Published in final edited form as:
Neural Netw. 2008 October ; 21(8): 1153–1163. doi:10.1016/j.neunet.2008.05.006.

Self-sustaining non-repetitive activity in a large scale neuronal-level model of the hippocampal circuit

Ruggero Scorcioni, **David J. Hamilton**, and **Giorgio A. Ascoli**
Center for Neural Informatics, Structure, and Plasticity (CN3) Krasnow Institute for Advanced Study, George Mason University Mail Stop 2A1, Fairfax, Virginia 22030−4444, Fax: (703)993−4325, Email: ascoli@gmu.edu

Abstract

The mammalian hippocampus is involved in spatial representation and memory storage and retrieval, and much research is ongoing to elucidate the cellular and system-level mechanisms underlying these cognitive functions. Modeling may be useful to link network-level activity patterns to the relevant features of hippocampal anatomy and electrophysiology. Investigating the effects of circuit connectivity requires simulations of a number of neurons close to real scale. Toward this end, we construct a model of the hippocampus with 16 distinct neuronal classes (including both local and projection cells) and 200,000 individual neurons. The number of neurons in each class and their interconnectivity are drawn from rat anatomy. Here we analyze the emergent network activity and how it is affected by reducing either the size or the connectivity diversity of the model. When the model is run with a simple variation of the McCulloch-Pitts formalism, self-sustaining non-repetitive activity patterns consistently emerge. Specific firing threshold values are narrowly constrained for each cell class upon multiple runs with different stochastic wiring and initial conditions, yet these values do not directly affect network stability. Analysis of the model at different network sizes demonstrates that a scale reduction of one order of magnitude drastically alters network dynamics, including the variability of the output range, the distribution of firing frequencies, and the duration of self-sustained activity. Moreover, comparing the model to a control condition with an equivalent number of (excitatory/inhibitory balanced) synapses, but removing all class-specific information (i.e. collapsing the network to homogeneous random connectivity) has surprisingly similar effects to downsizing the total number of neurons. The reduced-scale model is also compared directly with integrate-and-fire simulations, which capture considerably more physiological detail at the single-cell level, but still fail to reproduce the full behavioral complexity of the large-scale model. Thus network size, cell class diversity, and connectivity details may all be critical to generate self-sustained non-repetitive activity patterns.

Introduction

The hippocampus is one of the most intensively investigated brain regions, both because of its crucial involvement in memory representation, and its exquisite anatomical organization. Neocortical information enters and exits the hippocampus primarily through the entorhinal cortex. The hippocampal loop is often described as “trisynaptic”, as activity flows from the dentate gyrus (DG), through areas CA3, and to CA1. However, the structural complexity of this circuit defies oversimplification and has yet to be fully elucidated (Andersen, Morris,

Amaral, Bliss, & O'Keefe, 2007). Each of these regions features a large majority of principal cells (DG granule cells and CA3/CA1 pyramidal cells) characterized by excitatory projections, and a great variety of mostly local and inhibitory interneurons. The morphological, physiological, and functional specifications of these neuronal components constitute a major goal of recent and ongoing research.

The hippocampus is involved in spatial learning, and essential information for spatial navigation is present within the hippocampal circuit, such as in place cells (O'Keefe & Dostrovsky, 1971), head-direction cells (Ranck, 1985; Taube, Muller, & Ranck, 1990), and grid cells (Hafting, Fyhn, Molden, Moser, & Moser, 2005). The correlation of place cell firing to the body spatial position is a particularly resilient phenomenon to lesions (Miller & Best, 1980; Mizumori, McNaughton, Barnes, & Fox, 1989; McNaughton, Barnes, Meltzer, & Sutherland 1989; Brun et al., 2002), to environment modifications (O'Keefe & Conway, 1978), and to genetic manipulations (Rotenberg, Mayford, Hawkins, Kandel, & Muller, 1996; Wilson & Tonegawa, 1997). Head-direction cell firing is correlated to the direction of the head, and grid cell firing correlates to a spatial lattice regularly distributed in space. Understanding how the neurons in various classes and sub-regions can integrate across different spatial frames is not only essential to establish how spatial navigation functions, but also to determine how mammalian learning is achieved within the hippocampus (Redish, 1999).

If memories are constituted by temporal spiking patterns in cell assemblies, then characterizing the cellular bases of this activity is a critical step to elucidate the mechanism underlying mnemonic function. Computational modeling provides a means both to formalize and quantify current theories and state of knowledge, and to foster hypothesis generation and testing. The realism of a network model is reflected in the number, type, and specifications of its neuronal units. Simulating the quantitative details of membrane biophysics enables a fine-grained analysis of neuronal dynamics. At such a level, however, approaching the number of neurons of the whole hippocampus is computationally prohibitive. The correspondingly necessary downscaling implies a drastic loss of anatomical information. A complementary approach is to sacrifice most of the cellular-level biophysical detail to capture more realistic anatomical connectivity in larger-scale simulations (Ascoli & Atkeson, 2005). Identifying essential architectural characteristics of neuronal connectivity appears fundamental to correlate network structure and function. A real scale modeling effort could also help interpret information-theoretic approaches to the hippocampal circuit (Treves, Skaggs, & Barnes, 1996; Settanni & Treves, 2000), which usually adopt analytical models in the limit of infinitely large networks. Here we study the influence of cellular connectivity on network activity in isolation from other details such as fully accurate biophysical properties.

Earlier attempts to model the hippocampus with neural networks incorporating a high degree of anatomical realism include work related to the dentate gyrus (Patton & McNaughton, 1995), Ammon's Horn (Bernard & Wheal, 1994), and area CA1 (Senft & Ascoli, 1999). Only few studies, however, explicitly simulated network dynamics (Bernard, Cannon, Ben, & Wheal, 1997, Dyhrfjeld-Johnsen et al., 2007). Several studies addressed specific hippocampal phenomena and/or microcircuits with both computational and experimental approaches. These include the splitter cell phenomenon in CA1 cells (Katz, Kath, Spruston, & Hasselmo, 2007), gamma-frequency oscillations within the CA3 region (Mann, Suckling, Hajos, Greenfield, & Paulsen, 2005; Hocking & Levy, 2006), and modulation of place cell firing by the perforant pathway projection to CA1 (Brun et al., 2008). Others have studied the influence of dendritic structure on firing patterns in neocortical structures (Mainen & Sejnowski, 1996) and hippocampal regions (Krichmar, Nasuto, Scorcioni, Washington, & Ascoli, 2002), or the crucial role of NMDA receptors in the CA3 region for spatial learning (Nakazawa et al., 2003). However, none of these previous models were realistically large scale, nor did they

include accurate connectivity information from the full hippocampal circuit. Interestingly, past modeling efforts that generated self-sustaining network activity always involved the use of synaptic plasticity (e.g. Barak & Tsodyks, 2007; Izhikevich, 2008). For example, spike-timing-dependent plasticity (STDP) allows networks to operate at the border between randomness and synchrony: if network activity is random then STDP helps synchronize it, while if the network is synchronous, STDP tends to decouple it (Lubenov & Siapas, 2008).

Here we expand on some of these previous efforts and present a large-scale circuit model with 16 distinct cell classes and 200,000 individual neurons. To construct this model, we have begun to leverage the availability of rich data cataloging neuronal classes and associated connectivity for the rat hippocampus. These include connectivity patterns for numerous classes of interneurons (Freund & Buzsaki, 1996) and post-synaptic potential parameters for all principal cells (e.g., Urban, Henze, & Barrionuevo, 2001). To enable practical simulations near real-scale, we adopt simple McCullloch-Pitts-like (MP) neuronal computational units (McCulloch & Pitts, 1943). In essence, if the running sum of inputs received by a unit passes a given threshold at any time, that unit will fire and reset its inputs at the following time step. We show that the network achieves global, distributed, and non-repetitive self-sustaining activity in the absence of plasticity.

We also report the results of several variations of the model. First, we investigate the influence of the specific cell-class connectivity on network behavior by collapsing all excitatory and inhibitory neurons into two homogeneous classes, while maintaining the same number of cells and of excitatory-excitatory, excitatory-inhibitory, inhibitory-excitatory, and inhibitory-inhibitory synapses. Next, to study the effect of network size, we keep the network architecture (cell classes and connectivity) constant while varying the number of neurons. Although none of these models are anatomically accurate, each provides a unique perspective on network behavior and characterization.

Finally, network downscaling provides an opportunity to compare this model to simulations based on leaky Integrate-and-Fire (IF) neurons (Abbott, 1999). This approach captures considerably more biophysical parameters, such as membrane capacitance and time constant, resulting in more realistic synaptic currents and post-synaptic potentials with finite rise time. While this level of description is much more realistic than the MP formalism, our results suggest that the additional information is complementary to that provided by the connectivity details, and cannot directly compensate for network size.

