# Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer

M. Djurfeldt
M. Lundqvist
C. Johansson
M. Rehn
Ö. Ekeberg
A. Lansner

*Biologically detailed large-scale models of the brain can now be simulated thanks to increasingly powerful massively parallel supercomputers. We present an overview, for the general technical reader, of a neuronal network model of layers II/III of the neocortex built with biophysical model neurons. These simulations, carried out on an IBM Blue Gene/L™ supercomputer, comprise up to 22 million neurons and 11 billion synapses, which makes them the largest simulations of this type ever performed. Such model sizes correspond to the cortex of a small mammal. The SPLIT library, used for these simulations, runs on single-processor as well as massively parallel machines. Performance measurements show good scaling behavior on the Blue Gene/L supercomputer up to 8,192 processors. Several key phenomena seen in the living brain appear as emergent phenomena in the simulations. We discuss the role of this kind of model in neuroscience and note that full-scale models may be necessary to preserve natural dynamics. We also discuss the need for software tools for the specification of models as well as for analysis and visualization of output data. Combining models that range from abstract connectionist type to biophysically detailed will help us unravel the basic principles underlying neocortical function.*

## Introduction

Within a volume of 1,400 $cm^3$, the human brain provides a computing capacity vastly surpassing that of today's supercomputers while consuming only 20 W of power. The 20 billion neurons in our cerebral cortex, together with the connectivity supplied by the millions of kilometers of axons in the underlying white matter [1], are thought to subserve most of our higher functions, such as emotion, planning, thought, and memory. While a century of intense study has produced a huge body of hard scientific facts about neurons, their interactions, and the cortical architecture, the principles of cortical information processing remain enigmatic.

One of the reasons for this is that while experimental techniques have provided data at a brain-scale level and at the single-cell level, data at the intermediate network level is still scarce. Improvements in existing methodology and novel techniques, such as multielectrode recordings and calcium imaging [2], are beginning to fill this gap. This development must be paralleled by the strong development of hypotheses and quantitative models of cortical network function in order to guide experiments and the interpretation of experimental data.

The aim of this paper is to show the general technical reader how large-scale simulation of brain networks on supercomputers can bridge a range of spatial scales of brain organization and work as an interface between functional hypotheses of network computation and empirical data. The models used in the simulations we present are substantially upscaled versions of the model in Lundqvist et al. [3]. In order to be able to run simulations of this size, we have made a series of adaptations to our simulation software [4]. Novel results include the description of a dynamic ground state (the state when not



**0018-8646/08/$5.00 © 2008 IBM**

actively recalling a memory) showing alpha-like oscillations and the presentation of an artificial voltage-sensitive dye signal [5].

### Large-scale models of the brain

The structure of the nervous system has several spatial scales with substructures such as molecules, synapses, neurons, microcircuits, networks, regions, and systems (see, e.g., Churchland and Sejnowski [6]). Models are often formulated on a single or a few of these scales. In simulations of brain imaging data, the state variables of the model often correspond to regions or groups of neurons [7]. When simulating a neuron, the state variables can be variables gating ionic currents of the excitable membrane.

When addressing questions concerning areas where dynamics at the level of the neuron and at the synapse have an effect at the network level, the model must span all of these levels. As a consequence, we get a large set of state variables, i.e., a *large-scale model* [8]. In this paper, we report a model with 22 million neurons, 11 billion synapses, and on the order of 40 billion state variables. It should be noted, though, that a large number of state variables does not necessarily imply a large number of parameters. For example, members of a specific cell type may share a basic set of parameters across an entire neuronal population.

It is difficult to understand a global brain network from modeling only a local network of some hundred cells, such as a cortical minicolumn (see the section "Minicolumns in the cerebral cortex" below), or from dramatically subsampling the global network by letting one model neuron represent an entire cortical column or area. One reason is that using small or subsampled network models leads to unnatural connectivity and dynamics.

A network model comprises neurons modeled after empirical data. Just like their real counterparts, these model neurons require sufficient synaptic input current to become activated. In a small network, model neurons are bound to have very few presynaptic neurons. Thus, it is necessary to exaggerate either connection density or synaptic conductance, and most of the time, both are necessary.

This results in a network with a few but strong signals circulating—a stark contrast to the real cortical network in which many weak signals interact. Such differences tend to significantly distort the network dynamics. For example, artificial synchronization can easily arise, which is a problem, especially as synchronization is one of the more important phenomena one might want to study. By modeling the full network with a one-to-one correspondence between real and model neurons, such problems are avoided.

### Bottom-up and top-down

The *bottom-up* approach to modeling means using the level of physical realization as a starting point with the hope of capturing function as an emergent phenomenon. What does the anatomy of the cerebral cortex mean? If we can, from the physical level of synapses, dendrites, neurons, and networks, identify computational primitives of the cortex, such primitives can be abstracted, and we can move up one level of analysis. This strategy is currently pursued in ambitious modeling projects, such as the Blue Brain Project [9, 10], which has as one of its goals the building of a model cortical column on an IBM Blue Gene/L* (BG/L) supercomputer. The growing availability of detailed experimental data will make this type of model increasingly worthwhile, especially since we are now awaiting the results of an industrial-scale dissection of the cortical column at the nanoscale level [11].

However, when modeling the complex and intricate structure of the cortex, it turns out that we may need information from additional sources. Some model parameters are well constrained by experiment, while others (e.g., the structure of long-range connectivity) are still largely unknown. Hypotheses of cortical function, as expressed in more abstract models, can guide model development in selecting which elements to include, in giving additional constraints, and in filling in where empirical data is still missing. This is the *top-down* approach to modeling.

In practice, the approach of the modeler is usually neither purely top-down nor purely bottom-up. The model described here is the result of an integration of functional constraints given by a theoretical view of the neocortex as an associative attractor memory network and empirical constraints given by cortical anatomy and physiology. (*Attractor* is a concept in the theory of dynamic systems; see the next section.)

### Cell assemblies and attractor networks

The view of the cortex as an attractor network originated more than 50 years ago in Hebb's cell assembly theory [12] (see, e.g., Fuster [13] for a review). Hebb suggested that the functional unit of the cortex is a subset of neurons that are repeatedly active together, and that such a *cell assembly* is the basis of mental representation. The thought of an apple would invoke one cell assembly, and the thought of an orange another. The theory has since been mathematically instantiated in the form of the Willshaw–Palm [14, 15] and Little–Hopfield models [16] and has subsequently been elaborated on and analyzed in great detail [17, 18]. This has resulted in the view of the persistent firing of cell assemblies as attractors in a dynamical system.

