire groups and the number of neurons per group while keeping the number of incoming synapses per neuron constant (cf. Tab. 18). The fraction of inhibitory neurons always amounts to 20 %. Neuron and synapse parameters are not altered. Tab. 19 lists the combinations of group size and group count used for the synapse loss estimation in Fig. 16 **A**.

The background Poisson stimulus is scaled as follows. For the hardware implementation of the synfire chain we can not use one individual Poisson source for each neuron due to input bandwidth limitations. Instead, we assume one pool of 32 Poisson sources for each synfire group, and each neuron receives input from 8 random sources from that pool. The size of the background pool is then scaled with the number of neurons per synfire group, while always drawing 8 sources from the pool per neuron. This scaling of the background pool was chosen to make the total number of background sources proportional to the total number of neurons and independent of the group count.

Table 19: Scaling table for the synfire chain used for the synapse loss estimation in Fig. 16 **A**

groups	group size	total neurons
8	125	1000
16	125	2000
24	125	3000
20	200	4000
25	200	5000
15	400	6000
20	350	7000
20	400	8000
30	300	9000
25	400	10 000
20	500	10 000
40	500	20 000
60	500	30 000
40	1000	40 000
50	1000	50 000
30	2000	60 000
20	3500	70 000
20	4000	80 000
30	3000	90 000
25	4000	100 000

C.3 Additional simulation

C.3.1 All distortions

To check that the compensation mechanisms do not interfere with each other, all distortions were applied with weight noise values of 20 % and 50 % and synapse loss values of 30 % and 50 %, with an axonal delay of 1.0 ms. Without compensation no stable region exists in all four cases. Fig. 28 shows the result with all compensation methods applied. When several methods required modification of a network parameter, all modifications were applied. For instance, in the case of the synaptic weight which was scaled by both synapse loss and delay compensation methods, both scaling factors were multiplied. Fig. 28 shows the restoration of input selectivity in all four cases.

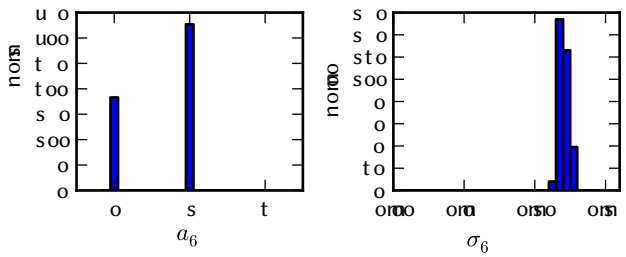

Fig. 27: Distribution of a_6 and σ_6 in the reference experiment for the synfire chain with feed-forward inhibition.

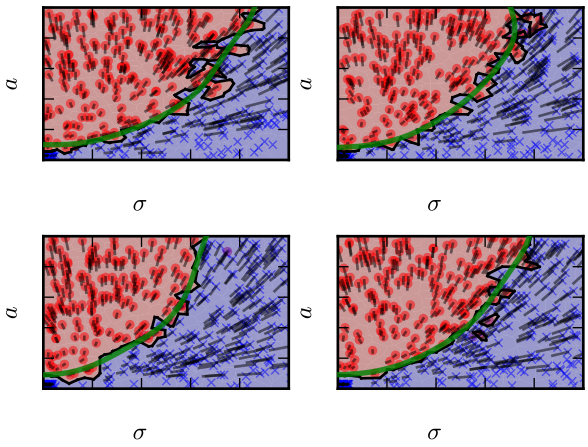

Fig. 28: (σ, a) state space of the synfire chain model with all compensation methods applied for four different values of distortion. **(A)** 30 % synapse loss, 20 % weight noise **(B)** 30 % synapse loss, 50 % weight noise **(C)** 50 % synapse loss, 20 % weight noise **(D)** 50 % synapse loss, 50 % weight noise

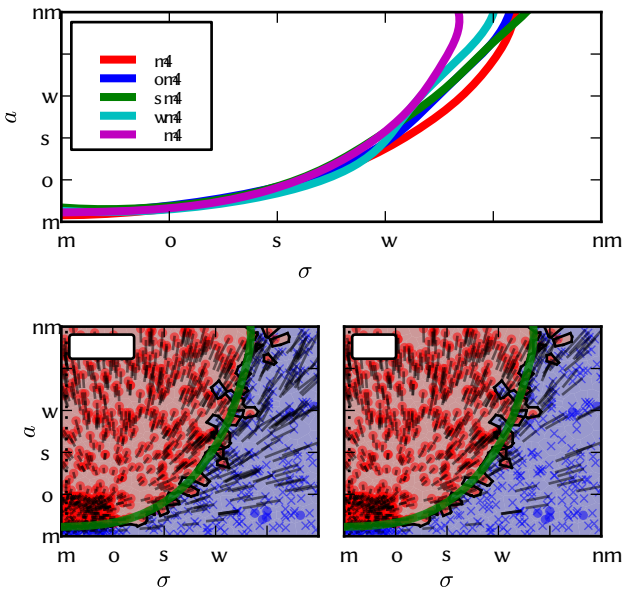

Fig. 29: Demonstration of spontaneous event filter in the weight noise compensation (Sec. C.3.4). **(A)** The same experiment as in Fig. 14 **C** (weight noise with active compensation) but without the filter for background spikes. The separatrix locations are comparable as the filter does not influence the result significantly in the compensated case. (**B**, **C**) Complete state space response for weight noise of 80%, once with, once without filter. This demonstrates that the applied filter does not affect the result in the compensated case.

C.3.2 Separatrix fit

To compare different separatrices, the a-values of the last group are characterized as successful (+1) or ex-

tinguished (−1) and the resulting values interpolated and smoothed by a gaussian kernel with a standard deviation (1.5 ms, 1.5) in the (σ, a) space. The iso-contour line of the resulting surface at a value of 0 is used as an approximation of the separatrix location, as shown in Fig. 12 **C** together with the individual simulation results. Data points with $\sigma \leq 0.2$ ms were not included in the fit to avoid bandwidth distortions in ESS simulations (Sec. 3.2.6) from affecting the fit quality. The data points are still shown individually as blue dots and regions, e.g., in Fig. 16. This modification was also included in the software simulations for consistency. Cases in which the separatrix does not capture the relevant behavior, e.g., if the separation is not reliable in a large region of the state space, are shown separately.

C.3.3 Weight noise compensation

Fig. 29 **A** shows the separatrix in the case of compensated weight noise.

C.3.4 Filtering of spontaneous activity

To prevent spontaneous background events from distorting the analysis, spikes are discarded as part of spontaneous activity if less then N spikes in the same excitatory group occur in a time window of $\pm T$. The utilized values for N and T are given at each point where the filter is applied; They are chosen such that authentic synchronous volleys with $a \geq 0.5$ (which would be counted as successful propagation, as defined above) are not removed. Fig. 29 **B** and **C** show that the influence of the filter for spontaneous activity is minimal in the compensated case.

C.3.5 Further ESS simulations

Distortion and compensation without synapse loss For the ESS simulation in Sec. 3.2.6 we enforced a certain amount of synapse loss by restricting the synfire chain network to very limited hardware resources. However, due to its feed-forward structure, the network can be easily mapped onto the BrainScaleS hardware without any synapse loss (Fig. 16 **A**). Thus, we also investigated the network without synapse loss, such that the active distortions in the ESS simulations were synaptic weight noise, non-configurable axonal delays as well as spike loss and jitter. The state space of the distorted network (Fig. 30 **A**) contains only a small and loosely connected region of sustained activity which indicates unreliable separation. Applying the compensation mechanism for synaptic weight noise and axonal delays fully restores the filter property of the synfire chain, as can be seen in Fig. 30 **B**, where different separatrices mimic different delay-dependent realizations. Compared to the compensation of all distortions, the compensated state space without synapse loss does not show any flaws (**C**).