Methods

The trisynaptic circuit of the rat hippocampus contains roughly two million neurons (Nakagami Saito, & Matsuki, 1997; Todtenkopf, 2000). The granule cells of the dentate gyrus (gcDG) are the first stage. Their inputs are primarily from entorhinal afferents (ecEC) via the perforant pathway. In the second stage, CA3 pyramidal cells (pcCA3) are fed by gcDG axons. CA3 pyramidal cells then drive CA1 pyramidal cells (pcCA1), which constitute the third and final stage of the trisynaptic circuit. CA1 pyramidal cells then project back to the ecEC (both directly and through the subiculum), completing the circuit. This general overview can be extended by the addition of interneurons into a schematic of the rat hippocampus consisting of 16 cellular classes (Figure 1).

Within the dentate gyrus there are five class types in the model besides granule cell: mossy cells (mcDG), GABAergic polymorphic cells (gpcDG), basket cells (bcDG), molecular layer perforant path cells (moppDG), and chandelier cells (ccDG). Area CA3 contains three class types in addition to pyramidal cells, namely radiatum interneuron (riCA3), oriens interneuron (oiCA3), and chandelier cell (ccCA3). Finally, area CA1 includes pyramidal cells and four

other class types, i.e. basket cells (bcCA1), lacunosum-moleculare interneurons (lmCA1), oriens-alveus interneurons (oaCA1), and chandelier cells (ccCA1). In this model, the entorhinal cortical columns (ecEC) are represented as a single class rather than as individual cell types.

This connectivity model sums up to 1.8M neurons and 10^{10} synapses before scaling (Table 1). All of the model variations studied here are derived from this anatomically accurate representation of network level connectivity, maintaining the same proportional numbers of both neurons and connections.

McCulloch-Pitts (MP) Models

The McCulloch-Pitts simulator is written in C++ running on an Intel Core 2 Quad 2.4GHz with 4GB RAM under Windows XP Pro. The simulator reads connectivity information from a text file in which each line specifies a connection between two classes, in the format of: <class1> <class2> <tot#1> <post#2> <type>. The first two fields label the pre- and post-synaptic classes, respectively. The third and fourth fields specify the total number of presynaptic neurons and the number of postsynaptic neurons contacted by each presynaptic neuron, thus determining the stoichiometry of the connections. The last field flags the synapse as either excitatory (+1) or inhibitory (−1).

The scale of the simulation is a free parameter determining the total number of neurons. The number of cells in each class and all connections are scaled proportionally. The largest network we simulate is composed of 200K neurons (16MP200), containing 300M synapses. Two reduced models are also tested with 100,000 (16MP100) and 17,000 neurons (16MP17), respectively. In each case, a “synaptically balanced” control of just two classes (excitatory and inhibitory) is also created to collapse all cell-class specific connectivity information (2MP200, 2MP100, and 2MP17, respectively).

The synaptic connections are formed following these steps: (1) Compute the ratio X of the desired network size divided by the total number of neurons; (2) Within one connection (the first or next line in the connectivity file), for every neuron in class1 randomly select X*post#2 neurons in class2, and store the connectivity information in a linked array where each element represents a contact from the given presynaptic cell (from class1) to the selected postsynaptic (class2) target; (3) Repeat step 2 for all connections (up to until the last line in the connectivity file).

Time is discretized in the MP simulations. Within each unitary time step, it is assumed that every axonal spike propagates to its targets without failure, and all synaptic signals travel along their dendrites and get integrated at the soma. In particular, all presynaptic inputs are summed (+/−1 for excitatory and inhibitory signals respectively) into a counter for each neuron. If the sum is higher than the class-specific threshold, the neuron will fire at the subsequent time step (t+1), and its synaptic counter is reset to zero at times t+1 and t+2 (corresponding to a relative refractory period of 2 steps).

All neurons within a given class have one and the same firing threshold, which assumes a constant value at steady-state. This steady-state threshold is determined for each class as the value attained at the end of an initial (transient) phase. During the transient phase, the threshold is computed at each time step as a fraction of the highest synaptic count value among the neurons of the same class. In this study, the fraction is set at 0.90 and 0.99 for excitatory and inhibitory neurons, respectively, and the transient phase lasts 200 time steps. To start the network, 70% of the neurons are randomly activated just at the first step of the transient phase.

Integrate-and-Fire (IF) Model

The NEST2 simulator (Diesmann & Gewaltig, 2002) is used for the integrate-and-fire neuronal network simulations utilizing the built-in IF neuron model. Integration is implemented with linear sub-threshold dynamics as described by the exact integration scheme (Rotter & Diesmann, 1999). The computational requirements of the greater represented complexity constrain network size to 17,880 neurons over all 16 classes (16IF17 model), scaled according to the same connectivity matrix described for the MP models (totaling 2,438,520 synapses). The IF simulations are run on an Intel Core 2 Duo 2GHz machine with 1.5GB RAM, under Fedora Core 5 Linux. The 3,231 lines of NEST2 syntax have a build time of approximately 3 minutes. Memory consumed to load the model is 100MB (breaking down to ~1KB per neuron and ~30 bytes per synapse). A 20K time step (each representing 0.1 ms) simulation generally has a run time of approximately 30 minutes and consumes about 1.1GB of RAM.

The IF parameter values are selected for all 16 classes as follows: resting membrane potential, 0 mV; threshold, random for each class between 0 and 10.88 mV (greater global threshold values cause a degradation of self-sustained network activity); reset potential, 0 mV; membrane capacity, 250 pF; membrane time constant, 10 ms; postsynaptic current time constant (both excitatory and inhibitory), 0.33 ms; refractory period, 0.5 ms. A Poisson stimulus is applied to all ecEC type neurons for the first 1000 time steps (100ms). All cell membrane potentials are monitored with the provided virtual voltage probes and logged to a file for each time step during simulation.

Results

All networks display sustained activity beyond the initial stimulation step. The large scale MP model with 16 cell classes and 200,000 neurons (16MP200) generates non-repetitive patterns that keep evolving past the transient phase. Shown in Figure 2A is the complete total activity trace of one such simulation which lasts 1164 steps. The underlying data from 900 to 1000 steps is further analyzed in the following panels of Figure 2. Figure 2B shows the time series of total activity for all principal cells as percent of neurons firing within their respective class. Figure 2C expands on the CA3 region by presenting activity traces for each of its neuronal classes. Figure 2D represents a scatter plot over all neuronal classes. For each class, 50 neurons are randomly selected and their spikes plotted over time. It is possible to notice how activity is transmitted across classes and regions. In addition the model is able to generate evolving patterns that over time are non-repetitive. Figure 2E shows a zoomed in view of the pcCA3 class and highlights the changing characteristic of the firing activity.

Differences between MP networks are shown in Figure 3. Figure 3A shows total activity for four different models, expressed as a percentage of active neurons at each time step. It is noticeable how both 16MP200 and the half-sized version, 16MP100, present qualitatively similar dynamical characteristics: presence of peaks with period of low activity. In contrast, networks either without cell-class specified connectivity details (2MP100), or with greatly reduced total neuron numbers (16MP17) present a narrowly constrained activity range and peaks are missing.

An alternative illustration of this information is provided in Figure 3B, where each trace of the same simulation is normalized in the 0−1 range relative to its minimum and maximum values reached over the total lifespan of self-sustained activity in that network (only 100 time steps are presented). Despite the “zoom in” effect produced for the narrowly constrained waveforms of 2MP100 and 16MP17, traces in these reduced network configurations are still strikingly missing extended low activity periods between peaks. It is important to notice that besides the numbers of neurons and/or of cell classes (and, proportionally, of synaptic connections), all

other parameters associated with the models (i.e. initialization mechanism, stimulation method, etc.) remain constant.

To analyze network differences further among the model variants, we compute the frequency distribution and power spectra of the total activity by Fast Fourier Transform (Seewave R-package: Sueur, Aubin, & Simonis-Sueur, 2007). A representative example for each of the networks is reported in Figure 4. In these illustrations, the more pronounced the striping of the frequency spectrum, the more repetitive is the corresponding activity. In particular, Figure 4A presents the 16MP200 network, which generates the least repetitive activity of all models. As the network size is reduced, the repetitive behavior increases (e.g. model 16MP100 in Figure 4B). Removal of class-specific connectivity in the 2MP200 and 2MP100 models (panels 4C and 4D, respectively) yields striping effects with harmonics behavior, highlighting even more pronounced repetitive patterns. The 16MP17 model (Figure 4E) has a sparse frequency spectrum and low power density, indicating a narrowly constrained output activity. Some, but by no means all of the behavioral complexity is recovered in this network configuration by switching from MP units to the biophysically more detailed integrate-and-fire formalism. The 16IF17 model (Figure 4F) in fact also presents multiple striping and harmonics that indicate strong presence of repetitive patterns.