**Table 1** Estimates of the number of minicolumns and hypercolumns for a selection of mammals.

| | *Human* | *Macaque* | *Cat* | *Rat* | *Mouse* |
|---|---|---|---|---|---|
| *Minicolumns* | $2.0 \cdot 10^8$ | $2.0 \cdot 10^7$ | $6.1 \cdot 10^6$ | $5.0 \cdot 10^5$ | $1.6 \cdot 10^5$ |
| *Hypercolumns* | $2.0 \cdot 10^6$ | $2.0 \cdot 10^5$ | $6.1 \cdot 10^4$ | $5.0 \cdot 10^3$ | $1.6 \cdot 10^3$ |

An *attractor* is a set of points in the phase space such that a trajectory that gets close enough will remain close. Activating enough cells of a cell assembly will cause more and more of the assembly to become activated—the system evolves toward the attractor. Because of nonlinearity, a memory network can have multiple attractors, each representing one memory. The connections in the network form the landscape of attractors. When the state of the system evolves into the attractor, the memory has been recalled. If a memory is a conjunction of attributes, activating a subset of attributes will put the system close to the attractor representing the memory, leading to the recall of the full conjunction—the network has made an *association*.

The olfactory cortex [19] and the hippocampal CA3 field [20] have previously been perceived and modeled as prototypical neuronal autoassociative attractor memory networks. More recently, sustained activity in an attractor memory of a similar kind has been proposed to underlie prefrontal working memory [21].

### Minicolumns in the cerebral cortex

If cell assemblies or attractor states are the basis of cortical function, how do they relate to cortical anatomy? In a classic work on the visual system, Hubel and Wiesel [22] penetrated the primary visual cortex with a recording electrode. They found that cells responded most strongly to a specific orientation of an oblong bar in the visual field and that the preferred angle seemed to shift discretely as the electrode moved tangentially to the cortical surface while cells along a line normal to the surface tended to have similar response properties. They deduced that the basic unit of cortical organization must be what they called a *functional column* and suggested that the cortex is a lattice of such columns. They also suggested that sets of such columns are grouped into larger entities, *hypercolumns*, that together form a complete representation of all possible attribute values within each region of retinotopic space.

Later, Mountcastle [23] suggested the concept of *anatomical minicolumns*. These were described in detail by Peters and Sethares [24]. Could such a minicolumn, consisting of some hundred neurons, be the basic functional unit of the cortex? If so, a Hebbian cell assembly may consist of a set of such columns, and the activation of these columns would correspond to entering a dynamic attractor of the cortical network. Here, we note that the same cortical tissue can have multiple, overlapping maps of different stimulus attributes, such as orientation, direction, and ocular dominance. Also, the spatial organization of stimulus attributes in sensory cortices varies across species, and the fact that there is still no firm evidence connecting *functional* with *anatomical* columns means that the hypothesis of the column as the basic functional unit of the cortex remains controversial. An alternative possibility would be that a local cluster of interconnected anatomical minicolumns serves as a functional unit.

The organization of the visual cortex into minicolumns and hypercolumns has inspired our view of cortical associative memory, which has been expressed in the form of an abstract neural network model [25–27] and in biophysically detailed models [3, 28].

The work of Johansson and Lansner [29] presents hypothetical estimates of the number of available minicolumns and hypercolumns in the brains of a number of mammals (**Table 1**). Their data is based on the assumption of an average minicolumn diameter of 40 $\mu$m and a hypercolumn diameter of 400 $\mu$m.

## The model

On the basis of anatomical criteria, the cerebral cortex is divided into six layers ranging from the cortical surface to the gray matter to white matter border. The pattern of connections suggests different roles for these layers. For example, incoming connections typically contact layer IV, while layers II/III send outgoing connections to areas further down a processing stream. The simulations in this paper are based upon a model of layers II/III of the association cortex of the rat. It is an upscaled version of the model presented by Lundqvist et al. [3].

The overall architecture of our model is shown in **Figure 1**. Figure 1(a) illustrates the geometric layout of a subset of 100 hypercolumns in the plane of the cortical sheet, each marked with a distinct color. Each hypercolumn consists of 100 minicolumns. Figure 1(b) shows the schematic connectivity of the model. Each minicolumn contains 30 pyramidal cells that excite each other through short-range axons. Pyramidal cells project locally as well as to pyramidal cells in other minicolumns that belong to the same cell assembly and to regular-spiking nonpyramidal (RSNP) cells in minicolumns

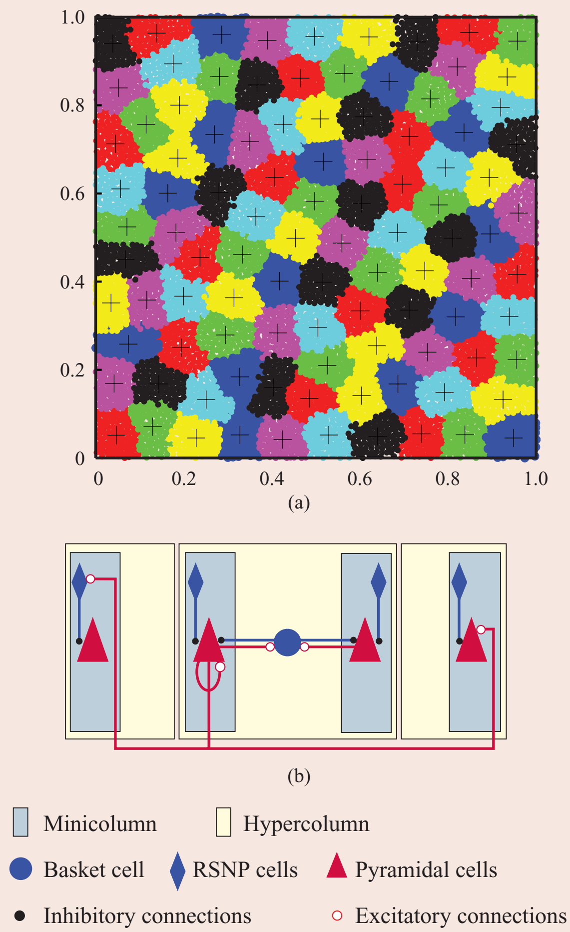


**Figure 1**

Model architecture: (a) the geometric layout of 100 hypercolumns consisting of 100 minicolumns each; (b) schematic connectivity of the model.

belonging to other assemblies. The basket cell normalizes activity in the local hypercolumn. RSNP cells provide local inhibition of pyramidal cells.

