Dissertation

submitted to the

Combined Faculties for the Natural Sciences and for Mathematics

of the Ruperto-Carola University of Heidelberg, Germany

for the degree of

Doctor of Natural Sciences

presented by

Grant Sutcliffe, B.Sc., M.Sc.

Born in Rochdale, United Kingdom

Oral-examination: 08.05.2017

Stochastic dynamics and delta-band oscillations in clustered spiking networks

Referees: Prof. Dr. Daniel Durstewitz

Dr. Kevin Allen

Statement of Originality

Declarations according to §8 (3) b) and c) of the doctoral degree regulations: a) I hereby declare that I have written the submitted dissertation myself and in this process have used no other sources or materials than those expressly indicated, b) I hereby declare that I have not applied to be examined at any other institution, nor have I used the dissertation in this or any other form at any other institution as an examination paper, nor submitted it to any other faculty as a dissertation.

Abstract

Following experimental measurements of clustered connectivity in the cortex, recent studies have found that clustering connections in simulated spiking networks causes transitions between high and low firing-rate states in subgroups of neurons. An open question is to what extent the sequence of transitions in such networks can be related to existing statistical and mechanical models of sequence generation. In this thesis we present several studies of the relationship between connection structure and network dynamics in balanced spiking networks. We investigate which qualities of the network connection matrix support the generation of state sequences, and which properties determine the specific structure of transitions between states. We find that adding densely overlapping clusters with equal levels of recurrent connectivity to a network with dense inhibition can produce sequential winner-takes-all dynamics in which high-activity states pass between correlated clusters. This activity is reflected in the power spectrum of spiking activity as a peak in the low-frequency delta range. We describe and verify sequence dynamics with a Markov chain framework, and compare them mechanically to “latching” models of sequence generation. Additionally we quantify the chaos of clustered networks and find that minimally separated states diverge in distinct stages. The results clarify the computational capabilities of clustered spiking networks and their relationship to experimental findings. We conclude that the results provide a supporting intermediate link between abstract models and biological instances of sequence generation.

Zusammenfassung

Basierend auf experimentellen Untersuchungen der geclusterten Konnektivität im Kortex haben neuere Studien herausgefunden, dass geclusterte Verbindungen in simulierten, gepulsten neuronalen Netzwerken Übergänge zwischen hohen und niedrigen Feuerraten in Untergruppen von Neuronen erzeugen. Eine offene Frage ist, inwieweit die Schaltaktivierung in solchen Netzwerken auf existierende statistische und mechanische Model der Sequenzerzeugung bezogen werden kann. In dieser Arbeit präsentieren wir eine Reihe von Studien über den Zusammenhang zwischen der Verbindungsstruktur und Netzwerkdynamik von balancierten, gepulsten neuronalen Netzwerken. Wir untersuchen welche Eigenschaften der Verbindungsmatrix eines Netzwerks zur Generierung von Zustandssequenzen führen, und welche Eigenschaften die spezifische Struktur der Übergänge zwischen einzelnen Zuständen bedingen. Wir konnten zeigen, dass das Hinzufügen dichter, überlappender Cluster mit dem gleichen Ausmaß an rezidivierender Aktivität zu einem Netzwerk mit starker Inhibition, sequentielle winner-takes-all Dynamiken erzeugen können, bei denen Zustände hoher Aktivität zwischen korrelierten Clustern hin- und herwechseln. Diese Aktivität spiegelt sich im Leistungsspektrum der Spiking-Aktivität als Peak im niederfrequenten Delta-Bereich wider. Wir beschreiben und verifizieren diese Sequenzdynamiken im Rahmen eines Markov-Ketten-Prozesses und vergleichen diese mit der Mechanik von „Latching“ Modellen der Sequenzgenerierung. Darüber hinaus quantifizieren wir das Chaos in geclusterten Netzwerken und finden, dass minimal separierte Zustände in bestimmten Stadien voneinander abweichen. Die Ergebnisse geben Einblick in die rechnerischen Fähigkeiten von geclusterten Spiking-Netzwerken und setzten diese in Bezug mit experimentellen Befunden. Wir schließen daraus das die Ergebnisse einen intermediären Zusammenhang zwischen abstrakten Modellen und biologischen Beispielen von Sequenzgenerierung herstellen.

Table of Contents

List of Figures

List of Tables

1 Introduction

Activity in neural systems is characterized by dynamic changes of state across multiple levels of space and time. Measuring and understanding the nature of these dynamic changes and their relationships to neural structure is one of the greatest technical and theoretical challenges in science.

A core element of neural activity at the single neuron level is the action potential, a transient change in an individual neuron's state that is the basis of inter-neuronal communication. At the level of neural populations, transient patterns of increases and decreases in the rate of action potential firing are associated with specific perceptions, motor actions and behavioral sequences (Balaguer-Ballester et al. 2011; Mante et al. 2013). The physiological bases of these rate variations and their relationships to thought and behavior are uncertain but under intense investigation. On the experimental side, the effort includes the development of imaging technology and statistical methods. The neural information gathered with these methods is used to develop and test hypotheses regarding activity-cognition relationships and computational principles.

1.1 In-vivo neural dynamics

In some lower brain areas action potential firing is closely tied to sensory or motor events, such as the stimulation of a sensory receptor or the activation of muscle fibres. Further up the brain hierarchy firing is irregular but typically varies in rate over time. In some cases, the relationship between firing patterns and cognition appears to have a relatively straightforward interpretation. One example is delay period activity during experimental working memory tasks. In-vivo recordings of primates have found raised firing rates in stimulus-specific neuron populations associated with recall of a stimulus a short period after its presentation. This activity is particularly prominent in the prefrontal and parietal cortex (Fuster and Alexander 1971; D'Esposito and Postle 2015; Constantinidis and Steinmetz 1996). Comparable activity has been observed using fMRI techniques in humans (D'Esposito and Postle 2015). The activity in this case has a clear theoretical interpretation in terms of attractor states in systems of recurrently connected units

(Durstewitz, Seamans, and Sejnowski 2000b).

In most cases, neural activity displays a more complex relationship to cognition, involving increases and decreases of firing rates across groups of neurons that change over time. Coordinated sequences of activity have been detected in-vivo in response to apparently simple sensory features, such as in rodent gustatory cortex in response to taste stimuli (Jones et al. 2007), and in rodent visual cortex in response to natural scenes and grating stimuli (Carrillo-Reid et al. 2015). Spatio-temporal sequences of activity are also observed in anesthetized animals and in in-vitro slice preparations, strongly suggesting that the spatio-temporal dynamics result from intrinsic local circuit properties, rather than being directly driven by external input (Ikegaya et al. 2004; Luczak et al. 2007; Cossart, Aronov, and Yuste 2003).

As might be expected, sequences of neural activity have also been experimentally associated with sequences of observable behavior. Response-predictive transitions between apparently discrete states have been observed during a delayed localization task in primates (Seidemann et al. 1996). In the hippocampus, internally-generated sequences have been associated with the storage of behavioral sequences and the generation of prospective future behavioral sequences during the planning of goal-directed behavior (Pezzulo et al. 2014).

Choice-specific sequences of activity have been observed in rodent prefrontal and parietal cortex during a maze task requiring perceptual judgement and behavioral choice (Harvey, Coen, and Tank 2012; Fujisawa et al. 2008), and stage-specific temporal patterns of activity have been found in the rodent anterior cingulate cortex during a maze task (Balaguer-Ballester et al. 2011).

Song production in songbirds is a particularly clear subject for the study of behavioral sequences. This is most often studied in finches. In finches, neurons in the forebrain premotor nucleus HVC show bursting activity which can be associated with specific segments in the birds song, such as phrases, syllables and notes (Fujimoto, Hasegawa, and Watanabe 2011; Long, Jin, and Fee 2010). This has been modeled in various ways, including statistical models of the song sequence's underlying syntax (Katahira et al. 2011) and mechanistic models of the neural circuits involved (Long, Jin, and Fee 2010).

Several of the experimental examples just mentioned explicitly describe sequences

in terms of sequential activations of groups of neurons in which each group forms a functional unit whose members activate more or less simultaneously (Cossart, Aronov, and Yuste 2003; Jones et al. 2007; Carrillo-Reid et al. 2015; Seidemann et al. 1996). This is in accordance with the hypothesis that the neural assembly is a fundamental unit in the functional organization of the brain (Yuste 2015). This hypothesis will be discussed further in section 1.3.3.

1.2 Oscillatory activity in the brain

The power spectral density of summed brain activity measures such as the local field potential shows peaks in several frequency bands. These can be related to various dynamic processes, cognitive states and computational functions. Delta oscillations in the sub 4 Hz band are particularly prominent in sleep states, but are present in all cognitive states and have been reported to increase during executive cognitive processes such as working memory, attention and response inhibition (Harmony 2013). Source-localization of electroencephalographic (EEG) data indicates that sources of delta oscillations in the awake brain include frontal cortical areas associated with executive functions, which supports this association (Harmony 2013; Guntekin and Basar 2015). Reduced power of delta oscillations during performance of experimental cognitive tasks has been found in subjects with various neurological and psychiatric diseases, including Alzheimer's disease, schizophrenia and bipolar disorder (Guntekin and Basar 2015). In contrast, the same patient groups tend to show *increased* spontaneous delta oscillations when not engaged in a task (Guntekin and Basar 2015).

The associations between delta oscillations and cognitive states suggests that they have a computational function. Harmony (2013) suggested that they could help synchronize large-scale cortical networks with control systems in the frontal lobes during the application of internally and externally directed attention. Another suggestion is that delta oscillations increase brain synchronization with autonomic processes, potentially including processes related to motivation (Knyazev 2012).

One relatively untested possibility is that delta oscillations influence neural dynamics in a way that alters their computational qualities. This could support effective processing in general, and oscillations could also be selectively introduced to a system in order to adapt computational qualities to behavioral demands. Animals

face a variety of different environments and behavioral challenges, and computationally optimum dynamics in a neural system could differ depending on context. An example of contrasting cognitive strategies is the contrast between exploration and exploitation when seeking resources (Hills et al. 2014). This can be related to the contrast between cognitive flexibility and stability, which has been related to neural dynamics and their adaptation through modulatory neurotransmitters such as dopamine (Durstewitz and Seamans 2008). Neuromodulatory neurotransmitters modulate oscillatory activity in vitro, which simulation analysis indicates occurs due to modulation of local neuron properties (Durstewitz 2009). The specific effect of neuromodulator-induced oscillations on behaviorally-relevant system properties is largely unexplored.

1.3 Models of neural dynamics

1.3.1 Balanced spiking networks

The firing of cortical neurons is often temporally irregular. The input to cortical neurons is largely from other cortical neurons, and an influential hypothesis is that irregular spiking activity emerges from intrinsic network-level dynamics when large numbers of otherwise noiseless neurons are sparsely interconnected with strong synapses, with parameters appropriately chosen so that recurrent inhibition within the network is strong enough to counteract excitation. Neurons in such networks have approximately equal mean levels of excitatory and inhibitory input, and so are termed “balanced” networks. The irregular spiking in this case arises as a result of fluctuations in the mean input to each neuron, which drive the neuron's membrane potential over the spiking threshold at irregular intervals. This phenomenon has been well demonstrated and explored theoretically and in simulation (van Vreeswijk and Sompolinsky 1996). Excitatory neurons in the cortex are sparsely connected (Perin, Berger, and Markram 2011), which supports this hypothesis.

A practical and commonly used assumption in the simulation and analysis of balanced networks is to assume that all neurons in the population being simulated send and receive connections to and from all other neurons with equal probability, effectively being randomly connected. This assumption has the advantage of facilitating analytical tractability, as the numbers of connections, and the mean and

variance of synaptic input, can be predicted in terms of probability distributions.

The computational capabilities of randomly connected balanced networks are uncertain but under continuing investigation. One simple but useful property is that the mean firing rate of such a network increases approximately linearly with the input to the network, with a timescale several times smaller than the integration time constant of a single neuron. A balanced network could therefore potentially function as a fast integrating relay of its inputs (van Vreeswijk and Sompolinsky 1996).

A further emergent property of a randomly-connected balanced network is that the firing rate of individual units can display a kind of amplifying selectivity to different patterns of input to the network, even when the connections are random rather than specifically designed to produce selective responses. The reason for this property is that the recurrent inhibition in a balanced network dampens the mean of excitation to individual units to a greater extent than the variance of input, which results in responses of individual units to different input patterns being relatively large even if the absolute input to each unit is similar (Pehlevan and Sompolinsky 2014).

1.3.2 Reservoir computing

A particularly notable computational capability of artificial randomly connected networks is to be found within the theoretical framework of *reservoir computing*. Reservoir computers are a class of computing system which perform computations, in the sense of mapping inputs to outputs, on spatial and / or temporal patterns of input. This is performed with the use of a dynamic “reservoir” of randomly connected units, which are “read” as a weighted sum by a layer of linear output units. A suitable reservoir has two particularly useful qualities. First, it can produce non-linear responses to linear combinations of patterns, therefore enabling complex non-linear computations. Second, because the state of a dynamic reservoir is influenced by both its past states and current input, it has a form of memory that can be used to perform computations on temporal series of input.

The set of mappings performed by reservoir computers can be powerful, but because the only computation-specific connection-setting they require in order to produce the desired input-output responses is the relatively simple training required to set the connections from the reservoir to the linear output units, they are relatively simple to create in comparison to neural computers in which connections between

hidden, recurrently-connected units must also be trained, as this form of training is difficult to design, parameter-sensitive and computationally comparatively expensive (Lukoševičius and Jaeger 2009).

According to several analyses, the neural reservoir in a reservoir computing system should have certain properties to be computationally useful. These properties are closely related to the mathematical concepts of stability and chaos in dynamic systems (Strogatz 2014). The degree of stability and chaos in systems such as neural networks is dependent on their parameters, and optimal computational power is hypothesized to be found in systems which are at the so-called “edge of chaos”, close to the boundary between stable and chaotic dynamics. This is because both stabilty and chaos have computationally useful properties. When considering inputs as initial states of a dynamic system and outputs as the state of the system at some later point, stable systems generate similar outputs for similar inputs, thereby supporting generalization of input-output transformations to novel inputs. On the other hand, as the current state of a chaotic system is sensitive to its past states, a chaotic system possesses a form of memory for its past inputs. This allows for computation on temporally-extended inputs. Systems at the edge of chaos potentially possess both qualities to a useful degree, and so support generalizing computation on temporally-extended inputs.

Several investigations have been made into the computational and chaotic properties of various types of simulated neural network and their relationships to network parameters. One particularly sharp distinction between network classes in this respect is the distinction between networks in which units are connected with continuously-valued connections, such as networks of firing rate units, and networks in which units are connected with binary-valued connections or pulses, such as spiking networks (Büsing, Schrauwen, and Legenstein 2010). Systems with binary connections show a particularly strong parameter-dependence of the degree of chaos, with a sharp transition between ordered and chaotic dynamics and a considerably smaller range of parameters which produce the useful edge of chaos state (Büsing, Schrauwen, and Legenstein 2010). Relevantly to the subject of in-vivo computation, spiking and binary-connected networks have been found to be produce chaotic dynamics at relatively low values of two important system parameters, those controlling recurrent connection strengths and the numbers of

connections to each unit. Of these, the parameter controlling the number of within-reservoir connections to each reservoir unit, or “in-degree”, is the parameter most problematic to the hypothesis that reservoir computing principles are used in cortical neural networks. Simulation, theoretical and in-vivo studies have provided evidence that fast chaotic dynamics occur when reservoir units have more than a small number of connections, typically less than 10 (Wallace, Maei, and Latham 2013; London et al. 2010; Büsing, Schrauwen, and Legenstein 2010; Monteforte and Wolf 2010). As cortical neurons typically have many more connections that this, this suggests that reservoir computing principles do not operate in the cortex. Similarly, the transition from order to chaos with increased connection weights has been found to be sharp in binary and pulse-connected networks. Although this finding has not been explicitly related to the balanced state, simulation studies indicate that this transition occurs at connection strengths lower than that required to spontaneously produce balanced-state irregular spiking (Natschläger, Bertschinger, and Legenstein 2005; Legenstein and Maass 2007). This suggests that balanced state networks in particular are not good reservoirs, agreeing with analysis that finds that balanced networks are highly chaotic (Jahnke, Memmesheimer, and Timme 2009; Monteforte and Wolf 2010). Compounding this issue is the finding that synaptic transmission is unreliable, with synaptic failures preventing a high proportion of action potentials from being transmitted to their postsynaptic targets (Branco and Staras 2009). Network dynamics are highly sensitive to the omission of spikes (Wallace, Maei, and Latham 2013).

A notable exception to the predominant finding of chaos in highly-connected spiking networks is the case of networks of deterministic spiking neurons connected with inhibitory delta synapses, where delta synapses are non-temporal synapses which transfer charge instantly. Such networks display order rather than chaos when differences in initial conditions are limited to small differences in spike times. (Jahnke, Memmesheimer, and Timme 2009; Monteforte and Wolf 2012).

1.3.3 Clustered networks and assemblies

It is very likely that computation in cortical neural networks requires some degree of non-random connectivity. As it is not yet possible to measure more than a tiny fraction of the synapses in any one biological neural circuit, the nature of this

connectivity requires inference from limited data and theoretical speculation (Yuste 2015). A popular hypothesis is that neurons act together as groups with correlated targets. A particularly influential further hypothesis is the existence of “assemblies” of excitatory neurons which project excitatory connections to each other with raised probabilities or potencies, thereby providing a mechanism by which they can mutually excite each other into a state of raised spike firing rates. As the raised firing rate state in all assembly members can be initiated by similar states of raised or partially raised firing rates of a subset of assembly neurons, the high-activity state of the assembly constitutes an attractor, into which nearby states are drawn. This assembly attractor therefore provides local stability in the neural dynamics at the level of firing rates.

The ability of an assembly to maintain a specific state for extended periods, in conjunction with the possibility that this state can be engaged and disengaged by external mechanisms, provides a potential mechanism for cognitive phenomena that require the maintenance of information for extended periods in the absence of external cues, such as working memory or sustained attention. The hypothesis that this is employed in the brain is supported by evidence of sustained firing activity in parts of the brain associated with state maintenance, such as the prefrontal cortex, during experimental tasks that require temporary storage of information, such as working memory tasks (Durstewitz, Seamans, and Sejnowski 2000b).

The attractor property of the recurrently-connected assembly, whereby states in the “basin of attraction” of the attractor lead to the system moving to the attractor state, has been proposed as a model of decision processes. In these, different attractors correspond to different choices in a decision task, with evidence for a specific decision causing excitatory input to the corresponding assembly that pushes it towards the basin of attraction of the high-activity state. Such systems have been modeled at various levels of realism and complexity, and have some experimental support from in-vivo electrophysiology in primates performing somatosensory and visual discrimination tasks (Braun and Mattia 2010; Durstewitz and Deco 2008), and in rodents performing a rule-learning task (Durstewitz et al. 2010).

1.3.4 Attractor assemblies and firing rate variability

In addition to the previously-mentioned potential for recurrently-connected neural

assemblies to create a kind of stability in neural activity, there has recently been increased interest in the somewhat paradoxical potential for recurrent assemblies to generate variability in large-scale neural dynamics. The variability in this case is at the level of variability of firing rates over longer timescales of 100s of milliseconds or more, which has been observed experimentally in the cortex (Churchland et al. 2011; Litwin-Kumar and Doiron 2012).

The core principle here is the existence of assembly attractors similar to those implemented in working memory models, but in which both the low and high firing rate states have transient or reduced stability, so that the assembles switch between their low and high activity states purely as a consequence of dynamics internal to the network. A central cause of state switching in existing models of this phenomenon is the irregularity of spike firing inherent to balanced-type spiking networks, which, due to the finite size of the assemblies, causes fluctuations in the cumulative recurrent excitation of the assembly which causes the assembly to “jump” between the basins of attractions of the low- and high-firing-rate states. The dependence of state transitions on the inherent irregularity of spike times results in stochastic transitions between states in specific assemblies, and a corresponding stochastic fluctuation of the firing rates of individual neurons. The combination of stochastic variation in firing rates along with spike times has correspondingly been referred to as a doubly stochastic process (Churchland et al. 2011; Litwin-Kumar and Doiron 2012).

Existing studies have mainly focused on the ability of spiking networks of multiple transiently-stable attractors to reproduce broad statistical measures of in-vivo activity, such as firing rate variability and the Fano factor. However, little attention has been focused on the question of how the activity of such networks relates to more complicated phenomena seen in biological networks, and in particular how it relates to the tendency for biological neural networks to show predictable activity trajectories.

The relationship between assemblies and firing rate variability suggests a relationship with theoretical neural mechanisms of sequence generation, particularly the “synfire chain”. Synfire chains are a theoretical form of sequence storage and generation in which assemblies of neurons are connected in feed-forward chains, and so produce sequential chains of high-firing rate activity. Synfire chains have

been well-explored theoretically and can be generated with various naturalistic neural plasticity mechanisms (Zheng and Triesch 2014; Abeles 2009), and have also received experimental confirmation (Ikegaya et al. 2004; Long, Jin, and Fee 2010). However, the groups in theoretical and simulated synfire chain models do not usually include recurrent connections within groups, and so do not necessarily fit the definition of recurrently connected assemblies (Abeles 2009).

Aside from synfire chains, the spontaneous activity in clustered neural networks suggests potential relationships with other computational models. These include models of spontaneous sequence generation, with functional applications including language production (Russo and Treves 2011) and memory search (Tsuda 2001).

1.4 Thesis motivation and organization

Although dynamic activity in clustered spiking networks has been reasonably well-explored in respect to their formation and effect on spiking statistics, their computational capacities remain largely untested. This study explores the effect of clustering on the dynamics of balanced networks, with an emphasis on implications for computation. The state-switching dynamics described in previous studies suggests a potential link to existing models of sequence generation, and particular effort will be given to interpret output in this context. The study uses computational neuroscience methods, particularly simulations and analysis of representative systems.

After this introduction, the first part of the thesis will describe a process of producing a clustered network with stochastic attractor dynamics, by measuring the dynamics of a balanced network in terms of assembly attractor states and low-frequency oscillations, and systematically determining which qualities determine dynamics in the case of dense overlapping clusters. In the second part we will analyze the output of a subset of the previously-determined network configurations and statistically test the hypothesis that activity can be described in terms of a Markov chain. We examine furthermore which dynamical phenomena induce transitions between states, and which features of the network connection matrix determine probabilities of transitions between specific pairs of states. In the final part we will test and quantify the order or chaos of the previously created clustered and unclustered networks, and test the hypothesis that clustering qualitatively or quantitatively

changes the form of chaos in the network.

2 Defining a sparse spiking network model with overlapping assemblies

In this chapter we describe the specification and preliminary analysis of the neural network that is the main subject of this thesis. The work described was performed to test the hypothesis that clustered spiking networks produce computationally noteworthy dynamics, and also to find potentially interesting parameter regimes for more in-depth analysis.

We specify a spiking neural network, and a “base” balanced network parameter configuration that produces irregular firing at low physiologically plausible rates. We then specify both an existing and novel method of clustering connections, and several statistical measures to be applied to the spiking output of the simulated networks.

Using these measures to describe the output, we then simulate networks created using the two clustering methods, systematically varying the clustering parameters in order to quantify their effect on network output. We furthermore measure the effect of clustering on network connection structure, using measures which include the clustering coefficient and the eigendecomposition of the connection matrix. Relationships between network output and connection structure are noted.

2.1 Spiking network model

2.1.1 Neuron and synapse model definitions

Neurons were modeled as adapting leaky integrate and fire units. The evolution of the voltage V of these units was determined by

$$C_m \frac{dV(t)}{dt} = \left(-g_L \left(V(t) - E_L \right) + I_{inj}(t) + I_{syn}(t) - w(t) \right) \quad (1)$$

where C_m is the capacitance of the membrane of the unit, g_L the membrane conductance, and E_L the reversal potential of the leakage current. I_{inj} and I_{syn} are externally-injected and net synaptic currents,

respectively. w is a hyperpolarising variable which obeys

$$\tau_w \frac{dw(t)}{dt} = a(V(t) - E_L) - w(t) \tag{2}$$

where τ_w is the time constant and a is a scalar which controls the subthreshold growth rate of w. The definition of the w adaptation current follows that of the equivalent current in the adaptive exponential integrate and fire neuron proposed by Brette and Gerstner (2005).

The action potential firing condition is triggered when V reaches a threshold value θ from below. When this happens, the membrane potential V is set to the reset value V_r and held at this value for a fixed refractory period V_{ref}, the adaptation variable w is augmented by a fixed value b, and a synaptic event is triggered at all downstream connected units.