The lifespan of each network's self-sustained activity changes between different simulations. To analyze this characteristic, every model is run 200 times with varying random seeds, each producing a network with a unique microscopic pattern of connections and different selection of neurons externally activated at the first time step. Figure 5A shows the network lifespan distribution, which is separated into two qualitatively similar groups. In both the 16MP200 and the 16MP100 models, the vast majority of simulation runs result in a bell-shaped distribution of the network lifespan. In particular, the large scale 16MP200 model displays self-sustained activity lasting 2785 time steps on average (standard error: 502.4, low quartile: 605,median: 958, high quartile: 1891). In clear contrast, the 16MP17 and 2MP100 models are characterized by a sharply bimodal distribution of life spans, where activity either damps out almost immediately after the end of the transient phase, or self-sustains for as long as the simulations are run (>50,000 steps).

Each MP unit has a threshold that is constant within the same class, but the values are set at the end of the transient phase, and thus generally vary in different simulation runs. It can be expected that a lower threshold for excitatory neurons extends the network lifespan. To test this hypothesis, we plot threshold values for each of the 16 classes over the network lifespan achieved in each run of the large-scale model 16MP200 (Figure 5B). No significant correlation (in fact, no trend whatsoever) is found between the duration of self-sustained activity and any of the plotlines. Figure 5C shows how narrowly constrained the threshold values are (note standard deviation bars within circles). Cell classes are sorted from left to right by the coefficient of variation of their threshold, plotted on the secondary ordinate axis (square boxes). Most classes settle in a very limited range of values (all but two displaying a coefficient of variation of less than 5%), while others are more pronounced in their threshold variability (in particular the only class of excitatory interneurons, mcDG). This suggests that connectivity information influences both average value and coefficient of variation of each class threshold. Moreover, to test whether using one and the same steady-state threshold value for all units in a class is a limiting factor for network dynamics, two additional sets of 200 simulations each were performed with variable thresholds. These networks were based on the same structure as 16MP17, except that the threshold of each individual unit was modified with uniform noise. The noise in the two new sets of simulations was set at ±5% and ±10%, respectively, of the class threshold value. Neither condition displayed any substantial difference compared to 16MP17 in the quantity and quality of self-sustaining and non-repetitive patterns (data not shown).

The observed network behavior in the full model appears unlikely to be fully understood in terms of relatively simple dynamics. For example, an activity wave moving from one neuronal class to another could trigger repetitive population spikes, but would be inconsistent with irregular patterns. Moreover, this mechanism could not explain the relationships among neuronal classes (e.g. Figure 2B). Thus it may be instructive to examine the temporal relations among the relative activity of the neuronal classes, each with respect to the total activity of the overall network. Figure 6A shows these data for the principal neurons, again displaying non-trivial dynamics. Each line represents the relative activity of a neuronal class normalized to the total network activity for the same simulation illustrated in Figure 2B. For instance, at t=950, 13.14% of pcCA1 units are firing, while at the same time 4.53% of all units in the network are active, yielding a normalized pcCA1 activity of 2.90 (13.14/4.53). All cells classes within the CA3 region are similarly analyzed in Figure 6B, but no obvious pattern seems to emerge.

In order to elucidate the mechanisms underlying the more periodic elements of network behavior, we attempted to separate the non-repetitive pattern features by averaging the population spikes of each class. Figure 6C shows the population spike of gcDG within ±15 steps computed over 45,000 activity peaks from all two hundred 16MP200 simulations, as well as the corresponding time-locked population activity of the other principal cell classes. The mean gcDG population spike has a peak of ~55% and a periodicity of ~11 steps. On average, gcDG spikes are preceded by spikes in both the ecEC class at t=−1 and the pcCA1 class at t=−2. At each step of this sequence, the population spikes of subsequent classes progressively increase their peaks (pcCA1 < ecEC < gcDG). Figure 6D shows the average population spikes of pcCA1. In this case there is no clear indication of previous population spikes other than a minimal alignment of pcCA3 spikes at t=−1. In contrast, pcCA1 activity is followed by increasing peaks in ecEC at t=+1 and in gcDG at t=+2, consistent with panel 6C. These findings suggest a role of CA3 pyramidal cells as a source of non-repetitive behavior. To further investigate this possibility, we analyzed the relationship between the normalized activity of pcCA1 and pcCA3. Displaying both traces from Figure 6A in the same scatterplot (Figure 6E) reveals two distinct and contrasting trends. In most cases (highlighted by empty circles), pcCA1 activity is proportional to pcCA3, but the opposite relationship (identified with filled symbols) is also prominent. These different trends tend to correspond with opposite phases of the overall network activity oscillations (Figure 6F), namely network spikes and lower activity.

Table 2 summarizes all MP model information. The total number of simulated neuronal units, the number of units in the smallest cell class (chandelier cells for 16MP models and inhibitory neurons for 2MP models), and the total number of synapses all characterize structural network properties. In contrast, the average activity over the network lifespan, the minimum and maximum activity, and the “confidence interval” of the activity range representing 95% of time steps, all characterize global behavioral dynamics. It is possible to notice that a reduction in the number of neurons leads to a lowered confidence range as well as a reduced average activity. Similar results are obtained by reducing the number of classes.

Discussion

This work embodies the first approximation of a realistically large-scale neural network model of the rat hippocampus based on cellular-level connectivity information. One of the main objectives and results of this study is the demonstration that a purely anatomical model endowed with a minimalist binary integration mechanism in its neuronal units is capable of producing plausible, non-repetitive patterns that are self-sustaining. This result, however, only holds as long as both the network size (in terms of total number of neurons) and complexity (in terms of cell-class specific connectivity details) approximate the real values.

There are several limitations to the modeling techniques described here at this early stage of development. (A) The models are still not to real scale, even the largest one (16MP200) only representing 1/9th of the number of neurons reported in the rat hippocampus. The smallest cell class in the 16MP200 model contains 112 neurons, which may still be relatively low for the emergence and maintenance of complex dynamics. (B) Some interneuron classes are not represented or are not segregated into the appropriate subclasses. For example, recent modeling work on the dentate gyrus (Dyhrfjeld-Johnsen et al., 2007) overviews and references the following cell classes in addition to granule, basket, mossy, chandelier, and MOPP cells: HIPP and HICAP hilar interneurons (possibly corresponding to our combined gpcDG cells), and interneuron selective cells. As more research is carried out, it is conceivable, in fact likely, that even more new classes of interneurons will be discovered and reported. (C) The connectivity matrix is derived from an early draft (Ascoli & Atkeson, 2005), and has yet to be updated. Several connections might be missing (most notably the temporo-ammonic pathway from EC to CA1) and new estimates for the numbers of cells and synaptic contacts should be gathered from the literature and archived in digitally accessible databases. This is a key requirement for future neuroinformatics research in computational neuroanatomy of the hippocampus and other brain regions. (D) In the model, delay and transmission times are uniform across cell types, while in actual hippocampal neurons significant differences are likely to exist between different classes. (E) The model is designed to identify the specific influence of connectivity and realistic scale on network activity. Toward this goal, all other details are kept to a minimum. However, simulating neurons as simple MP units is far from representative of plausible electrophysiological dynamics. Greater biological insights could be gained by using more accurate neuronal models when sufficient computational power becomes available. (F) Synaptic strengths do not vary among individual contacts, neurons, cell classes, and especially not even between excitatory and inhibitory neurons. (G) Last but not least, local geometry is not accounted for in establishing interconnections between neurons, when it is instead known that hippocampal interneurons project locally.

Despite these limitations, the networks produce self-sustaining activity that grows in complexity with an increase in the numbers of neurons and cell classes. Of the different networks we compared, the larger network model with detailed connectivity (16MP200) tended to exhibit more non-repetitive behaviors with greater dynamic range. If enough neurons and classes are present, the distribution of frequency components is also enriched. Conversely, models that are limited in the size or diversity of cell classes yield a sparser frequency distribution and reduced power spectra density. These differences are not affected by the variability or complexity of the individual computational units. For example, the use of variable thresholds (up to ±10% uniform noise) did not substantially alter behavior of the 16MP17 network. Likewise, the NEST2 model had an additional element of variability, in that the synaptic strength was normally distributed within each class. Nevertheless, similarly to the 16MP17 model, the NEST2 simulations failed to exhibit the same level of non-repetitive self-sustaining behavior observed in the larger networks.