The long-range projection between distant minicolumns constitutes the memory matrix of the attractor memory and defines the cell assemblies: Only minicolumns belonging to the same memory pattern or cell assembly excite each other.

Each hypercolumn also contains a population of 100 inhibitory basket cells, which are excited by the pyramidal cells of that hypercolumn. The basket cells, in turn, inhibit the pyramidal cells of that hypercolumn, thereby providing a mechanism for normalization of activity. This enables the hypercolumn to operate like a winner-take-all module in which different patterns can compete. Each minicolumn also contains two inhibitory RSNP cells, which contact the local pyramidal cells. The abstract neural network model upon which the long-range connectivity of the present model is based suggests an additional way in which cell assemblies can compete. This competition has been realized through long-range axons from pyramidal cells to RSNP cells of minicolumns that belong to other cell assemblies. Since RSNP cells inhibit pyramidal cells within the local minicolumn, the activity of targeted assemblies will be suppressed. Such connections have not yet been identified anatomically. Connectivity is otherwise compatible with experimental data, to the extent that it is available. However, for simplicity, we have made the borders of minicolumns and hypercolumns sharp, in contrast to the local approximately Gaussian structure observed experimentally (see, e.g., Buzás et al. [30]). For details, see Lundqvist et al. [3]. A separate set of cells model cortical layer IV, which provides input to the pyramidal cells described above. External input to the attractor memory is provided as simulated synaptic events in these cells.

Cells are modeled using the Hodgkin–Huxley formalism [31, 32] in which the cell membrane potential $V(t)$ of a neural compartment is expressed as the differential equation

$$C_{\mathrm{m}} \frac{\mathrm{d}V}{\mathrm{d}t} = I_{\mathrm{comp}} + I_{\mathrm{cond}} + I_{\mathrm{syn}}, \tag{1}$$

where $C_{\mathrm{m}}$ is the membrane capacitance, $I_{\mathrm{comp}}$ the sum of currents from adjacent compartments, $I_{\mathrm{cond}}$ the sum of ionic currents through channels in the cell membrane, and $I_{\mathrm{syn}}$ the sum of synaptic currents. The electrical behavior of the cell is determined by the ionic currents that are described through activation and inactivation variables. For example, the delayed rectifier current, carried by potassium ions, is described by

$$I_{\mathrm{Kr}} = (E_{\mathrm{Kr}} - V(t))\, G_{\mathrm{Kr}} n^4. \tag{2}$$

Here, $n$ is an activation function described by

$$\frac{\mathrm{d}n}{\mathrm{d}t} = \alpha_n(1-n) - \beta_n n, \tag{3}$$

where $\alpha_n$ and $\beta_n$ depend nonlinearly on $V(t)$. A pyramidal cell in our model consists of six compartments. Each compartment has one state variable representing the membrane potential and carries up to five ionic currents, with one to two state variables per current. Some compartments have a flow of calcium into an intracellular store, which is represented by an additional state variable. Furthermore, some synapses carry a separate flow of calcium with yet another associated state variable. Synapses are generally governed by three state variables, one for the degree of opening and two for short-term changes in synapse strength (facilitation and depression).

For the simulations in Lundqvist et al. [3] and here, an orthogonal set of nonoverlapping memory patterns was formed. One minicolumn was selected from each hypercolumn to form one pattern. A long-range connection between two distant minicolumns was formed stochastically if the two minicolumns belonged to the same pattern. Thus, each pyramidal cell received long-range excitation from only a subset of the cells in the pattern. Similarly, each RSNP cell received excitation from a subset of pyramidal cells belonging to minicolumns of foreign patterns.

## SPLIT simulator

The development of parallel simulation in computational neuroscience has been relatively slow. Today, there are few publicly available parallel simulators, and they are not nearly as general, flexible, and well documented as the more commonly used serial simulators, such as Neuron [33, 34] and Genesis [35, 36]. However, there is some ongoing development. For Genesis there is PGenesis, and the parallel version of Neuron has just been released. In addition, there exists simulators such as NeoCortical Simulation (NCS) [37], Neural Simulation Toolbox (NEST) [38], and our own parallelizing simulator SPLIT [39]. However, they are in many ways still in the experimental and developmental stage.

The SPLIT simulator [39] was developed in the mid-1990s with the aim of exploring how to efficiently use the resources of various parallel computer architectures for large-scale biophysically detailed neuronal-network simulations. The simulator has also served as a platform for experiments with communication algorithms.

SPLIT takes the form of a C++ library that is linked into the user program. The SPLIT application programming interface is provided by an object of the class `split`, which is the only means of communicating with the library. The user program specifies the model using method calls on the `split` object. The user program is serial and can be linked with a serial or parallel version of the library. Parallelism is thus completely hidden from the user. In the parallel case, the serial user program runs in a master process that communicates, through mechanisms internal to the SPLIT library, with a set of slave processes. On clusters, SPLIT uses Message Passing Interface (MPI).

The library exploits data locality for better cache-based performance. To benefit from vector architectures, state variables are stored in sequence. It uses adjacency lists for compact representation of the neural projections and address event representation (AER) for spike events [40].

The neurons in the model can be distributed arbitrarily over the set of slaves. This gives great freedom in optimizing communication so that densely connected neurons reside on the same CPU and so that axonal delays between neurons simulated on different slaves are maximized. This way, CPUs do not have to communicate as often, giving higher efficiency.

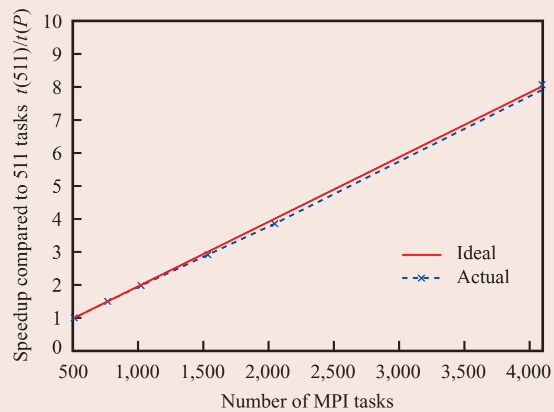


**Figure 2**

Speedup for a simulation of a model with 4 million cells and 2 billion synapses on the BG/L platform. Data points up to 2,048 processors were collected on the Rochester BG/L system, while the last data point was obtained on the Watson Research BG/L system. ($t$: time; $P$: number of tasks.)