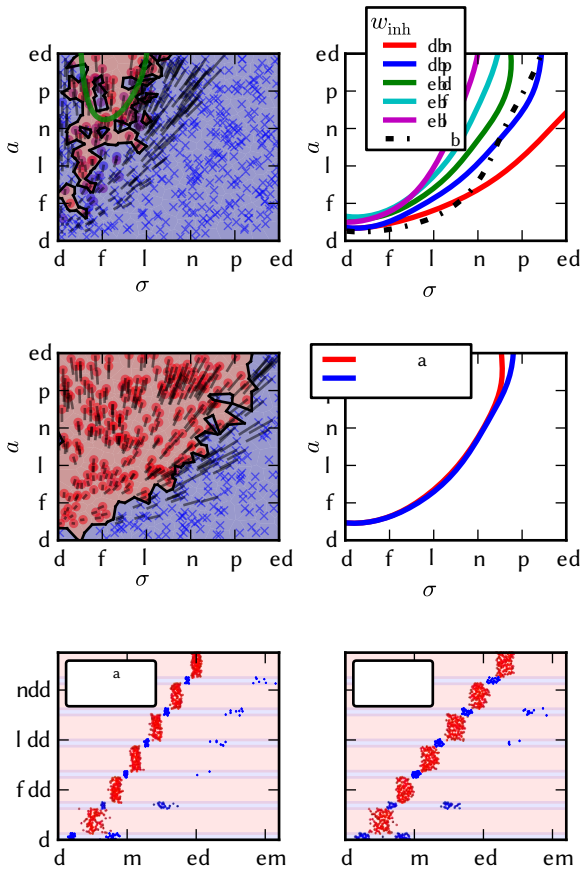

Fig. 30: Additional simulations of the feed-forward synfire chain on the ESS without synapse loss: (**A**) (σ,a) state space on the ESS with default parameters and 20% weight noise. (**B**) After compensation of all distortions, different separatrices are possible by setting different values of the inhibitory weight. (**C**) Compensated state space belonging to the blue separatrix in **B**. w refers to the synaptic weight of local inhibition. (**D**-**F**) Investigation of effects of spike loss and jitter by using two different approaches for neuron placement. (**D**) Separatrices for round-robin and sequential neuron placement with parameters as for the green curve in **B**. Raster plots for round-robin (**E**) and sequential (**F**) neuron placement. Stimulus parameters: $a_0 = 1$ and $\sigma_0 = 1$ ms.

Effect of spike loss and jitter We investigated the effect of spike loss and jitter in the HICANN, where the spikes of the neurons connected to the same on-wafer routing bus are processed subsequently (Sec. 2.1.2), which can

lead to spike time jitter and in rare cases to spike loss when firing is highly synchronized.

Which 64 neurons inject their spikes into a routing bus is determined by the placement of the neurons on the HICANN. Hence, in order to study the effect of spike loss and jitter, we simulated the synfire chain network in two different placement setups: First, neurons of the same synfire group were placed sequentially onto the same routing bus, and second, neurons were distributed in a round-robin manner over different routing buses, such that neurons of different groups injected their spikes into one routing bus. Hence, we expect the spiking activity on each routing bus to be more synchronous in the first case than in the second. In both setups, the utilized hardware and the number of neurons per routing bus was equal, allowing a fair competition between both. The separatrices for the two different placement strategies with otherwise identical parameters are indistinguishable except for a small region in the top right of the (σ, a) space (Fig. 30 **D**). Nevertheless, the raster plots (Fig. 30 **E** and **F**) reveal the effect of the introduced jitter: For sequential placement, the spread of spike times within a group is roughly double than for round-robin placement and also the onset of the volley in the last group comes 1.5 ms later. In contrast to the reference simulation (cf. Fig. 27), the fixed point of succesful propagation is not (0.12 ms,1) but (0.20 ms,1) for round-robin and (0.35 ms,1) for sequential placement.

We conclude that, especially for dense pulses, the subsequent processing of spikes in the hardware leads to a temporal spread of the pulse volley, which however has only a minor influence on the filter properties of the synfire chain.

Appendix D Self-sustained asynchronous irregular activity

D.1 Network simulation setup

The default model consists of 3920 neurons (80 % pyramidal and 20 % inhibitory) equally distributed on a two-dimensional lattice of 1×1 mm^2 folded to a torus. The connection probability is distance-dependent and is normalized such that each neuron receives synaptic input from 200 excitatory and 50 inhibitory neurons. All simulations run for 10 s. 2 % of all neurons in the network are initially stimulated by one individual Poisson source for 100 ms in order to induce initial network activity. The default size was chosen such that the model can be fully realized on the BrainScaleS hardware without losing any synaptic connections in the mapping step (Sec. 2.2), thereby allowing us to compare topologically equivalent software simulations, with the only remaining difference lying in the non-configurable delays and dynamic constraints on the ESS.

Table 20: AdEx Neuron parameters used in the AI network

Parameter	Pyramidal	Inhibitory	Unit
C_m	0.25	0.25	nF
τ_refrac	5	5	ms
E^spike	-40	-40	mV
E^r	-70	-70	mV
E_L	-70	-70	mV
τ_m	15	15	ms
a	1	1	nS
b	0.005	0	nA
Δ_T	2.5	2.5	mV
τ_w	600	600	ms
E_T	-50	-50	mV
$E^\mathrm{rev,e}$	0	0	mV
$E^\mathrm{rev,i}$	-80	-80	mV
$\tau^\mathrm{syn,e}$	5	5	ms
$\tau^\mathrm{syn,i}$	5	5	ms

D.1.1 Model parameters

The neuron parameters of the AdEx model used in this benchmark are listed in Tab. 20 and are equal to those in (Muller and Destexhe, 2012) with the only difference being that excitatory pyramidal cells have neuronal spike-triggered adaptation while inhibitory cells do not. Sweeps are performed over the two-dimensional $(g_\mathrm{exc}, g_\mathrm{inh})$ parameter space, with the ranges being 3 nS to 11 nS for g_exc and 50 nS to 130 nS for g_inh. The Poisson sources for the initial network stimulation have a mean rate of 100 Hz and project onto the network's neurons with a synaptic weight of 100 nS. The distance-dependent connection probability has a Gaussian profile with a spatial width of $\sigma = 0.2$ mm. Synaptic delays depend on the distance according to the following equation: $t_\mathrm{delay} = 0.3\,\mathrm{ms} + \frac{d}{v_\mathrm{prop}}$, with d being the distance between two cells and $v_\mathrm{prop} = 0.2\,\mathrm{mm\,ms^{-1}}$ the spike propagation velocity. The distribution of delays is shown in Fig. 31, the average delay in the network amounts to 1.55 ms.

D.1.2 Network scaling

When the network is scaled up in size, we only increase the number of neurons while keeping the number of afferent synapses per neuron constant. All other parameters concerning the connectivity do not change, including the size of the cortical sheet, the distance-dependent delays and connection probability, as well as the ratio of excitatory to inhibitory cells. Neuron and synapse parameters remain unaltered.

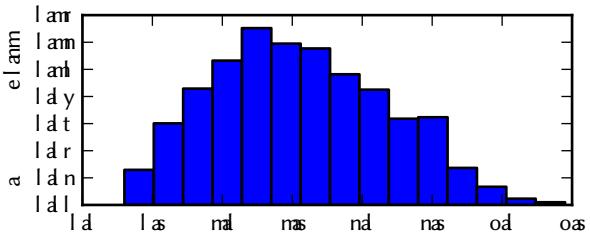

Fig. 31: Histogram of delays in the AI network. The mean delay is 1.55 ms.

D.2 Functionality criteria

The survival time is defined as the last spike time in the network. If the network survives until the end of the simulation, we consider it as self-sustaining. Additionally, several criteria are employed to characterize the network's activity, regularity and synchrony.

The mean firing rate of all pyramidal neurons is used to classify the overall activity of the network. The variance of the firing rates across the pyramidal neurons measures the homogeneity of their response. For a better comparison, we look at the relative variance, i.e., the coefficient of variation of the firing rates $\mathrm{CV}_{\mathrm{rate}} = \frac{\sigma(\nu)}{\overline{\nu}}$, where $\overline{\nu}$ and $\sigma(\nu)$ are the mean and standard deviation of the average firing rates ν of the inidividual neurons.