Total synaptic currents I_{syn} are given by

$$I_{syn}(t) = I^{GABA}(t) + I^{NMDA}(t) + I^{AMPA}(t) \tag{3}$$

Synapses were modeled as conductance synapses, in which synaptic currents are dependent on postsynaptic unit membrane potential V_m. Synapses were either excitatory or inhibitory, with inhibitory synapses defined entirely by GABA conductances, and excitatory synapses consisting of both NMDA (N-Methyl-D-aspartic acid) and AMPA (α-Amino-3-hydroxy-5-methylisoxazole-4-propionic acid) conductances in each synapse. These obeyed

$$I^{GABA}(t) = (V(t) - E^{GABA}) \sum_{j=1}^{K_j^{inh}} G_j^{GABA}(t) \tag{4}$$

$$I^{AMPA}(t) = (V(t) - E^{AMPA}) \sum_{j=1}^{K_j^{exc}} G_j^{AMPA}(t) \tag{5}$$

$$I^{NMDA}(t) = (V(t) - E^{NMDA}) s \sum_{j=1}^{K_j^{exc}} G_j^{NMDA}(t) \tag{6}$$

In which E is the reversal potential of the synapse type in superscript, K_j is the in-

degree (number of input connections) of neuron j, and G_j is the respective synaptic conductance. s is a dynamic coefficient which models the effect of the voltage-dependent Mg^{2+} block found in NMDA synapses, according to the formulation of Mel (1993)

$$s = \frac{1.50265}{\left(1 + 0.33\,\mathrm{e}^{-0.0625V(t)}\right)} \tag{7}$$

Individual synaptic spike responses were modelled as a temporal difference-of-exponentials conductance change according to

$$G_j(t) = \sum_{l=1}^{L_t} \left(F_l\, g_l B \left(e^{\frac{t0_l - t}{\tau_{off}}} - e^{\frac{t0_l - t}{\tau_{on}}} \right) \right) \tag{8}$$

where l is the index of presynaptic spikes prior to t and L_t is the total number of these spikes, $t0_l$ is the time of postsynaptic arrival of spike l, determined by the respective spike time of the presynaptic unit plus a transmission delay D, and τ_{on} and τ_{off} are the time constants of the two exponentials, corresponding physiologically to the time constants of neurotransmitter-receptor binding and of the clearance of transmitter from the synaptic cleft, respectively. B is a normalization scalar obeying (9) which normalizes the maximum value of the double exponential term to 1, so that the maximum conductance is directly set by the scalar g_l (Dayan and Abbott 2001):

$$B = \left(\left(\frac{\tau_{on}}{\tau_{off}}\right)^{\tau_{rise}/\tau_{off}} - \left(\frac{\tau_{on}}{\tau_{off}}\right)^{\tau_{rise}/\tau_{on}} \right)^{-1}, \quad \tau_{rise} = \frac{\tau_{off}\,\tau_{on}}{\left(\tau_{off} - \tau_{on}\right)} \tag{9}$$

Excitatory synapses also modelled short-term synaptic depression, according to the formulation of Tsodyks and Markram (1997). According to this formulation, synaptic efficacy is produced by utilization of a finite synaptic resource. Each synapse has a pool of usable synaptic resource, and each presynaptic action potential results in the use and inactivation of a proportion U of the total available synaptic resource R, with this utilized proportion immediately being added to the pool I of the resource in the inactivated state. The resource is gradually recovered from the inactivated to the available state R with first order dynamics:

$$\frac{dR(t)}{dt} = \frac{I(t)}{\tau_{rec}}, \quad I = 1 - R \tag{10}$$

where τ_{rec} is the time constant of recovery. As the magnitude of the postsynaptic conductance change (PSC) produced by a presynaptic spike is scaled by the proportion of resource utilized UR, depletion of R by recent previous action potentials causes a cumulative reduction in PSC magnitude, i.e. synaptic depression. Using this formulation, the current fraction F of the total potential postsynaptic conductance change was calculated with the iterative expression

$$F_{n+1} = F_n(1-U)e^{-\Delta/\tau_{rec}} + U\left(1 - e^{-\Delta/\tau_{rec}}\right) \tag{11}$$

where Δ is the time interval between action potentials n and $n+1$ (Tsodyks and Markram 1997). The computed fraction F scaled the corresponding synaptic conductance change in equation (8).

2.1.2 Network composition

The form of network activity selected as the base of our study was a balanced network of the general type described in van Vreeswijk and Sompolinsky (van Vreeswijk and Sompolinsky 1996), in which network excitation is dynamically balanced by inhibition from a pool of recurrently connected inhibitory neurons. The resulting small-magnitude net mean input to units in such networks results in a moderately sub-threshold mean membrane potential V, while irregularities in each neuron's input cause occasional deviations of V away from the mean to the spike threshold θ, resulting in irregular spiking. While some balanced network models have used irregular network-external input as a source of this irregularity (e.g. Brunel 2000), the system used here was deterministic and the irregularity originated from the pseudo-random sparse connectivity of the network, as external input to the network was held constant unless otherwise indicated.

The network was composed of excitatory and inhibitory units in a 4:1 ratio, a ratio which corresponds to the ratio of pyramidal cells to interneurons in the cortex (Brunel and Wang 2001). Unless otherwise specified, in the simulations described here the number of excitatory neurons $N^E = 1000$ and the number of inhibitory neurons $N^I = 250$. Excitatory and inhibitory units projected exclusively excitatory or inhibitory connections to other units, as well as differing in their neural model

parameters. The parameters used for simulations are displayed in Table 1. Connection probabilities p^{xy} from neuron population y to population x were $p^{II} = p^{IE} = p^{IE} = 0.5$ and $p^{EE} = 0.1$. The higher connectivity to and from inhibitory neurons here comes from the finding of dense inhibitory connectivity in the cortex and the corresponding principle of nonspecific inhibition (Fino and Yuste 2011; Hofer et al. 2011). The connection algorithm set the total number of received connections K^x of each individual unit of population x from population y to $floor\left(p^{xy} N^x\right)$, and these K^{xy} connections were randomly chosen from the N^x possible connections.

Neuron parameter values were chosen mainly according to standard values used in previous works (Renart, Brunel, and Wang 2004; Durstewitz, Seamans, and Sejnowski 2000a), which were intended to reflect typical physiological values. Some neuron parameter values were modified in order to reproduce experimental findings, such as a shorter refractory period in inhibitory units in order to facilitate fast-spiking interneuron-like behavior (Compte, Constantinidis, et al. 2003).

Synaptic coupling strengths were set with the aim of producing a desired level and form of spontaneous activity, using a combination of manual setting of parameters to according to previously-outlined principles of balanced network dynamics (Durstewitz and Seamans 2002; Brunel 2000), and automated parameter search by iterating simulations over ranges of parameter values.

2.1.3 Clustering definitions and algorithms

We used two clustering methods, both of which were based on the principle of defining fixed-sized groups of excitatory units as clusters and then increasing recurrent connectivity within these groups in terms of both connection probability and connection weight. This definition has been used in earlier studies which investigate clustering in simulated networks (e.g. Litwin-Kumar and Doiron 2012; Klinshov et al. 2014), and is also similar to methods used in neural network paradigms which employ bistable assemblies as a form of memory (e.g. Latham and Nirenberg 2004; Brunel and Wang 2001).

The first clustering definition employed was that specified by Litwin-Kumar and Doiron (2012), in which connectivity within each group is determined by a ratio of

connection probabilities between within-group connections and without-group connections. For brevity we will refer to this as the ratio method.

A central aim of our study was to embed multiple overlapping groups in a network in order to observe possible dynamics that result from the existence of multiple nearby quasi-attractors. Preliminary simulations using the ratio method indicated that attractor dynamics were reduced when networks with multiple overlapping clusters were created with this method. Previous work on assembly attractors indicated that intra-cluster recurrent excitation is an important determinant of activity in clustered networks (Durstewitz, Seamans, and Sejnowski 2000a; Latham and Nirenberg 2004), and so we created a definition of connection clustering that included within-cluster recurrent connectivity as as a parameter. In this section we define both methods and sketch an algorithm to produce connection matrices in accordance with the new method.

For both cluster methods, first a number of clusters were defined, each comprising a fixed number $N^{E}_{assembly}$ of excitatory units. For the ratio method, for each unit the overall excitatory connection probability p^{EE} is divided into within and without-group connection probabilities p^{EE}_{in} and p^{EE}_{out}, with a ratio parameter $R^{EE}=\frac{p^{EE}_{in}}{p^{EE}_{out}}$ determining the relative value of these two parameters and hence the degree of clustering. A normalization term is used for each unit so that the overall connection probability p^{EE} remains the same for all units regardless of the number of potential clustered connections. Additionally, the synaptic strength g was set to a higher value for within-group connections.

For the second clustering method, clustering was defined not by the ratio of within-group to without-group connection probabilities in individual units, but by a parameter p^{EE}_{in} which set the overall connection probability within individual clusters. As the number of units in each assembly was also fixed by the parameter $N^{E}_{assembly}$, clusters in this method had a fixed total recurrent connection weight and number of connections. Because this method produced networks containing assemblies with uniform total recurrent excitation, we termed this method the “equal assemblies” method, or EQASS method for short. As with the ratio clustering model,

the fixed overall connection probability p^{EE} remained, while the without-group connection probability p_{out}^{EE} and therefore the ratio of connection probabilities were left as free parameters. As with the ratio model, the synaptic connection strength was higher for within-group connections, but the total input synaptic connection strength was fixed per-neuron, with without-group input connection strength left as a free parameter.

In order to not exceed the fixed connection probability and total connection weight of each unit and group, the algorithm which created network connections according to the EQASS definition functioned according to a general principle of first randomly generating within-cluster connections for clustered units one by one until the required within-cluster connectivity p_{in}^{EE} was achieved, checking total per-unit and per-cluster connection numbers after each connection was created, and blocking units from receiving any further within-cluster connections if they reached the maximum number of incoming cluster connections possible. The total number of within-cluster connections per-unit was limited so as not to exceed the fixed uniform connection probability and cumulative connection weight. After all within-cluster connections were generated, non-cluster connections were then generated for each unit with the required number and connection strength for the unit to reach the pre-defined connection probability and total weight values.

Specifically, assuming a total set of N^E excitatory units, a parameter $N_{assembly}^E$ was set, determining the number of units in a clustered assembly. A parameter g^{EE} was set, determining the average weight of an excitatory synapse, and also a parameter $g_{in}^{EE} > g^{EE}$, which determined the weight of within-assembly synapses. Similarly, the parameter p^{EE} set the overall and p_{in}^{EE} the within-assembly connection probabilities. From these values the maximum numbers of per-assembly within-assembly connections, $K_{assembly}^{EE} = \mathrm{floor}\left(p^{EE}\left(N_{assembly}^E\right)^2\right)$ and per-neuron assembly connections $K_{in,max}^{EE} = \mathrm{floor}\left(\frac{\left(g^{EE} p^{EE} N^E\right)}{g_{\mathrm{in}}^{EE}}\right)$ were calculated, the latter being the maximum number of assembly connections that a neuron can receive before exceeding the fixed total presynaptic weight $g^{EE} p^{EE} N^E$.

All assemblies were defined by choosing $N^{E}_{assembly}$ units for each assemblies, either as disjoint sets or randomly (i.e. without regard as to whether units were members of multiple assemblies). After assemblies were defined, connections were created by randomly selecting from all possible intra-assembly connections. After each connection was created, the total number of connections to each unit and within each assembly checked, and all further potential connections to the unit or assembly were removed from the pool of available connections if the number of connections was equal to $K^{EE}_{in,max}$ or $K^{EE}_{assembly}$, respectively. This process continued until either all assemblies had the requisite number of within-assembly connections $K^{EE}_{assembly}$, or until no further possible within-assembly connections were available, in which case the algorithm terminated and returned an error. As will be shown in the results, the likelihood of this error condition increased with increasing numbers of assemblies and higher values of p^{EE}_{in}.

After all within-assembly connections were defined, without-assembly connections were created for each unit by randomly selecting from available possible non-assembly connections, until each unit had $K^{EE} = p^{EE} N^{E}$ connections in total. The connection weight of these connections g_{out} was set so that the total input connection weight was equal to $g^{EE} p^{EE} N^{E}$, i.e. so that the mean synaptic weight was g^{EE}.

For both the ratio and equal assembly clustering methods, all other connection matrices (i.e. excitatory-inhibitory and inhibitory-inhibitory) were chosen randomly so that each unit received $K^{xy} = p^{xy} N^{x}$ connections, each of equal strength g^{xy}.

2.1.4 Iteration routine

In order to explore the effect of clustering, simulations were systematically run with progressively increasing levels of key clustering parameters. Networks were generated with either 10, 20 or 30 embedded assemblies $N^{assembly}$, each comprising a proportion 0.1 of $N^{E} = 1000$ total excitatory neurons. In the networks with 10 embedded assemblies, networks were created with both overlapping and non-overlapping assemblies. In networks with 20 or 30 assemblies, assemblies were

necessarily overlapping as $N^E < 0.1 N^E N^{assembly}$. In networks with overlapping assemblies, assembly members were selected randomly.

Simulations were run with progressively greater values of the respective clustering parameters R^{EE} and p_{in}^{EE}. In the case of the networks created with the ratio clustering method, R^{EE} was varied from 1 to 3.9 in increments of 0.145, and in the case of the networks created with the equal assemblies method, p_{in}^{EE} was varied from 0.1 to 0.3 in increments of 0.01. In both of these cases, the lowest level of the parameter range corresponded to no additional clustering above that produced by random connectivity. The ranges and steps were chosen to produce equivalent levels of within-cluster recurrent connectivity for both methods in the case of non-overlapping clusters.

For each combination of parameter levels, 24 different and independent network instances were created using different random number generator seeds. Simulations had a duration of 32000ms (32 s), of which the first 2000ms was discarded in subsequent analysis to remove initial transients.

2.1.5 Computational framework

Simulations were performed using both an in-house system written in C and MATLAB (The MathWorks Inc., Natick, MA, 2000)., and the NEST spiking neural network simulator with Python and C++ (Gewaltig and Diesmann 2007). Analysis was performed in MATLAB and Python. Pseudorandom numbers (referred to as random numbers in this thesis for convenience) were generated by the Mersenne Twister algorithm implemented in both MATLAB and Python.

2.2 Analysis methods and results

The output of the network was analyzed in order to compare it in a general sense to in-vivo data, and to characterize the dynamics introduced by the clustering, under the assumption that clustering would cause cluster-specific variations in firing rate. Analysis was primarily performed on the spiking output, where individual spikes were represented as single timepoints. We specify here the measures used.

2.2.1 Methods

The simplest descriptive measure of spike firing is the firing rate of a neuron or group of neurons. Following Dayan and Abbott (2001, 9), this can be formally defined by representing spikes as Dirac δ functions, and the spiking response of a neuron at a given time *t* as $\rho(t)$. The firing rate *r* can then be defined as

$$r(t)=\frac{1}{\Delta t}\int_{t}^{t+\Delta t} d\tau\,\langle\rho(\tau)\rangle \tag{12}$$

Another measurement used to characterize spiking output was the interspike interval (ISI), the time difference between two consecutive spiking events

$$ISI_j=t_j-t_{j-1} \tag{13}$$

where j is the spike index and t_j is the time of the *j*th spike.

In addition to the mean and standard deviation of the ISI, the coefficient of variation C_V of the ISI distribution, the ratio of the standard deviation to the mean, is valuable as measure of spiking irregularity.

$$C_V=\frac{\sigma_{ISI}}{\langle ISI\rangle} \tag{14}$$

Here, a C_V of 0 corresponds to completely regular firing. The Poisson distribution, in which ISIs have an exponential distribution, has a C_V of 1.

The Fano factor F was also computed on spike trains. This is a measure of the variability of the spike rate over time, computed as the ratio of the variance to the mean of binned spike counts N, where bin sizes are equal:

$$F=\frac{\sigma_N^2}{\langle N\rangle} \tag{15}$$

2.2.1.1 Spike power spectrum analysis

The power spectrum of the population spike firing rate was calculated using Welch's

method (Welch 1967), by first binning the spike output of all neurons then using the resulting time series as input to the *welch* function in the SciPy signal processing toolbox. Briefly, Welch's method creates an estimate of the power spectral density by dividing the data into overlapping segments, computing a periodogram for each segment using a fast Fourier transform, and averaging the periodograms. Compared to the use of a single periodogram calculated on the whole time series, this method has the advantage of producing a less noisy estimate, at the expense of a loss of frequency resolution at the lower end.

Low frequency periodic activity in the delta range of 0 – 4 Hz is of particular physiological interest, and has been associated with disease states (Broyd et al. 2009; Ford et al. 2002) and different forms of neural coding and cognition (Lakatos et al. 2008). Accordingly, we used a simple measure of delta band power by taking the root mean square (the square root of the mean of the square of the values) of the 0 – 4 Hz range of the spiking power spectrum calculated with Welch's method.

2.2.1.2 Spike train autocorrelation function

Related to the power spectrum is the autocorrelation function of spiking activity when calculated on a population of neurons. This measures the distribution of times between all pairs of spikes from all sources in a population. This was calculated by binning the spikes, subtracting the mean value from this binned spike train, and calculating its autocorrelation, normalized by the variance of the spike rate r

$$Autocorrelation_{\rho\rho}(\tau) = \frac{\langle (\rho(t) - \langle r \rangle)(\rho(t+\tau) - \langle r \rangle) \rangle}{\sigma_r^2} \tag{16}$$

where $\rho(t)$ in this case is the binned spike response at time t.

2.2.1.3 Assembly-specific measures

We define also basic measures of the effect of the clustering definitions, on both the connection matrix and activity during simulations, in order to verify and quantify the effect of the clustering methods, particularly on assembly-specific high-activity states and winner-takes-all dynamics.

First we measure the assembly-specific recurrent connectivity, the summed

projection of a group to itself. By representing a group as a binary vector g of size N_E in which entries are 1 when the corresponding unit is a group member and 0 otherwise, the summed recurrent group projection rec can be calculated as $rec = g^T Wg$ where W is the excitatory weight matrix.

Following Russo et al. (2008) and Hopfield (1982), we define an activation measure that quantifies the state of the system in terms of the attractor states which are embedded in it. Following models of distributed memory storage such as that of Hopfield (1982), in which retrieval success of a specific memory is measured as the summed total of binary units which are in the memory state, we set a firing rate threshold of $\theta = 10Hz$, and define neurons firing above this threshold as being active. The activation state A of assembly i is then defined as the proportion of active neurons within its members.

$$A_i = \frac{1}{N} \sum_{0}^{N} H(r_i - \theta) \tag{17}$$

where H is the Heaviside step function. This measure was calculated on the spiking output of the network binned with a bin size of 100ms.

Unlike the memory retrieval definitions of Russo et al. (2008) and Hopfield (1982), the definition used here does not include both activity and inactivity of units across the whole system in its state retrieval definitions. Instead, only activity within specific assemblies is counted, similarly to spiking neuron memory state networks (e.g. Brunel and Wang 2001; Durstewitz, Seamans, and Sejnowski 2000a). In principle, this allows multiple assembly states to be active simultaneously.

As a measure of winner-takes-all behavior, and also the degree to which the system state can be described in terms of a single assembly, we also define two further simple measures. First, we measure the ratio of the mean firing rate of neurons within the most activated assembly in the network to the mean firing rate of the excitatory population as a whole. Second, we measure the “separation” of the activation state in terms of the measure A. The separation measure is defined as the difference in the activation state A between the most activated assembly and the second most activated. Finally, we calculate the “transition rate”, the rate at which the identity of most activated assembly changes.

2.2.2 Output of the network

In this section we systematically describe the output of the network qualitatively and in terms of the measures described above. First we report the output of the unclustered network, which displays typical balanced network activity. Then we report separately the output of clustered networks with non-overlapping and overlapping clusters. In doing so we highlight the emergence of transient high-activity states with increasing clustering, and contrast the sensitivity of the output of the two clustering mechanisms to overlaps between embedded clusters. Output regarding the adaptation parameter w and delta band oscillations are reported in separate sections.

2.2.2.1 Without clustering

Parameters were determined at which the network spontaneously produced irregular low-rate spiking activity (Figures 2.1 and 2.2). Averaging over 24 independently generated networks, the mean (standard deviation) steady state firing rate of the excitatory and inhibitory populations was 1.62(±0.26) Hz and 6.11(±0.55) Hz respectively. The mean C_V and Fano Factor of the excitatory and inhibitory populations were both close to 1, indicating approximately Poisson spiking with a stationary firing rate. In addition to the low rate asynchronous spiking of individual neurons, both the population averaged autocovariance and power spectrum indicated high frequency population oscillations with peaks at approximately 50Hz and 200Hz.

When interpreted through the analytical framework of Brunel and Wang (2003), which explored mechanisms of population oscillations in spiking simulations with parameters similar to those used here, these oscillations likely result from inhibitory feedback loops, with the 50 Hz oscillation generated by the excitatory – inhibitory – excitatory loop, and the faster 200 Hz oscillation coming from the more direct inhibitory – inhibitory loop. One notable difference between the two oscillatory peaks produced here and the results explored in Brunel and Wang (2003) is that Brunel and Wang described only simulations with a single oscillatory peak, with differences in the intrinsic frequencies of the excitatory-inhibitory and inhibitory-inhibitory loops resulting in a frequency intermediate to the two. It is possible that the dual peaks

observed here were possible because the two intrinsic frequencies are well separated, and because the higher frequency is approximately an integer multiple of the lower. Confirming this, test simulations run with a slight alteration to one of the synaptic latencies (mean excitatory latency D^{ex} = 5ms rather than 3ms) produced output with approximately the same excitatory firing rate but two different population oscillation peaks. Example output is shown in Figure 2.3B, which shows population frequency peaks at approximately 30Hz and 150Hz.

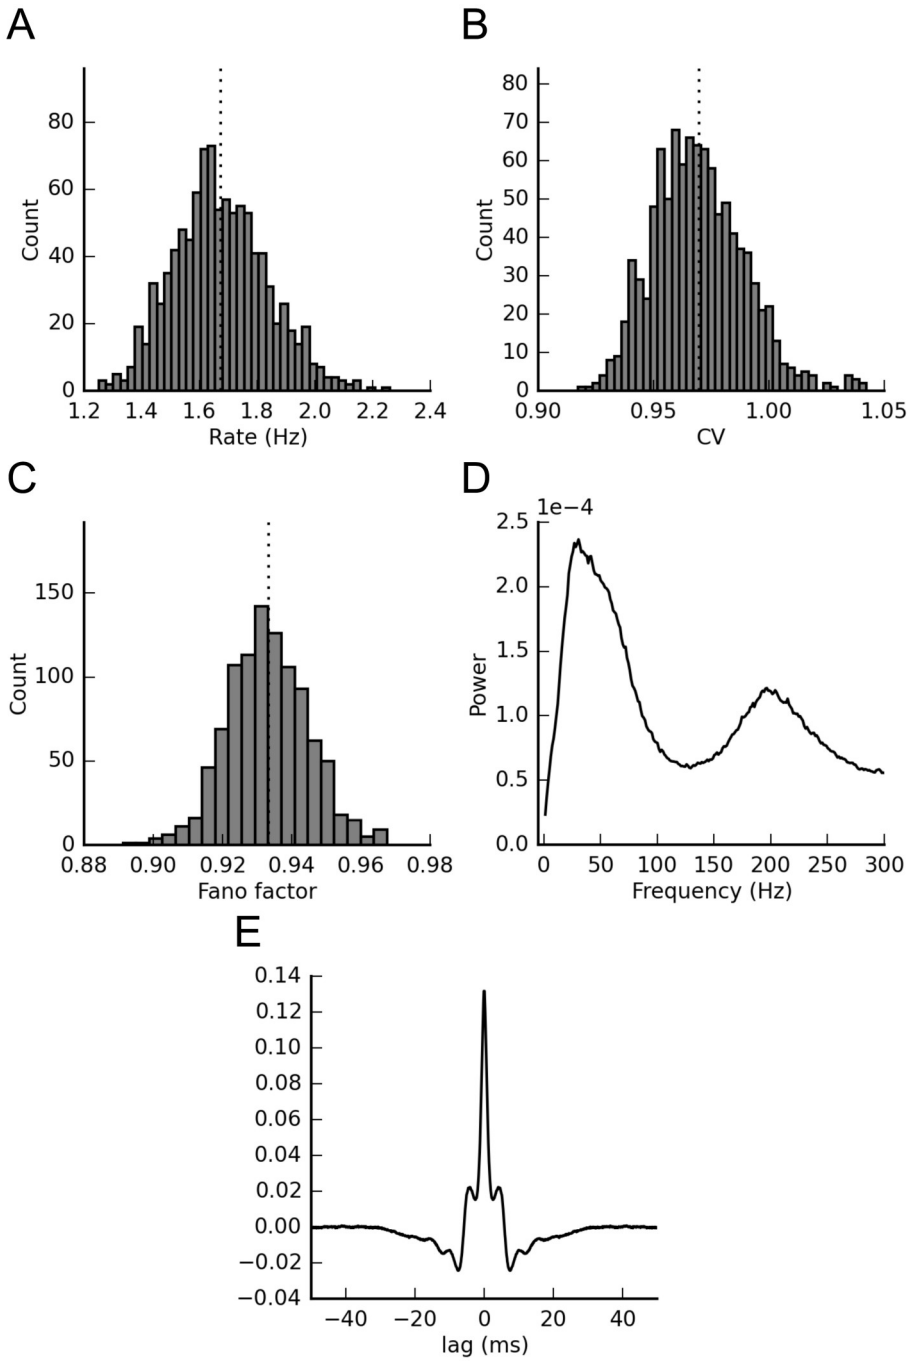

Figure 2.1: Spiking statistics of excitatory neurons in an unclustered balanced network.