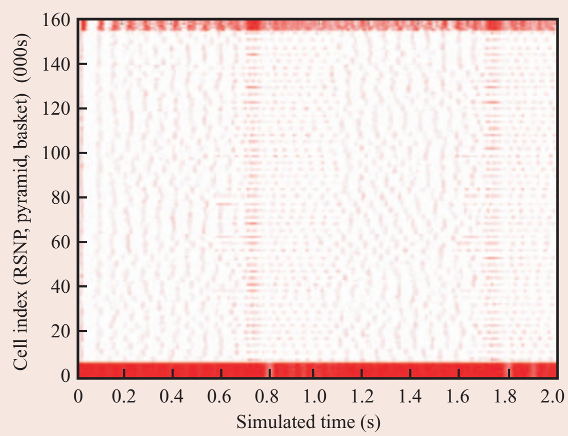


**Figure 3**

Raster plot for a simulation of 49 hypercolumns with 100 minicolumns per hypercolumn.

SPLIT also makes use of a novel abstraction, the connection-set algebra, which implements an efficient domain decomposition of the connectivity metadata. With connection-set algebra, network connectivity structure can be described in a modular way. It provides a set of basic types of connectivity structure and a set of operators by which it is possible to describe new types of connectivity as combinations of existing types [4].

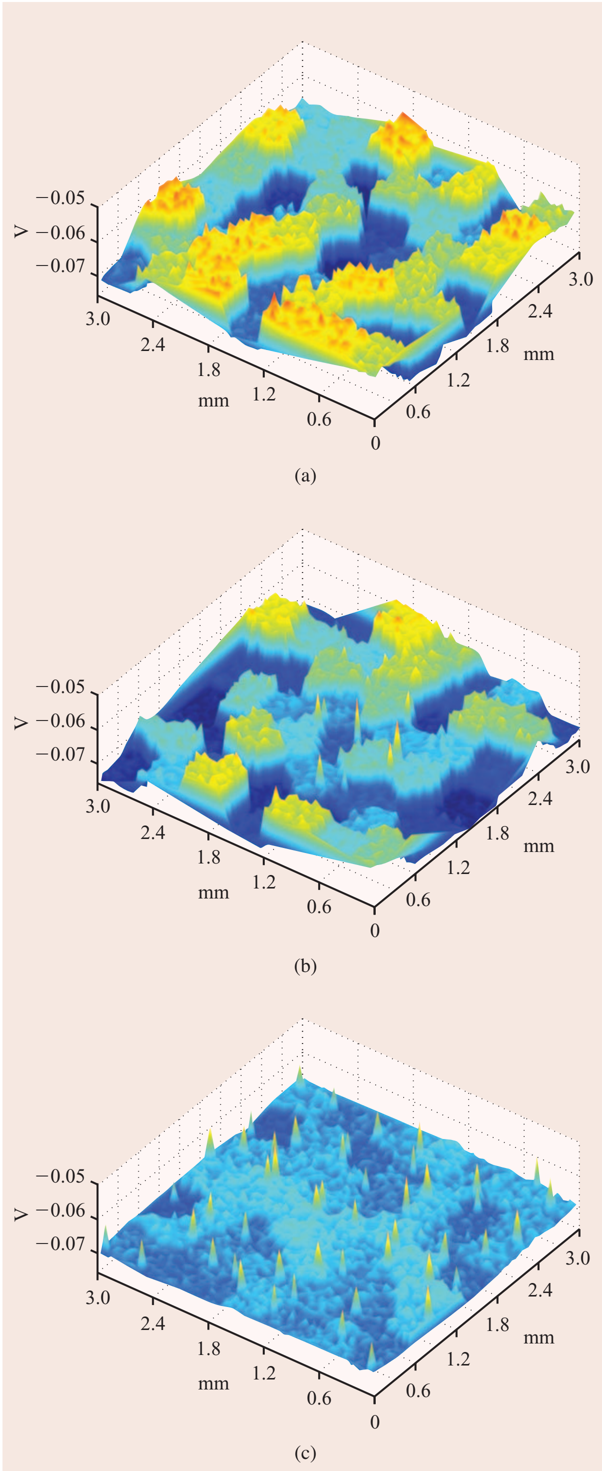


Figure 4

Simulated VSD signal of 4,900 minicolumns in a simulated cortical patch of size 3 × 3 mm: (a) the ground state with waves of hypercolumnar activity; (b) a part of a memory pattern is stimulated through layer IV; (c) the network has attained an attractor memory state.

We ran our simulations on the BG/L system installations at IBM Rochester, Minnesota, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and at The Royal Institute of Technology (KTH), Stockholm. **Figure 2** shows the speedup for a model with 4 million cells and 2 billion synapses on the BG/L platform up to 4,096 processors. Data points up to 2,048 processors were collected on the Rochester BG/L platform, while the last data point was obtained on the Watson Research BG/L platform. Results are given for 1 second of simulated time. The last run (4,096 processors) took 1,205 seconds.

## Simulation results

**Figure 3** shows the spiking activity of a simulation of 49 hypercolumns (100 minicolumns each). Each row of the raster plot shows the spikes of one neuron. The lower portion of the raster (the first 9,800 cells) shows activity in all RSNP cells, the mid-portion (147,000 cells) shows the pyramidal cells, and the upper portion (4,900 cells) shows the basket cells. The long-range pyramidal–pyramidal and pyramidal–RSNP synapses store orthogonal memories. As a consequence of stimulation of the cells in layer IV between 0.5 seconds and 0.64 seconds and between 1.5 seconds and 1.64 seconds, the network state can be seen switching from a ground state to an active memory state. Only cells in layer IV representing a part of one of the memory patterns stored in the inter-hypercolumnar memory matrix are stimulated. This partial pattern is quickly completed to the full memory pattern. This behavior was robust for all memories stored, and it shows that on a population level (although each pyramidal cell connects only to a random subset of cells in the pattern), the cells have formed a cell assembly corresponding to the pattern. Apart from pattern completion, the model is capable of all the functionality usually ascribed to attractor networks, such as noise reduction and resolution of ambiguity.

One of the experimental techniques used to study activity in populations of real neurons is to record changes in the color of a voltage-sensitive dye (VSD). **Figure 4** shows a synthesized VSD signal for a model network with 49 hypercolumns during three phases of activity. The signal was computed as the low-pass filtered sum of the membrane potentials of all cells in each minicolumn. Figure 4(a) shows the ground state condition, and Figure 4(b) shows the VSD signal just after stimulation of a partial pattern. In Figure 4(c), the network has completed the shift to the active memory state.