The coefficient of variation of interspike intervals ($\mathrm{CV}_{\mathrm{ISI}}$) serves as an indicator of spiking regularity. It is calculated via

$$\mathrm{CV}_{\mathrm{ISI}} = \frac{1}{N}\sum_{i=1}^{N}\frac{\sigma_i(\mathrm{ISI})}{\overline{\mathrm{ISI}_i}} \tag{31}$$

$\sigma_i(\mathrm{ISI})$ is the standard deviation of interspike intervals in the i-th spike train, while $\overline{\mathrm{ISI}_i}$ is the mean interspike interval in the same spike train. N is the number of averaged spike trains which is set to the number of pyramidal cells for each simulation. $\mathrm{CV}_{\mathrm{ISI}}$ is 0 for a regular spike train and approaches 1 for a sufficiently long Poisson spike train.

The correlation coefficient CC is defined via

$$\mathrm{CC} = \frac{1}{P}\sum_{j,k}^{P}\frac{\mathrm{Cov}(S_j, S_k)}{\sigma(S_j)\sigma(S_k)} \tag{32}$$

The sum runs over $P = 5000$ randomly chosen pairs of spike trains (j, k) from the excitatory population. S_i is the time-binned spike count in the i-th spike train with a bin width of $\Delta = 5$ ms. $\sigma(S_i)$ denotes the standard deviation of S_i, and $\mathrm{Cov}(S_j, S_k)$ the covariance of S_j and S_k. CC approaches 0 for sufficiently long independent spike trains and is 1 for linearly dependent (S_j, S_k). The simulation results were cross-checked with a bin width of $\Delta = 2$ ms.

The power spectrum $S(\omega)$ of a spike train is calculated via

$$A_k = \sum_{m=0}^{N-1} r_m \exp\left(-2\pi i \frac{mk}{N}\right) \qquad k = 0, \ldots, N-1 \tag{33}$$

$$\omega_k := \frac{2\pi k}{N\Delta} \tag{34}$$

$$S(\omega_k) := |A_k|^2 N\Delta \tag{35}$$

using the time-binned population firing rate r_i with $i \in \{0, ..., N-1\}$ with a bin width of Δ for a spike train of length $N\Delta$ (see, e.g. Rieke et al, 1997, 3.1.4). For the AI network we used a bin width of $\Delta = 1$ ms for calculating the raw power spectra, and a $\sigma = 5$ Hz for the Gauss-filtered versions which where then used to determine the peak frequency (i.e. the first non-zero peak in the power spectrum).

In case of the L2/3 model, the power spectra were calculated from Gauss-filtered ($\sigma = 5$ ms) spike data with a bin width of $\Delta = 0.1$ ms and (unless otherwise stated) smoothed with a $\sigma = 0.3$ ms Gauss-filter.

For all statistics, the first second of the simulation is left out, i.e. only the 9 seconds from 1 s to 10 s are considered. If the network did not survive until the end of the simulation, the firing rate was calculated between 1 s and the survival time, or between 0.1 s and the survival time for the case when the latter was smaller than 1 s.

D.3 Iterative compensation

In the so-called iterative compensation, we sequentially modify individual parameters such that the response of each neuron is modified to match its target response. In our case, we iteratively change the spike detection voltage E_{T} such that the firing rate of each neuron is shifted towards the target rate. At each step, the threshold voltage is adapted as follows:

$$E_{\mathrm{T}}^{\mathrm{n+1,i}} = E_{\mathrm{T}}^{\mathrm{n,i}} + (\nu^{\mathrm{tgt}} - \nu^{\mathrm{n,i}})c_{\mathrm{comp}} \tag{36}$$

where $E_{\mathrm{T}}^{\mathrm{n,i}}$ and $\nu^{\mathrm{n,i}}$ are the threshold voltage and firing rate of neuron i of the n-th step, ν^{tgt} is the target rate for all neurons of a population and c_{comp} is a compensation factor that links the firing response and the threshold voltage. The target rate ν^{tgt} is computed separately for the excitatory and inhibitory population from the reference simulations (Sec. 3.3.2). We choose the compensation factor for each $(g_{\mathrm{exc}}, g_{\mathrm{inh}})$ state in the following manner: Similar to the mean-field approach in

Sec. 3.3.6, we consider the response rate of an excitatory neuron given a network firing rate of ν^{tgt}, that is, the neuron is stimulated by 200 excitatory and 50 inhibitory Poisson sources with rate ν^{tgt}. We then vary the threshold voltage of said neuron between $-54\,\text{mV}$ and $-46\,\text{mV}$ and thereby determine the dependency of the response rate on the threshold voltage. From a linear fit of this dependency, we extract the slope m, and set the compensation factor to $c_{\text{comp}} = \frac{0.5}{m}$ (Fig. 32). The factor of 0.5 was chosen to limit the change of the mean rate in each step in order to avoid oscillations in the compensation procedure. Whenever we changed the spike initiation voltage E_{T}, we shifted the spike detection voltage E^{spike} equally.

We remark that this compensation method requires the parameters for every individual neuron to be fine-tunable. This is the case for the BrainScaleS wafer-scale hardware, where the AdEx parameters of every hardware neuron are independently configurable with sufficient precision by means of analog floating gate memories (Sec. 2.1), in contrast to the synaptic weights which are restricted to a 4-bit precision in typical operation mode.

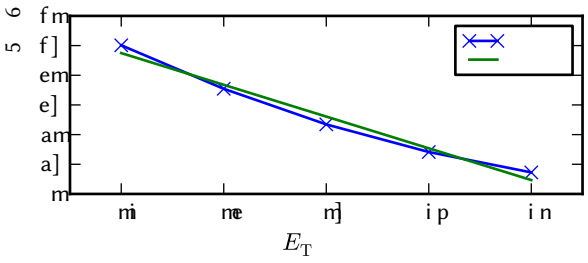

Fig. 32: Example compensation factor assertion for the state ($g_{\text{exc}} = 9\,\text{nS}$, $g_{\text{inh}} = 90\,\text{nS}$) of the AI network: The Figure shows the response rate of an excitatory neuron stimulated by 200 excitatory and 50 inhibitory Poisson sources with rate 12.38 Hz depending on its spike initiation threshold E_{T}. The slope $m = -2.6745\,\frac{\text{Hz}}{\text{mV}}$ of the linear fit is then used to calculate the compensation factor $c_{\text{comp}} = \frac{0.5}{m} = -0.18695\,\frac{\text{mV}}{\text{Hz}}$.

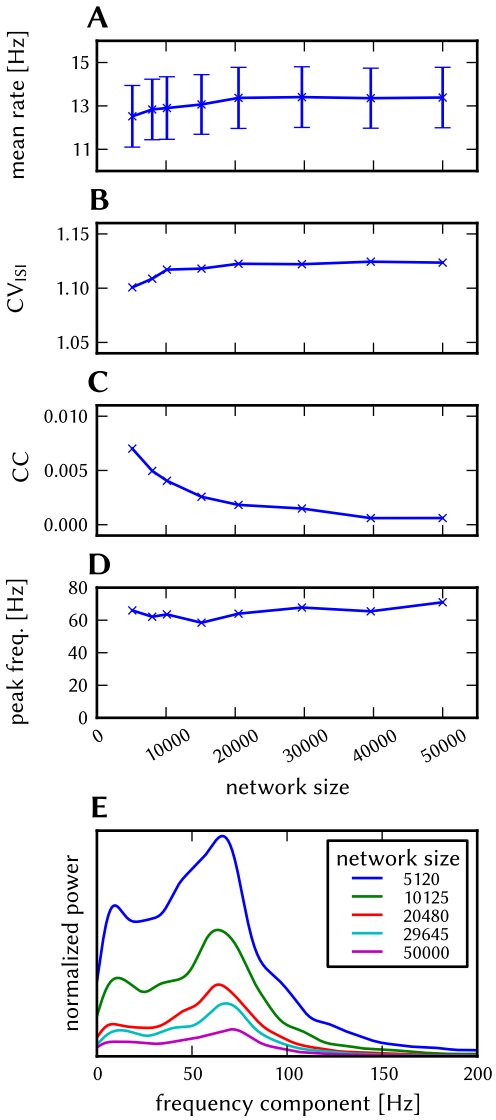

Fig. 33: Network size scaling behavior of the AI network for the (9 nS, 90 nS) state. Mean and variance of firing rate across the PY neurons (**A**), coefficient of variance of inter-spike intervals (**B**), coefficient of pairwise cross-correlation (**C**), and power spectrum of global activity: full spectra (**D**) and peak frequencies (**E**).