Dotted vertical lines indicate mean values. (A) Histogram of neuron firing rates, mean firing rate 1.67 (0.16) Hz. (B) Histogram of coefficients of variation of inter-spike intervals, mean (SD) C_V 0.97 (0.02). (C) Histogram of Fano factors of spike counts, computed with a 100ms window, mean (SD) Fano factor 0.93 (0.01). (D) Power spectrum of spiking across all excitatory spikes, showing peaks at approximately 50 and 200 Hz. (E) Population spiking autocovariance.

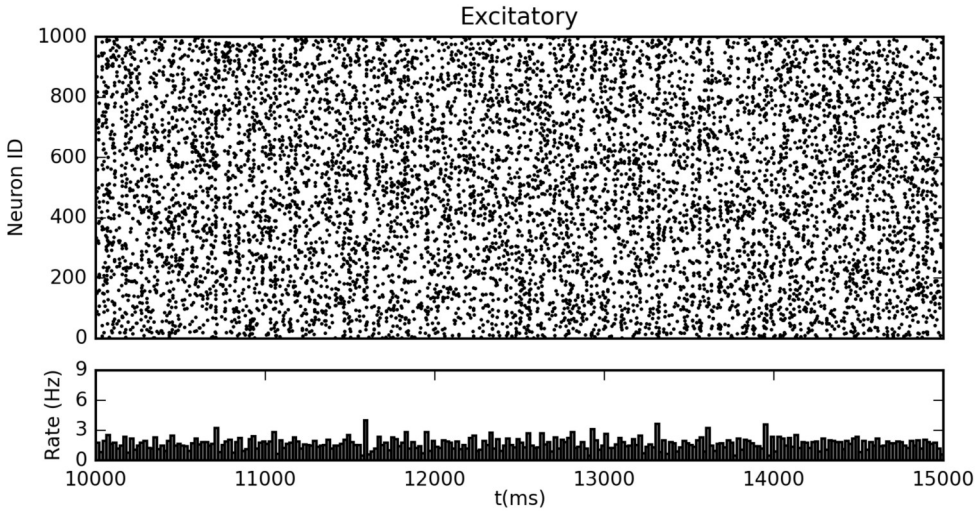

Figure 2.2: Example spiking output of an unclustered balanced network over a 5 second period.
(Top) Raster of spike output of 1000 excitatory neurons. (Bottom) Corresponding mean population firing rate.

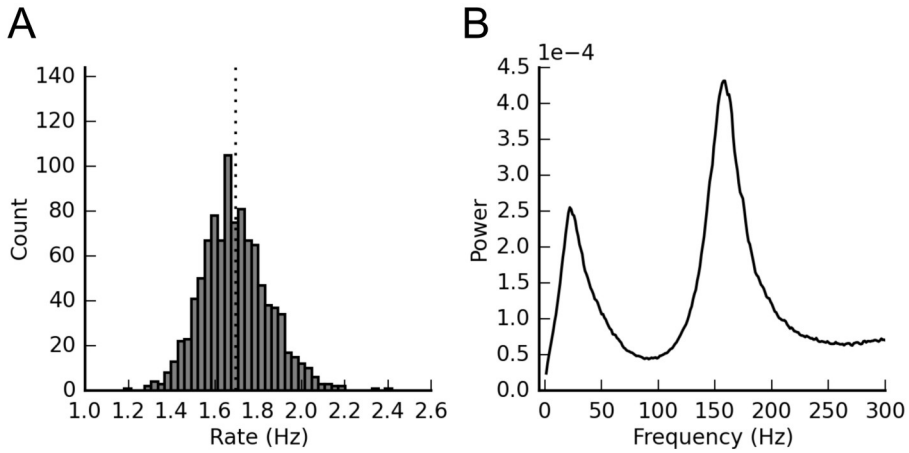

Figure 2.3: Statistics of spiking activity of excitatory neurons in a simulated network with an increased excitatory synaptic latency of 5ms rather than 3ms.
(A) Histogram of neuron firing rates, showing similar mean and distribution as Figure 2A. (B) Power spectrum of spiking across all excitatory spikes, showing peaks at lower values of approximately 30 and 150 Hz, and increased amplitude of the higher-frequency peak.

2.2.2.2 Non-overlapping assemblies

When creating non-overlapping assemblies, both the ratio and equal assemblies methods produced qualitatively similar effects on clustering strength and network activity Figure 2.4, Figure 2.5. The clustering parameter of each method had a mainly linear relationship with total within-assembly connections $K^{EE}_{assembly}$ and corresponding within-assembly weights, with a slight deviation in the case of the ratio method (Figure 2.4A) due to rounding effects arising from the constraint that the number of inward assembly connections per unit K^{EE}_{in}, which was determined internally by $K^{EE}_{in} = \mathrm{round}\left(p^{EE}_{in} N^{E}\right)$, was restricted to integer values.

The relationship between clustering and the quantitative measures of spiking activity was strongly non-linear, particularly in regards to a form of soft phase transition that occurred at values of approximately $R^{EE} = 1.9$ and $p^{EE}_{in} = 0.165$. Around these values the individual assemblies started to display bimodality in their firing rates, in that the increased within-assembly connectivity supported transient periods of higher firing rates within individual assemblies. The mean rate of firing over the whole network increased roughly linearly with the clustering parameter after this phase transition point, up to the maximum clustering values tested (Figure 2.4B).

These assembly-specific higher-firing-rate states, or “activations” as we will refer to them, occured initially only in one assembly at any one time, indicating competition via mutual inhibition. The disparity in firing rates between the activated assembly and the rest of the network is reflected in the firing rate ratio and activation separation plots (Figure 2.4C and Figure 2.5A), which show a steep increase in these measures after the phase transition point. Visual inspection of plots of spiking activity and the assembly activation output confirm this interpretation (Figure 2.6).

With further increases in clustering, the firing rate ratio and separation measures do not increase monotonically but rather decline after reaching a peak (Figure 2.4C and Figure 2.5A). Inspection of simulation output confirms that this occurs as a result of increasingly frequent activations in multiple assemblies simultaneously. At the highest tested level of the clustering parameters $R^{EE} = 3.9$ and $p^{EE}_{in} = 3.0$ the

separation measure returns to a mean value of approximately 0, as at least two assemblies are equally active at most times for the duration of the simulation (Figure 2.5A).

The transition rate measure (Figure 2.5B) initially shows a rapid increase in transitions with the onset of bistability, up to a mean rate of slightly under 1 per second, before falling back to near 0. Inspection of simulation output indicates that this corresponds to the initial emergence of relatively unstable activated states at moderate levels of clustering, which then become more stable and therefore longer in duration as recurrent excitation increases with increased clustering.

The behavior of the networks created by the two different routines differs slightly in respect to the transition rate at higher levels of clustering. In the case of the equal assemblies routine, the emergence of simultaneous activated assemblies is accompanied by an increase in the transition rate, as activation switches between different assemblies, while in the case of the ratio routine the transition rate remains low, as multiple assemblies remain largely stable in the activated state. These differences in dynamics can be attributed to the two differences between networks created by the two clustering methods In the case of non-overlapping assemblies. First, total input connection weights are kept equal for all units by the EQASS method but allowed to be higher in the ratio method. Second, all assembly units have a uniform number of within-assembly connections in ratio networks, but this is allowed to vary in EQASS networks. Assembly units in non-overlapping ratio networks therefore have greater and more uniform levels of total synaptic input than those in EQASS networks, and one or both of these factors will have caused the differences in behavior.

The effect of increased clustering on the Fano factor appears qualitatively similar to the effect on the transition rate (Figure 2.5C), rising from its baseline value of approximately 1 with the onset of transient activations, then falling as activated states become stable. As the Fano factor measures variations in neuron firing rates over 100ms+ periods, there is a direct association between the Fano factor and the tendency for assemblies to transition between low- and high-firing rate states.

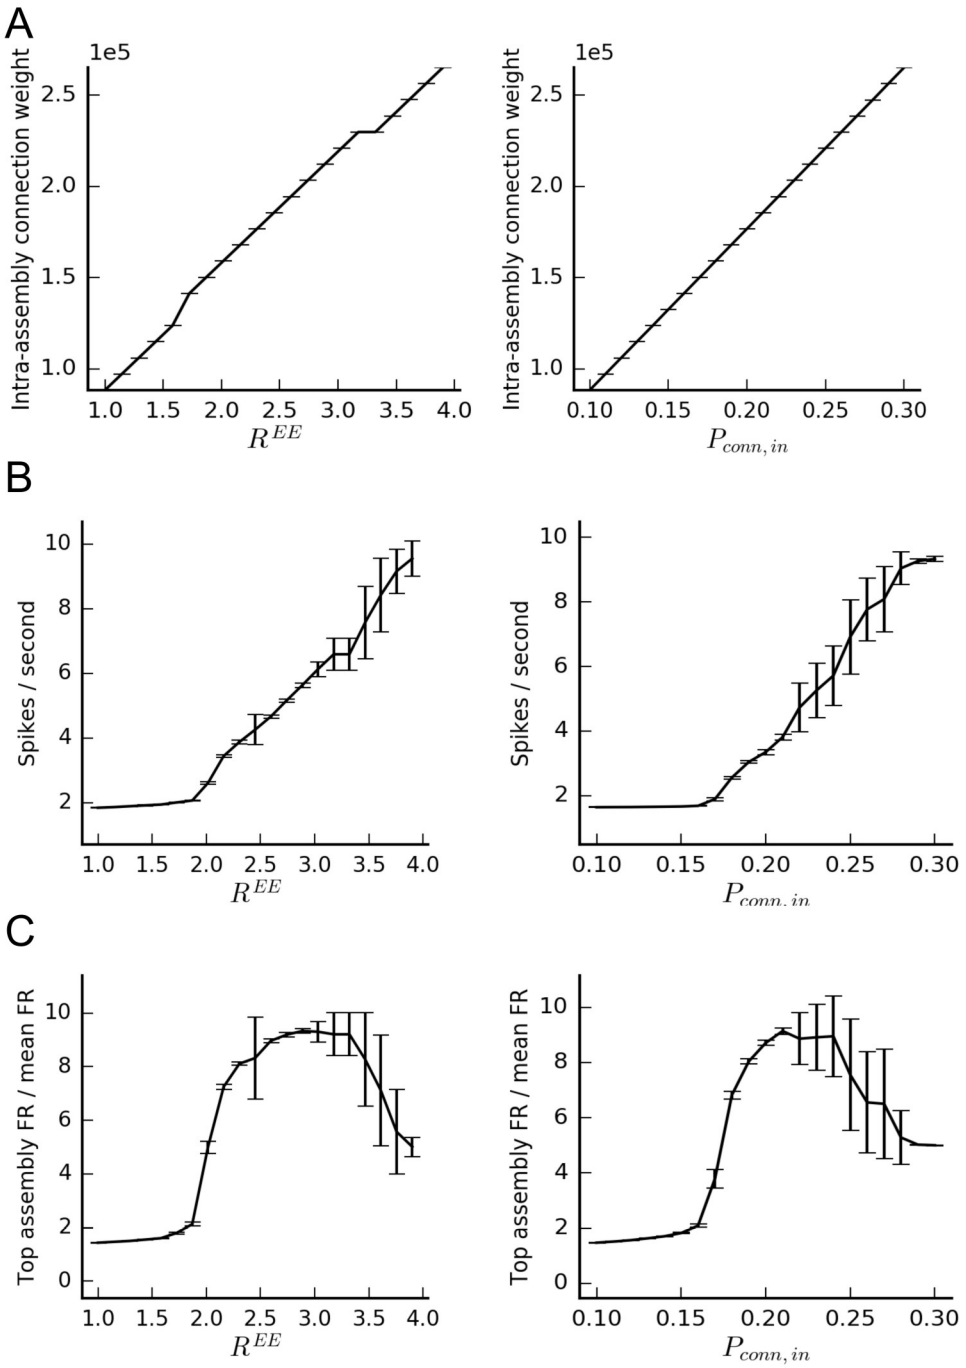

Figure 2.4: Effects of progressively increased clustering within non-overlapping assemblies, part 1.

Results from networks created with the ratio (left) and EQASS (right) definitions. Mean (SD) values over 24 randomly generated networks. (A) Intra-assembly connection weight. (B) Mean excitatory neuron firing rate. (C) Ratio of mean firing rates between the most activated assembly and all excitatory neurons.

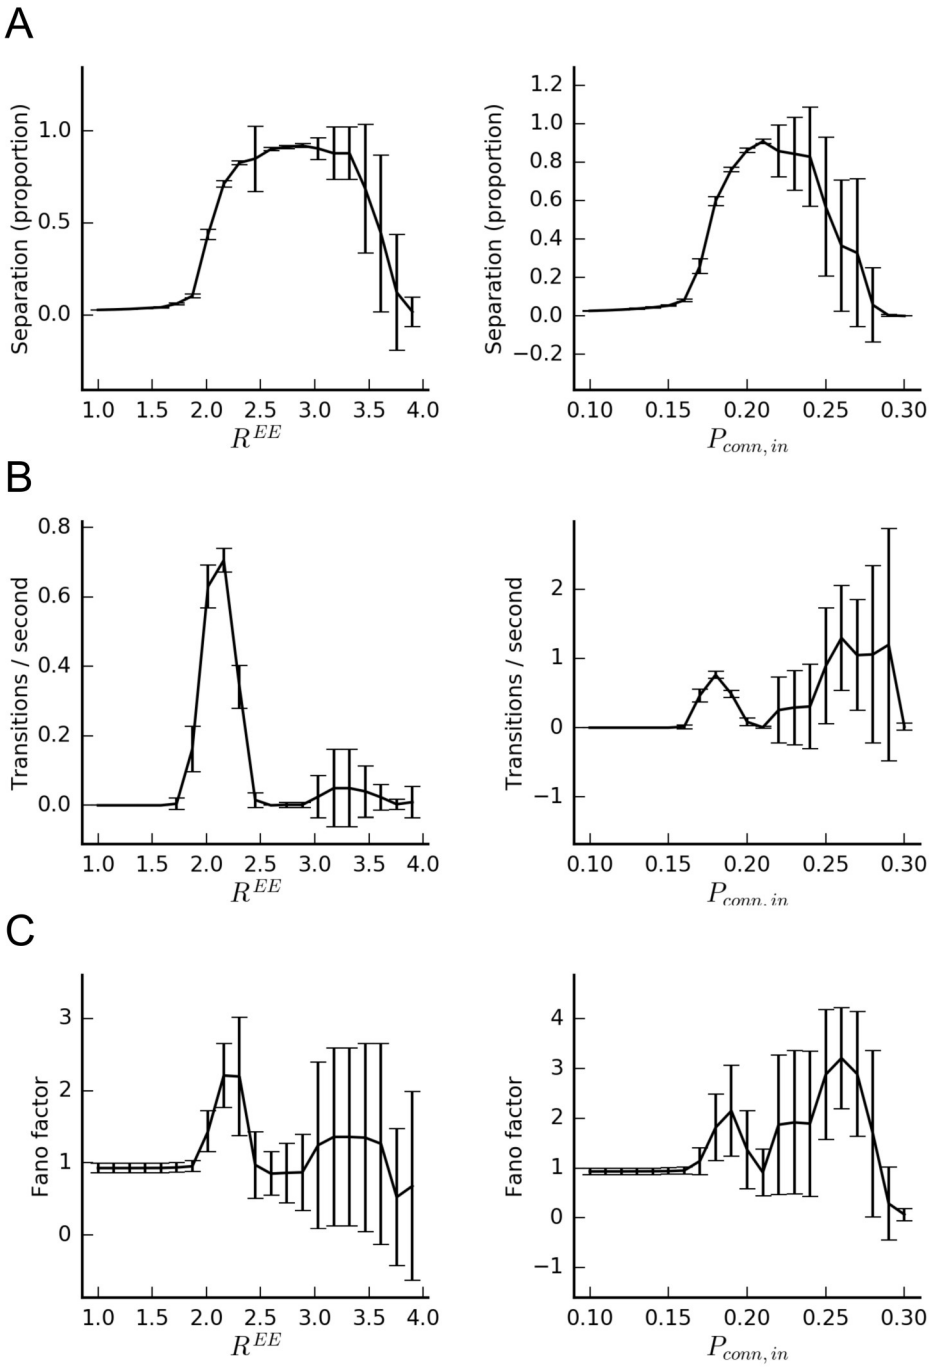

Figure 2.5: Effects of progressively increased clustering within non-overlapping assemblies, part 2.

Results from networks created with the ratio (left) and EQASS (right) definitions, part 2. Mean (SD) values over 24 randomly generated networks. (A) Separation between first and second most activated assemblies in terms of activations measure A. (B) Rate of transition between the identity of the most-activated assembly. (C) Fano factor of excitatory neuron firing rate computed over 100ms windows.

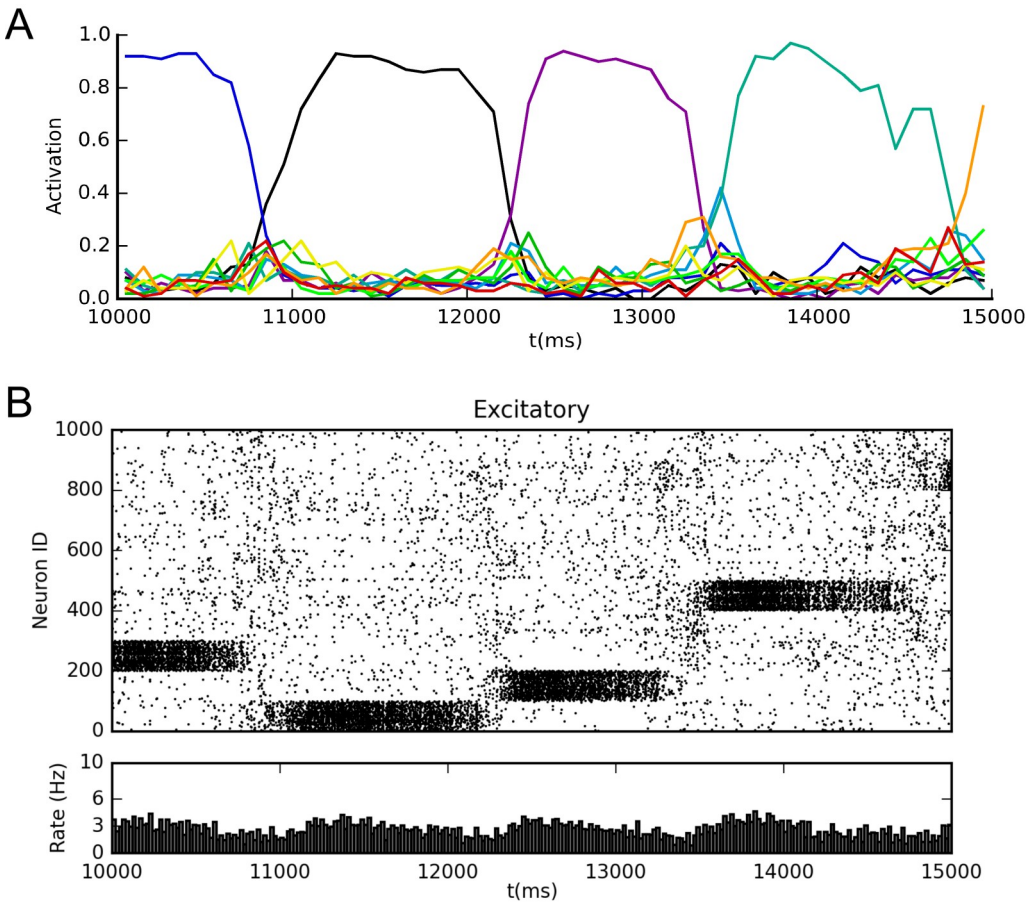

Figure 2.6: Example output of a balanced network with 10 non-overlapping assemblies created with the equal assemblies method.

Example network created with $p_{in}^{EE} = 1.8$ *, plots correspond to the same simulation period. (A) Assembly activation measure* A*, with colored lines corresponding to activation of different assemblies. (B)(Top) Raster of spike output of 1000 excitatory neurons. (Bottom) Mean population firing rate.*

2.2.2.3 Overlapping assemblies

Clustered connection matrices were defined according to both clustering definitions with 10, 20 and 30 overlapping assemblies and 21 equally-spaced levels of the clustering parameters R^{EE} and P_{in}^{EE}. The equal assemblies algorithm failed to complete creation of connection sets with 30 assemblies and a clustering level of greater than $P_{in}^{EE}=0.23$, as it was unable to create enough connections to reach the specified total recurrent connectivity for each assembly without exceeding the connection limits for individual units.

In the networks with overlapping assemblies created with the ratio assembly definition the clustering parameter R^{EE} had a non-linear relationship to total intra-assembly connection weight rec, which was also dependent on the number of assemblies in the network $N^{assembly}$ (Figure 2.7A). The increase in rec was sub-linear, and slower with greater $N^{assembly}$. Networks created with the equal assemblies definition had a linear relationship with P_{in}^{EE} which was independent of $N^{assembly}$ and whether or not assemblies overlapped. As the equal assemblies definition was created specifically to maintain this independence, this result simply confirmed that the algorithm had functioned correctly.

For the overlapping assembly networks created with the equal assemblies routine, the effect of increased clustering on the mean firing rate and the top assembly / mean firing rate ratio was similar to the effect of increased clustering on networks with non-overlapping assemblies, and was approximately equal for all 3 levels of $N^{assembly}$ (Figure 2.7). The mean separation measure (Figure 2.8A) appears to marginally decrease at higher levels of P_{in}^{EE} with increasing $N^{assembly}$, while the mean transition rate and Fano factor are consistently slightly higher with increasing $N^{assembly}$. This will be discussed in more depth later.

The response of the ratio networks in terms of the various activity measures was strongly dependent on the number of embedded assemblies $N^{assembly}$. At the $N^{assembly}=10$ level, the relationship between the measures and clustering parameter

R^{EE} was qualitatively similar to the corresponding relationship in the EQASS networks, although slightly weaker. At the higher levels $N^{assembly}=20$ and $N^{assembly}=30$ the relationship between clustering and the measures is considerably weaker in most cases, as might be predicted from the lower intra-assembly connectivity rec. The relationship also differed qualitatively in some respects, which can be seen most clearly in the results at the $N^{assembly}=30$ level. Salient examples of these qualitative differences include the considerable increase in overall network firing rate but relatively small increases in the firing rate ratio and separation measures (Figure 2.7B, C and Figure 2.8A), and sharp large increase in the transition rate measure with a relatively small increase in the mean Fano factor (Figure 2.8B, C).

The primary reason for the loss of assembly activation dynamics here was the decrease in recurrent connectivity with increasing numbers of embedded clusters (Figure 2.7A). The increase in the total number of clusters in the network causes an increase in the mean number of clusters of which each unit is a member and therefore an increase in the mean number of assembly co-members. As the ratio of within-cluster to without-cluster connection probabilities is determined by $R^{EE}=\frac{p_{in}^{EE}}{p_{out}^{EE}}$ with the cumulative probability of any connection held constant at p^{EE} by a normalization factor, the inwards recurrent connection probability p_{in}^{EE} is determined individually for each unit as a function of the number of assembly co-members according to $p_{in}^{EE}=\frac{P^{EE}R^{EE}N^{E}}{N_{in}^{E}\left(R^{EE}-1\right)+N^{E}}$, where N_{in}^{E} is the number of assembly co-members for the unit being considered. The increases in N_{in}^{E} caused by the increase in overlaps therefore cause a lower within-cluster connection probability p_{in}^{EE} and correspondingly a decrease in recurrent excitation.

The higher mean firing rates, despite reduced assembly-specific activations, were partly a consequence of the effect of increased general excitatory transmission with increases in R^{EE} and $N^{assembly}$, both of which increase the proportion of assembly connections in the network, which have a fixed higher strength than non-assembly

connections. As total connection weight is not normalized in the ratio definition, this leads to higher overall excitatory-excitatory connection strengths, and higher mean firing rates. The increase in total excitatory-excitatory connection strength as a function of $N^{assembly}$ and R^{EE} is shown in Figure 2.9A. The higher transition rates in absence of assembly activation dynamics (Figure 2.8B) were largely caused by these higher firing rates and an artifact of the method used to determine assembly activations, as the higher mean firing rates caused the firing rates of all neurons to approach the activation threshold of 10Hz (2.2.1.3) and the incorrect detection of assembly-specific activations. Figures 2.9B and C illustrate this behavior.