Analysis of network activity shows the robust presence of global rhythms as well as specific patterns within individual classes. Overall, activity is clearly differentiated among cell classes. Self-sustained activity is consistent over multiple runs initiated with different random seeds. Although thresholds are narrowly constrained (and cell-class specific), their values do not influence network lifespan. Correlations exist both among assemblies of neurons in the same class as well as across multiple cell classes.

It seems evident from Figures 2A,B, and 6A,B that network activity follows non trivial dynamics. Averaging the population spikes of a given class of principal neurons (e.g., gcDG or pcCA1, shown in Figures 6C and 6D, respectively) reveals clear sequences of activation that “travel” from pcCA1, through ecEC, and to gcDG. At each move, the percentage of active

neurons in the class increases from ~12% in pcCA1, to ~28% in ecEC, and finally ~55% in gcDG, suggesting an avalanche effect within the sub-circuit pcCA1 → ecEC → gcDG. The end point of this cascade, the dentate gyrus granule cells, also constitutes the cellular class responsible for network spike activity, as indicated by the coincidence of gcDG and overall network spikes (Figure 2A,D). This is due the large number of units composing the gcDG class, both in the model and in the rat hippocampus, where more than half of all neurons are granule cells.

In contrast, CA3 pyramidal cells appear to exercise a more subtle effect on these dynamics (Figure 6C,D). A direct influence of pcCA3 would imply network oscillations with a period of 4 time steps, following the loop pcCA3 → pcCA1 → ecEC → gcDG and back to pcCA3. Instead, we observe a weak periodicity of 11, characterized by reduced average peaks at ±11 steps (Figure 6C,D). Thus the pcCA3 influence on the activity of pcCA1 is more complex, and takes on average 7 time steps to generate a pcCA1 spike. Figure 6E suggests a potential explanation by highlighting two distinct relationships between the activity patterns of pcCA3 and pcCA1. In correspondence to gcDG population spikes (empty circles), pcCA1 and pcCA3 follow similar (either high or low) activity. In contrast, during quiet periods for the gcDG class (filled circles), pcCA3 and pcCA1 are in opposite phases or alternating behavior. This interaction eventually generates a population spike in pcCA1, which triggers the avalanche effect ending with a gcDG (and overall network) spike.

The 200MP16 network model displays dynamic behaviors that qualitatively match experimental multiunit recordings and local field potentials. The two most striking features are the presence of non-repetitive patterns and the alternation of quiet periods with epochs of high activity. Non-repetitive and rhythmic patterns were reported in all primary hippocampal regions in early single unit recordings from unrestrained rats (Ranck, 1973). The relationship between the periodicity of the population activity and the interspike intervals of individual neurons in our simulations also strikingly reflects the observations of numerous more recent investigations (Geisler, Robbe, Zugaro, Sirota, & Buzsaki, 2007; Csicsvari, Jamieson, Wise, & Buzsaki, 2003; Isomura et al., 2006; Rice et al., 1996). Similarly, several studies based on tetrode recordings reported periods of high activity followed by quiescence in sequences that are rhythmic and non-repetitive. For example, both multiunit raster plots and field potentials continuously recorded for 5 seconds from the hippocampus during slow-wave sleep (Ji & Wilson, 2007) closely resemble the spiking and population activities generated in our model (cf. their Fig. 2a with our Figure 2A). Similar experimental pairings of hippocampal local field potentials and neuronal activity showed high activity peaks overlapping with slow-wave ripple (Siapas & Wilson, 1998), again extremely similar to the dynamics observed in our simulations.

In contrast, numerous experimental studies demonstrated higher firing rates in interneurons than in principal cells (e.g. Whittington, Traub, Kopell, Ementrout, & Buhl, 2000; Klausberger et al., 2003). This feature is not appropriately captured in our model, where interneurons fire sparsely and at low frequencies. It is tempting to speculate that a more complete and accurate review of hippocampal connectivity will ultimately result in even more realistic complex activity, thus providing additional insights. However, much more neuroinformatics work is still required with respect to both anatomical and physiological details. Future refinement plans include addition of positioning information within the connectivity constraints and plausible variation of temporal delays as well as synaptic weights among neuronal units.

Acknowledgements

The authors like to thank Dr. Duncan Donohue and Mr. Deepak Ropireddy for valuable feedback on the manuscript. This work was supported by NIH grants NS39600 and AG025633, and by NSF grant SGER 0747864.

References

Abbott LF. Lapicque's introduction of the integrate-and-fire model neuron. Brain Research Bulletin 1999;50:303–304. [PubMed: 10643408]

Andersen, P.; Morris, R.; Amaral, D.; Bliss, T.; O'Keefe, J. The Hippocampus Book. Oxford University Press; 2007.

Ascoli GA, Atkeson JC. Incorporating anatomically realistic cellular-level connectivity in neural network models of the rat hippocampus. BioSystems 2005;79:173–181. [PubMed: 15649602]

Barak O, Tsodyks M. Persistent activity in neural networks with dynamic synapses. PLoS Computational Biology 2007;3(2):323–332.

Bernard C, Cannon RC, Ben Ari Y, Wheal HV. Model of spatio-temporal propagation of action potentials in the Schaffer collateral pathway of the CA1 area of the rat hippocampus. Hippocampus 1997;7:58–72. [PubMed: 9138669]

Bernard C, Wheal HV. Model of local connectivity patterns in CA3 and CA1 areas of the hippocampus. Hippocampus 1994;4:497–529. [PubMed: 7889123]

Brun VH, Otnass MK, Molden S, Steffenach HA, Witter MP, Moser M-B, Moser EI. Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. Science 2002;296:2243–2246. [PubMed: 12077421]

Brun VH, Leutgeb S, Wu HQ, Schwarcz R, Witter MP, Moser EI, Moser MB. Impaired spatial representation in CA1 after lesion of direct input from entorhinal cortex. Neuron 2008;57(2):290–302. [PubMed: 18215625]

Csicsvari J, Jamieson B, Wise KD, Buzsaki G. Mechanisms of gamma oscillations in the hippocampus of the behaving rat. Neuron 2003;37(2):311–322. [PubMed: 12546825]

Diesmann, M.; Gewaltig, MO. NEST: An Environment for Neural Systems Simulations.. In: Plesser, T.; Macho, V., editors. Forschung und wisschenschaftliches Rechnen, Beiträge zum Heinz-Billing-Preis 2001. 58. Ges. für Wiss. Datenverarbeitung; Gottingen: 2002. p. 43-70.

Dyhrfjeld-Johnsen J, Santhakumar V, Morgan RJ, Huerta R, Tsimring L, Soltesz I. Topological determinants of epileptogenesis in large-scale structural and functional models of the dentate gyrus derived from experimental data. Journal of Neurophysiology 2007;97:1566–1587. [PubMed: 17093119]

Freund TF, Buzsaki G. Interneurons of the hippocampus. Hippocampus 1996;6:347–470. [PubMed: 8915675]

Geisler C, Robbe D, Zugaro M, Sirota A, Buzsaki G. Hippocampal place cell assemblies are speed-controlled oscillators. Proceedings of the National Academy of Sciences 2007;104(19):8149–8154.

Hafting T, Fyhn M, Molden S, Moser M-B, Moser EI. Microstructure of a spatial map in the entorhinal cortex. Nature 2005;436:801–806. [PubMed: 15965463]

Hocking AB, Levy WB. Gamma scillations in a minimal CA3 model. Neurocomputing 2006;69:1244–1248.

Isomura Y, Sirota A, Ozen S, Montgomery S, Mizuseki K, Henze DA, Buzsaki G. Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. Neuron 2006;52(5):871–882. [PubMed: 17145507]

Izhikevich EM. Polychronization: computation with spikes. Neural Computation 2008;18(2):245–282. [PubMed: 16378515]

Ji D, Wilson MA. Coordinated memory replay in the visual cortex and hippocampus during sleep. Nature Neuroscience 2007;10(1):100–107.

Katz Y, Kath WL, Spruston N, Hasselmo ME. Coincidence detection of place and temporal context in a network model of spiking hippocampal neurons. PLoS Computational Biology 2007;3(12):2432–2445.