D.4 Further simulations

D.4.1 Network size scaling behavior

To investigate what happens when the network is scaled according to the rules given in Sec. D.1.2, we pick one state of the (g_{exc}, g_{inh}) space and vary the network size between 5000 and 50 000 neurons. The results for the (9 nS, 90 nS) state can be seen in Fig. 33: The mean firing rate slightly increases with size until its saturates, while the variance of the firing rate across neurons remains approximately constant (**A**). Like the firing rate, the irregularity (CV_{ISI}) increases and saturates with size (**B**). The synchronicity (CC) decreases with size, as one would expect (**C**). The power spectrum of global activity exhibits the same profile for all sizes, however the power is scaled inversely to the network size (**D** and **E**).

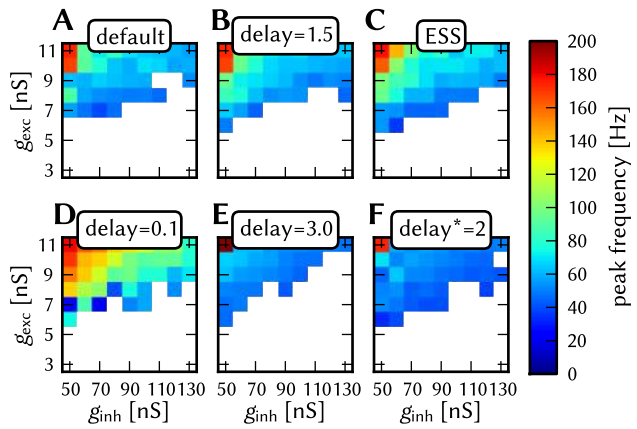

Fig. 34: Effects of axonal delays on the AI network. (g_{exc}, g_{inh}) spaces with the peak frequency of the global pyramidal activity for different axonal delay setups: default with distance-dependent delays (**A**), constant delay of 1.5 ms (**B**), simulation on the ESS where delay is not configurable (**C**), constant delay of 0.1 ms (**D**), constant delay of 3.0 ms (**E**), distance-dependent delays scaled by factor of 2 with respect to default setup (**F**).

D.4.2 Non-configurable axonal delays

In Sec. 3.3.3 we argue that non-configurable delays on the BrainScaleS hardware only have a minimal effect on the AI network because the average delay in the model matches the estimated average delay on the hardware. Here, we provide the simulation results and further investigations on the influence of the delay on the network dynamics. For the analysis of the effects of non-configurable delay we repeated the (g_{exc}, g_{inh}) sweep with all synaptic delays set to 1.5 ms, cf. Sec. 2.4. The delay distortion only affected the power spectrum of global actitivity but not the other criteria such that we show only the peak frequency parameter spaces in Fig. 34. The distorted network (**B**) with a constant delay of 1.5 ms is not significantly different from the default network with distance-dependent delays (**A**), both the region of sustained activity and the position of the peak in the power spectrum are in good match. The same holds for ESS simulations (**C**), where non-configurable delays were the only active distortion.

To further investigate the influence of the delays, we ran additional simulations where all delays in the network were set to 0.1 ms (**D**), and 3 ms (**E**), respectively. Lowering delays increases the speed of activity propagation such that the position of the peak in the power spectrum is shifted towards higher frequencies. For higher delays the peak frequency decreases analogously, but also the region of sustained activity diminishes significantly. (**F**) shows simulations with distance-dependent delays scaled by a factor of 2 with respect to the baseline model, thus having an average delay of 3.1 ms (cf. Fig. 31). While the peak frequency is in good agreement with the 3 ms simulations, the region of sustained states is extended and even larger than in the baseline setup. Herewith our simulations affirm that distance-dependent delays in fact do expand the region of self-sustained states in the (g_{exc}, g_{inh}) space (cf. Sec. 3.3.1).

D.4.3 Combining distortions: synapse loss and synaptic weight noise

We also investigate what happens when both synapse loss and synaptic weight noise are active at the same time. Additionally, we test up to which extent we can compensate both distortions. To do so we scaled both distortions up to 90 % and tried to restore the original behavior for two states: (9 nS, 90 nS) and (10 nS, 70 nS).

The relative change of the mean rate and CV_{rate} are shown in Fig. 35 for the (9 nS, 90 nS) state. For this state, synapse loss compensation works fine up to a level of 50 %: the relative change of the mean rate and CV_{rate} are close to 0. The compensation fails for synapse losses at 70 % distortion and above: when the original firing rate is recovered, the network is unstable, i.e. it does not survive until the end of the experiment. The amount of synaptic weight noise has no effect on this behavior. We remark that, during the iterative compensation, there are stable networks with a slightly higher firing rate than the target rate: the network becomes unstable when approaching its target rate. This is in accordance with observations from the 50 % loss parameter space compensation in Fig. 20, where the region of sustained activity is smaller than before, i.e. requiring a higher frequency for fewer synapses. We also note that our iterative compensation algorithm does not recover distorted networks that die out shortly after initial stimulation (cf. the 90 % synapse loss column in Fig. 35 **A**). Synaptic weight noise does not pose a problem to the iterative compensation: In all cases the mean rate could be fully recovered and the variance of firing rates close to the original level, with the relative difference of CV_{rate} being smaller than 1.5.

For the (10 nS, 70 nS) state, compensation was capable of restoring a synapse loss including 70 %, cf. Fig. 36. Interestingly, the reduction of the variation of firing rates after 10 compensation steps performed slightly better when starting with a higher synapse loss.

We summarize that the iterative method effectively compensates for combined distortions with both synapse loss and synaptic weight noise, at least when the synapse loss does not exceed 50 %. Furthermore, we expect these results to hold also for a large area in the

(g_{exc}, g_{inh}) space where the network is in the asynchronous regime.

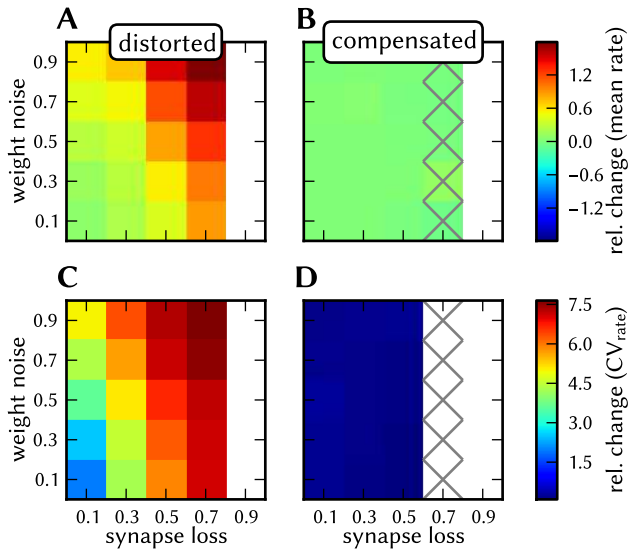

Fig. 35: Compensation of combined distortions in the AI network with the iterative method. Sweep over synapse loss and synaptic weight noise for the ($g_{\mathrm{exc}} = 9\,\mathrm{nS}$, $g_{\mathrm{inh}} = 90\,\mathrm{nS}$) state. Relative change of the firing rate with respect to the undistorted network for distorted (**A**) and compensated (**B**) simulations. Relative change of $\mathrm{CV}_{\mathrm{rate}}$ with respect to the undistorted network for distorted (**C**) and compensated (**D**) simulations. The compensated simulations refer to the 10th step of iterative compensation. White data points stand for networks where the distorted network did not survive. Data points marked with a cross denote cases where the compensated network did not survive.