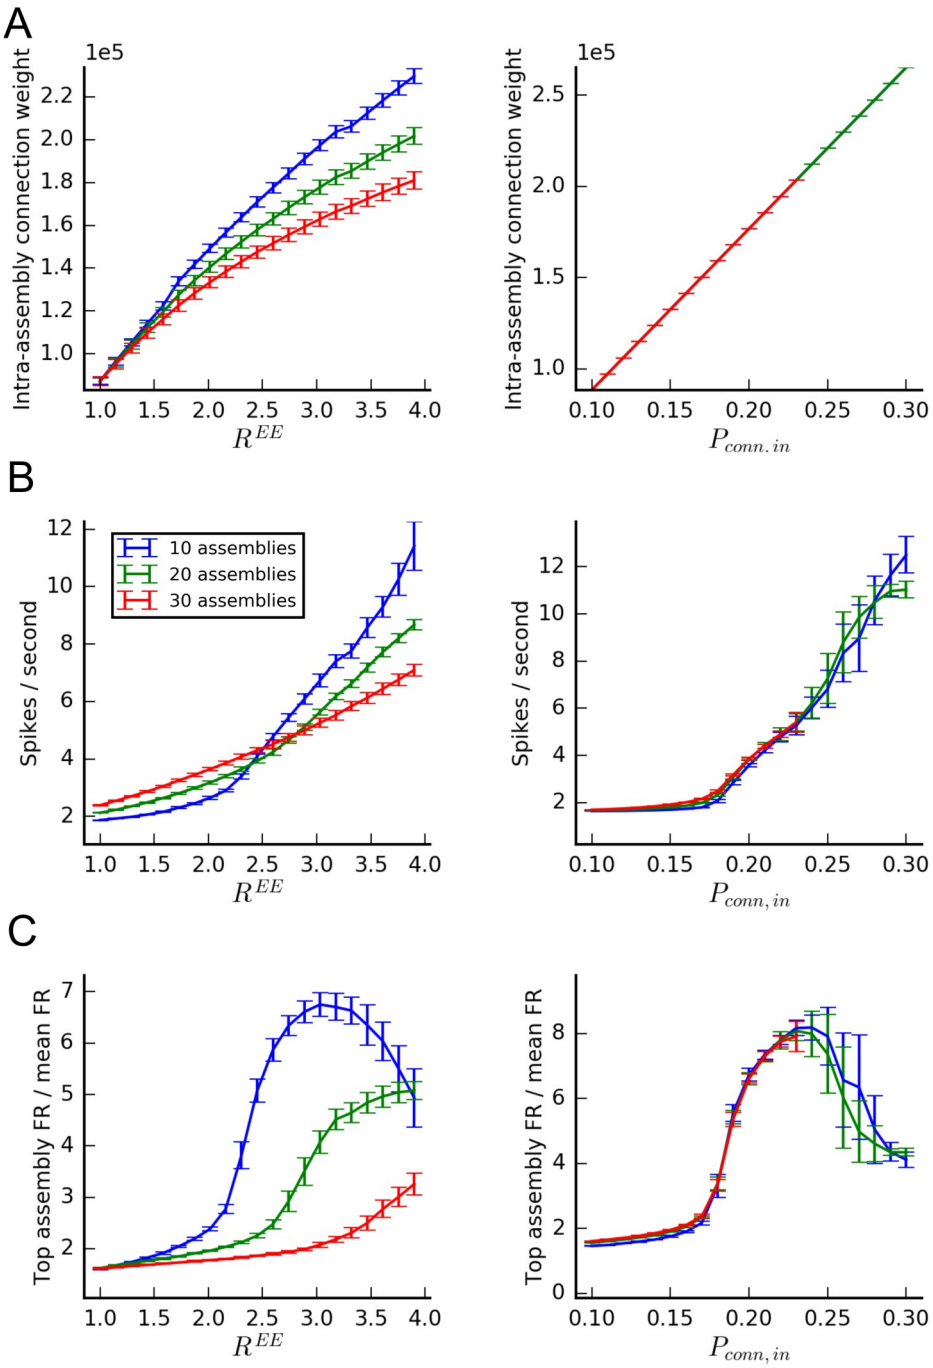

Figure 2.7: Effects of progressively increased clustering in networks within randomly overlapping assemblies, part 1.
Results from networks created with the ratio (left) and EQASS (right) definitions. Mean (SD) values over 24 randomly generated networks. Blue: 10 assemblies, green: 20 assemblies, red: 30 assemblies. (A) Intra-assembly connection weight. (B) Mean excitatory neuron firing rate. (C) Ratio of mean firing rates between the most activated assembly and all excitatory neurons.

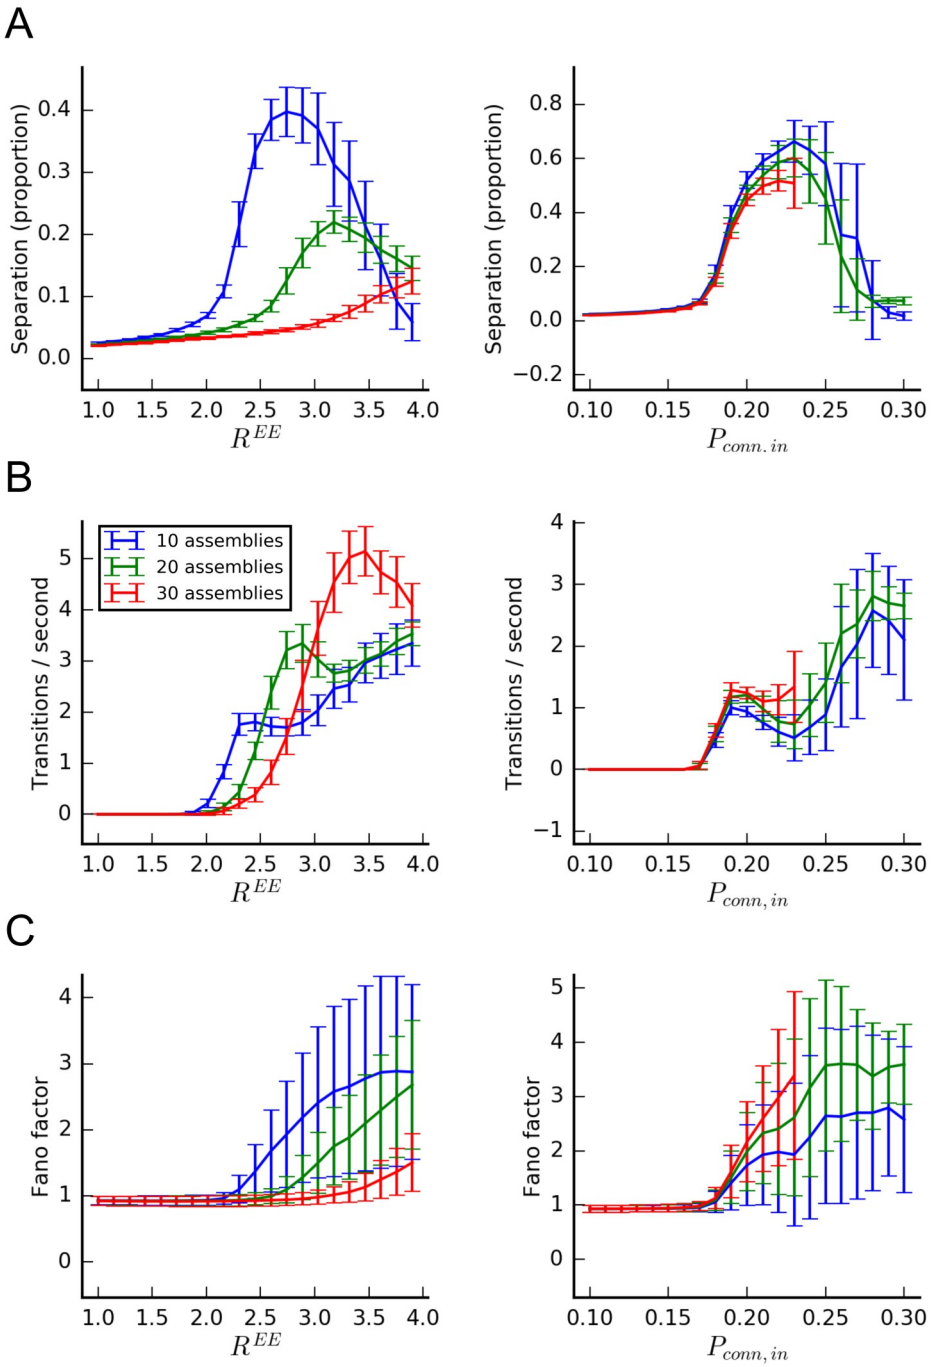

Figure 2.8: Effects of progressively increased clustering in networks within randomly overlapping assemblies, part 2.
Results from networks created with the ratio (left) and EQASS (right) definitions. Mean (SD) values over 24 randomly generated networks. Blue: 10 assemblies, green: 20 assemblies, red: 30 assemblies. (A) Separation between first and second most activated assemblies in terms of activations measure A. (B) Rate of transition between the identity of the most-activated assembly. (C) Fano factor of excitatory neuron firing rate computed over 100ms windows.

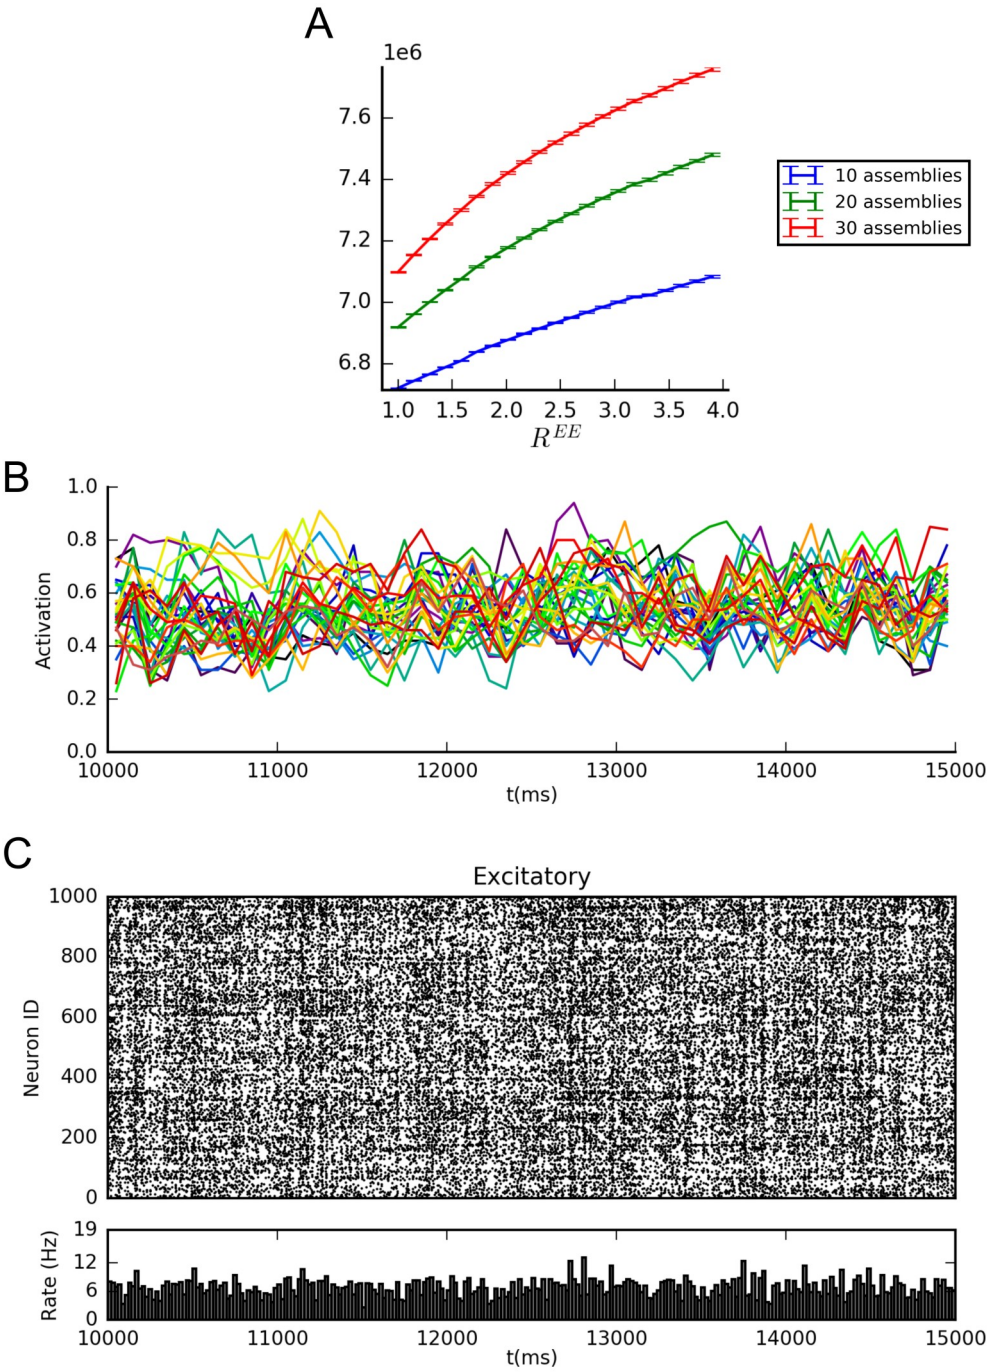

Figure 2.9: Plots related to the effects of embedding multiple overlapping assemblies with the ratio routine.

(A) Sum of all excitatory-excitatory connection weights as relation of clustering ratio and number of assemblies. Blue: 10 assemblies, green: 20 assemblies, red: 30 assemblies.. (B) and (C) example output from a highly clustered network with $N^{assembly}=30$ *and* $R^{EE}=3.465$ *. (B) Assembly activation measure* A *for all assemblies in network. (C) (Top) Raster plot of excitatory spiking output of excitatory neurons. (Bottom) Mean excitatory firing rate.*

2.2.3 Switching activation dynamics and phase transitions in networks with overlapping clusters

We now take a more detailed look at spiking dynamics of clustered networks, highlighting 3 qualitatively different states that are found with different levels of clustering. For brevity, we restrict our analysis to networks with overlapping assemblies created by the equal assemblies algorithm.

With fixed assembly size but progressively increasing degree of clustering within assemblies, spontaneous network activity transitions through 3 qualitatively different stages (Figure 2.10). At low levels of clustering, activity is essentially identical to that of unclustered networks, with approximately Poisson firing and fast oscillations at the population level, as shown in figures 2.1 and 2.2. Above a threshold level of clustering, firing rates within individual assemblies become bimodal, showing transient periods of activation which become higher-rate and longer in duration with increased clustering, before reaching a peak (Figure 2.8). In this state, up-states occur only in one assembly at a time, with brief periods of overlap, indicating mutual inhibition through the inhibitory population (Figure 2.11). The mean firing rate, C_V and Fano factor of excitatory spiking is considerably higher than the baseline state, indicating “bursting” spiking at the neuron level (Figure 2.12). The power spectrum of spiking shows a peak in the <2Hz range, in addition to the peaks at approximately 50Hz and 200Hz that are seen in the baseline state. This is a result of sequential switching of activation between assemblies, which the transition rate metric shows as occurring at an approximately <2Hz rate at moderate clustering levels (Figure 2.8B). This subject is covered further in section 2.2.5.

As within-cluster connectivity increases, activation states start to occur in multiple assemblies simultaneously. Additionally, the networks more frequently enter a state of highly synchronous and regular spiking, similar to that described by Brunel (2000) in strongly-connected sparse random networks of integrate and fire neurons. The distributions of neuron firing rates, C_V and Fano factors are shifted still further to the right (Figure 2.14B, C and D). The power spectrum and autocorrelation of spiking in the highly synchronous state are dominated by the effect of the oscillations (Figure 2.14E and F). The peaks in the power spectrum (Figure 2.14E)

at harmonics of the oscillation frequency of ~50Hz are due to the pulse-like signal produced by the binned spikes.

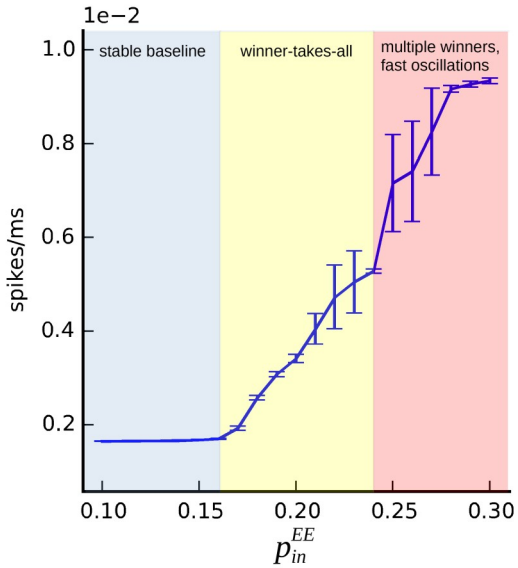

Figure 2.10: Mean population firing rate as a function of clustering, showing approximate boundaries of state changes

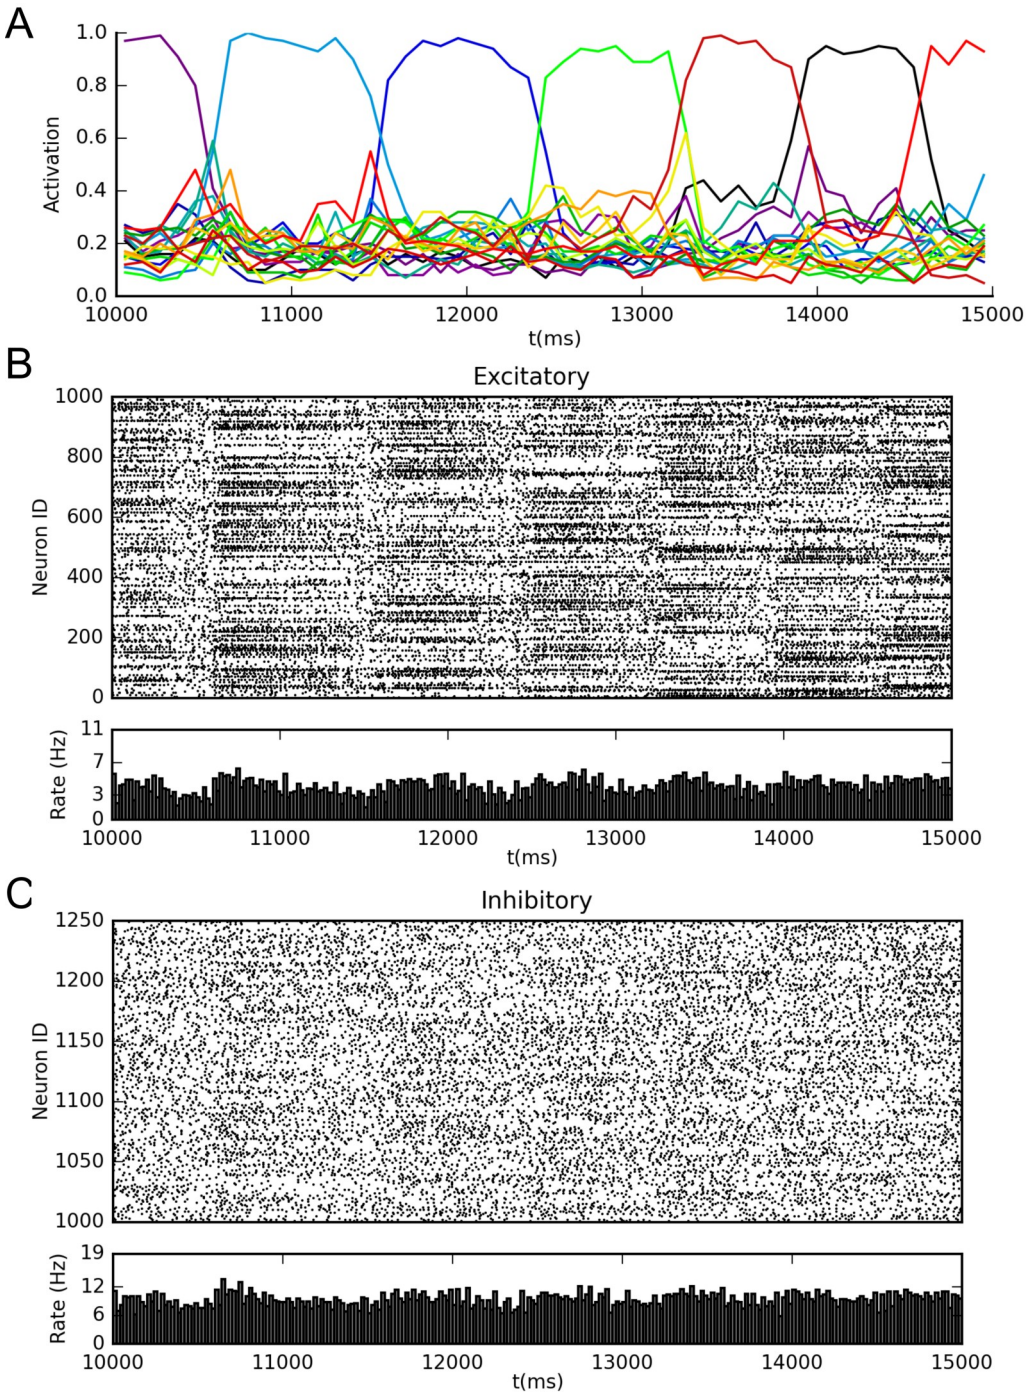

Figure 2.11: Example output of a balanced network with 20 overlapping assemblies created with the equal assemblies method with moderate clustering.
Networks created with $p_{in}^{EE}=2$*, plots correspond to the same simulation period. (A) Assembly activation measure* A*, with colored lines corresponding to activation of different assemblies. (B)(Top) Raster of spike output of 1000 excitatory neurons. (Bottom) Mean population firing rate.*

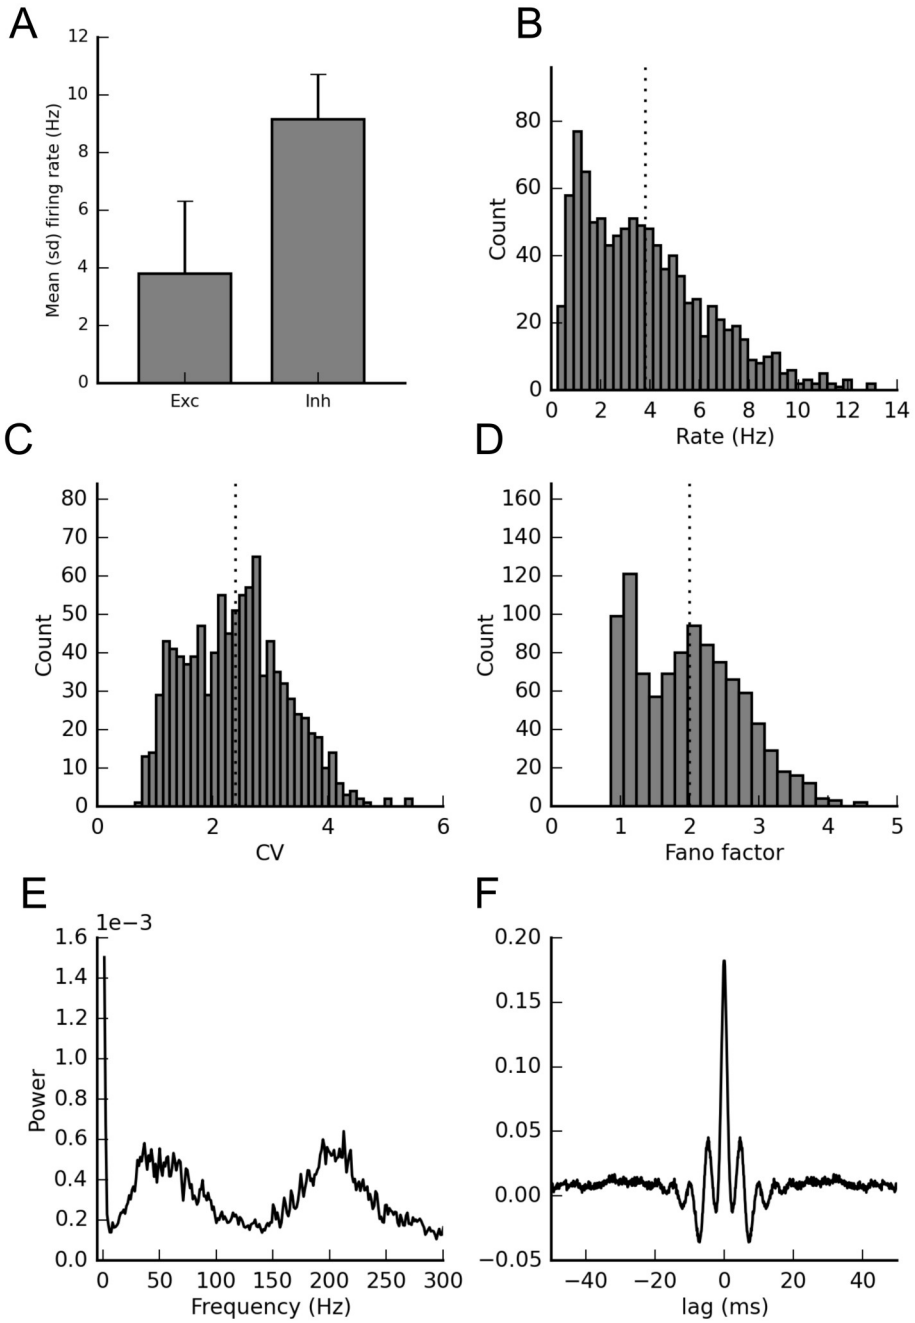

Figure 2.12: Statistics for an example balanced network with 20 overlapping assemblies created with the equal assemblies method with moderate clustering. *Networks created with* $p_{in}^{EE}=2$ *. Dotted vertical lines indicate mean values. (A) Mean (SD) of excitatory and inhibitory firing rates. (B) Histogram of excitatory neuron firing rates. (C) Histogram of coefficients of variation of inter-spike intervals. (D) Histogram of Fano factors of spike counts, computed with a 100ms window. (D) Power spectrum of spiking across all excitatory spikes, showing peaks at approximately 1, 50 and 200 Hz. (E) Population spiking autocovariance.*