In the present version of the model, the ground state is characterized by oscillations at a frequency of approximately 15 Hz [**Figure 5(a)**], where the oscillations of individual minicolumns are phase-locked to other minicolumns in the hypercolumn. The coexistence of a stable ground state with active memory states was first

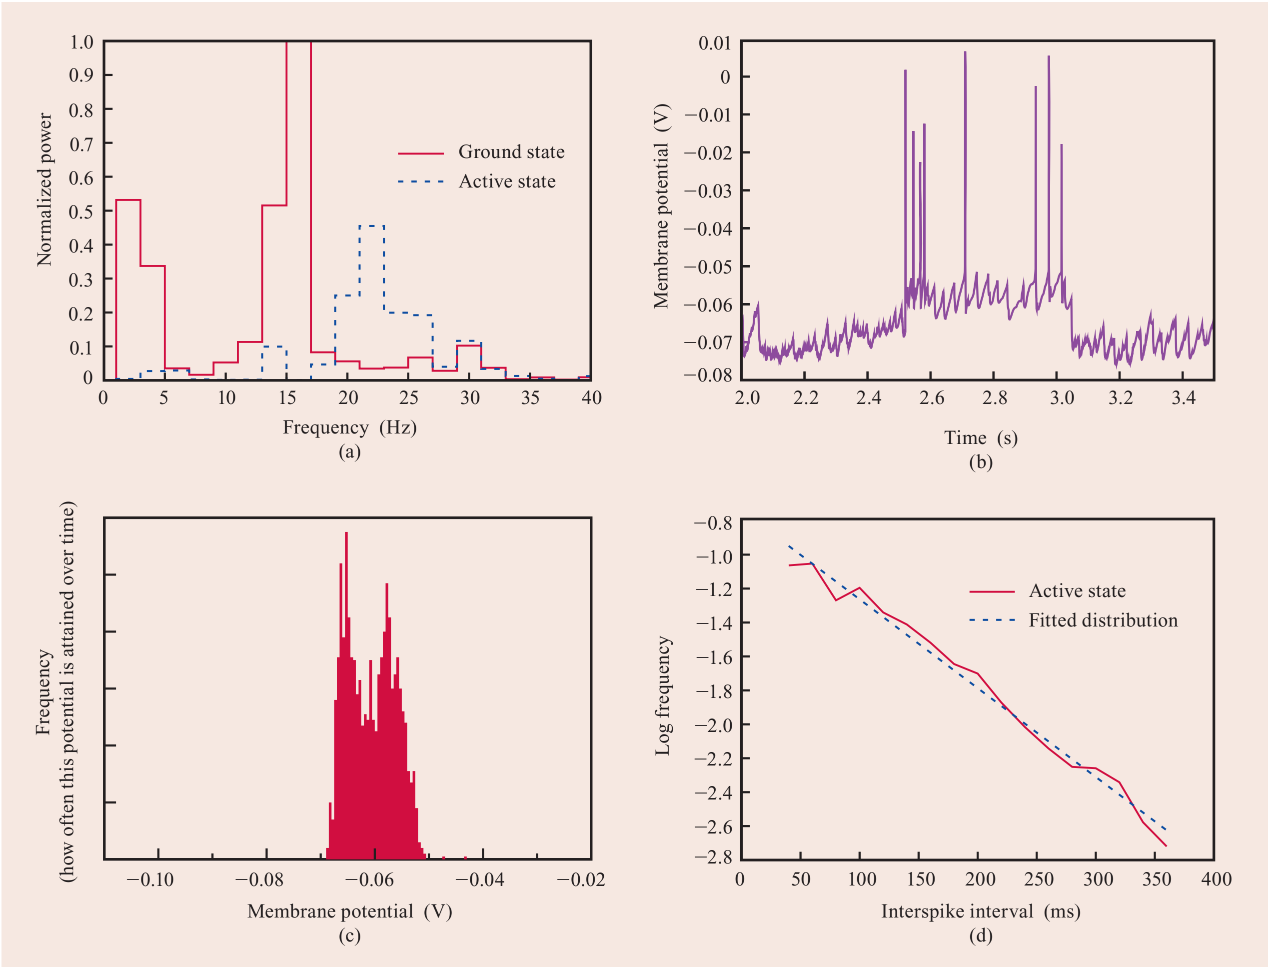


**Figure 5**

Characterization of dynamic states of the model: (a) normalized power spectrum for a VSD signal from one minicolumn; (b) a model pyramidal neuron switching from a hyperpolarized state (similar to a down state) to a more depolarized state (similar to an up state) and back; (c) membrane potential histogram; (d) exponentially distributed activity of pyramidal cells in the depolarized state in a network with 25 hypercolumns.

shown in a model of delay period activity in the prefrontal cortex [41] and has been further analyzed in, e.g., [42]. In the ground state of our model, pyramidal cells fire at about 0.1 Hz. The 15-Hz rhythm thus emerges as a collective network-level phenomenon. The fact that it appears only when there is no input to the network and the network is not in one of its active memory states is suggestive of the class of alpha rhythms, which has been proposed to reflect cortical idling [43]. This frequency lies close to the alpha band and is consistent with the cat mu rhythm [44].

In the active state, only the pyramidal cells of a single minicolumn are active in each hypercolumn. In this state, pyramidal cells fire at 10–15 Hz, basket cells at 50 Hz, and RSNP cells at 25–35 Hz. One particularly interesting phenomenon, which consistently arises in our simulations for a broad range of parameters and all model sizes, is a rhythmic modulation of pyramidal cell activity with a frequency of 25–40 Hz [Figure 5(a)]. It occurs only during active states and is reminiscent of the gamma-band activity observed in working memory tasks [45].

Our model exhibits some emergent phenomena that have also been observed in the brain. When the network attains an active memory state, pyramidal cells participating in the active cell assembly are bombarded with synaptic events. This elevates their membrane potential so that the global shift of dynamic state is reflected in a shift of their membrane potential from a hyperpolarized state to a more depolarized state. **Figure 5(b)** shows a model pyramidal neuron undergoing such a shift. The mean membrane potential becomes elevated at $t = 2.5$ s as a result of network activity when the neuron starts participating in an active attractor state. The attractor state was activated through simulated electrical stimulation of other member neurons.