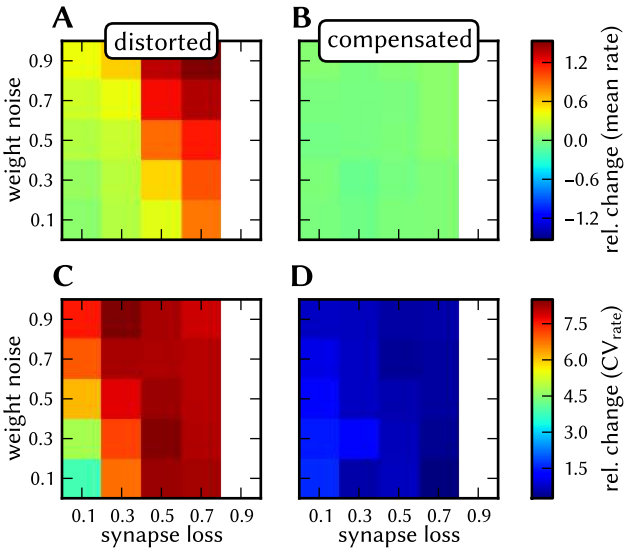

Fig. 36: Compensation of combined distortions in the AI network with the iterative method. Sweep over synapse loss and synaptic weight noise for the ($g_{\mathrm{exc}} = 10\,\mathrm{nS}$, $g_{\mathrm{inh}} = 70\,\mathrm{nS}$) state. Relative change of the firing rate with respect to the undistorted network for distorted (**A**) and compensated (**B**) simulations. Relative change of $\mathrm{CV}_{\mathrm{rate}}$ with respect to the undistorted network for distorted (**C**) and compensated (**D**) simulations. The compensated simulations refer to the 10th step of iterative compensation. White data points stand for cases where the distorted network did not survive.

References

Abeles M, Hayon G, Lehmann D (2004) Modeling compositionality by dynamic binding of synfire chains. *Journal of computational neuroscience* 17(2):179–201

Aertsen A, Diesmann M, Gewaltig MO (1996) Propagation of synchronous spiking activity in feedforward neural networks. *J Physiol Paris* 90(3-4):243–247

Berge HKO, Häfliger P (2007) High-speed serial AER on FPGA. In: ISCAS, IEEE, pp 857–860

Bergman K, Borkar S, Campbell D, Carlson W, Dally W, Denneau M, Franzon P, Harrod W, Hill K, Hiller J, et al (2008) Exascale computing study: Technology challenges in achieving exascale systems.

Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 18(24):10,464–10,472

Bill J, Schuch K, Brüderle D, Schemmel J, Maass W, Meier K (2010) Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity. *Front Comp Neurosci* 4(129)

Bontorin G, Renaud S, Garenne A, Alvado L, Le Masson G, Tomas J (2007) A real-time closed-loop setup for hybrid neural networks. In: Proceedings of the 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS2007)

Brette R, Gerstner W (2005) Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J Neurophysiol* 94:3637 – 3642

Brette R, Rudolph M, Carnevale T, Hines M, Beeman D, Bower JM, Diesmann M, Morrison A, Goodman PH, Harris Jr FC, Zirpe M, Natschlager T, Pecevski D, Ermentrout B, Djurfeldt M, Lansner A, Rochel O, Vieville T, Muller E, Davison AP, Boustani SE, Destexhe A (2007) Simulation of networks of spiking neurons: A review of tools and strategies. *Journal of Computational Neuroscience* 23(3):349–398

Bringuier V, Chavane F, Glaeser L, Frégnac Y (1999) Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. *Science* 283(5402):695–699

Brüderle D, Müller E, Davison A, Muller E, Schemmel J, Meier K (2009) Establishing a novel modeling tool: A python-based interface for a neuromorphic hardware system. *Front Neuroinform* 3(17)

Brüderle D, Petrovici M, Vogginger B, Ehrlich M, Pfeil T, Millner S, Grübl A, Wendt K, Müller E, Schwartz MO, de Oliveira D, Jeltsch S, Fieres J, Schilling M, Müller P, Breitwieser O, Petkov V, Muller L, Davison A, Krishnamurthy P, Kremkow J, Lundqvist M, Muller E, Partzsch J, Scholze S, Zühl L, Mayr C, Destexhe A, Diesmann M, Potjans T, Lansner A, Schüffny R, Schemmel J, Meier K (2011) A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems. *Biological Cybernetics* 104:263–296

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience* 8(3):183–208

Buxhoeveden D, Casanova M (2002) The minicolumn and evolution of the brain. *Brain Behav Evol* 60:125–151

Connors B, Gutnick M (1990) Intrinsic firing patterns of diverse neocortical neurons. *Trends Neurosci* 13:99–104

Cossart R, Aronov D, Yuste R (2003) Attractor dynamics of network up states in the neocortex. *Nature* 423:238–283

Costas-Santos J, Serrano-Gotarredona T, Serrano-Gotarredona R, Linares-Barranco B (2007) A spatial contrast retina with on-chip calibration for neuromorphic spike-based AER vision systems. *IEEE Transactions on Circuits and Systems* 54(7):1444–1458

Davison AP, Brüderle D, Eppler J, Kremkow J, Muller E, Pecevski D, Perrinet L, Yger P (2008) PyNN: a common interface for neuronal network simulators. *Front Neuroinform* 2(11)

Delbrück T, Liu SC (2004) A silicon early visual system as a model animal. *Vision Res* 44(17):2083–2089

Destexhe A (2009) Self-sustained asynchronous irregular states and Up/Down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. *Journal of Computational Neuroscience* 3:493 – 506

Destexhe A, Contreras D (2006) Neuronal computations with stochastic network states. *Science* 314(5796):85–90

Destexhe A, Pare D (1999) Impact of Network Activity on the Integrative Properties of Neocortical Pyramidal Neurons In Vivo. *J Neurophysiol* 81(4):1531–1547

Destexhe A, Rudolph M, Pare D (2003) The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience* 4:739–751

Diesmann M (2002) Conditions for stable propagation of synchronous spiking in cortical neural networks: Single neuron dynamics and network properties. PhD thesis, Ruhr-Universität Bochum

Diesmann M, Gewaltig MO (2002) NEST: An environment for neural systems simulations. In: Plesser T, Macho V (eds) Forschung und wisschenschaftliches Rechnen, Beiträge zum Heinz-Billing-Preis 2001, GWDG-Bericht, vol 58, Ges. für Wiss. Datenverarbeitung, Göttingen, pp 43–70

Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533

Diesmann M, Gewaltig MO, Rotter S, Aertsen A (2001) State space analysis of synchronous spiking in cortical neural networks. *Neurocomputing* 38:565–571

Djurfeldt M, Lundqvist M, Johansson C, Rehn M, Ekeberg O, Lansner A (2008) Brain-scale simulation of the neocortex on the ibm blue gene/l supercomputer. *IBM Journal of Research and Development* 52(1.2):31–41

Ehrlich M, Mayr C, Eisenreich H, Henker S, Srowig A, Grübl A, Schemmel J, Schüffny R (2007) Wafer-scale VLSI implementations of pulse coupled neural networks. In: Proceedings of the International Conference on Sensors, Circuits and Instrumentation Systems (SSD-07)

Ehrlich M, Wendt K, Zühl L, Schüffny R, Brüderle D, Müller E, Vogginger B (2010) A software framework for mapping neural networks to a wafer-scale neuro-

morphic hardware system. In: Proceedings of the Artificial Neural Networks and Intelligent Information Processing Conference (ANNIIP) 2010, pp 43–52

El Boustani S, Destexhe A (2009) A master equation formalism for macroscopic modeling of asynchronous irregular activity states. *Neural Computation* 21(1):46–100

Eppler JM, Helias M, Muller E, Diesmann M, Gewaltig MO (2008) PyNEST: a convenient interface to the NEST simulator. *Front Neuroinform* 2(12)

Fieres J, Schemmel J, Meier K (2008) Realizing biological spiking network models in a configurable wafer-scale hardware system. In: Proceedings of the 2008 International Joint Conference on Neural Networks (IJCNN)

Furber SB, Lester DR, Plana LA, Garside JD, Painkras E, Temple S, Brown AD (2012) Overview of the SpiNNaker system architecture. *IEEE Transactions on Computers* 99(PrePrints)

Galluppi F, Rast A, Davies S, Furber S (2010) A general-purpose model translation system for a universal neural chip. In: Wong K, Mendis B, Bouzerdoum A (eds) Neural Information Processing. Theory and Algorithms, Lecture Notes in Computer Science, vol 6443, Springer Berlin / Heidelberg, pp 58–65

González-Burgos G, Barrionuevo G, Lewis DA (2000) Horizontal synaptic connections in monkey prefrontal cortex: An in vitro electrophysiological study. *Cerebral Cortex* 10(1):82–92

Häfliger P (2007) Adaptive WTA with an analog VLSI neuromorphic learning chip. *IEEE Transactions on Neural Networks* 18(2):551–72