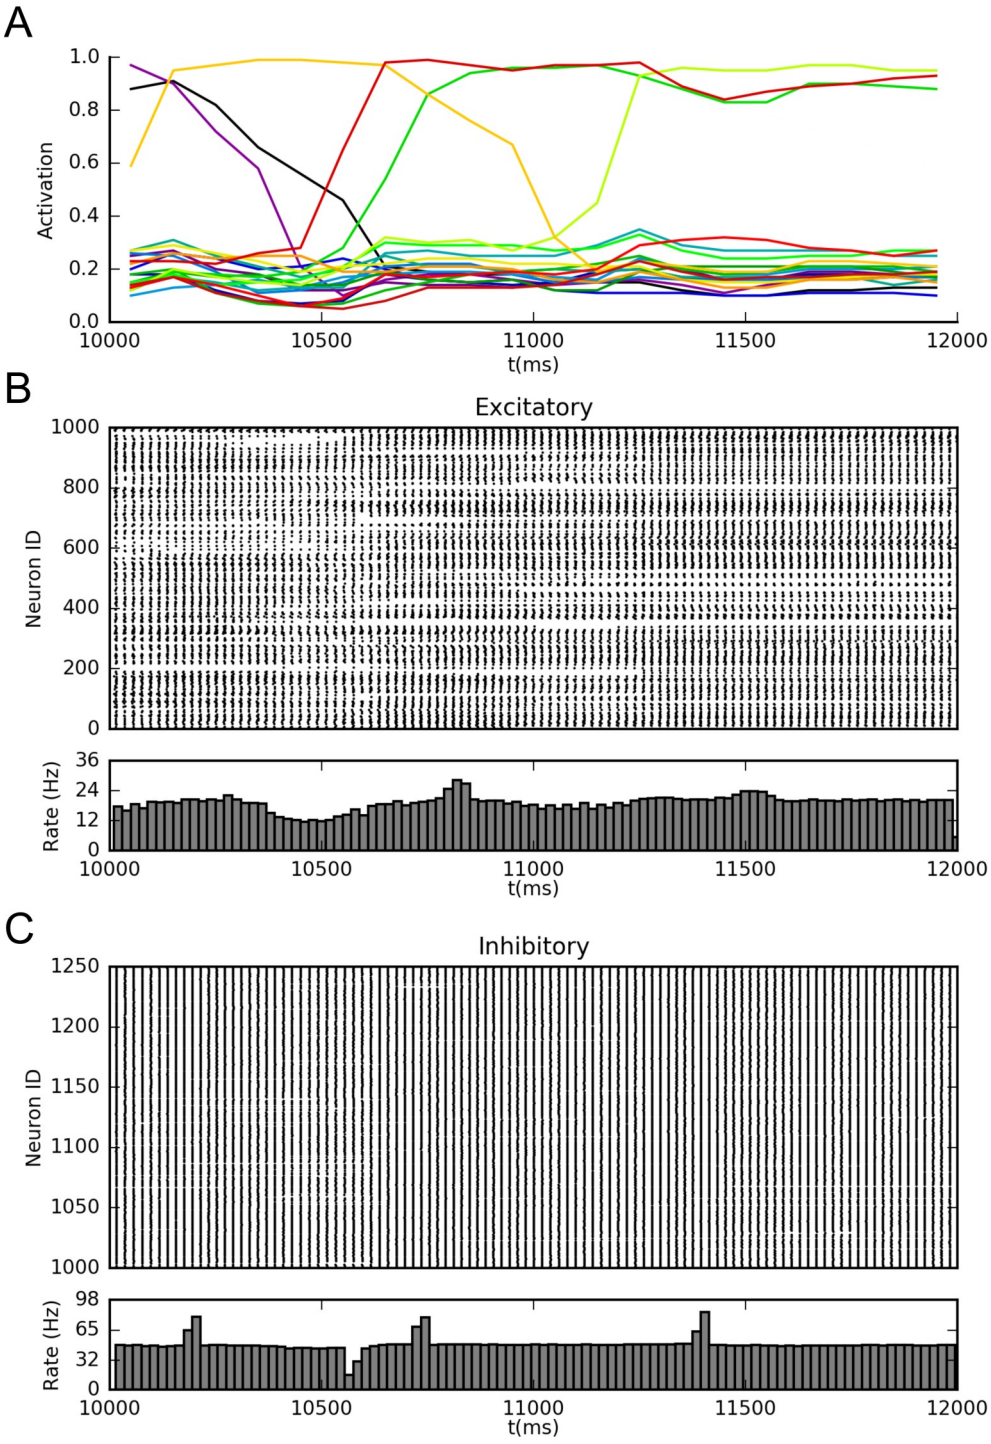

Figure 2.13: Example output of a balanced network with 20 overlapping assemblies created with the equal assemblies method with high clustering.

Networks created with $p_{in}^{EE} = 2.8$ *, plots correspond to the same simulation period. (A) Assembly activation measure* A*, with colored lines corresponding to activation of different assemblies. (B)(Top) Raster of spike output of 1000 excitatory neurons. (Bottom) Mean population firing rate.*

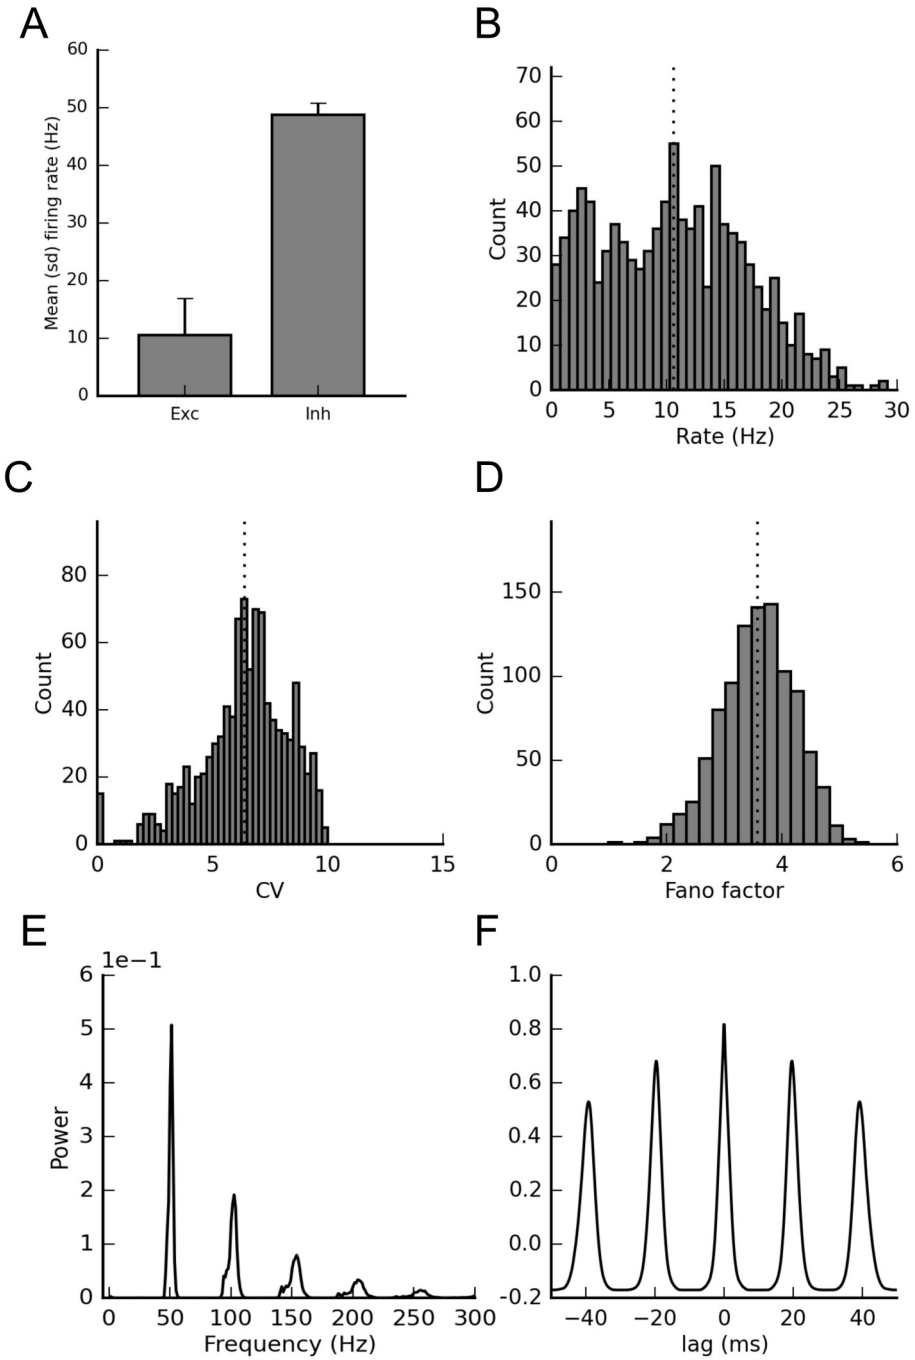

Figure 2.14: Statistics for an example balanced network with 20 overlapping assemblies created with the equal assemblies method with high clustering.

Networks created with $p_{in}^{EE} = 2.8$ *. Dotted vertical lines indicate mean values. (A) Mean (SD) of excitatory and inhibitory firing rates. (B) Histogram of excitatory neuron firing rates. (C) Histogram of coefficients of variation of inter-spike intervals. (D) Histogram of Fano factors of spike counts, computed with a 100ms window. (D) Power spectrum of spiking across all excitatory spikes, showing peaks at multiples of 50Hz. (E) Excitatory population spiking autocovariance.*

2.2.4 Adaptation current *w* during switching activation dynamics

As shown in Figure 2.15, the adaptation current w in excitatory units was strongly modulated by the activation state of the assemblies to which the unit belonged. By progressively reducing excitability with time in the activated state, this adaptation current increased the probability that a network would fall out of the activated state as a function of time. Test simulations (results not shown) found that the rate of activation transitions increased and became more regular with increases of the adaptation parameters a and b, although transitions still occurred when adaptation was removed.

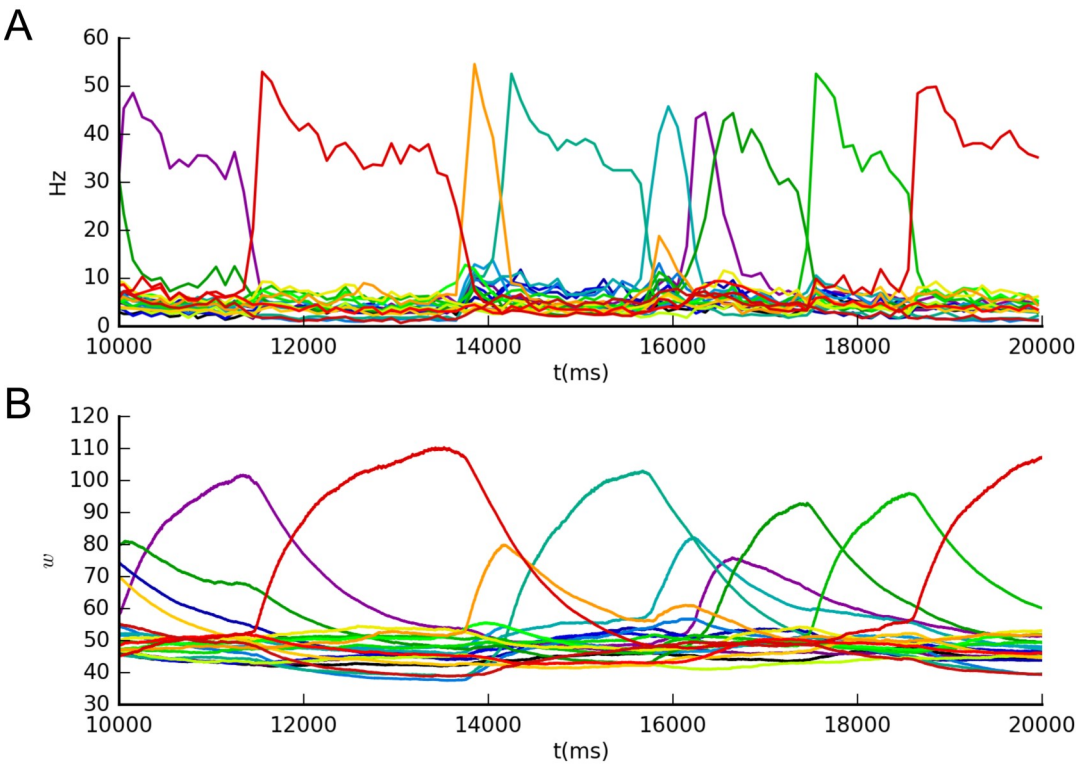

Figure 2.15: Mean values of adaptation current w in relation to assembly firing rates and switching activation dynamics.
Network created with $p_{in}^{EE}=0.22$*, plots and line colours correspond to the same simulation period and assemblies. (A) Mean firing rates of assemblies. (B) Mean values of hyperpolarising conductance* w *within specific assemblies.*

2.2.5 Clustering and delta oscillations

The clustering parameters modulated the power of delta-band oscillations in network spiking, which appears to have been mediated through the switching activation dynamics. This can be inferred by comparing the measured delta oscillation power with the measured Fano factor, shown in Figure 2.16A. While the correspondences are not perfect, they do show an emergence of delta oscillations with the emergence of >1 Fano factor values at approximately $P_{conn,in}=0.17$, which corresponds to the emergence of switching behavior. The power in the delta band increases approximately with firing rate in networks with overlapping clusters (Figure 2.16B). The rate at which activation switches between assemblies is in the delta band of 0-4Hz (Figure 2.8B), suggesting a causal link between the two.

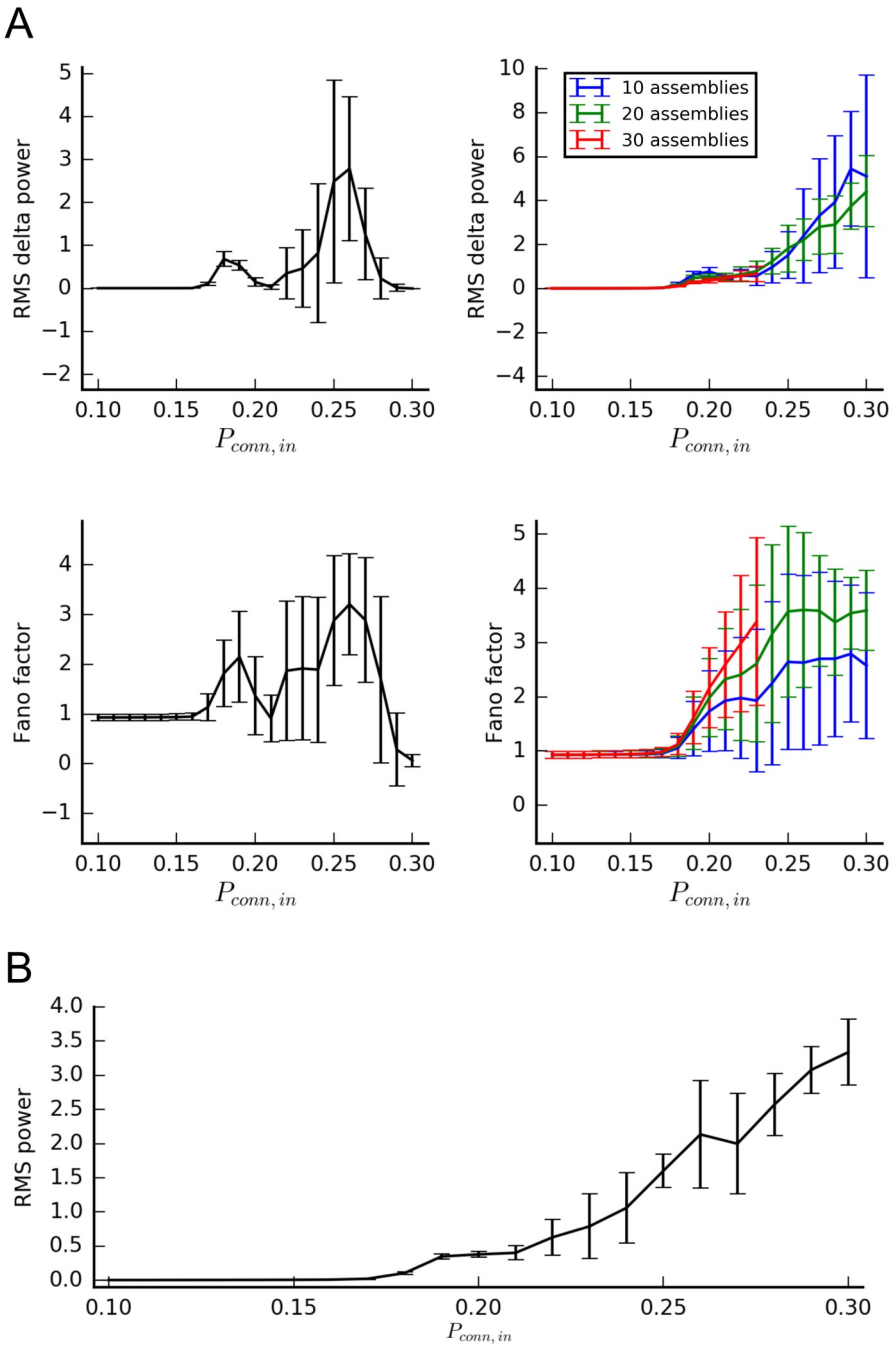

Figure 2.16: Mean (SD) root mean square delta frequency band (0-4Hz) power and Fano factor as a function of clustering in networks created with the equal assemblies definition.

(A) Left: non-overlapping assemblies, right: overlapping assemblies. Legend for overlapping plots: Blue: 10 assemblies, green: 20 assemblies, red: 30 assemblies. (B) Mean (SD) RMS delta power measurements from a set of 8 networks with 20 assemblies run for an extended period of 30 minutes of simulated time. See section 3.2 for details of the simulations.

2.3 Effects of clustering on network connectivity structure

The variations in spiking activity just described are due to variations in the excitatory connection matrix. The question arises of which qualities of the connection matrix cause the switching activation dynamics. The connection matrix itself is highly complex, with $(N^E)^2$ possible connections, and therefore requires descriptive measures to be understood. The clustering parameters R^{EE} and P_{in}^{EE} are themselves descriptive measures, and have a direct effect on the intra-assembly connection weight rec. These have a clear but simple effect on dynamics, with activation occurring after a threshold value and cumulative spiking in the activated state increasing with further increases of these values. In this section we preform exploratory analysis of the excitatory connection matrix using the clustering coefficient measure and with the eigendecomposition of the matrix, in order to find properties that determine the switching activation dynamics described previously.

2.3.1 Clustering coefficient

The connection matrix can be described as a directed graph or network, and can be analyzed with graph and network analysis measures. The *clustering coefficient* (CC) is one such measure which could potentially be informative in the present context. The original definition of this quantifies clustering by counting the instances of two mutually-connected nodes both being directly connected to a third node (thereby forming a “triangle”), and normalizing this by the maximum number of such triangles that could be formed given the number of connections possessed by the two nodes. As this original definition of the CC was for binary undirected graphs, we use an extension of the clustering coefficient for networks with weighted, directed connections (Fagiolo 2007), which considers all 8 potential configurations of directed triangles between three units and their connection weights:

$$\widetilde{C}_i^D(W) = \frac{\widetilde{t}_i^D}{\widetilde{T}_i^D} = \frac{\left[W^{[1/3]} + (W^T)^{[1/3]}\right]_{ii}^3}{2\left[d_i^{tot}(d_i^{tot}-1) - 2d_i^{\leftrightarrow}\right]} \tag{18}$$

where $\widetilde{C}_i^D(W)$ is the weighted directed clustering coefficient of unit i of weighted directed connection matrix W, $\widetilde{t}_i^D$ is the number of triangles which include unit i and $\widetilde{T}_i^D$ is the maximum possible number of triangles given the degree of the unit, d_i^{tot} is the total degree (in + out connections) of the unit, $2d_i^{\leftrightarrow}$ is the number of bilateral sets of connections of i, i.e. the number of other units which a connection is both sent to and received from. The subscript ii denotes the value at location (i,i) of the corresponding matrix, i.e. the diagonal entry corresponding to unit i. The implementation used was taken from the Python conversion of the Brain Connectivity Toolbox (Rubinov and Sporns 2010).

2.3.1.1 Results

The clustering coefficient measure $\widetilde{C}_i^D$ as described above was measured for each unit in each of the 24 randomly generated connection matrices for both clustering routines at each level of clustering. This was only calculated in respect to excitatory-excitatory connections, as other connection types remained random in all simulations. The mean and SD of these measurements in respect to clustering parameters R^{EE} and P_{in}^{EE} is plotted in Figure 2.17.

For the ratio networks, mean clustering coefficient $\widetilde{C}_i^D$ as described above rises monotonically and supralinearly with increasing clustering R^{EE}. In the case of overlapping assemblies, $\widetilde{C}_i^D$ increases with number of assemblies $N^{assembly}$ in ratio networks (Figure 2.17, bottom left).

The relationship between parameters and $\widetilde{C}_i^D$ is more complex in the equal assembly networks, as the mean value does not increase monotonically for either $N^{assembly}$ and P_{in}^{EE} (Figure 2.17). For increasing P_{in}^{EE}, mean $\widetilde{C}_i^D$ first decreases before rising, with this pattern being more pronounced with higher $N^{assembly}$. Increased P_{in}^{EE} caused instead a monotonic increase in the standard deviation of $\widetilde{C}_i^D$, with closer inspection revealing the development of a broader and bimodal

distribution. Overall, the mean weighted directed clustering coefficient does not show a clear correspondence to any of the measures of network activity used in this study. This indicates that this measure should be interpreted carefully when applied to neural networks (Saramäki et al. 2007), and in particular that the distribution of values should be considered in addition to the mean. This is in agreement with the general conclusions of the analysis of Saramäki et al. (2007), that clustering coefficient measures have subtle and specific relationships with network structure.

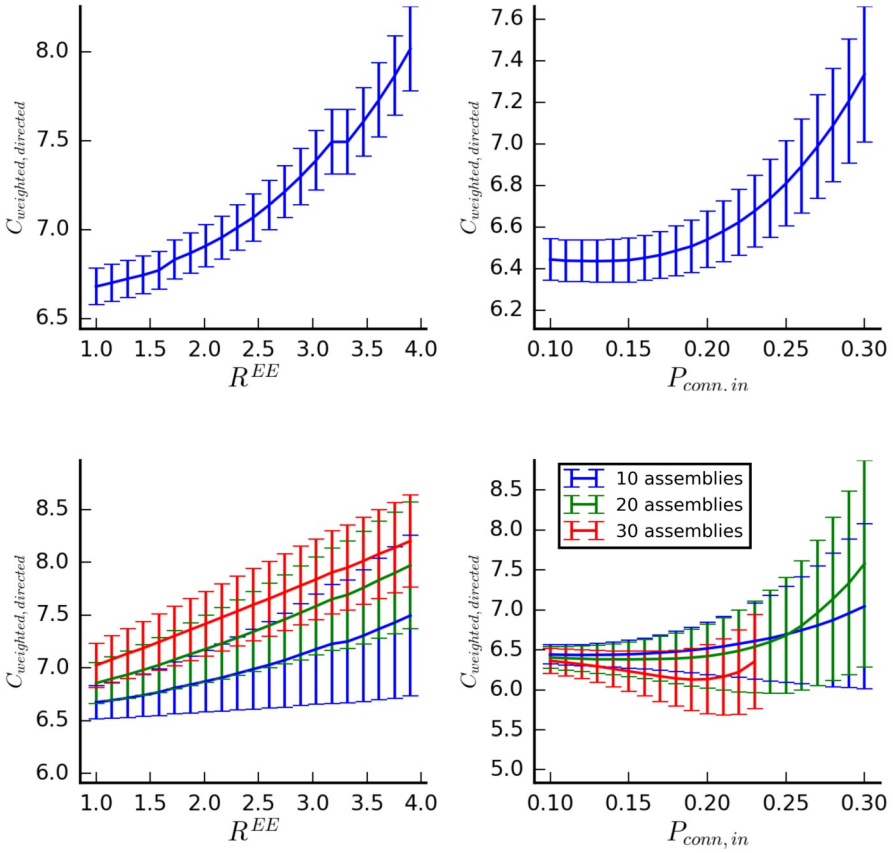

Figure 2.17: Clustering coefficients of connection matrices as a function of clustering.
Left: ratio networks, Right: EQASS networks, Top: non-overlapping, Bottom: overlapping.