These states are very similar, respectively, to the *down* and *up* states that have been observed physiologically [46, 47]. These shifts cause many cells to have a bimodal distribution of membrane potentials [**Figure 5(c)**], also consistent with physiology [48]. Despite the regularities seen on a network level, the firing of individual pyramidal cells is Poisson distributed in the active state [**Figure 5(d)**], which shows exponentially distributed activity of pyramidal cells in the depolarized state in a network with 25 hypercolumns. Inter-spike intervals (ISIs) (9,443 spikes) were collected from all pyramidal cells during a total of 2 seconds of simulated time when the network was in a globally active state, i.e., one in which a subset of cells shows up-state activity. The logarithm of the distribution of ISI was plotted as a function of ISI length. An exponential distribution was fitted to the data and is shown as a straight line ($r^2$ was 0.98 for the exponential fit and 0.86 for a power-law distribution [not shown]). Again, this result is consistent with physiology [see, e.g., Bédard et al. [49]). One particularly attractive feature of the model is that it is robust to the perturbation of parameters [3], which is to be expected from a biological system.

At the Thomas J. Watson Research Center, we had the opportunity to perform a run on 8,192 processors on the BG/L system. The model chosen for that run is the largest full-scale Hodgkin–Huxley type of model of a cortical patch ever simulated. The simulation, running in coprocessor mode, occupied 336 MB of memory at each node, giving a total of 2.8 TB. We simulated 22 million neurons and 11 billon synapses, which corresponds to a cortical surface area of 16 cm$^2$, comparable to the cortex of a small mammal. While real pyramidal cells have 10,000 synapses, the average number of synapses per neuron is only 500 in our model because of both the orthogonality of the memory matrix and the lack of connections to other cortical areas. One answer we sought was whether such a large patch of cortex could maintain a stable attractor state despite the significant propagation delays caused by axonal conduction time. The simulation showed that this is indeed the case.

## Discussion

Real vertebrate neuronal networks typically comprise millions of neurons and billions of synaptic connections. They have a complex and intricate structure, including a modular and layered layout of the neurons at several levels (e.g., cortical minicolumns, hypercolumns, and areas). It can still be useful to model a single module or microcircuit in isolation (e.g., in relation to a correspondingly reduced experimental *in vitro* preparation, such as a cortical slice), but such a module is, by definition, a component in a much larger network. In a real nervous system, it is embedded in a mosaic of other modules and receives afferent input from multiple other sources. We may therefore need to consider an entire network of modules.

However, partly because of limited computing resources, computational studies have often focused on problems that can be approached at the cellular and small-network levels (e.g., the emergence and dynamics of local receptive fields in the primary visual cortex [50]). Modeling at the large-network level poses specific challenges, but ultimately we believe that it is necessary to build models that take global and dynamic network interactions into account if we are to understand the functioning of a neuronal system. We have demonstrated that today's computing technology makes such simulations possible.

One application of biologically detailed, large-scale neuronal network models is to study the network effects of pharmacological manipulations. This can be done in great detail to answer such questions as "Does the addition of the drug alter the firing patterns of a particular neuron?" and "What happens to the activity of a large ensemble of neurons?" It might even be possible to study effects on functional properties, such as memory function. Other types of manipulations are also possible, including the addition or removal of different cell types or the rearrangement of connectivity patterns.

Large-scale network models will help to bridge the gap between the neuronal and synaptic phenomena and the phenomenology of global brain states and dynamics, as observed by such techniques as extracellular recording, local field potential, electroencephalogram, magnetoencephalography, VSD, and blood-oxygenation-level-dependent signal.

Models of working memory and persistent activity (see, e.g., Tegnér et al. [51] and Compte et al. [21]) usually rely on local, purely Gaussian connectivity. The model presented here is unique in that it incorporates the more long-range, patchy connectivity that is observed anatomically. In our model, this connectivity stems from pyramidal cells contacting distant minicolumns. The model makes use of two types of inhibitory interneurons to achieve natural firing rates of the ground state and active states of the cortical network.

Some care is required when comparing the results in Figure 4 to a real VSD signal. First, the real VSD signal usually has a somewhat lower spatial resolution. Second, the visualization in Figure 4 is shaped by some artificial features of connectivity structure in the present model. Minicolumn and hypercolumn connectivity is delimited by sharp boundaries, which is not the case in the neocortex. While it is likely that this difference does not cause a major change in network dynamics, a more diffuse connectivity would cause a smoother VSD signal. We will address this issue in future work. Third, topographical representations of attributes, present in many areas of the brain, would give rise to denser connections with nearby minicolumns than those far away. This is likely to give a more wavelike organization to the alpha-like oscillations of the ground state shown in Figure 4(a). In addition, the thalamus, not included in our model, might be even more important than the cortex in shaping such rhythms [44].

We are currently working on extending the present model with a layer V and a more complete model of layer IV. These extensions will enable modeling of systems consisting of more than one cortical area. An exciting possibility is, then, to use information from connectivity databases, such as the collation of connectivity data on the macaque brain (CoCoMac) [52], in determining the architecture of the larger network.

## Conclusion

We have presented full-scale, biologically detailed models of layers II/III of the cerebral cortex at a scale corresponding to the cortex of a small mammal. These models combine the functional hypothesis of the cortex as an attractor memory network with empirical constraints of known anatomy and physiology. The models do indeed have the capabilities of an abstract attractor memory network, such as pattern completion and resolution of ambiguities, while at the same time they display some of the phenomenology of the living brain.

Computational neuroscience will benefit greatly from the current development of new, affordable, massively parallel computers likely to enter the market in the next few years. This development constitutes an enabling technology for the modeling of complex and large-scale neuronal networks representing multiple cortical areas and systems. This kind of model can be used to link dynamic phenomena at the cellular and synaptic levels to behavior.

## Acknowledgments

We would like to thank Carl G. Tengwall and Erling Weibust, Blue Gene Solutions, IBM Deep Computing EMEA, IBM Svenska AB, and Cindy Mestad, James C. Pischke, Steven M. Westerbeck, and Michael Hennecke for support when running on the Blue Gene/L installation at the Deep Computing Capacity on Demand Center, IBM, Rochester, Minnesota. Fred Mintzer and Bob Walkup provided support during our runs at the Thomas J. Watson Research Center, Yorktown Heights, New York. Many thanks also to Gert Svensson and Mattias Claesson for support during runs on the Hebb Blue Gene/L installation at PDC, KTH. This work was partly supported by grants from the Swedish Science Council (Vetenskapsrådet, VR-621-2004-3807), the European Union (FACETS project, FP6-2004-IST-FETPI-015879), and the Swedish Foundation for Strategic Research (through the Stockholm Brain Institute).