Klausberger T, Magill PJ, Márton LF, Roberts JD, Cobden PM, Buzsaki G, Somogyi P. Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. Nature 2003;421(6925):844–848. [PubMed: 12594513]

Krichmar J, Nasuto S, Scorcioni R, Washington S, Ascoli G. Influence of dendritic morphology on CA3 pyramidal cell electrophysiology. Brain Research 2002;941:11–28. [PubMed: 12031543]

Lubenov EV, Siapas AG. Decoupling through Synchrony in Neuronal Circuits with Propagation Delays. Neuron 2008;58(1):118–31. [PubMed: 18400168]

Mainen ZF, Sejnowski TJ. Influence of dendritic structure on firing pattern in model neocortical neurons. Nature 1996;382(6589):363–366. [PubMed: 8684467]

Mann EO, Suckling JM, Hajos N, Greenfield SA, Paulsen O. Perisomatic feedback inhibition underlies cholinergically induced fast network oscillations in the rat hippocampus in vitro. Neuron 2005;45 (1):105–117. [PubMed: 15629706]

Miller VM, Best PJ. Spatial correlates of hippocampal unit activity are altered by lesions of the fornix and endorhinal cortex. Brain Research 1980;194(2):311–323. [PubMed: 7388617]

Mizumori SJ, McNaughton BL, Barnes CA, Fox KB. Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of CA3c output: evidence for pattern completion in hippocampus. Journal of Neuroscience 1989;9(11):3915–3928. [PubMed: 2585060]

McCulloch W, Pitts W. A Logical Calculus of Ideas Immanent in Nervous Activity. Bulletin of Mathematical Biophysics 1943;5:115–133.

McNaughton BL, Barnes CA, Meltzer J, Sutherland RJ. Hippocampal granule cells are necessary for normal spatial learning but not for spatially-selective pyramidal cell discharge. Experimental Brain Research 1989;76(3):485–496.

Nakagami Y, Saito H, Matsuki N. Optical recording of trisynaptic pathway in rat hippocampal slices with a voltage-sensitive dye. Neuroscience 1997;81:1–8. [PubMed: 9300395]

Nakazawa K, Sun LD, Quirk MC, Rondi-Reig L, Wilson MA, Tonegawa S. Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. Neuron 2003;38(2):305–315. [PubMed: 12718863]

O'Keefe J, Conway DH. Hippocampal place units in the freely moving rat: why they fire where they fire. Experimental Brain Research 1978;31(4):573–590.

O'Keefe J, Dostrovsky J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Research 1971;34:171–175. [PubMed: 5124915]

Patton PE, McNaughton B. Connection matrix of the hippocampal formation. Part I. The dentate gyrus. Hippocampus 1995;5:245–586. [PubMed: 8589792]

Ranck JB. Studies on Single Neurons in Dorsal Hippocampal Formation and Septum in Unrestrained Rats, Part I. Behavioral correlates and firing repertoires. Experimental Neurology 1973;41(2):462–531.

Ranck, JB. Head direction cells in the deep cell layer of dorsolateral presubiculum in freely moving rats.. In: Busaki, G.; Vanderwolf, CH., editors. Electrical Activity of the Archicortex. Akademia Kiado; Budapest: 1985. p. 217-220.

Redish, AD. Beyond cognitive map: From place cells to episodic memory. The MIT Press; Cambridge: 1999.

Rice A, Rafiq A, Shapiro SM, Jakoi ER, Coulter DA, DeLorenzo RJ. Long-lasting reduction of inhibitory function and gamma-aminobutyric acid type A receptor subunit mRNA expression in a model of temporal lobe epilepsy. Proceedings of the National Academy of Sciences 1996;93(18):9665–9669.

Rotenberg A, Mayford M, Hawkins RD, Kandel ER, Muller RU. Mice expressing activated CaMKII lack low frequency LTP and do not form stable place cells in the CA1 region of the hippocampus. Cell 1996;87(7):1351–1361. [PubMed: 8980240]

Rotter S, Diesmann M. Exact simulation of time-invariant linear systems with applications to neuronal modeling. Biological Cybernetics 1999;81:381–402. [PubMed: 10592015]

Senft SL, Ascoli GA. Reconstruction of brain networks by algorithmic amplification of morphometry data. Lecture Notes in Computer Science 1999;1606:25–33.

Settanni G, Treves A. Analytical model for the effects of learning on spike count distributions. Neural Computation 2000;12(8):1773–1787. [PubMed: 10953238]

Siapas AG, Wilson MA. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. Neuron 1998;21(5):123–1128. [PubMed: 9697857]

Sueur, J.; Aubin, T.; Simonis-Sueur, C. Seewave 1.4.4. Université Paris XI; 2007. MNHN. http://cran.r-project.org/src/contrib/Descriptions/seewave.html

Taube JS, Muller RU, Ranck JB Jr. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. Journal of Neuroscience 1990;10:420–435. [PubMed: 2303851]

Todtenkopf, MS. Mark's Hippocampus Web Site. Northeastern University; Boston, MA: 2000. http://www.psych.neu.edu/stellar/lab/people/MST/Hipp/Index.html

Treves A, Skaggs WE, Barnes CA. How much of the hippocampus can be explained by functional constraints? Hippocampus 1996;6(6):666–674. [PubMed: 9034853]

Urban NN, Henze DA, Barrionuevo G. Revisiting the role of the hippocampal mossy fiber synapse. Hippocampus 2001;11:408–417. [PubMed: 11530845]

Whittington MA, Traub RD, Kopell N, Ementrout B, Buhl EH. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. Journal of Psychophysiology 2000;38:315–336.

Wilson MA, Tonegawa S. Synaptic plasticity, place cells and spatial memory: study with second generation knockouts. Trends in Neuroscience 1997;20(3):102–106.

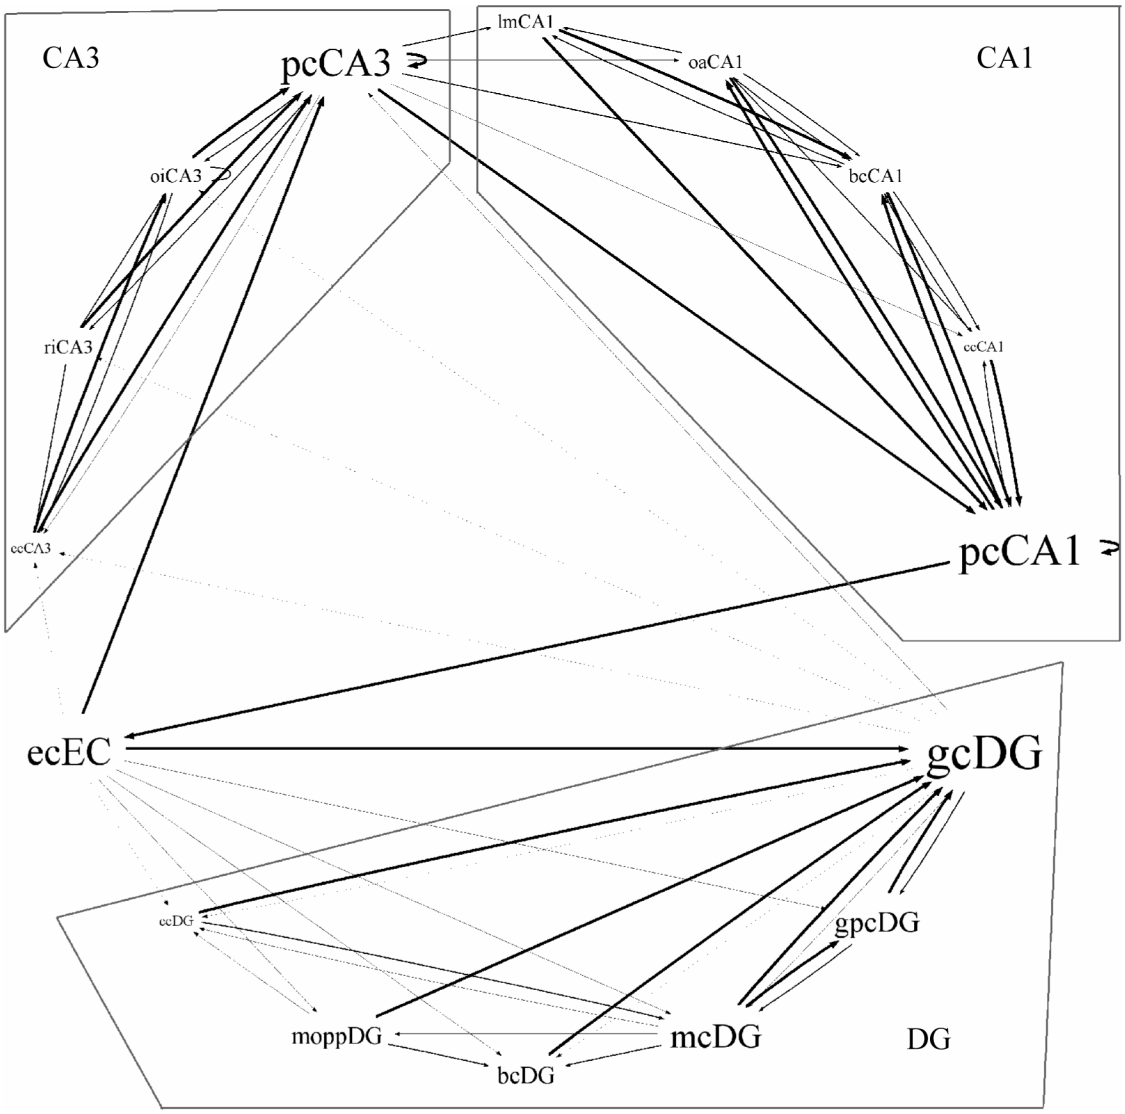

Figure 1.