2.3.2 Connection matrix eigenvalue analysis

The connections between neurons can be represented as a connection weight matrix W, which is amenable to analysis with linear algebra methods. This can be used to give insight into the qualities of the connectivity matrix that lead to bistability and switching behavior, although the nonlinearities of the spiking neuron and synapse models make a full analysis of the structure-dynamics relationship intractable (Schaub et al. 2015). Following previous analyses, we calculated the eigendecomposition of each of the 24 randomly generated connection matrices produced by each of the clustering routines at each level of the clustering parameters R^{EE} and p_{in}^{EE}. This was done as an exploratory analysis to find

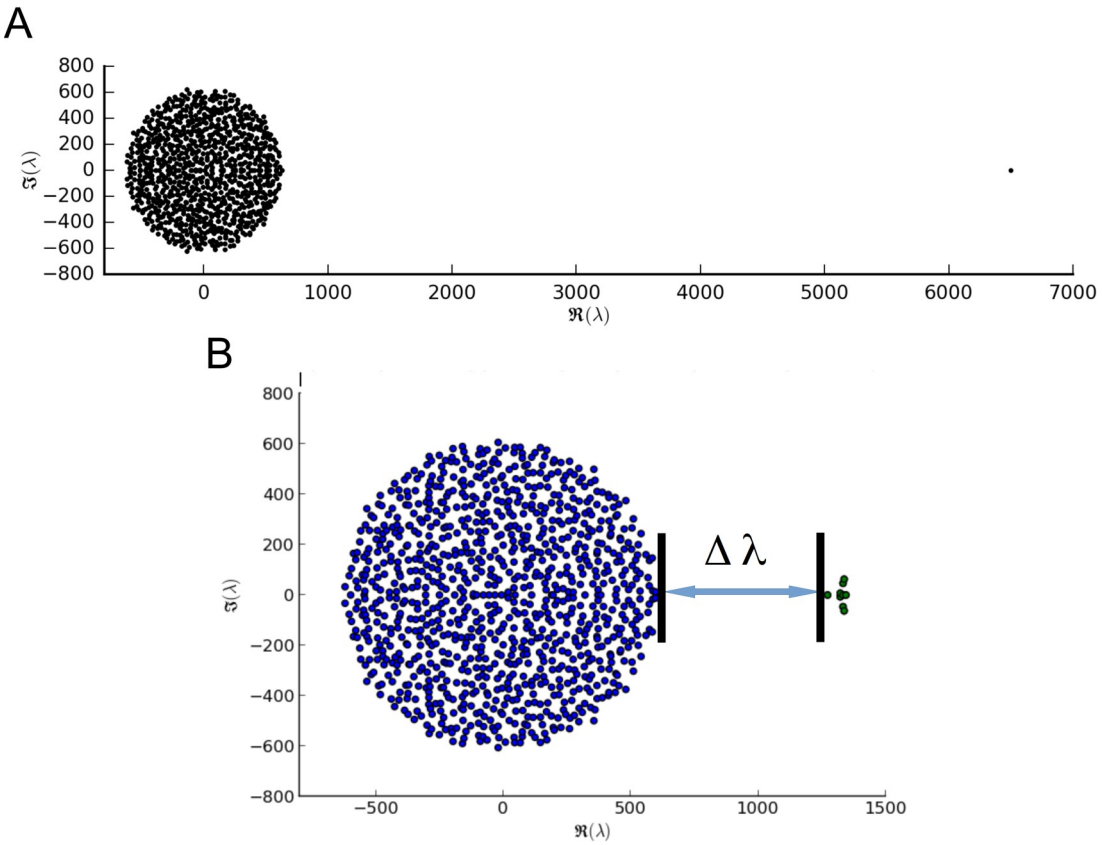

Figure 2.18: Eigenvalues of excitatory-excitatory connection matrices in the complex plane.
(A) All eigenvalues of an unclustered random excitatory connection matrix. Note the single large positive real value on the right. (B) Illustration of the value measured by the eigenvalue gap measure $\Delta\lambda$, the real part of the distance between the main circular group of eigenvalues and the assembly-associated group.

features of the eigenvalue spectrum that are associated with switching activation dynamics or other measured qualities, with the expectation that clustering sufficient to change activity regimes would be associated with the emergence of a subgroup of eigenvalues that are separated from the main group (Schaub et al. 2015).

As with the clustering coefficient, the eigendecomposition was only calculated in respect to excitatory-excitatory connections, as other connection types remained random in all simulations, and because the differences between the excitatory and inhibitory neuron and synapse parameters would complicate the interpretation of the full connection matrix.

The eigenvalue spectrum of a random graph such as a randomly-connected connection matrix has, with a few exceptions, a circular distribution in the complex plane (Rajan and Abbott 2006). This holds perfectly when matrix values sum to 0 and positive and negative elements have the same variance, which is not the case here. Here, the excitatory-excitatory connection matrix has an isolated large real eigenvalue which corresponds to the recurrent excitatory drive of the network, with a value equal to $K_{in}^{EE}\mu^{EE}$, the number of input connections per unit multiplied by the mean excitatory connection weight. The remaining values in an unclustered matrix lie within a circle with a radius of $\sqrt{1+K_{in}^{EE}(\mu^{EE})^2}$ (Rajan and Abbott 2006), as shown in Figure 2.18.

Above a threshold level of the clustering parameter p_{in}^{EE}, the addition of clustered assemblies to a connection matrix with the equal assemblies routine produces a subset of eigenvalues with real parts that are outside this circle, which correspond to the assembly-specific recurrent excitation within each assembly. The assemblies in this case can be conceived as vector directions which are amplified when multiplied by the connection matrix. Schaub et al. (2015) stress the importance of an "eigenvalue gap" in the onset of switching activation dynamics. An eigenvalue gap is the distance in the real plane between the smallest of the group of eigenvalues corresponding to the assemblies and the main group of eigenvalues lying within the circle (Figure 2.18). We measured the eigenvalue gap as the largest gap in the ordered list of the real parts of all eigenvalues excluding the large eigenvalue corresponding to the overall recurrent drive. As an additional measure, we also measured the mean of the "assembly eigenvalues", those eigenvalues with a real

value greater than the radius of the circle of the large group of eigenvalues defined by $\sqrt{1+K_{in}^{EE}(\mu^{EE})^2}$. Finally, we measured the standard deviation of the real parts of the assembly eigenvalues.

2.3.2.1 Results

The relationship of these metrics to clustering parameters is shown in Figures 2.20 and 2.21. Qualitative differences are evident between the matrices with non-overlapping and overlapping assemblies, with notable differences between the overlapping assembly matrices generated by the ratio and equal assembly routines. In this respect the eigenvalues of the connection matrices resemble the spiking dynamics of the simulated networks.

The eigenvalue metric output is qualitatively similar for non-overlapping assembly matrices created by both the ratio and equal assembly definitions, with an “assembly” group of eigenvalues emerging after a threshold level of clustering, and both the mean of this group and the eigenvalue gap increasing linearly as a function of clustering after (Figure 2.20A and B). The distribution of assembly eigenvalues is initially relatively large, before becoming more compact as clustering increases (Figure 2.20C).

The eigenvalue gap measurement has looser relationships with clustering parameters and spiking dynamics when assemblies are allowed to overlap (Figure 2.21), with the relationships becoming looser with increasing number of assemblies $N^{assembly}$. This is also true for networks created with the equal assemblies algorithm, in which switching activation dynamics were largely independent of $N^{assembly}$. However, the mean of the emergent group of assembly eigenvalues was also virtually independent of $N^{assembly}$ in matrices created by the equal assembly algorithm(Figure 2.21B). As with the non-overlapping assembly matrices, the emergence of a group of eigenvalues with a mean real value greater than a threshold value of $\sqrt{1+K_{in}^{EE}(\mu^{EE})^2}$ slightly preceded the emergence of assembly activation dynamics in simulations. The measurements of the standard deviation of the assembly eigenvalues revealed a wider distribution in comparison to the non-overlapping assembly matrices (Figure 2.19 and Figure 2.21B), and that the

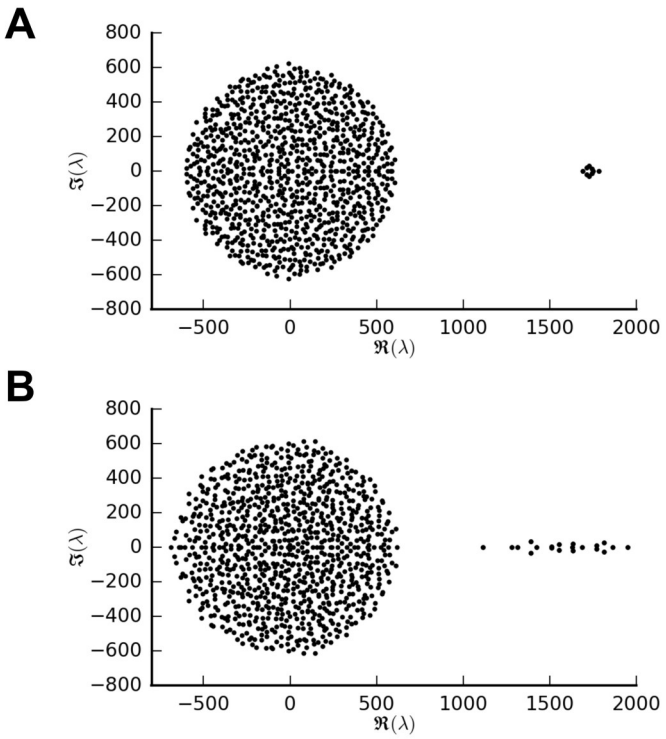

Figure 2.19: Representative eigenvalues in matrices with non-overlapping versus overlapping clusters

Eigenvalues of a matrix with clusters created with $p_{in}^{EE} = 0.25$ *. (A) 10 non-overlapping clusters (B) 20 overlapping clusters. Note wider distribution of values in the case of overlapping clusters.*

distribution became wider with increased clustering and number of assemblies. This increased variance of eigenvalues appears to account for the discrepancy between the eigenvalue gap and the assembly eigenvalue mean. The finding of weak association of the eigenvalue gap measure with switching activation and the strong association of the mean assembly eigenvalue measure can be considered a minor correction to the conclusions of Schaub et al. (2015).

The difference between matrices with overlapping and non-overlapping assemblies in this respect can be explained in terms of the eigenvectors associated with the assembly eigenvalues. While non-overlapping recurrent assemblies can be described in terms of orthogonal directions in the N^E -dimensional feature space, the overlapping assemblies correspond to non-orthogonal directions. As recurrent excitation within assemblies can be considered as matrix amplification of activity in the "direction" of the assembly, orthogonal assemblies with equal recurrent

excitation correspond to a set of approximately equal eigenvalues, with some variance due to the semi-randomness of the connectivity. Equal amplification in non-orthogonal directions, on the other hand, will correspond to a more varied set of eigenvalues, whose associated eigenvectors can be related to commonalities and differences between the assemblies.

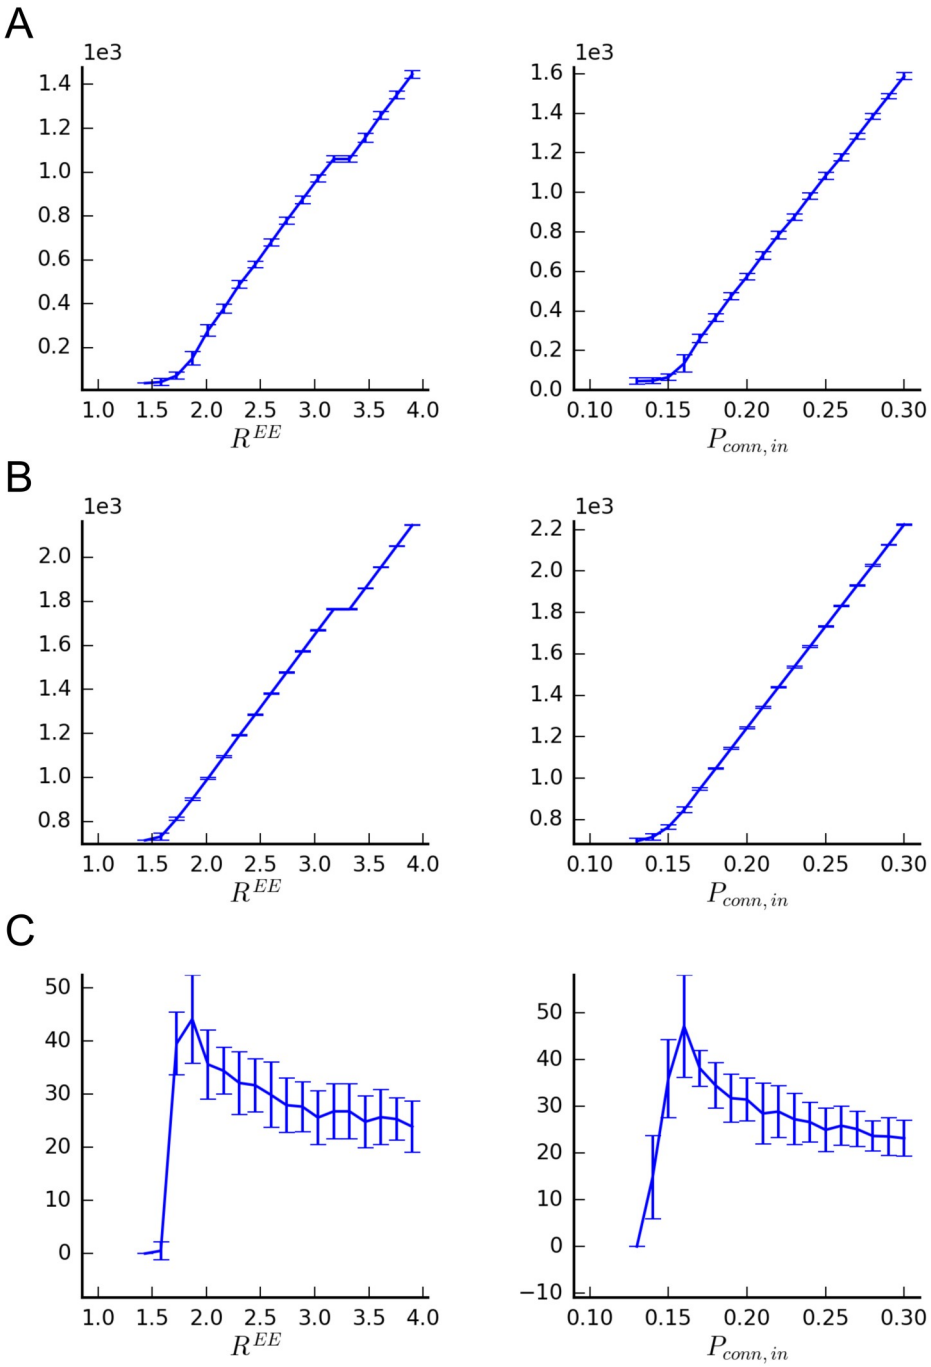

Figure 2.20: Effects of progressively increased clustering on eigenvalue metrics of clustered connection matrices with non-overlapping assemblies.
Results from matrices created with the ratio (left) and EQASS (right) definitions. Mean (SD) values over 24 randomly generated networks. See main text for definition of metrics. Note that left plots for (A) and (B) show a rounding artifact as described in section 2.2.2.2. (A) Eigenvalue gap. (B) Mean of real part of assembly eigenvalues. (C) SD of real part of assembly eigenvalues.

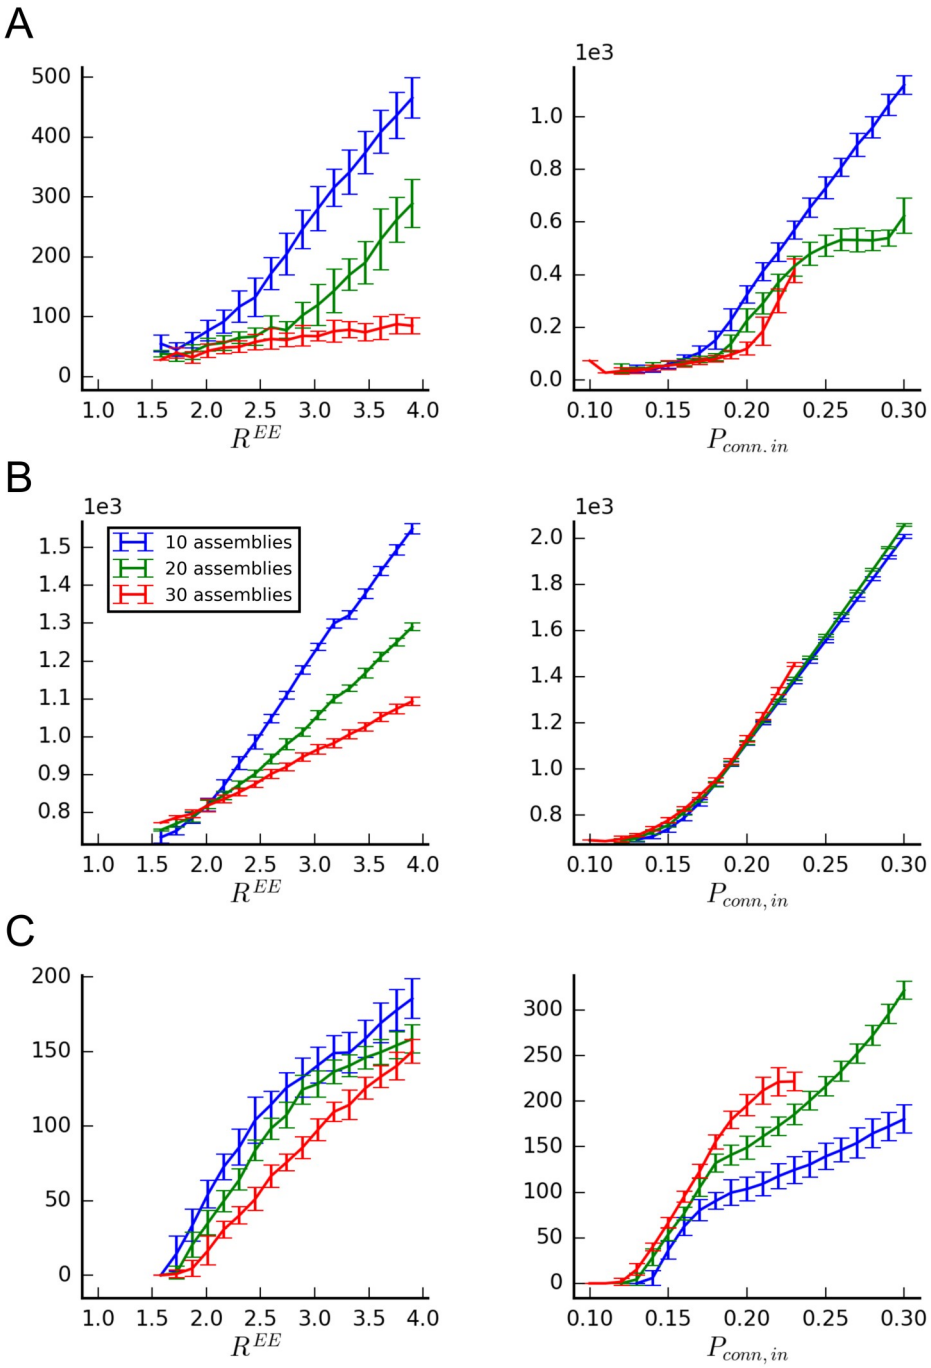

Figure 2.21: Effects of progressively increased clustering on eigenvalue metrics of clustered connection matrices with overlapping assemblies.
Results from matrices created with the ratio (left) and EQASS (right) definitions. Mean (SD) values over 24 randomly generated networks with overlapping assemblies. See main text for definition of metrics. Blue: 10 assemblies, green: 20 assemblies, red: 30 assemblies. (A) Eigenvalue gap. (B) Mean of real part of assembly eigenvalues. (C) SD of real part of assembly eigenvalues.

2.4 Discussion

The simulation results presented demonstrate that embedding clustered groups into balanced networks of spiking neurons fundamentally changes spiking dynamics. This replicates and extends previous work in this area (Litwin-Kumar and Doiron 2012; Schaub et al. 2015; McDonnell and Ward 2014).

The results indicate that intra-assembly recurrent excitatory connectivity is a key determinant of activation dynamics, and that this knowledge can be used to predict specific types of spiking dynamics when setting connections in a neural network. We have also specified an algorithm which generates sparse connection matrices with embedded assemblies with specific levels of recurrent connectivity. This was compared to the ratio-based clustering definition specified by Litwin-Kumar and Doiron (2012), in which intra-assembly excitation and therefore dynamics is reduced when assemblies overlap. When comparing the two, it should be stated for the sake of fairness that the clustering definition of Litwin-Kumar and Doiron (2012) was not designed to fulfill the measurement criteria used here (assembly-specific activations in networks with densely overlapping assemblies), and so should not be considered inferior in terms of design.

The association between delta band oscillations, variability of firing rates and switching activation dynamics in this study suggests that these could be related in the cortex. We found that the distribution of assemblies (i.e. whether or not units can be members of multiple assemblies) has a direct effect on statistical descriptions of network activity. Specifically, we found that overlapping assemblies were associated with more frequent state transitions, with a corresponding effect on delta oscillations and firing rate variance. This effect will be further explored later in this thesis.

Compared to other computational studies, the network model we used had an intermediate level of biophysical realism. The network used an adapting spiking neuron model, a conductance synaptic model with multiple channel types and synaptic depression, and sparse excitatory connectivity and moderately dense nonspecific inhibitory connectivity. This is more realistic than models with binary, firing rate or other simple spiking units and all-to-all network connectivity, but less realistic than models with multiple neuron types, multi-compartment neuron models,

or anatomically accurate network topology and connectivity. While greater biophysical realism generally increases biological plausibility, all else being equal, increased complexity also decreases analytical tractability. The relatively complex model used here builds upon work reported in previous studies, and we refer interested readers to these studies for further detail and theoretical justification of model elements, particularly in regards to the function of NMDA synapses and synaptic biophysics in supporting bistability of firing rates at physiologically realistic rates (Brunel and Wang 2001; Durstewitz, Seamans, and Sejnowski 2000a; Wang 1999). Neuron adaptation currents are relatively unusual in balanced network models, perhaps because analysis has often focused on steady-state activity, where adaptation is held constant. Their inclusion here was partly inspired by their inclusion in models of neural computation which use rate neuron models, in which adaptation drives transitions between attractor states (Wennekers and Palm 2009; Russo et al. 2008; Gros 2007). Adaptation has a similar function of driving transitions between states in the model presented here. In previous spiking models of firing rate transitions, transitions have been driven by firing rate fluctuations due to irregular spiking and finite-size effects (Litwin-Kumar and Doiron 2012; Schaub et al. 2015; Moreno-Bote, Rinzel, and Rubin 2007). While irregular spiking and finite size effects certainly contribute to transition dynamics in the model shown here, the inclusion of an adaptation current increases the regularity and frequency of state transitions. Neuron adaptation currents are certainly biophysically realistic and are a well-explored inclusion in neuron and network models (Hertäg et al. 2012; Durstewitz 2007; Brette and Gerstner 2005). At least two previous spiking network studies have linked adaptation with firing-rate transitions and delta band oscillations, although not necessarily in the context of assembly attractor states (Compte, Sanchez-Vives, et al. 2003; Durstewitz 2007).

The network showed winner-take-all dynamics at intermediate levels of clustering, in which the neurons of single assembly fired at a high rate, and neurons not belonging to this assembly fired at a lower rate. This dynamical regime allows a simple interpretation of the system state in terms of attractor states corresponding to a single assembly. The sequential transitions of the network between these states also potentially allows a macroscopic description of dynamics in terms of activity trajectories and transitions. This will be explored in the next section.

We found that the network produced highly synchronized firing at higher levels of clustering. Comparable states have been observed in simpler random spiking networks at high levels of connectivity (Brunel 2000), but are not reported in-vivo as far as we are aware. The biological plausibility and relevance of these fast regular spiking states is uncertain. The occurrence of such states is reduced by increased non-uniformity of synaptic delays (Brunel 2000), and it may be that biological networks do not possess the regularity and reliability to sustain such states.

Higher clustering rates produced states in which multiple assemblies were simultaneously activated. Transitions in and out of activation states also occurred in this regime, with the transitions largely asynchronous between the different assemblies. These dynamics would be more challenging to describe and analyze in comparison to the single assembly attractor states that occur with moderate clustering.

Worthy of note is the distributed nature of assemblies when they are allowed to randomly overlap. Each neuron is on average a member of 3 assemblies in a network with 30 embedded assemblies, and so knowledge of the state of multiple neurons is generally necessary to determine the attractor state of the network. Similar phenomena have been reported during in-vivo experimentation. Determination of system state from multicellular experimental recordings requires statistical techniques for estimating the system state from the subset of neurons for which data are available (Seidemann et al. 1996; Balaguer-Ballester et al. 2011). Such techniques are not necessary for analysis of the simulated network, as the output of all neurons is available.

2.5 Standard simulation parameters

Table 1: Standard simulation parameters

Neuron and synapse parameters	
Symbol	**Value**
C_m	200 pF
g_L^{ex}	10 nS
g_L^{inh}	15 nS
E_L	-70 mV
θ	-50 mV
V_{reset}	-80 mV
τ_{ref}^{ex}	5 ms
τ_{ref}^{inh}	2 ms
a	4 nS
b	2.5 pA
τ_w	1000 ms
I_{inj}	350 mV
τ_{on}^{AMPA}	0.1 ms
τ_{off}^{AMPA}	1.0 ms
g_{peak}^{AMPA}	1 nS
τ_{on}^{NMDA}	2.0 ms
τ_{off}^{NMDA}	100.0 ms
g_{peak}^{NMDA}	0.09 nS
τ_{on}^{GABA}	0.1 ms
τ_{off}^{GABA}	3.0 ms
g_{peak}^{GABA}	1 nS
E_L^{ex}	0 mV
E_L^{inh}	-80 mV
D^{ex}	1 – 5 ms (uniform distribution)
D^{inh}	1 ms

Symbol	Value
U	0.1
τ_{rec}	200 ms

Network parameters	
Symbol	**Value**
N^E	1000
N^I	250
p^{EE}	0.1
p^{EI}	0.5
p^{IE}	0.5
p^{II}	0.5
J^{EE}	65
J^{EI}	6
J^{IE}	18
J^{II}	5

3 Markov transition dynamics and sequences as emergent properties of clustered networks

3.1 Introduction

In the simulations described in the previous section, the state of the neural network is readily interpretable in terms of the activation state of the embedded assemblies. This is particularly the case for those parameter combinations in which a single assembly is in the activated state at most times of the simulation. The sequential switching nature of these states suggests an interpretation of network dynamics in terms of trajectories through network state space, and particularly in terms of transitions between discrete states.