## References

1. J. Karbowski, "How Does Connectivity Between Cortical Areas Depend on Brain Size? Implications for Efficient Computation," *J. Comput. Neurosci.* **15**, No. 3, 347–356 (2003).
2. R. Cossart, Y. Ikegaya, and R. Yuste, "Calcium Imaging of Cortical Networks Dynamics," *Cell Calcium* **37**, No. 5, 451–457 (2005).
3. M. Lundqvist, M. Rehn, M. Djurfeldt, and A. Lansner, "Attractor Dynamics in a Modular Network Model of Neocortex," *Network: Computation in Neural Systems* **17**, No. 3, 253–276 (2006).
4. M. Djurfeldt, C. Johansson, Ö. Ekeberg, M. Rehn, M. Lundqvist, and A. Lansner, "Massively Parallel Simulation of Brain-Scale Neuronal Network Models," *Technical Report TRITA-NA-P0513*, School of Computer Science and Communication (CSC), The Royal Institute of Technology (KTH), Stockholm, 2005.
5. A. Grinvald, E. E. Lieke, R. D. Frostig, and R. Hildesheim, "Cortical Point-Spread Function and Long-Range Lateral Interactions Revealed by Real-Time Optical Imaging of Macaque Monkey Primary Visual Cortex," *J. Neurosci.* **14**, No. 5, 2545–2568 (1994).
6. P. S. Churchland and T. J. Sejnowski, *The Computational Brain*, The MIT Press, Cambridge, MA, 1992.
7. B. Horwitz, K. J. Friston, and J. G. Taylor, "Neural Modeling and Functional Brain Imaging: An Overview," *Neural Networks* **13**, No. 8/9, 829–846 (2000).
8. M. Djurfeldt and A. Lansner, "Workshop Report: 1st INCF Workshop on Large-Scale Modeling of the Nervous System," *Nature Precedings* (2007); see *http://dx.doi.org/10.1038/npre.2007.262.1*.
9. Blue Brain Project, École Polytechnique Fédérale de Lausanne (EPFL); see *http://bluebrain.epfl.ch/page17871.html*.
10. H. Markram, "The Blue Brain Project," *Nature Rev. Neurosci.* **7**, 153–160 (2006).
11. W. Denk and H. Horstmann, "Serial Block-Face Scanning Electron Microscopy to Reconstruct Three-Dimensional Tissue Nanostructure," *PLoS Biol.* **2**, No. 11, 1900–1909 (2004).
12. D. Hebb, *The Organization of Behavior*, John Wiley, New York, 1949.
13. J. M. Fuster, *Memory in the Cereral Cortex*, MIT Press, Cambridge, MA, 1995.
14. D. Willshaw and H. Longuet-Higgins, "Associative Memory Models," *Machine Learning*, Vol. 5, B. Meltzer and O. Michie, Eds., Edinburgh University Press, Edinburgh, Scotland, 1970.
15. G. Palm, *Neural Assemblies, an Alternative Approach to Artificial Intelligence*, Springer-Verlag, New York, 1982.
16. J. J. Hopfield, "Neural Networks and Physical Systems with Emergent Collective Computational Abilities," *Proc. Natl. Acad. Sci. USA* **79**, No. 8, 2554–2558 (1982).
17. D. J. Amit, *Modeling Brain Function: The World of Attractor Neural Networks*, Cambridge University Press, New York, 1989.
18. J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computation*, Westview Press, Boulder, CO, 1991.