Connectivity circuit of the hippocampal network. Font size and arrow size are approximately proportional to the number of neurons in each class and to the number of postsynaptic connection made by each presynaptic cell, respectively. Gray boxes indicate the major subdivisions: DG, CA3, and CA1.

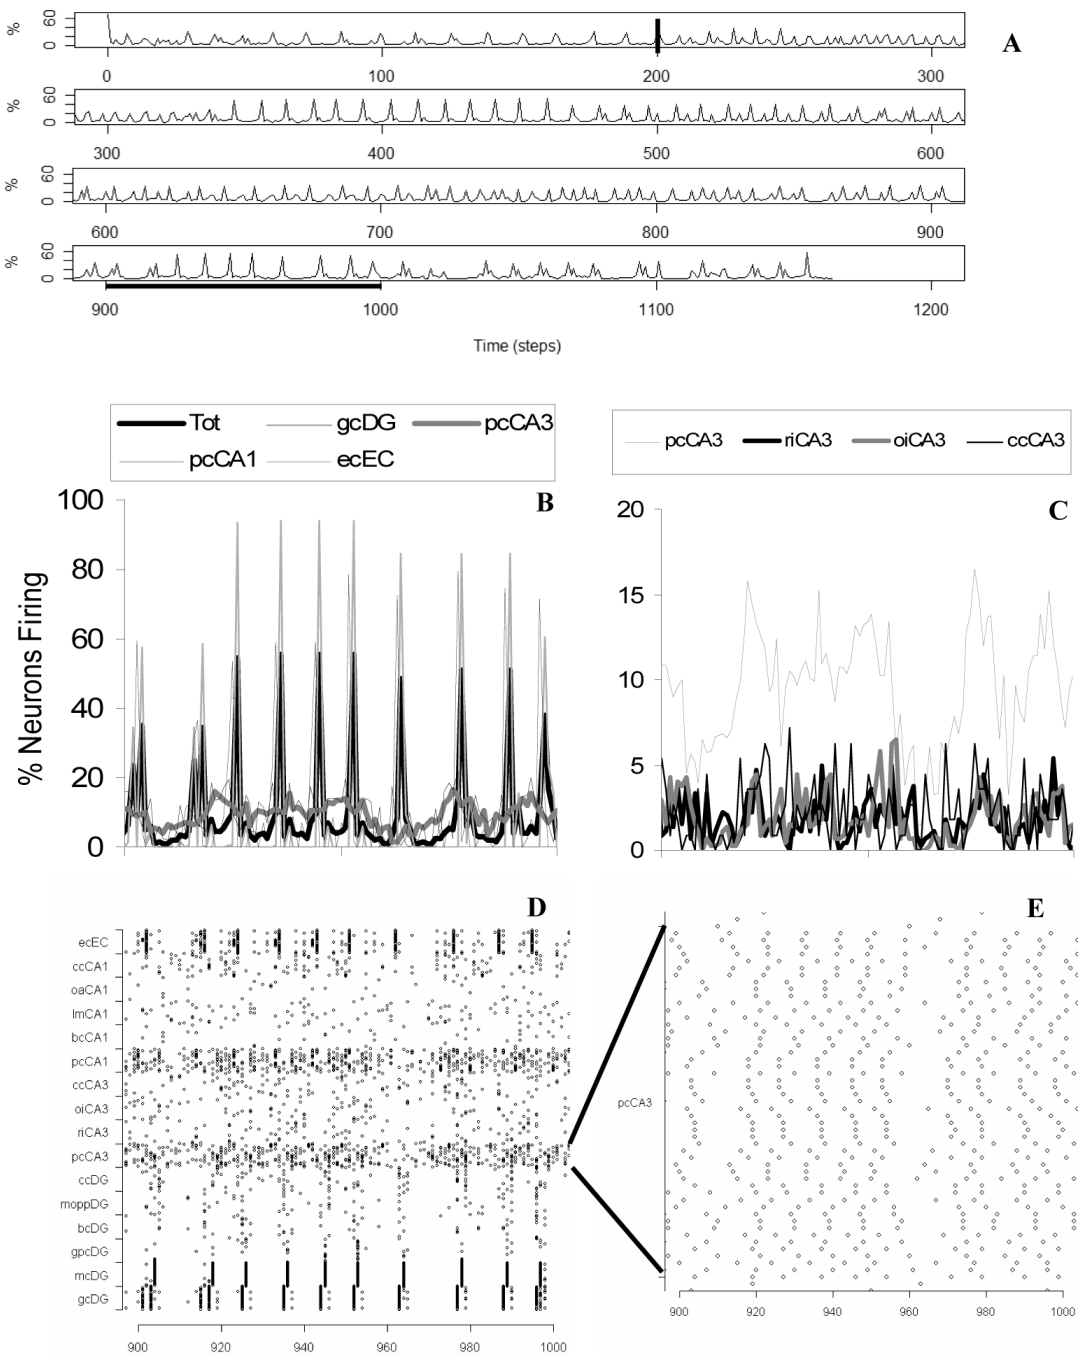

Figure 2.

Activity in a representative simulation of the 16MP200 model. (A) Overall activity trace (expressed as percentage of spiking neurons over all cell classes) for the complete simulation lasting 1164 steps. At t=0 the network is initiated by randomly activating 70% of all neurons. Thresholds become constant at t=200 (vertical black bar). The horizontal black bar from 900 to 1000 steps marks the temporal period illustrated in all subsequent panels (B-E). (B) Activity traces for the overall network (Tot) and for each of the principal neurons, expressed as percentage of spiking neuron within each class. (C) Activity traces for every neuronal class within the CA3 region. (D) Raster plot of 50 random neurons per class. (E) Zoom-in on the pcCA3 units.

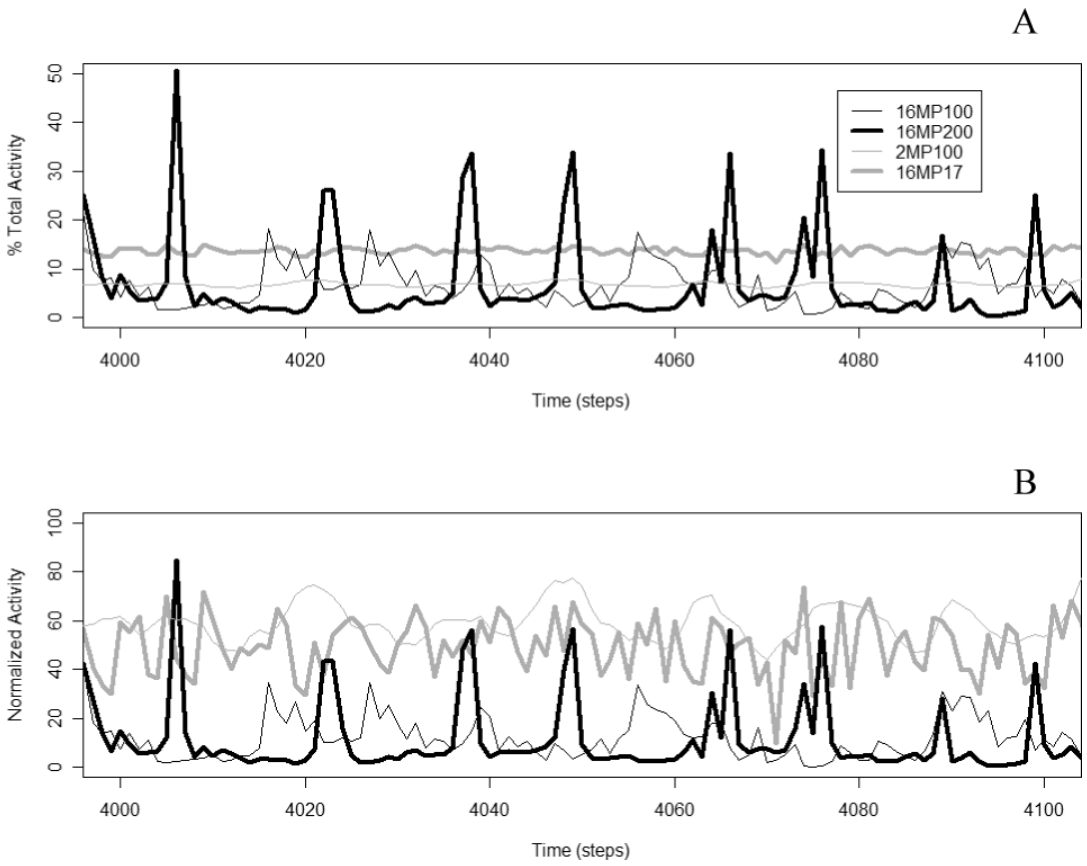

Figure 3.