In the latter case, the dynamics of transitions between discrete states can be analyzed within the framework of Markov chains, in which a system is described in terms of a set of states $S=\{s_{1,}s_{2,}\ldots,s_r\}$ and a set of transition probabilities, such that probability p_{ij} is the probability of a transition from state s_i to state s_j. These can be conveniently represented in the form of a transition matrix, which facilitates analysis (Grinstead and Snell 2010).

In this section we describe a process by which the activity in clustered balanced networks is converted to a set of discrete states, and dynamics described as a sequence of states. We test statistically the hypothesis that transition probabilities are state-dependent rather than uniform, and quantify the extent to which probabilities are non-uniform. We furthermore examine the relationship of this non-uniformity to the degree of clustering in the network, and quantify the relationship between individual transition probabilities and structural features of networks, particularly the correlations between clusters. Finally, we propose a mechanism by which transitions are determined, with reference to previous neural models of sequence generation (Russo et al. 2008; Rabinovich et al. 2001).

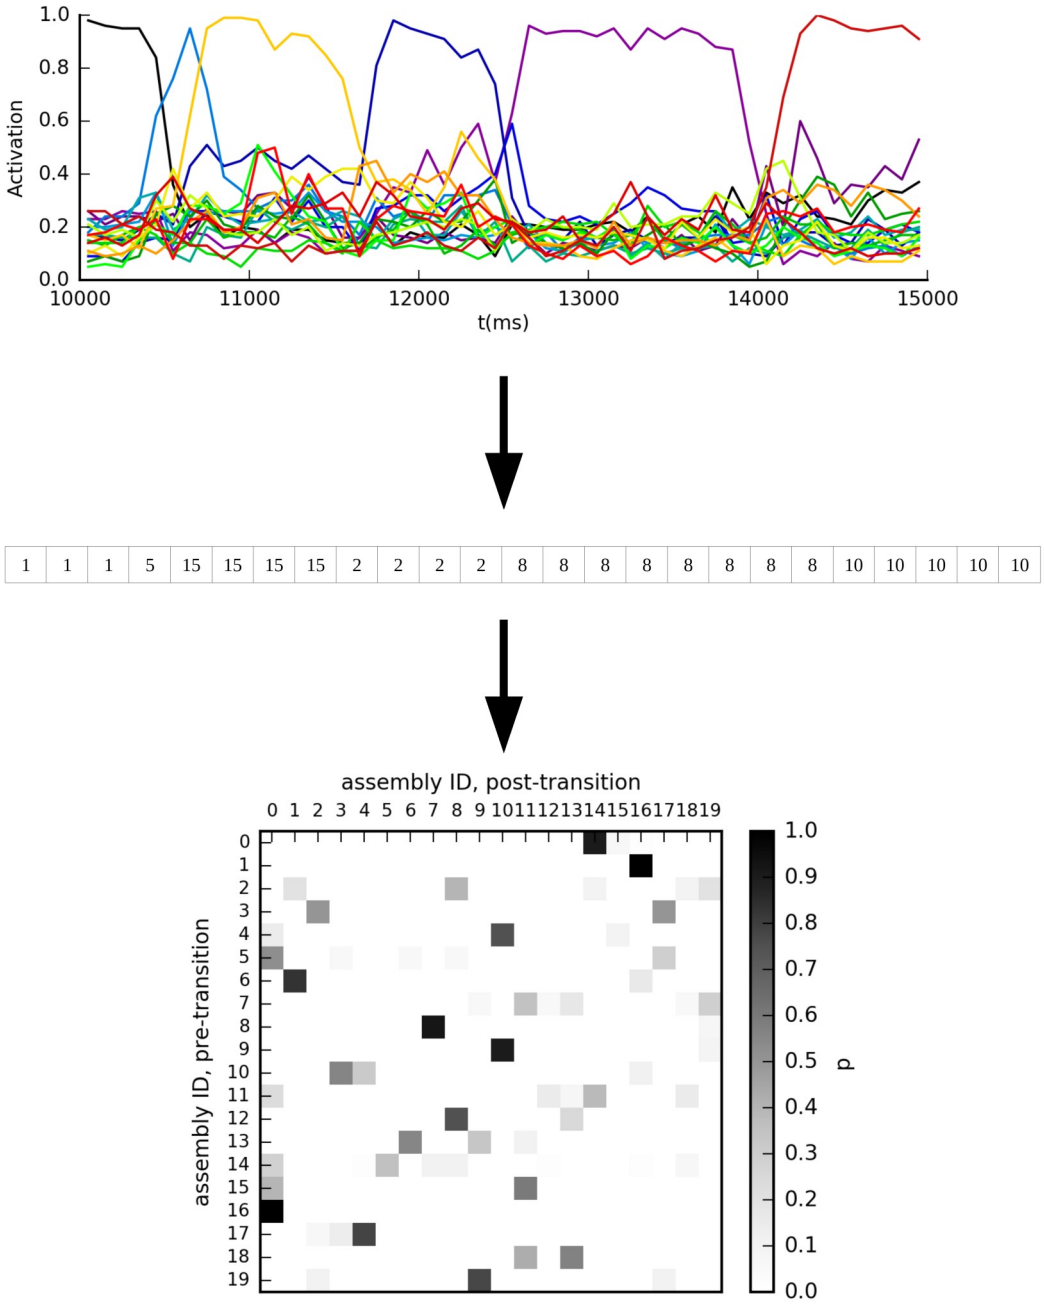

Figure 3.1: Schematic of the production of a transition matrix from simulation data.
(Top) Clustered networks are simulated for extended periods and the activation state of the embedded assemblies calculated. (Middle) A time series of most-activated assembly identities is generated from the activation measurements. (Bottom) Transitions between different states are counted, and a transition probability matrix is calculated. Probabilities are represented here as a heat plot.

3.2 Methods

We created representations of network dynamics as a transition matrix by sampling network activity during extended simulations. This process is depicted schematically in Figure 3.1. For this, we restricted simulations to networks with $N^{assembly}=20$ overlapping assemblies created with the equal assemblies algorithm as before, as these had previously been found to show switching activation dynamics over a wide range of clustering levels. Simulations were run with the clustering parameter p_{in}^{EE} systematically varied from 0.1 to 0.3 in increments of 0.01. For each level of p_{in}^{EE}, 8 independent network instances were created using different random number generator seeds. Simulations were run for 30 minutes (1800000 ms) of simulated time, plus an additional 5000ms which was discarded in subsequent analysis so as to remove initial transients.

To produce a simplified characterization of network state in terms of the identity of a single assembly, the spiking activity was binned with a resolution of 100ms and the activation state A of each assembly was calculated as described in section 2.2.1.3. The resulting set of time series was then further processed to define the system state at each time by determining the most activated assembly at each time point (i.e. that with the highest value of A) and setting the state correspondingly, except those periods where no assembly had an activation value greater than a threshold value of $A=0.7$, in which case the state was set to the same state as the previous time step. This was to allow for brief low-activity periods during transitions between states.

The resulting time series was then further processed to produce a transition matrix by counting the number of times that the time series changed from state x to state y and recording this at t_{xy} of transition count matrix T. This was then converted to a matrix of transition probabilities P by dividing each entry t_{xy} by the sum of all transitions from state x, i.e. $p_{xy}=\frac{t_{xy}}{\sum_{j=1}^{N^{assembly}} t_{xj}}$.

We chose only to count and estimate transition probabilities for transitions between

different states, and not for the probability that the system remains in the same state for successive time steps. This was because the assembly state of the network could be reasonably assumed to be a time inhomogenous rather than a time homogenous process, and as such would not be well-represented by a transition matrix which represents the probability of a transition as independent of the number of steps in which the system had already been in its current state. One reason for this inhomogeneity was that transitions between states were partly driven by the depolarizing conductance w, which increases in magnitude while assemblies are in the activated state, thereby increasing the likelihood of a switch to a different assembly state with time (Figure 2.15).

3.3 Analysis methods and results

As before, switching activation dynamics were produced by networks with $p_{in}^{EE} \geq 0.17$. Transition matrices were calculated for all networks at these levels, but were discarded for $p_{in}^{EE} = 0.17$ due to insufficient numbers of transitions. Mean number of transitions as a function of clustering is shown in Figure 3.2A.

3.3.1 Quantification of deterministic transition dynamics

Transition matrices showed large disparities in transition probabilities and appeared to clearly be the result of systems with inherent dynamical structures. We attempted to quantify the extent that this was the case, in order to obtain statistical confirmation that the system produces predictable sequences at the level of assembly activations, to compare the degree to which different systems are deterministic rather than random, and to be able to make correspondences between the non-randomness of dynamics and features of the systems that produce them.

The activation statistics suggested a potentially simple form of non-random temporal structure that could be considered trivial. The counts of activations per assembly were highly inhomogeneous, which suggests that some states simply have larger basins of attraction, and so are more active and 'attractive' independent of current activation state. An example plot of typical activation counts during a 30 minute simulation is shown in Figure 3.2B. Transition dynamics could therefore be

determined by a kind of unbiased competition between sub-active assemblies to become the new most-active assembly.

We defined multiple measures and tests of the inhomogeneity and predictability of transition dynamics that are based on Pearson's chi-squared test. The first was a test of the inhomogeneity of activation counts for each assembly. This is done by calculating the chi-squared statistic X^2, with the null hypothesis of an equal number of activations per assembly:

$$\mathrm{X}^2 = \sum_{i=0}^{N_{assembly}} \frac{(\mathrm{O}_i - E_i)^2}{E_i} \tag{19}$$

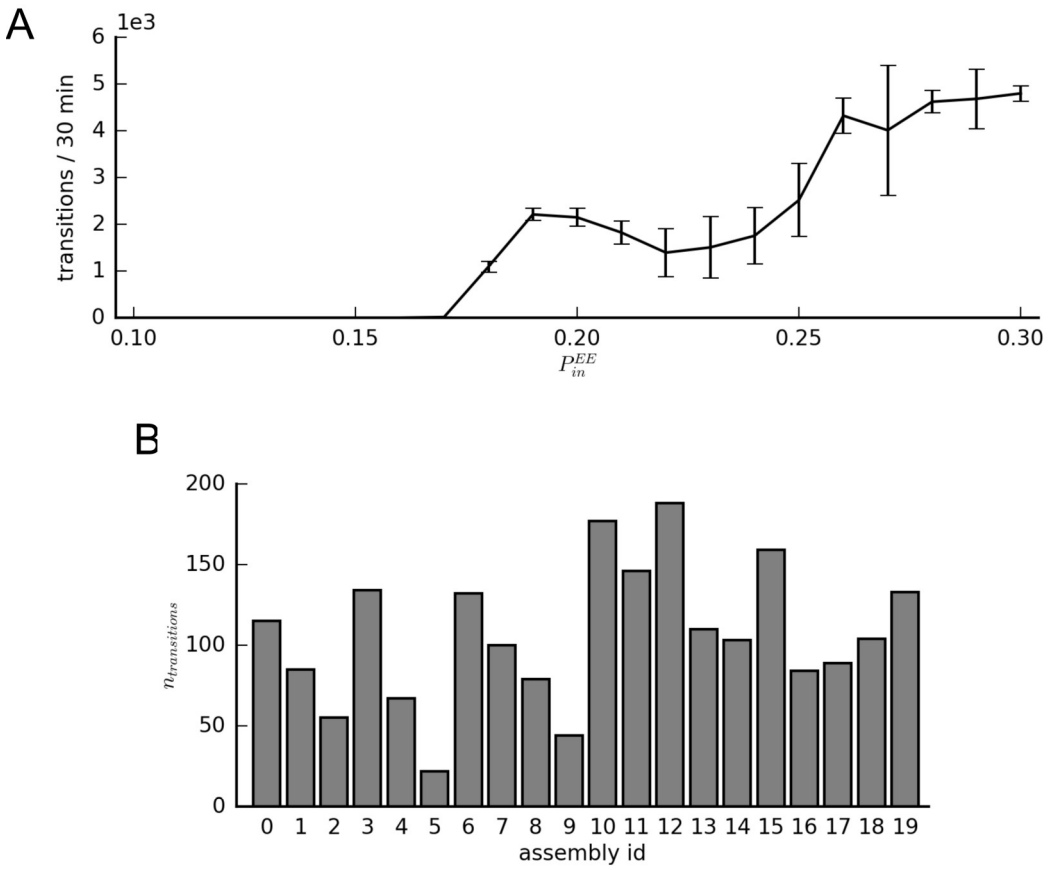

Figure 3.2: Frequency of changes of assembly activation state as a function of clustering.
(A) Number of separate assembly state transitions in 30 minutes of simulated time. (B) Example counts of separate assembly state occurrences in one 30 minute simulation with $p_{\mathrm{in}}^{EE} = 0.21$ *, demonstrating unevenness of activation counts.*

where O_i is the observed frequency of activations of assembly i, E_i is the expected value of O_i, which in this case is $N^{assembly}/n$, where n is the total number of activations.

Additionally, we also defined a normalized chi-squared statistic that we use to compare homogeneity of distributions as a function of parameters. As a test of deviation from a given distribution the chi-squared statistic is sensitive to the number of samples available. It is only exact in the case of an infinite number of samples, and tends to be conservative when sample numbers are low, particularly when the expected frequencies of categories are low (Cressie and Read 1989). Total transition numbers in each 30 minute simulation had a wide, parameter-dependent distribution in the simulations, which had a corresponding effect on expected frequency counts (Figure 3.2B). To compensate for this, we therefore calculated a normalized chi-squared statistic by dividing counts by the total number of activations n, so that the total number of transitions was equal for all simulations. The resulting statistic can be considered a measure of the *degree* of inhomogeneity of activation counts.

The results of both the normalized and unnormalized chi-squared statistics are shown in Figure 3.3A and B. The unnormalized chi-squared statistic X^2 (d.f.=19) had a mean value of >500 ($p<0.0001$) for all values of $p_{\text{in}}^{EE} \geq 0.17$, confirming that activation probabilities were unequal. The normalized statistic indicated peaks in activation inhomogeneity at $p_{\text{in}}^{EE} = 0.18$ and in the range $p_{\text{in}}^{EE} = 0.22 - 0.25$, suggesting a loose association with low total transition counts (compare Figure 3.2A).

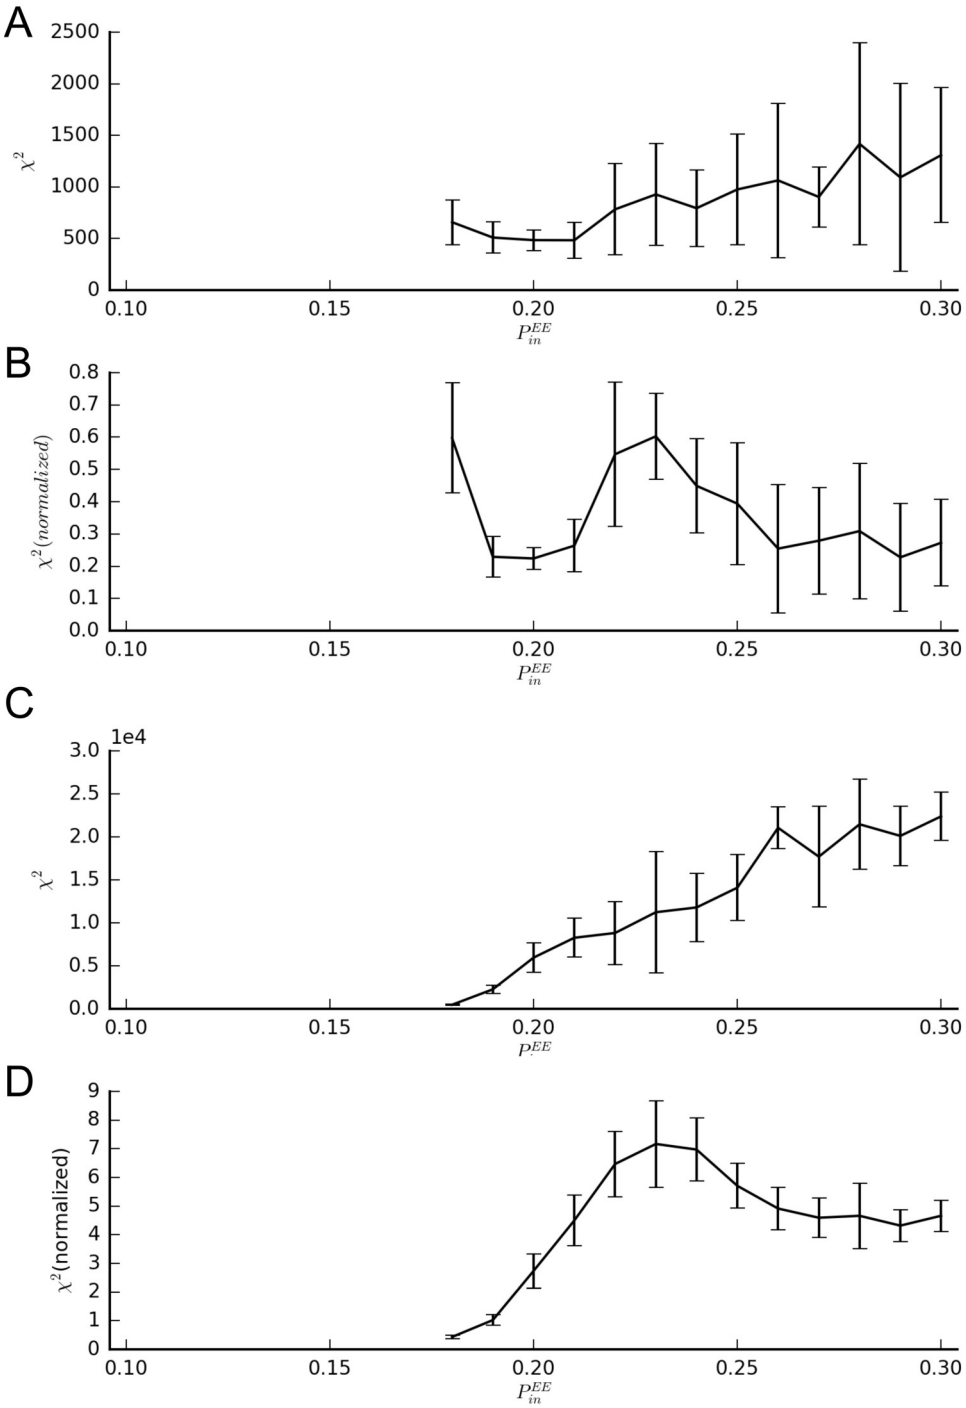

Figure 3.3: Assembly activation inhomogeneity and state-dependence as a function of clustering.

(A) Chi-squared statistic of state activation count inhomogeneity. (B) Normalized chi-squared statistic of state activation count inhomogeneity. (C) Chi-squared statistic of the fit of the transition matrix to a model of state-independent relative activation probabilities. (D) Normalized chi-squared statistic of the fit of the transition matrix to a model of state-independent relative activation probabilities.

3.3.2 State-dependent relative transition probabilities

The strong hypothesis that we are interested in is that relative transition probabilities are state dependent, such that the relative probabilities of the system moving to states C or D are different depending on whether the system is currently in state A or B. As the transition matrix we use is constructed such that transitions from a state to the same state do not occur, the probability of a transition of a state to itself is always zero, and so absolute transition probabilities are always to some extent state dependent. However, it could be possible that the relative activation probabilities of the other assemblies are state independent. If this is the case, and the probabilities of the transitions to the other states depend only on the intrinsic general excitability of each state, then this is a simpler and computationally less interesting form of dynamics than the alternative hypothesis that the relative probabilities and levels of excitability differ depending on current state.

We create an expected distribution which corresponds to this null hypothesis that relative transition probabilities are state-independent.

$$E_{ij} = \frac{n_i n_j}{\sum_{j=0,\, j\neq i}^{N_{assembly}} n_j} \tag{20}$$

where E_{ij} is the expected number of transitions from state i to j and n_i is the number of occurrences of state i.

The chi-squared measure for the test of state-dependent transition probabilities is therefore

$$\mathrm{X}^2 = \sum_{i=0}^{N_{assembly}} \sum_{j=0,\, j\neq i}^{N_{assembly}} \frac{(\mathrm{O}_{ij} - E_{ij})^2}{E_{ij}} \tag{21}$$

As with the assembly activation counts, we defined a normalized chi-squared statistic in order to compare the state-dependence of different parameter levels. As before, we calculated a normalized chi-squared statistic by first dividing transition counts by the total number of transitions n, so that the total number of transitions is equal for all simulations. Similarly to the normalized chi-squared measure of

activation count inhomogeneity, this can be considered a measure of the degree of state-dependence, i.e. the degree to which probabilities of transitions to specific states differ depending on the current state.

The results of these measures are shown in Figure 3.3C and D. The mean (SD) chi-squared statistic X^2 (d.f.=324) rises from 472 (91) at $p_{in}^{EE}=0.18$ to 22384 (2811) at $p_{in}^{EE}=0.3$, which corresponds to p<0.05 for $p_{in}^{EE}=0.18$ and p<0.0001 for higher values. This indicates that relative transition probabilities were state-dependent, with high probability for higher clustering values. The normalized chi-squared value shows a non-monotonic relationship to clustering level which resembles that of of the top assembly / mean firing rate ratio metric, rising from a low level at $p_{in}^{EE}\approx0.18$, reaching a peak at $p_{in}^{EE}\approx0.23$ then falling to an intermediate level. It is notable that the form of the normalized chi-squared state-dependence graph has a close correspondence to the firing rate ratio / P_{in}^{EE} graph (Figure 2.7C).

The drawback of the normalization measure that we apply is the proportionally greater amplification of chance variation in simulations with lower transition counts, which could potentially inflate values. As a control against this potential bias, we repeated the analysis with an equalized dataset using only the first 900 transitions from each simulation, which produced results which were virtually identical to those produced using the full dataset (results shown in appendix in Figure 6.1). The results of the simulations at $P_{in}^{EE}=0.18$ and $P_{in}^{EE}=0.22$ also demonstrate that there is not a strong relationship between total transition count and normalized chi-squared value, as the simulations at these parameter levels have similar mean numbers of total transitions (1091 and 1392 respectively, Figure 3.2A) and levels of activation inequality (Figure 3.3A and B), but are at opposite ends of the range of normalized chi-squared values for state-dependence of transitions (Figure 3.3D).

3.3.3 Network structure and transition probability

The finding of deterministic transitions raises the questions of how activation is passed between assemblies and which features of network structure determine transition probability.

The activity of the spiking network here is similar to that of the networks of Potts

units created by Treves and others (Treves 2005; Russo et al. 2008). The Potts unit, or model, is a generalization of the Ising spin-state model, which in this case is taken to represent a local neural network which has multiple attracting states. In latching models, Potts units are connected together in a larger network in which multiple attractor states are embedded, each of which corresponds to a subset of units being in unit-specific attractor states, so that a global attractor state consists of the co-occurrence of a set of sub-attractors. An adapting term is added to the Potts units, so that the system falls out of each attracting state after the state is reached, so allowing the system to move to a different attractor.

While the network units of the Potts associative memory network are very different from the spiking units presented here, the qualitative activity of the Potts network for some parameter combinations is notably similar to the switching activation regime (Figure 3.4). The transitions between states in the Potts network are also probabilistic, and a mechanism underlying these transitions has been termed “latching”, a process involving activity being passed between correlated attractors, such that activity in one attractor is initiated by a subset of units of that attractor which are also active in the previous attractor state. This activity is associated with correlations between successively active attractor states and higher levels of activity within two attractors at the time at which the second attractor becomes more activated that the first.

We note here that our use of the term “latching” is slightly different from that used by Treves and colleagues, who described all transitions in their networks as latching transitions, making a distinction between latching transitions between correlated attractors, which proceeded as described in the previous paragraph, and latching transitions between non-correlated attractors, which occurred via a different mechanism. We use the term latching to denote only the first transition type, while we characterize other transitions as resulting from a process in which the currently-active assembly spontaneously falls out of the high-activity state due to fatigue or fluctuation of the recurrent synaptic drive, and another assembly spontaneously entering the high-activity state in its place.

Transitions between in attractor states in the spiking network can also be characterized by the activation state of the two attractors at the point of transition, suggesting that a similar latching process could determine transitions (Figure 3.4). In

order to test this hypothesis we measured various aspects of network structure and activity during the extended simulations. We hypothesized that, similarly to latching Potts networks (Russo et al. 2008), transition probabilities and transition cross-over levels would be higher for positively correlated assemblies.