19. L. B. Haberly and J. M. Bower, "Olfactory Cortex: Model Circuit for Study of Associative Memory?" *Trends Neurosci.* **12**, No. 7, 258–264 (1989).
20. A. Treves and E. T. Rolls, "Computational Analysis of the Role of the Hippocampus in Memory," *Hippocampus* **4**, No. 3, 374–391 (1994).
21. A. Compte, N. Brunel, P. S. Goldman-Rakic, and X.-J. Wang, "Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model," *Cereb. Cortex* **10**, No. 9, 910–923 (2002).
22. D. H. Hubel and T. N. Wiesel, "Ferrier Lecture. Functional Architecture of Macaque Monkey Visual Cortex," *Proc. R. Soc. Lond. B Biol. Sci.* **198**, No. 1130, 1–59 (1977).
23. V. B. Mountcastle, "An Organizing Principle for Cerebral Function: The Unit Module and the Distributed System," *The Mindful Brain: Cortical Organization and the Group-Selective Theory of Higher Brain Function*, G. M. Edelman and V. B. Mountcastle, Eds., MIT Press, Cambridge, MA, 1978, pp. 7–50.
24. A. Peters and C. Sethares, "Organization of Pyramidal Neurons in Area 17 of Monkey Visual Cortex," *J. Comp. Neurol.* **306**, No. 1, 1–23 (1991).
25. A. Lansner, E. Fransén, and A. Sandberg, "Cell Assembly Dynamics in Detailed and Abstract Attractor Models of Cortical Associative Memory," *Theory Biosci.* **122**, No. 1, 19–36 (2003).
26. A. Sandberg, A. Lansner, K. M. Petersson, and Ö. Ekeberg, "A Bayesian Attractor Network with Incremental Learning," *Network: Computation in Neural Systems* **13**, No. 2, 179–194 (2002).
27. A. Sandberg, J. Tegnér, and A. Lansner, "A Working Memory Model Based on Fast Hebbian Learning," *Network* **14**, No. 4, 789–802 (2003).
28. E. Fransén and A. Lansner, "A Model of Cortical Associative Memory Based on a Horizontal Network of Connected Columns," *Network: Computation in Neural Systems* **9**, No. 2, 235–264 (1998).
29. C. Johansson and A. Lansner, "Towards Cortex Sized Artificial Neural Systems," *Neural Networks* **20**, No. 1, 48–61 (2006).
30. P. Buzás, K. Kovács, A. S Ferecskó, J. M. Budd, U. T. Eysel, and Z. F. Kisvárday, "Model-Based Analysis of Excitatory Lateral Connections in the Visual Cortex," *J. Comp. Neurol.* **499**, No. 6, 861–881 (2006).
31. A. L. Hodgkin and A. F. Huxley, "A Quantitative Description of Membrane Current and its Application to Conduction and Excitation in Nerve," *J. Physiol.* **117**, No. 4, 500–544 (1952).
32. Ö. Ekeberg, P. Wallén, A. Lansner, H. Tråvén, L. Brodin, and S. Grillner, "A Computer Based Model for Realistic Simulations of Neural Networks. I: The Single Neuron and Synaptic Interaction," *Biol. Cybern.* **65**, No. 2, 81–90 (1991).
33. M. L. Hines and N. T. Carnevale, "The Neuron Simulation Environment," *Neural Comput.* **9**, No. 6, 1179–1209 (1997).
34. NEURON for Computer Simulations of Neurons and Neural Networks, Duke University, *http://neuron.duke.edu/*.
35. J. M. Bower and D. Beeman, *The Book of GENESIS: Exploring Realistic Neural Models with the General Neural Simulation System*, Second Edition, Springer-Verlag, New York, 1998.
36. GEneral NEural SImulation System (GENESIS); see *http://www.genesis-sim.org/GENESIS/*.
37. J. Frye, NCS Documentation, University of Nevada, Goodman Brain Computation Lab; see *http://brain.cse.unr.edu/ncsDocs/*.
38. A. Morrison, C. Mehring, T. Geisel, A. Aertsen, and M. Diesmann, "Advancing the Boundaries of High-Connectivity Network Simulation with Distributed Computing," *Neural Computation* **17**, No. 8, 1776–1801 (2005).
39. P. Hammarlund and Ö. Ekeberg, "Large Neural Network Simulations on Multiple Hardware Platforms," *J. Comput. Neurosci.* **5**, No. 4, 443–459 (1998).
40. J. Bailey and D. Hammerstrom, "Why VLSI Implementations of Associative VLCNs Require Connection Multiplexing," *Proceedings of the International Conference on Neural Networks*, San Diego, CA, 1988, pp. 173–180.
41. D. J. Amit and N. Brunel, "Model of Global Spontaneous Activity and Local Structured Activity During Delay Periods in Cerebral Cortex," *Cereb. Cortex* **7**, No. 3, 237–252 (1997).
42. E. Latham and S. Nirenberg, "Computing and Stability in Cortical Networks," *Neural Comput.* **16**, 1385–1412 (2004).
43. G. Pfurtscheller, A. Stancák, Jr., and C. Neuper, "Event-Related Synchronization (ERS) in the Alpha Band—An Electrophysiological Correlate of Cortical Idling: A Review," *Int. J. Psychophysiol.* **24**, No. 1/2, 39–46 (1996).
44. S. W. Hughes and V. Crunelli, "Thalamic Mechanisms of EEG Alpha Rhythms and Their Pathological Implications," *The Neuroscientist* **11**, No. 4, 357–372 (2005).
45. B. Pesaran, J. S. Pezaris, M. Sahani, P. P. Mitra, and R. A. Andersen, "Temporal Structure in Neuronal Activity During Working Memory in Macaque Parietal Cortex," *Nat. Neurosci.* **5**, No. 8, 805–811 (2002).
46. M. Steriade, F. Amzica, and D. Contreras, "Synchronization of Fast (30–40 Hz) Spontaneous Cortical Rhythms During Brain Activation," *J. Neurosci.* **16**, No. 1, 392–417 (1996).
47. R. Cossart, D. Aronov, and R. Yuste, "Attractor Dynamics of Network UP States in the Neocortex," *Nature* **423**, No. 6937, 283–288 (2003).
48. J. Anderson, I. Lampl, I. Reichova, M. Carandini, and D. Ferster, "Stimulus Dependence of Two-State Fluctuations of Membrane Potential in Cat Visual Cortex," *Nat. Neurosci.* **3**, No. 6, 617–621 (2000).
49. C. Bédard, H. Kröger, and A. Destexhe, "Does the $1/f$ Frequency Scaling of Brain Signals Reflect Self-organized Critical States?" *Phys. Rev. Lett.* **97**, No. 11, 118102 (2006).
50. V. Dragoi and M. Sur, "Dynamic Properties of Recurrent Inhibition in Primary Visual Cortex: Contrast and Orientation Dependence of Contextual Effects," *J. Neurophysiol.* **83**, No. 2, 1019–1030 (2000).
51. J. Tegnér, A. Compte, and X. J. Wang, "The Dynamical Stability of Reverberatory Neural Circuits," *Biol. Cybern.* **87**, No. 5/6, 471–481 (2002).
52. K. E. Stephan, L. Kamper, A. Bozkurt, G. A. Burns, M. P. Young, and R. Kötter, "Advanced Database Methodology for the Collation of Connectivity Data on the Macaque Brain (CoCoMac)," *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **356**, No. 1412, 1159–1186 (2001).


*Received March 23, 2007; accepted for publication June 26, 2007; Internet publication December 5, 2007*

**Mikael Djurfeldt** *Computational Biology and Neurocomputing Group, Royal Institute of Technology, S-100 44, Stockholm, Sweden (mikael@djurfeldt.com)*. Mr. Djurfeldt is a Research Engineer at the Computational Biology and Neurocomputing (CBN) group at the Royal Institute of Technology (KTH), Stockholm. His research interests include information processing and function in the cortex and the basal ganglia. He also has experience in software tool development.

**Mikael Lundqvist** *Computational Biology and Neurocomputing Group, Royal Institute of Technology, S-100 44, Stockholm, Sweden*. Mr. Lundqvist studied theoretical physics at Stockholm University, and is now a research engineer. His research interest is in working memory and the generic circuitry of the cortex.

**Christopher Johansson** *Computational Biology and Neurocomputing Group, Royal Institute of Technology, S-100 44, Stockholm, Sweden*. Dr. Johansson holds an M.S. degree in electrical engineering and a Ph.D. degree in computer science from the Royal Institute of Technology, Stockholm, Sweden. His areas of research cover parallel implementation of neural networks, modeling of cortical circuits, reinforcement learning, and brain-inspired computing.

**Martin Rehn** *Computational Biology and Neurocomputing Group, Royal Institute of Technology, S-100 44, Stockholm, Sweden*. Dr. Rehn is a postdoctoral researcher. His research interests include sparse sensory coding in the visual and auditory domains and models of computation in cortical circuits.

**Örjan Ekeberg** *Computational Biology and Neurocomputing Group, Royal Institute of Technology, S-100 44, Stockholm, Sweden*. Dr. Ekeberg is an Associate Professor in computer science. His research focus is on simulation models and computational techniques for studying the interaction between neuronal circuits and musculo-skeletal mechanics in the context of animal movement control. During 2001–2002, Dr. Ekeberg was a Resident Fellow of the Institute of Advanced Study, Berlin.

**Anders Lansner** *Computational Biology and Neurocomputing Group (CBN), Royal Institute of Technology, S-100 44, Stockholm, Sweden*. Dr. Lansner is Professor of computer science at the Royal Institute of Technology (KTH) and Stockholm University, and he is head of the CBN group. He holds an M.S. degree in engineering chemistry and a Ph.D. degree in computer science from KTH. His research interests range from neural computation and control to biophysically detailed computational models of specific brain systems and functions.