Hartmann S, Schiefer S, Scholze S, Partzsch J, Mayr C, Henker S, Schuffny R (2010) Highly integrated packet-based aer communication infrastructure with 3gevent/s throughput. In: Electronics, Circuits, and Systems (ICECS), 2010 17th IEEE International Conference on, pp 950 –953

Hasler J, Marr HB (2013) Finding a roadmap to achieve large neuromorphic hardware systems. *Frontiers in Neuroscience* 7(118)

Helias M, Kunkel S, Masumoto G, Igarashi J, Eppler JM, Ishii S, Fukai T, Morrison A, Diesmann M (2012) Supercomputers ready for use as discovery machines for neuroscience. *Frontiers in Neuroinformatics* 6(26)

Hellwig B (2000) A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological Cybernetics* 82:111–121

Hines M, Carnevale N (2003) The NEURON simulation environment., M.A. Arbib, pp 769–773

Hines ML, Davison AP, Muller E (2009) NEURON and Python. *Front Neuroinform*

Hirsch J, Gilbert C (1991) Synaptic physiology of horizontal connections in the cat's visual cortex. *The Journal of Neuroscience* 11(6):1800–1809

Indiveri G (2008) Neuromorphic vlsi models of selective attention: From single chip vision sensors to multichip systems. *Sensors* 8(9):5352–5375

Indiveri G, Chicca E, Douglas R (2006) A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity. *IEEE Transactions on Neural Networks* 17(1):211–221

Indiveri G, Chicca E, Douglas R (2009) Artificial cognitive systems: From VLSI networks of spiking neurons to neuromorphic cognition. *Cognitive Computation* 1(2):119–127

Indiveri G, Linares-Barranco B, Hamilton TJ, van Schaik A, Etienne-Cummings R, Delbruck T, Liu SC, Dudek P, Häfliger P, Renaud S, Schemmel J, Cauwenberghs G, Arthur J, Hynna K, Folowosele F, Saighi S, Serrano-Gotarredona T, Wijekoon J, Wang Y, Boahen K (2011) Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience* 5(0)

Kampa BM, Letzkus JJ, Stuart GJ (2006) Cortical feed-forward networks for binding different streams of sensory information. *Nature Neuroscience* 9(12):1472–1473

Kremkow J, Aertsen A, Kumar A (2010a) Gating of signal propagation in spiking neural networks by balanced and correlated excitation and inhibition. *The Journal of neuroscience* 30(47):15,760–15,768

Kremkow J, Perrinet L, Masson G, Aertsen A (2010b) Functional consequences of correlated excitatory and inhibitory conductances in cortical networks. *J Comput Neurosci* 28:579–594

Kumar A, Schrader S, Aertsen A, Rotter S (2008) The high-conductance state of cortical networks. *Neural Computation* 20(1):1–43

Laing C, Lord GJ (2009) *Stochastic Methods in Neuroscience*. Oxford University Press

Lande T, Ranjbar H, Ismail M, Berg Y (1996) An analog floating-gate memory in a standard digital technology. In: Microelectronics for Neural Networks, 1996., Proceedings of Fifth International Conference on, pp 271 –276

Lewis MA, Etienne-Cummings R, Cohen AH, Hartmann M (2000) Toward biomorphic control using custom aVLSI chips. In: Proceedings of the International conference on robotics and automation, IEEE Press

Lundqvist M, Rehn M, Djurfeldt M, Lansner A (2006) Attractor dynamics in a modular network of neocortex. *Network: Computation in Neural Systems* 17:3:253–276

Lundqvist M, Compte A, Lansner A (2010) Bistable, irregular firing and population oscillations in a modular attractor memory network. *PLoS Comput Biol* 6(6)

Markram H, Gupta A, Uziel A, Wang Y, Tsodyks M (1998) Information processing with frequency-dependent synaptic connections. *Neurobiol Learn Mem* 70(1-2):101–112

Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G, Wu C (2004) Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci* 5(10):793–807

Mead CA (1989) *Analog VLSI and Neural Systems*. Addison Wesley, Reading, MA

Mead CA (1990) Neuromorphic electronic systems. *Proceedings of the IEEE* 78:1629–1636

Mead CA, Mahowald MA (1988) A silicon model of early visual processing. *Neural Networks* 1(1):91–97

Merolla PA, Boahen K (2006) Dynamic computation in a recurrent network of heterogeneous silicon neurons. In: Proceedings of the 2006 IEEE International Symposium on Circuits and Systems (ISCAS 2006)

Millner S, Grübl A, Meier K, Schemmel J, Schwartz MO (2010) A VLSI implementation of the adaptive exponential integrate-and-fire neuron model. In: Lafferty J, Williams CKI, Shawe-Taylor J, Zemel R, Culotta A (eds) Advances in Neural Information Processing Systems 23, pp 1642–1650

Mitra S, Fusi S, Indiveri G (2009) Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI. *IEEE Transactions on Biomedical Circuits and Systems* 3:(1):32–42

Morrison A, Mehring C, Geisel T, Aertsen A, Diesmann M (2005) Advancing the boundaries of high connectivity network simulation with distributed computing. *Neural Comput* 17(8):1776–1801

Morrison A, Diesmann M, Gerstner W (2008) Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics* 98(6):459–478

Mountcastle VB (1997) The columnar organization of the neocortex. *Brain* 120(4):701–722

Muller L, Destexhe A (2011) A model of propagating waves in cerebral cortex across network states. *BMC Neuroscience* 12

Muller L, Destexhe A (2012) Propagating waves in thalamus, cortex and the thalamocortical system: Experiments and models. *Journal of Physiology-Paris* 106(5–6):222 – 238

Murakoshi T, Guo JZ, Ichinose T (1993) Electrophysiological identification of horizontal synaptic connections in rat visual cortex in vitro. *Neuroscience Letters* 163(2):211 – 214

Netter T, Franceschini N (2002) A robotic aircraft that follows terrain using a neuromorphic eye. In: Conf. Intelligent Robots and System, pp 129–134

Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *PNAS* pp 5419–5424

Peters A, Sethares C (1997) The organization of double bouquet cells in monkey striate cortex. *Journal of Neurocytology* 26(12):779 – 797

Pfeil T, Potjans TC, Schrader S, Potjans W, Schemmel J, Diesmann M, Meier K (2012) Is a 4-bit synaptic weight resolution enough? - constraints on enabling spike-timing dependent plasticity in neuromorphic hardware. *Frontiers in Neuroscience* 6(90)

Pfeil T, Grübl A, Jeltsch S, Müller E, Müller P, Petrovici MA, Schmuker M, Brüderle D, Schemmel J, Meier K (2013) Six networks on a universal neuromorphic computing substrate. *Frontiers in Neuroscience* 7:11

Renaud S, Tomas J, Bornat Y, Daouzli A, Saighi S (2007) Neuromimetic ICs with analog cores: an alternative for simulating spiking neural networks. In: Proceedings of the 2007 IEEE Symposium on Circuits and Systems (ISCAS2007)

Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W (1997) *Spikes - Exploring the neural code.* MIT Press, Cambridge, MA.