Comparison among network configurations. (A): Total activity demonstrating a wider dynamic range for 16MP100 and 16MP200 (black lines) compared to both 2MP100 and 16MP17 (grey lines). (B) Normalized activity within the minimum and maximum values of each simulation from the end of the transient phase through the network lifespan. Only 16MP100 and 16MP200 present peaks followed by periods with lower activity.

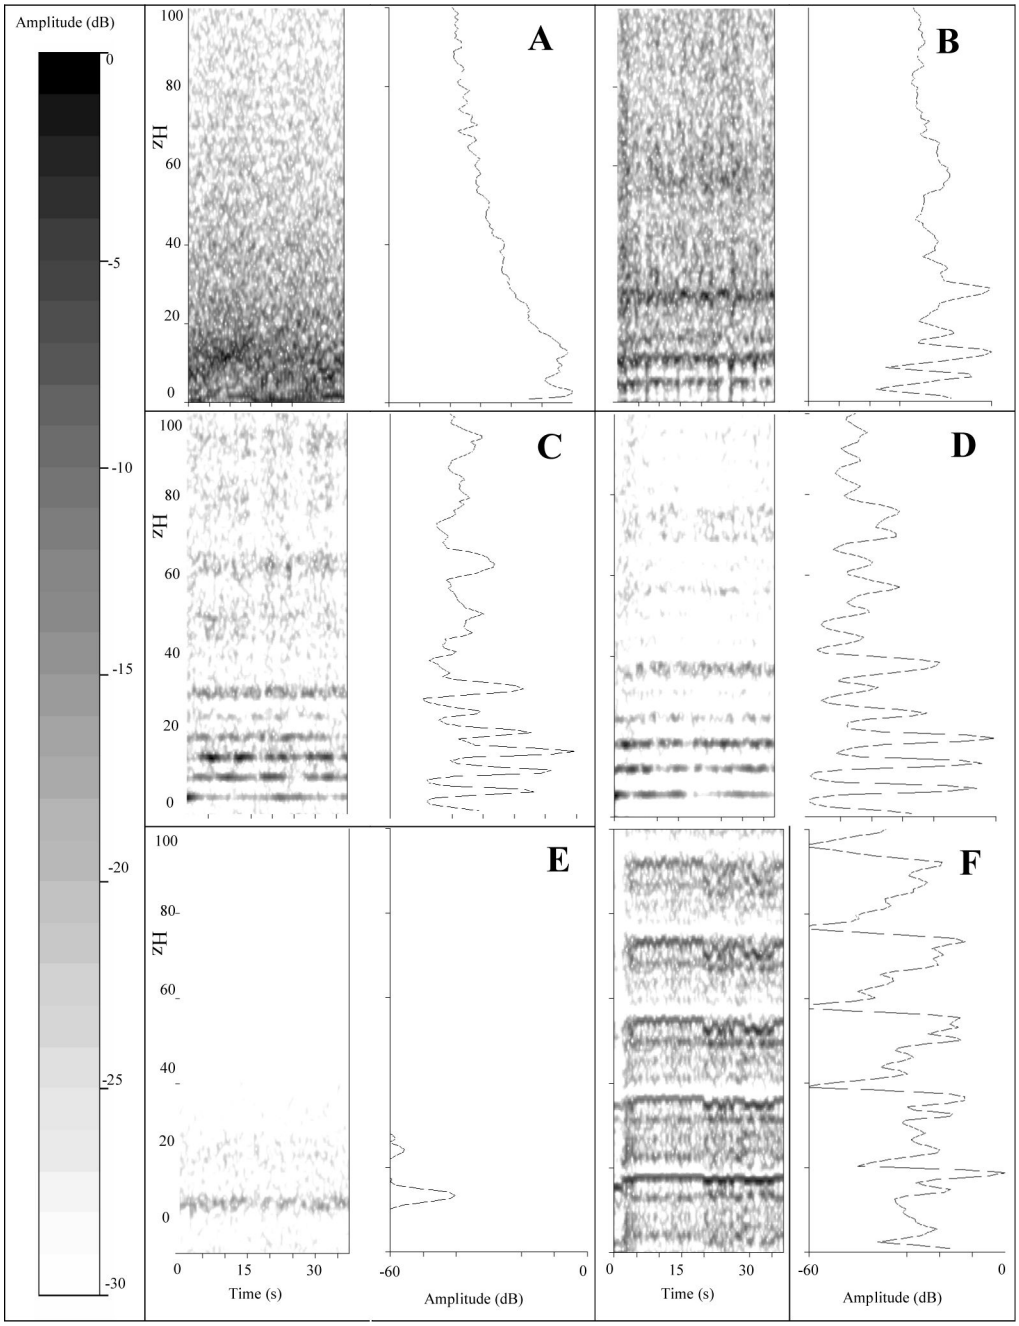

Figure 4.

Frequency distribution over time and associated power spectra density of the overall activity (computed after subtracting the average activity for each model) in representative simulations of each network configuration. For these plots one time step is equivalent to 2.5ms. (A) 16MP200, (B) 16MP100, (C) 2MP200, (D) 2MP100, (E) 16MP17, (F) 16IF17. Amplitude gray scale on left panel applies to all frequency distributions.

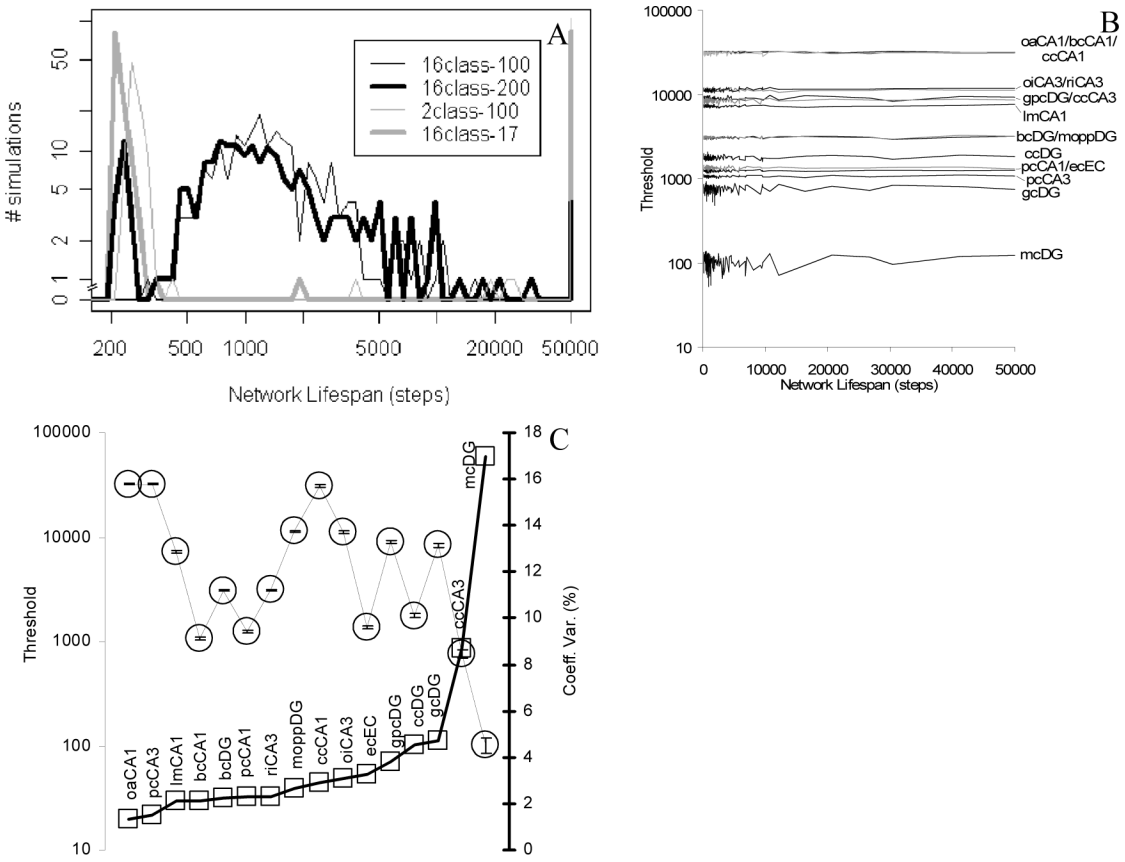

Figure 5.