Rocke P, McGinley B, Maher J, Morgan F, Harkin J (2008) Investigating the suitability of fpaas for evolved hardware spiking neural networks. In: Hornby G, Sekanina L, Haddow P (eds) Evolvable Systems: From Biology to Hardware, Lecture Notes in Computer Science, vol 5216, Springer Berlin / Heidelberg, pp 118–129

Roxin A, Brunel N, Hansel D (2005) Role of delays in shaping spatiotemporal dynamics of neuronal activity in large networks. *Phys Rev Lett* 94:238,103

Schemmel J, Grübl A, Meier K, Muller E (2006) Implementing synaptic plasticity in a VLSI spiking neural network model. In: Proceedings of the 2006 International Joint Conference on Neural Networks (IJCNN), IEEE Press

Schemmel J, Brüderle D, Meier K, Ostendorf B (2007) Modeling synaptic plasticity within networks of highly accelerated I&F neurons. In: Proceedings of the 2007 IEEE International Symposium on Circuits and Systems (ISCAS), IEEE Press, pp 3367–3370

Schemmel J, Fieres J, Meier K (2008) Wafer-scale integration of analog neural networks. In: Proceedings of the 2008 International Joint Conference on Neural Networks (IJCNN)

Schemmel J, Brüderle D, Grübl A, Hock M, Meier K, Millner S (2010) A wafer-scale neuromorphic hard-

ware system for large-scale neural modeling. In: Proceedings of the 2010 IEEE International Symposium on Circuits and Systems (ISCAS), pp 1947–1950

Scholze S, Henker S, Partzsch J, Mayr C, Schuffny R (2010) Optimized queue based communication in vlsi using a weakly ordered binary heap. In: Mixed Design of Integrated Circuits and Systems (MIXDES), 2010 Proceedings of the 17th International Conference, pp 316 –320

Scholze S, Eisenreich H, Höppner S, Ellguth G, Henker S, Ander M, Hänzsche S, Partzsch J, Mayr C, Schüffny R (2011a) A 32 GBit/s communication SoC for a waferscale neuromorphic system. *Integration, the VLSI Journal*

Scholze S, Schiefer S, Partzsch J, Hartmann S, Mayr CG, Höppner S, Eisenreich H, Henker S, Vogginger B, Schüffny R (2011b) VLSI implementation of a 2.8GEvent/s packet based AER interface with routing and event sorting functionality. *Frontiers in Neuromorphic Engineering* 5(117):1–13

Schrader S, Diesmann M, Morrison A (2010) A compositionality machine realized by a hierarchic architecture of synfire chains. *Frontiers in Computational Neuroscience* 4

Serrano-Gotarredona R, Oster M, Lichtsteiner P, Linares-Barranco A, Paz-Vicente R, Gómez-Rodríguez F, Riis HK, Delbrück T, Liu SC, Zahnd S, Whatley AM, Douglas RJ, Häfliger P, Jimenez-Moreno G, Civit A, Serrano-Gotarredona T, Acosta-Jiménez A, Linares-Barranco B (2006) AER building blocks for multi-layer multi-chip neuromorphic vision systems. In: Weiss Y, Schölkopf B, Platt J (eds) Advances in Neural Information Processing Systems 18, MIT Press, Cambridge, MA, pp 1217–1224

Serrano-Gotarredona R, Oster M, Lichtsteiner P, Linares-Barranco A, Paz-Vicente R, Gomez-Rodriguez F, Camunas-Mesa L, Berner R, Rivas-Perez M, Delbruck T, Liu SC, Douglas R, Hafliger P, Jimenez-Moreno G, Ballcels A, Serrano-Gotarredona T, Acosta-Jimenez A, Linares-Barranco B (2009) Caviar: A 45k neuron, 5m synapse, 12g connects/s aer hardware sensory–processing–learning–actuating system for high-speed visual object recognition and tracking. *Neural Networks, IEEE Transactions on* 20(9):1417–1438

Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly nonrandom features of synaptic connectivity in cortical circuits. *PLOS Biology* 3(3):517–519

Telfeian AE, Connors BW (2003) Widely integrative properties of layer 5 pyramidal cells support a role for processing of extralaminar synaptic inputs in rat neocortex. *Neuroscience Letters* 343(2):121 – 124

Thomson AM, West DC, Wang Y, Bannister AP (2002) Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2-5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. *Cerebral Cortex* 12:936–953

Vogels TP, Abbott LF (2005) Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci* 25(46):10,786–95

Vogelstein RJ, Mallik U, Vogelstein JT, Cauwenberghs G (2007) Dynamically reconfigurable silicon array of spiking neuron with conductance-based synapses. *IEEE Transactions on Neural Networks* 18:253–265

Zou Q, Bornat Y, Tomas J, Renaud S, Destexhe A (2006) Real-time simulations of networks of hodgkin-huxley neurons using analog circuits. *Neurocomputing* 69:1137–1140

References

Alexander DM, Jurica P, Trengove C, Nikolaev AR, Gepshtein S, Zvyagintsev M, Mathiak K, Schulze-Bonhage A, Ruescher J, Ball T, van Leeuwen C (2013) Traveling waves and trial averaging: The nature of single-trial and averaged brain responses in large-scale cortical signals. *Neuroimage* 73: 95 – 112.

Ayzenshtat I, Meirovithz E, Edelman M, Werner-Reiss W, Bienenstock E, Abeles M, Slovin H (2010) Precise spatiotemporal patterns among visual cortical areas and their relation to visual stimulus processing. *J Neurosci* 30: 11232 – 11245.

Benucci A, Frazor R, Carandini M (2007) Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron* 55: 103 – 117.

Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU (2006) Complex networks: Structure and dynamics. *Physics reports* 424: 175 – 308.

Braitenberg V, Schüz A (1998) *Cortex: statistics and geometry of neuronal connectivity* Springer Berlin.

Bringuier V, Chavane F, Glaeser L, Frégnac Y (1999) Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. *Science* 283: 695 – 699.

Brüderle D, Petrovici MA, Vogginger B, Ehrlich M, Pfeil T, Millner S, Grübl A, Wendt K, Müller E, Schwartz MO et al. (2011) A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems. *Biological cybernetics* 104: 263 – 296.

Brunel N, Wang X (2003) What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *Journal of neurophysiology* 90: 415 – 430.

Brunel N, Hakim V (1999) Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural computation* 11: 1621 – 1671.

Bullier J (2001) Integrated model of visual processing. *Brain Res Rev* 36: 96 – 107.

Campbell PK, Jones KE, Normann RA (1993) Three-dimensional electrode device US Patent 5,215,088.

Chemla S, Chavane F (2010) A biophysical cortical column model to study the multi-component origin of the vsdi signal. *Neuroimage* 53: 420 – 438.

Chen Y, Anand S, Martinez-Conde S, Macknik S, Bereshpolova Y, Swadlow H, Alonso J (2009) The linearity and selectivity of neuronal responses in awake visual cortex. *J Vis* 9: 1 – 17.

Chen Y, Geisler W, Seidemann E (2006) Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat Neurosci* 9: 1412 – 1420.

Cruikshank S, Lewis T, Connors B (2007) Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nat Neurosci* 10: 462 – 468.

Destexhe A, Paré D (1999) Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol* 81: 1531 – 1547.

Destexhe A, Rudolph M, Fellous J, Sejnowski T (2001) Fluctuating synaptic conductances recreate in-vivo-like activity in neocortical neurons. *Neuroscience* 107: 13 – 24.

Destexhe A, Rudolph M, Paré D (2003) The high-conductance state of neocortical neurons *in vivo*. *Nat Rev Neurosci* 4: 739 – 751.

Destexhe A, Rudolph-Lilith M (2012) *Neuronal Noise* Springer, New York, 1 edition.

Destexhe A (2009) Self-sustained asynchronous irregular states and up–down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. *Journal of computational neuroscience* 27: 493 – 506.

Destexhe A (2011) Intracellular and computational evidence for a dominant role of internal network activity in cortical computations. *Current opinion in neurobiology* 21: 717 – 725.

Ermentrout B (1998) Neural networks as spatio-temporal pattern-forming systems. *Reports on progress in physics* 61:353.

Ermentrout G, Kleinfeld D (2001) Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. *Neuron* 29: 33 – 44.

Feldman M (2011) *Hilbert Transform Applications in Mechanical Vibration* Wiley, West Sussex, 1 edition.

Ferezou I, Bolea S, Petersen C (2006) Visualizing the cortical representation of whisker touch: Voltage-sensitive dye imaging in freely moving mice. *Neuron* 50: 617 – 629.

Ferezou I, Haiss F, Gentet L, Aronoff R, Weber B, Petersen C (2007) Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* 56: 907 – 923.

Gabor D (1946) Theory of communication. Part 1: The analysis of information. *J Inst EE - Part III: Radio and Communication Engineering* 93: 429 – 441.

Gabriel A, Eckhorn R (2003) A multi-channel correlation method detects traveling γ-waves in monkey visual cortex. *J Neurosci Meth* 131: 171 – 184.

Girard P, Hupé J, Bullier J (2001) Feedforward and feedback connections between areas v1 and v2 of the monkey have similar rapid conduction velocities. *J Neurophysiol* 85: 1328 – 1331.

González-Burgos G, Barrionuevo G, Lewis D (2000) Horizontal synaptic connections in monkey prefrontal cortex: an in vitro electrophysiological study. *Cereb Cortex* 10: 82 – 92.

Goswami J, Hoefel A (2004) Algorithms for estimating instantaneous frequency. *Signal Processing* 84: 1423 – 1427.

Goulet J, Ermentrout GB (2011) The mechanisms for compression and reflection of cortical waves. *Biological cybernetics* 105: 253 – 268.

Grienberger C, Konnerth A (2012) Imaging calcium in neurons. *Neuron* 73: 862 – 885.

Grinvald A, Hildesheim R (2004) VSDI: a new era in functional imaging of cortical dynamics. *Nat Rev Neurosci* 5: 874 – 885.

Grinvald A, Lieke E, Frostig R, Hilesheim R (1994) Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *J Neurosci* 14: 2545 – 2568.

Han F, Caporale N, Dan Y (2008) Reverberation of recent visual experience in spontaneous cortical waves. *Neuron* 60: 321 – 327.

Hubel D, Wiesel T (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160: 106 – 154.

Humphries MD, Gurney K (2008) Network 'small-world-ness': a quantitative method for determining canonical network equivalence. *PloS one* 3:e0002051.

Johansson M (1999) The Hilbert transform. Master's thesis, Växjö University.

Kuhn A, Aertsen A, Rotter S (2004) Neuronal integration of synaptic input in the fluctuation-driven regime. *J Neurosci* 24: 2345 – 2356.

Lubenov E, Siapas A (2009) Hippocampal theta oscillations are travelling waves. *Nature* 459: 534 – 539.

Marple S (1999) Computing the discrete-time "analytic" signal via FFT. *IEEE Trans Signal Process* 47: 2600 – 2603.

Marre O, Baudot P, Levy M, Frégnac Y (2005) High timing precision and reliability, low redundancy and low entropy code in v1 neurons during visual processing of natural scenes. In *Society for Neuroscience Abstracts*.

Marre O, Yger P, Davison AP, Frégnac Y (2009) Reliable recall of spontaneous activity patterns in cortical networks. *The Journal of Neuroscience* 29: 14596 – 14606.

Maynard E, Nordhausen C, Normann R (1997) The Utah intracortical electrode array: A recording structure for potential brain-computer interfaces. *Electroencephalography and Clinical Neurophysiology* 102: 228 – 239.

McCormick DA (1992) Neurotransmitter actions in the thalamus and cerebral cortex and their role in neuromodulation of thalamocortical activity. *Progress in neurobiology* 39: 337 – 388.

Meirovithz E, Ayzenshtat I, Bonneh Y, Itzhack R, Werner-Reiss U, Slovin H (2010) Population response to contextual influences in the primary visual cortex. *Cereb Cortex* 20: 1293 – 1304.

Metherate R, Cruikshank S (1999) Thalamocortical inputs trigger a propagating envelope of gamma-band activity in auditory cortex in vitro. *Exp Brain Res* 126: 160 – 174.

Mitra P, Pesaran B (1999) Analysis of dynamic brain imaging data. *Biophys J* 76: 691 – 708.

Mohajerani M, Chan A, Mohsenvand M, LeDue J, Liu R, McVea D, Boyd J, Wang Y, Reimers M, Murphy T (2013) Spontaneous cortical activity alternates between motifs defined by regional axonal projections. *Nat Neurosci* 16: 1426 – 1435.

Mohajernai M, McVea D, Fingas M, Murphy T (2010) Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *J Neurosci* 30: 3745 – 3751.

Mountcastle V (1957) Modality and topographic properties of single neurons of cat's somatic sensory cortex. *J Neurophysiol* 20: 408 – 434.

Muller L, Reynaud A, Chavane F, Destexhe A (2014) The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nature Communications* In press.

Muller L, Polack P, Contreras D, Destexhe A (2010) Analysis of voltage-sensitive dye imaging data during propagating cortical waves in mouse visual cortex reveals fine structure and state dependence. In *Society for Neuroscience Abstracts*.

Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. *Nature neuroscience* 12: 70 – 76.

Newman ME (2003) The structure and function of complex networks. *SIAM review* 45: 167 – 256.

Papoulis A (1977) *Signal Analysis* McGraw-Hill, New York, 1 edition.

Paré D, Shink E, Gaudreau H, Destexhe A, Lang EJ (1998) Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons in vivo. *Journal of Neurophysiology* 79: 1450 – 1460.

Peyrache A, Dehghani N, Eskandar EN, Madsen JR, Anderson WS, Donoghue JA, Hochberg LR, Halgren E, Cash SS, Destexhe A (2012) Spatiotemporal dynamics of neocortical excitation and inhibition during human sleep. *Proc Natl Acad Sci USA* 109: 1731 – 1736.

Pinto DJ, Ermentrout GB (2001) Spatially structured activity in synaptically coupled neuronal networks: I. traveling fronts and pulses. *SIAM journal on Applied Mathematics* 62: 206 – 225.

Poulet J, Petersen C (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454: 881 – 885.

Prechtl J, Cohen L, Pesaran B, Mitra P, Kleinfeld D (1997) Visual stimuli induce waves of electrical activity in turtle cortex. *Proc Natl Acad Sci U S A* 94: 7621 – 7626.

Reimer A, Hubka P, Engel AK, Kral A (2011) Fast propagating waves within the rodent auditory cortex. *Cerebral Cortex* 21: 166 – 177.

Reynaud A, Takerkart S, Masson G, Chavane F (2011) Linear model decomposition for voltage-sensitive dye imaging signals: Application in awake behaving monkey. *NeuroImage* 54: 1196 – 1210.

Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. *The Journal of Neuroscience* 22: 5639 – 5651.

Roelfsema PR, Engel AK, Konig P, Singer W (1997) Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385: 157 – 161.

Roland P, Hanazawa A, Undeman C, Eriksson D, Tompa T, Nakamura H, Valentiniene S, Ahmed B (2006) Cortical feedback depolarization waves: A mechanism of top-down influence on early visual areas. *Proc. Natl. Acad. Sci. U.S.A.* 103: 12586 – 12591.

Roxin A, Brunel N, Hansel D (2005) Role of delays in shaping spatiotemporal dynamics of neuronal activity in large networks. *Physical review letters* 94:238103.

Rubino D, Robbins K, Hatsopoulos N (2006) Propagating waves mediate information transfer in the motor cortex. *Nature Neurosci* 9: 1549 – 1557.

Rudolph M, Pospischil M, Timofeev I, Destexhe A (2007) Inhibition determines membrane potential dynamics and controls action potential generation in awake and sleeping cat cortex. *J Neurosci* 27: 5280 – 5290.

Sanchez-Vives M, McCormick D (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat Neurosci* 3: 1027 – 1034.

Slovin H, Arieli A, Hildesheim R, Grinvald A (2002) Long-term voltage-sensitive dye imaging reveals cortical dynamics in behaving monkeys. *J Neurophysiol* 88: 3421 – 3428.

Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *The Journal of Neuroscience* 28: 12591 – 12603.

Stroh A, Adelsberger H, Groh A, Rühlmann C, Fischer S, Schierloh A, Deisseroth K, Konnerth A (2013) Making waves: initiation and propagation of corticothalamic Ca^{2+} waves in vivo. *Neuron* 77: 1136 – 1150.

Tan AY, Chen Y, Scholl B, Seidemann E, Priebe NJ (2014) Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature.* In press.

Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393: 440 – 442.

Witte R, Rousche P, Kipke D (2007) Fast wave propagation in auditory cortex of an awake cat using a chronic microelectrode array. *J Neural Eng* 4: 68 – 78.

Wu J, Guan L, Bai L, Yang Q (2001) Spatiotemporal properties of an evoked population activity in rat sensory cortical slices. *J Neurophysiol* 86: 2461 – 2474.

Wu J, Guan L, Tsau Y (1999) Propagating activation during oscillations and evoked responses in neocortical slices. *J Neurosci* 19: 5005 – 5015.

Xing D, Ringach DL, Shapley R, Hawken MJ (2004) Correlation of local and global orientation and spatial frequency tuning in macaque V1. *The Journal of Physiology* 557: 923 – 933.

Xu W, Huang X, Takagaki K, Wu J (2007) Compression and reflection of visually evoked cortical waves. *Neuron* 55: 119 – 129.