(A) Distribution of network lifespan (the duration of self-sustained activity) over 200 simulations per model. 16MP200 and 16MP100 present different patterns from 16MP17 and 2MP100. (B) Threshold plot over network lifespan for 200 simulations of 16MP200. There are no trends, suggesting that the algorithmic determination of threshold values does not influence network stability. (C) Average threshold values (circles and dotted line, left axis) for each cell class, ordered by coefficient of variation (squares and full line, right axis). Error bars within circle represent standard deviation: thresholds are narrowly constrained across different simulations.

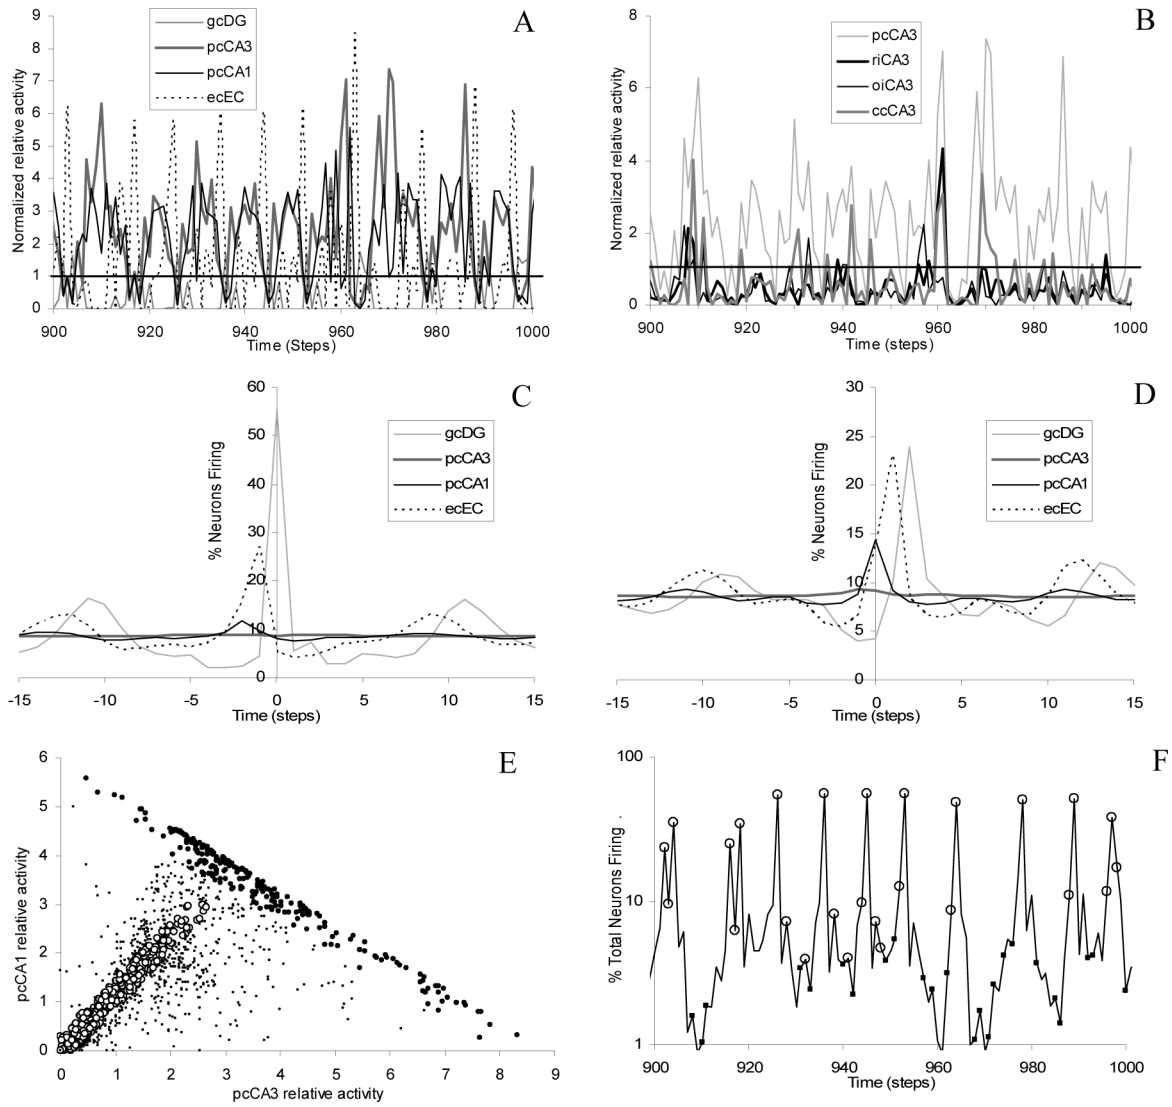

Figure 6.

Network dynamics. A) Relative activity of all principal neurons normalized to the relative activity of the overall network. For example, if at a given time step 10% of all neurons, 5% of gcDG, and 20% of pcCA3 cells are active, the normalized values are 0.5 and 2 for gcDG and pcCA3, respectively. In this plot, the overall network activity has the constant value of 1 (solid horizontal line). (B) Normalized relative activity for all CA3 neurons. (C) Average gcDG population spike and activity of all other principal neurons within ±15 steps. (D) Same as C for average pcCA1 population spike. (E) Scatter plot of pcCA1 vs pcCA3 activities, both relative to total network activity. Empty and filled circles identify the time steps in which pcCA1 and pcCA3 activities are positively and negatively related, respectively. (F) Temporal distribution of symbols highlighted in (E) over total network activity as percent of neurons firing.

Table 1

Connectivity matrix. Each element in row I and column J indicates the average number of pre-synaptic neurons of class J contacting one neuron of class I. The number of neurons in the real rat hippocampus is reported for each class in parenthesis.

	gcDG	pcCA3	pcCA1	ecEC	mcDG	gpcDG	bcDG	moppDG	ccDG
gcDG (1,000,000)				6800	1080	600	142	4733	6
mcDG (30,000)	800			160		100			6
gpcDG (15,000)	9000			400	3000				
bcDG (10,000)	800			2000	999			150	
moppDG (10,000)	800			2000	999				
ccDG (1,000)	800			2000	990			150	
					riCA3	**oiCA3**	**ccCA3**		
pcCA3 (200,000)	70	12000		1	2400	2400	300		
riCA3 (4,000)	250	20000		50					
oiCA3 (4,000)	250	20000			800	800	300		
ccCA3 (1,000)	1000	20000		200	800	800			
					bcCA1	**lmCA1**	**oaCA1**	**ccCA1**	
pcCA1 (300,000)		12000	2100		2400	2666	3600	133	
bcCA1 (4,000)		12000	90000			2000	800	200	
lmCA1 (4,000)		12000			800		800		
oaCA1 (4,000)		12000	187500		800				
ccCA1 (1,000)		12000	90000		800		800		
ecEC (200,000)			15000						

Table 2

Summary information comparing MP model among different network configurations. Averages and standard deviations are computed over 200 simulations from values collected from the end of the transient phase through the network lifespan.

(n=200)	**16MP200**	**16MP100**	**16MP17**	**2MP200**	**2MP100**	**2MP17**
# neurons	200k	100k	17k	200k	100k	17k
# neurons in smallest class	112	56	9	6488	3244	551
# synapses	300M	70M	2M	300M	70M	2M
Max Activity	49.84± 13.75	30.79± 11.16	14.86± 18.81	10.5± 1.80	9.20± 2.39	6.97± 1.47
Min Activity	0.014± 0.0988	0.0045± 0.0052	0.32± 0.48	0.04± 0.24	0.54± 0.87	0.23± 0.48
95% Range	39.64± 14.89	17.96± 6.07	3.32.± 2.76	7.5± 1.28	5.58± 1.17	3.97± 0.96
Avg Activity	7.73± 1.92	6.15 0.82±	4.23± 4.30	5.20± 1.62	4.20± 2.88	2.96± 1.37