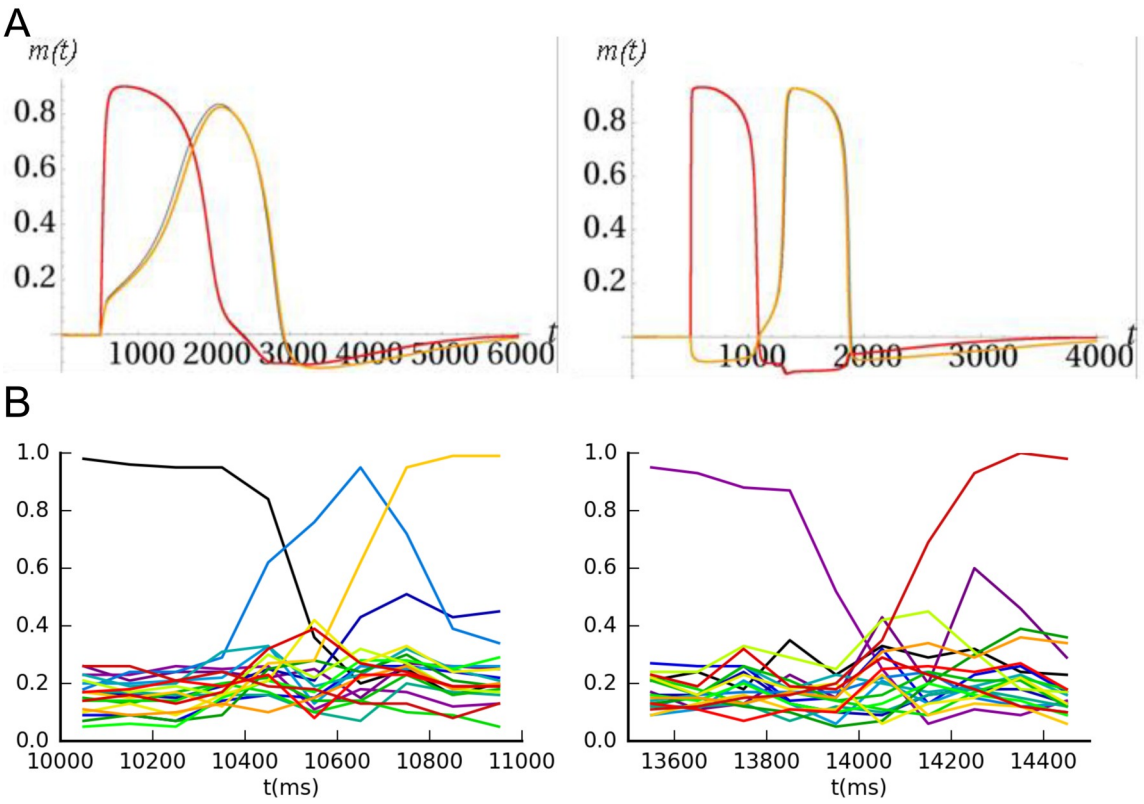

Figure 3.4: Latching vs non-latching transitions in an adapting Potts associative memory with comparison transitions in the clustered spiking network.
Latching transitions in the Potts network are driven by correlations between a currently active attractor state and second attractor state, which initiate a transition into the second state. (A) Activation state of two attractors undergoing a latching (left) and random / non-latching (right) transition in the Potts network. Figure taken from (Russo et al. 2008). (B) Comparable spontaneous transitions in the clustered spiking network.

3.3.3.1 Transition level

The "transition level", the activity level $A_i = A_j$ at which the identity of the most activated assembly transitions from assembly i to assembly j, was measured for each transition and the mean transition level for each transition identity t_{ij} was

calculated. If, between activations, a bin period occurred in which the activity of both assemblies dropped below $A = 0.7$ (such as in Figure 3.4B, right), both assemblies were assumed to have fallen to baseline activity levels during the transition. In these cases the transition level was recorded as the baseline activity level, which was set to be the mean activity level of sub-threshold assemblies over the duration of the simulation.

3.3.3.2 Assembly overlap, assembly projection, and Z-scores.

As a simple measure of the correlation between assemblies, the overlap between each assembly was calculated as the number of units which were members of both assemblies. The overlap between assembly a_i and a_j was therefore

$$overlap_{ij} = a_i^T a_j \tag{22}$$

where a_i is a binary vector of size N^E corresponding to all the excitatory units in the network, in which entries had a value of 1 when the corresponding unit was a member of assembly i and a value of 0 otherwise.

As the latching mechanism in the spiking network was hypothesized to function by transmitting excitation between assemblies, we also measured the summed excitatory projection weight from each assembly to each other assembly. The summed excitatory projection from assembly i to assembly j was calculated as

$$proj_{ij} = (Wa_i)^T a_j \tag{23}$$

where W is the excitatory weight matrix.

The motivation for measuring overlaps and projections was to make quantitative estimations of the extent to which transitions were determined by these structural features.

Transitions were hypothesized to be determined by some form of competition between sub-active assemblies. Assuming that only overlaps with or projections from the currently active assembly were relevant in determining transitions, we also standardized overlap and projection measurements so means and variances of values were equal in relation to each pre-transition assembly, i.e.:

$$z_{ij}^{overlap} = \frac{overlap_{ij} - \mu_i^{overlap}}{\sigma_i^{overlap}} \quad (24)$$

where μ_i is the mean and σ_i is the standard deviation of overlap values with assembly i, and

$$z_{ij}^{proj} = \frac{proj_{ij} - \mu_i^{proj}}{\sigma_i^{proj}} \quad (25)$$

where μ_i is the mean and σ_i is the standard deviation of summed excitatory projections from assembly i to each of the other assemblies.

3.3.3.3 Results

The relationships of interest were those between the structural measurements and the transition probabilities. Figure 3.6 shows plots of mean transition probabilities as a function of the $proj$ and $overlap$ metrics. Distributions of the corresponding structural metrics are shown underneath for comparison.

The correlation between the structural metrics and transition probability was also calculated at each level of p_{in}^{EE} at which switching activation activity occurred, using the set of measurements from all randomly-generated networks at each level as a single dataset. As initial analysis indicated that the distribution of transition probabilities was highly non-normal, the nonparametric Spearman rank correlation ρ between transition probability and the structural metrics was calculated at each level of clustering. The results are shown in Figure 3.5A.

It is clear from these measurements that transitions are strongly determined by network structure at higher clustering levels, with virtually zero transitions between assemblies with below-mean overlaps and projection strengths, and a supralinear rise in transition probability with higher overlap and projection strength values after the mean. The concentration of transition probabilities on the upper arms of the distribution of connectivity indexes indicates that transitions occur primarily across a small proportion of highly connected assembly conjunctions. The slightly steeper curve of the structure-probability plots at the $p_{in}^{EE} = 0.23$ level, which corresponds to

the clustering level at which transitions were most state-dependent as measured by the normalized chi-squared metric, suggests that the concentration of transitions at the upper end of inter-assembly connectivity scale (i.e. the supralinear non-linearity) is greater at intermediate levels of clustering. However it is notable that the rank correlation coefficient ρ does not reflect this difference in distribution or the difference in normalized chi-squared values (Figure 3.5A).

The relationship between structure and transition probability is notably weaker at the $p_{in}^{EE}=0.18$ level, as indicated by both the flatter structure-transition probability plots and the non-significant correlation coefficient (Figures 3.6 and 3.5). The mean transition level of $A=0.25$ at this level of clustering indicates that transitions were mostly the non-latching type, as compared to $A>0.7$ for levels of clustering of $p_{in}^{EE}\geq 0.21$ (Figure 3.5B).

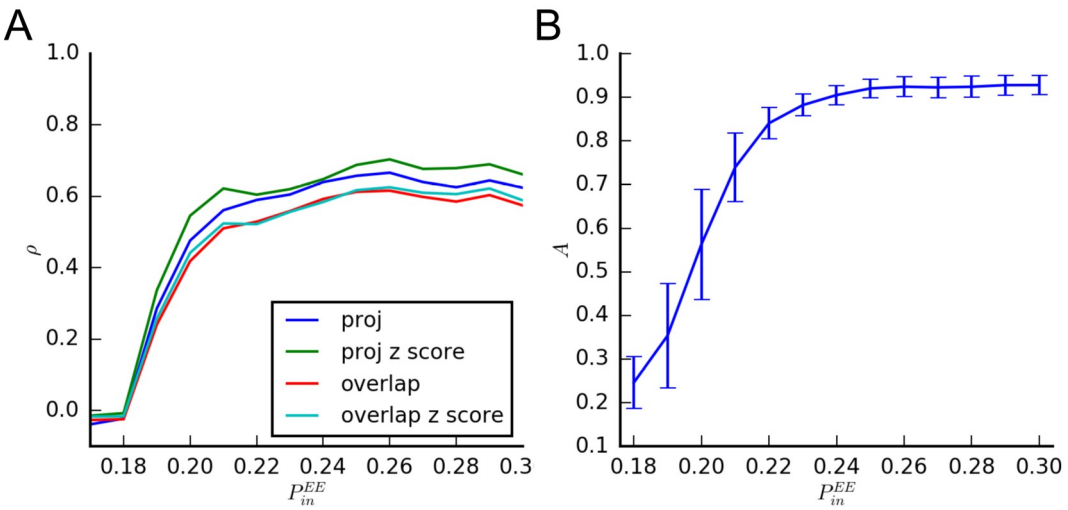

Figure 3.5: Structure dynamics-correlations as a function of clustering.
(A) Spearman ρ rank correlation between transition probability and network structure measures as a function of clustering. ρ is non-significant at $p_{in}^{EE}=0.17-0.18$ and significant at the P<0.0001 level otherwise. (B) Mean (SD) transition level as a function of clustering.

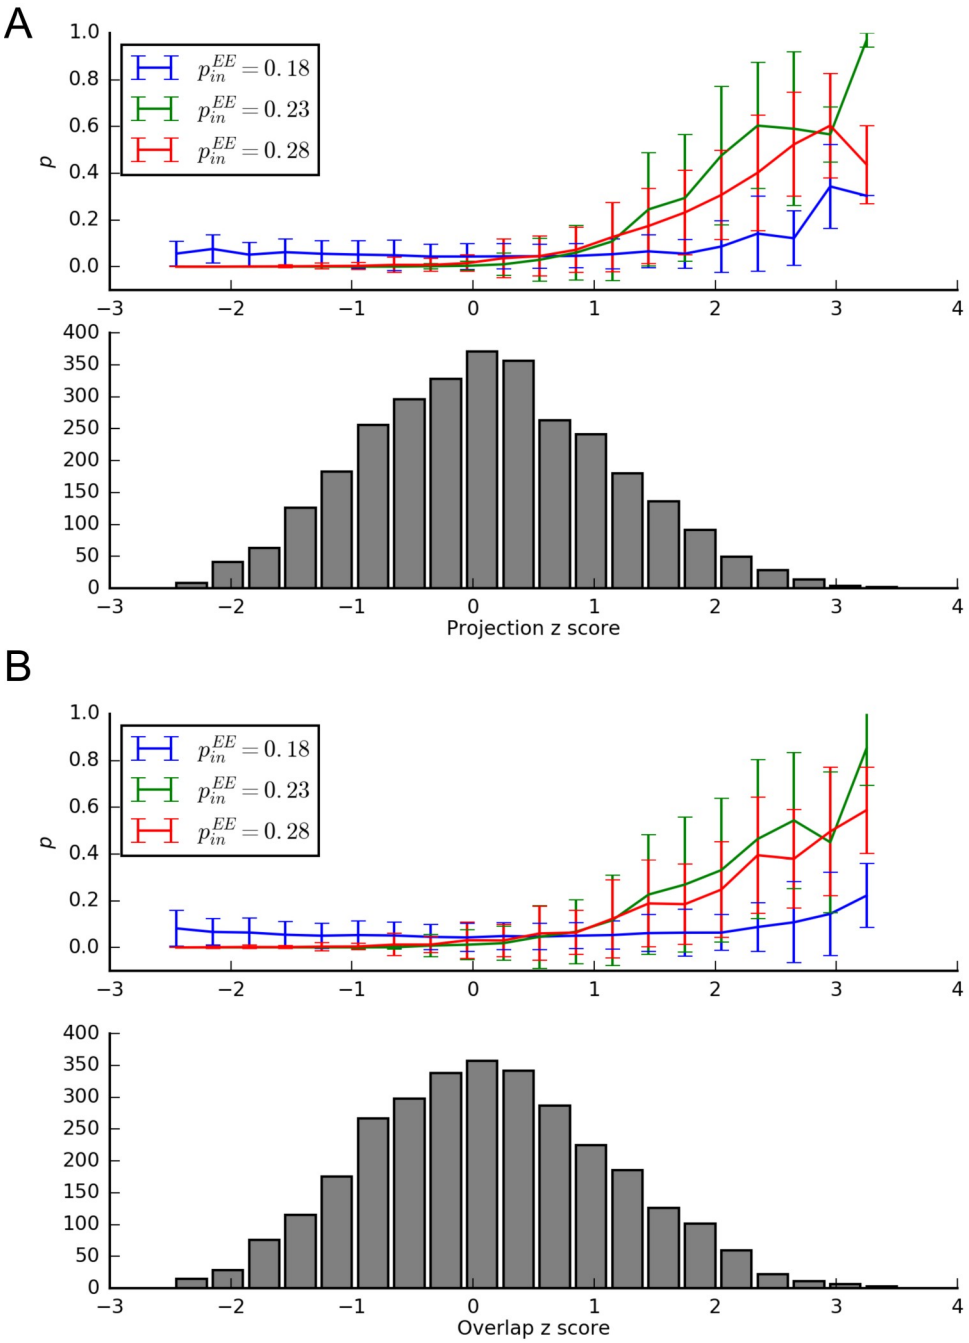

Figure 3.6: Transition probability as a function of network structure and distributions of structural metric values.
(A) (Top) Mean (SD) transition probability as a function of standardized excitatory projections. Values calculated by binning transition identities by excitatory projection values then calculating mean (SD) probabilities. (Bottom) Distribution of standardized excitatory projections at $p_{in}^{EE} = 0.23$ *. Distributions at other levels of* p_{in}^{EE} *are similar (data not shown). (B) (Top) Mean (SD) transition probability as a function of standardized overlap values. (Bottom) Distribution of standardized overlap values.*

3.4 Discussion

The results demonstrate that balanced spiking networks with connections clustered into overlapping assemblies produce predictable sequences of activity, at the level of firing rates of groups of neurons. These sequences appear to be generated by a relatively intuitive mechanism which has clear relationships with network connectivity. Transitions between states are stochastic and state-dependent, and can be described as a Markov chain of at least first-order.

Many functional applications have been proposed for systems that generate sequences of distinct states, and multiple examples of sequence generation are found in biology, as detailed in the introduction. The switching activation model presented here has face similarity with several of these, particularly those which have been associated with sequences of assembly-like firing-rate states (e.g. Seidemann et al. 1996; Jones et al. 2007).

The correlations between transition probability and both assembly correlations and excitatory connectivity, and the qualitative similarity of assembly activity during transitions, indicate that activity is propagated and transferred between assemblies by a mechanism similar to the latching mechanism studied by Treves and others (Treves 2005; Russo et al. 2008; Linkerhand and Gros 2013). There are some significant features of the latching Potts networks that are not reproduced in our study, in particular the specification that global attractor states are formed from the conjunction of multiple local attractor states in the subnetworks that are represented by the Potts units. Despite these structural differences, our results provide evidence that a latching transition mechanism is possible with relatively biophysically realistic spiking neurons in a network with sparse clustered connectivity.

The sequences produced by latching networks of Potts units have been specifically associated with language generation, in a framework in which attractors correspond to words, and the Markov chain quality of transitions corresponds to the restriction of word order according to syntactic rules. One particularly salient application of Markov chain modeling in the current context is that of song generation in songbirds (Katahira et al. 2011). Although custom neural models of birdsong generation have been proposed before (Long, Jin, and Fee 2010; Hanuschkin, Diesmann, and

Morrison 2011), the Markov sequence generation in this study suggests a generation model with relatively realistic spiking neurons and relatively few structural assumptions.

The latching mechanism can be considered in terms of the linking together of multiple correlated but discrete quasi-attractor states, such that each state is close to the basin of attraction of another attractor state, so that when the system starts to fall out of one attracting state due to an instability the system moves towards another semi-stable attracting point. This kind of dynamics, when transitions between multiple quasi-attractors are stochastic and dynamics in some parts of the system are chaotic, has been termed “chaotic itinerancy”, and has been proposed to be a form of dynamics found in the brain. In this context, the quasi-attractors could correspond to items of memory and the spontaneous transitions could function as a form of search mechanism, testing different behaviors for effectiveness or stored percepts for matches to sensory input (Tsuda 2001; Friston, Breakspear, and Deco 2012).

The overlap and connectivity measures had a Spearman correlation with transition probability of approximately $\rho = 0.6$ at higher clustering levels, and so did not account entirely for transitions and first-order transition probabilities. For various reasons, these measures are only a partial measurement of the structural factors which determine transition probabilities. For example, the relative 'advantage' that a specific assembly has due to structural connectivity with the currently active assembly is also dependent on the structural connectivity of other assemblies with the currently active assembly, which is not accounted for by the overlap and connectivity measures used here. Chance differences in indirect connectivity between assemblies via the inhibitory population are also likely to be relevant.

In addition to differences in transition probabilities resulting from differences in connectivity, it is also possible that temporally extended processes have a role. One example might be the lingering effect of differences in the slow adaptation current w resulting from previous assembly activations. This could cause an interaction between previously-active assemblies and current transition probabilities. In this case the system would be better described as a second order or higher Markov process, rather than the first order Markov process implicitly assumed by the transition matrices calculated here. Further analysis would be necessary to test

these hypotheses.

The inherent stochasticity of transitions is likely partly driven by fluctuations in synaptic drive due to the irregular and presumably chaotic spiking. This chaos will be explored further in the final section of this thesis. Finally, a proportion of the missing correlation will also result from chance effects due to finite sampling of the stochastic transitions.

3.4.1 Future outlook

The structure of the networks studied here was determined largely at random. Despite this, chance inequalities in connectivity and correlation structure produced systems with meaningfully reliable temporal dynamics. It should be possible to intentionally produce networks which produce an arbitrary pattern of activity by purposefully setting or modifying connection structure, such as by modifying excitatory connectivity between two assemblies in order to increase the probability of sequential activation of these assemblies. Further studies could test how this might be produced in-vivo by neural plasticity processes.

4 Chaos in clustered networks

As discussed in the introduction, the degree of chaos, roughly defined as the divergence rates of the trajectories of nearby system states, is a computationally important quality of a neural system. Analysis of the chaos of neural systems is challenging, and analytically rigorous investigations have only so far been performed on highly simplified systems (e.g. Wallace, Maei, and Latham 2013; Monteforte and Wolf 2010). Numerical methods can be used with more complex systems (Jahnke, Memmesheimer, and Timme 2009; Legenstein and Maass 2007), and there have even been efforts to measure the chaos of biological neural networks (London et al. 2010). The conclusions of these studies have been consistent, in that they conclude that networks with connectivity rates similar to those in the cortex are highly chaotic, apart from some special cases such as simulated purely inhibitory networks with delta synapses (Jahnke, Memmesheimer, and Timme 2008).

However it is also notable that the chaos of large artificial neural systems has only been measured for networks with random or near-random connectivity, to the best of our knowledge. In this section we therefore investigate the divergence or convergence rates of the neural network types so far discussed in this thesis, and so investigate the qualitative and quantitative effects of clustering on chaos in balanced networks. The neuron and synapse models that we use in this study are of the level of complexity that requires analysis with numerical techniques, and so we employ perturbation methods similar to those used in previous studies (Legenstein and Maass 2007). We hypothesized that, although spike times would likely remain chaotic, the switching activation dynamics in clustered networks could lead to a slower divergence of trajectories.

4.1 Trajectory divergence of adjacent states in clustered balanced networks

4.1.1 Methods

We numerically measured the sensitivity of the system to perturbations by running pairs of otherwise identical simulations in which one was minimally perturbed, and

measuring the differences in system states between the two systems after the occurrence of the perturbation.

Specifically, we randomly generated network instances and initial conditions according to the specifications and parameters in section 2, with both no clustering and 20 overlapping assemblies generated by the equal assemblies method with $p_{in}^{EE}=0.19$, $p_{in}^{EE}=0.23$ and $p_{in}^{EE}=0.28$. Simulations were run as matching pairs, with the difference that the occurrence of a spontaneously-occurring single spike in one of the simulations was delayed by the simulation step size (0.1ms) by "freezing" the state of the neuron producing this spike for one simulation step immediately before the occurrence of this spike. This perturbation was introduced at approximately $t=3000\,ms$, after initial transients had subsided. The difference between the states of the paired simulations was measured by summing the absolute differences of the values of the membrane potential variable V in the excitatory neurons of each simulation. This process was repeated for 50 iterations at each level of p_{in}^{EE}.

4.1.2 Results: three-stage divergence of trajectories

The states of the perturbed and unperturbed systems diverged to saturation values in 46/50 ($p_{in}^{EE}=0.1$), 47/50 ($p_{in}^{EE}=0.19$), 48/50 ($p_{in}^{EE}=0.23$) and 49/50 ($p_{in}^{EE}=0.28$) of the simulations, and in the remaining simulations the two systems quickly converged back to identical states after the initial perturbation. In those simulation pairs that did not diverge, it was found that no further spike times were shifted at all after the induced perturbation. This is likely to be a result of limited machine numerical resolution and of the spike time quantization inherent to the simulation software, which calculated the network state on a 'grid' of successive time steps. Each excitatory neuron projected to approximately 100 excitatory and 125 inhibitory neurons, so it is highly likely that a change in spike time in system of infinite resolution would lead to a change of the membrane V and therefore next spike time of at least one neuron which would spike at a later point in the simulation.

In those simulation pairs in which system states diverged, the divergence showed a distinctive form, an example of which is shown in Figure 4.1A. States diverged at a relatively gradual rate for a short period after the induced perturbation, during which

the order of neuron firing of the two simulations is approximately equal, but with some spikes occurring at slightly differing times, no greater than a few milliseconds. This continued until activity differed by the occurrence of a spike in one simulation which was not matched within a short time period by the occurrence of a spike from the corresponding neuron in the other simulation, an event we refer to as a “spike mismatch”. After this point, spiking activity rapidly diverged to a saturation level of dissimilarity in ~20ms. Mean V divergence as a function of clustering is shown in Figures 4.3B, C and D).

The times which elapsed between the perturbation and the first spike mismatch were relatively short, with a mean (SD) of 25.43(24.1)ms, and did not substantially differ as a function of clustering (Figure 4.3A). Spike mismatches were detected by identifying spikes in either simulation which were not matched by a spike from the corresponding neuron in the paired simulation within a ±5ms time window.

The membrane potential V of neurons in a balanced network has a low effective resistance and consequently a low effective time constant, and fluctuates rapidly as a function of spike input (Renart, Brunel, and Wang 2004). The state difference value ΔV therefore mainly reflects the difference of spike times when calculated from the raw membrane potentials V. To proportionally reduce the influence of individual spike times and increase the influence of firing rates on ΔV, we therefore additionally calculated ΔV after convolving the V of each excitatory neuron with a Gaussian filter with $\sigma = 30\,ms$.

When calculated on the raw, unconvolved membrane potentials, the state difference measure ΔV showed a similar response to perturbation at all tested levels of clustering p_{in}^{EE} apart from $p_{in}^{EE} = 0.28$ (Figure 4.3B). The occurrence of a spike mismatch caused a fast and linear divergence of system states to near-saturation level in approximately 20ms in networks with p_{in}^{EE} values of 0.1, 0.19 and 0.23 (Figure 4.3C), indicating that clustering does not *per se* alter the degree of chaos in spike times. The divergence of membrane potentials in networks with $p_{in}^{EE} = 0.28$ showed a slower, qualitatively different form. As described in section 2.2.3, at $p_{in}^{EE} = 0.28$ excitatory neurons across the network fire synchronously at approximately 50Hz, which appears to both retard the divergence of states after a

perturbation and exhibit a large effect on ΔV, in that ΔV is greater when firing in the perturbed and unperturbed networks is out of phase (Figure 4.1B). The relatively gradual divergence of system states shown in Figure 4.3B and C is therefore partially a result of the gradual desynchronization of the fast oscillations of the two networks. Convolving the V of each neuron removed the fast oscillation and also the effect of its desynchronization. At all levels of clustering, convolution reduced fast fluctuations caused by spike firing and reduced the effect of these on ΔV (Figure 4.3D).

In the clustered networks, this proportional reduction of the effect of individual spikes made clearer the slower divergence of trajectories at the level of neuron firing rates, which appeared to reflect the divergence of the assembly attractor state sequences of the paired simulations. A representative example of the effect of perturbation on a pair of clustered network simulations is shown in Figure 4.2. In the aggregate, the divergence of neuron firing rates in highly-clustered networks occurred over a timescale of 3000-5000ms, 2 orders of magnitude greater than the timescale of the divergence of spike patterns after the occurrence of a spike mismatch (compare Figure 4.3D with Figure 4.3C).

The chaotic divergence of nearby trajectories in clustered spiking networks therefore seems to pass through 3 distinct stages: first the gradual divergence of spike times but approximate preservation of the sequence of spikes across the network until the first spike mismatch, then a fast divergence of spike sequences, then a slower but stochastic divergence of assembly activations and assembly activation sequences.

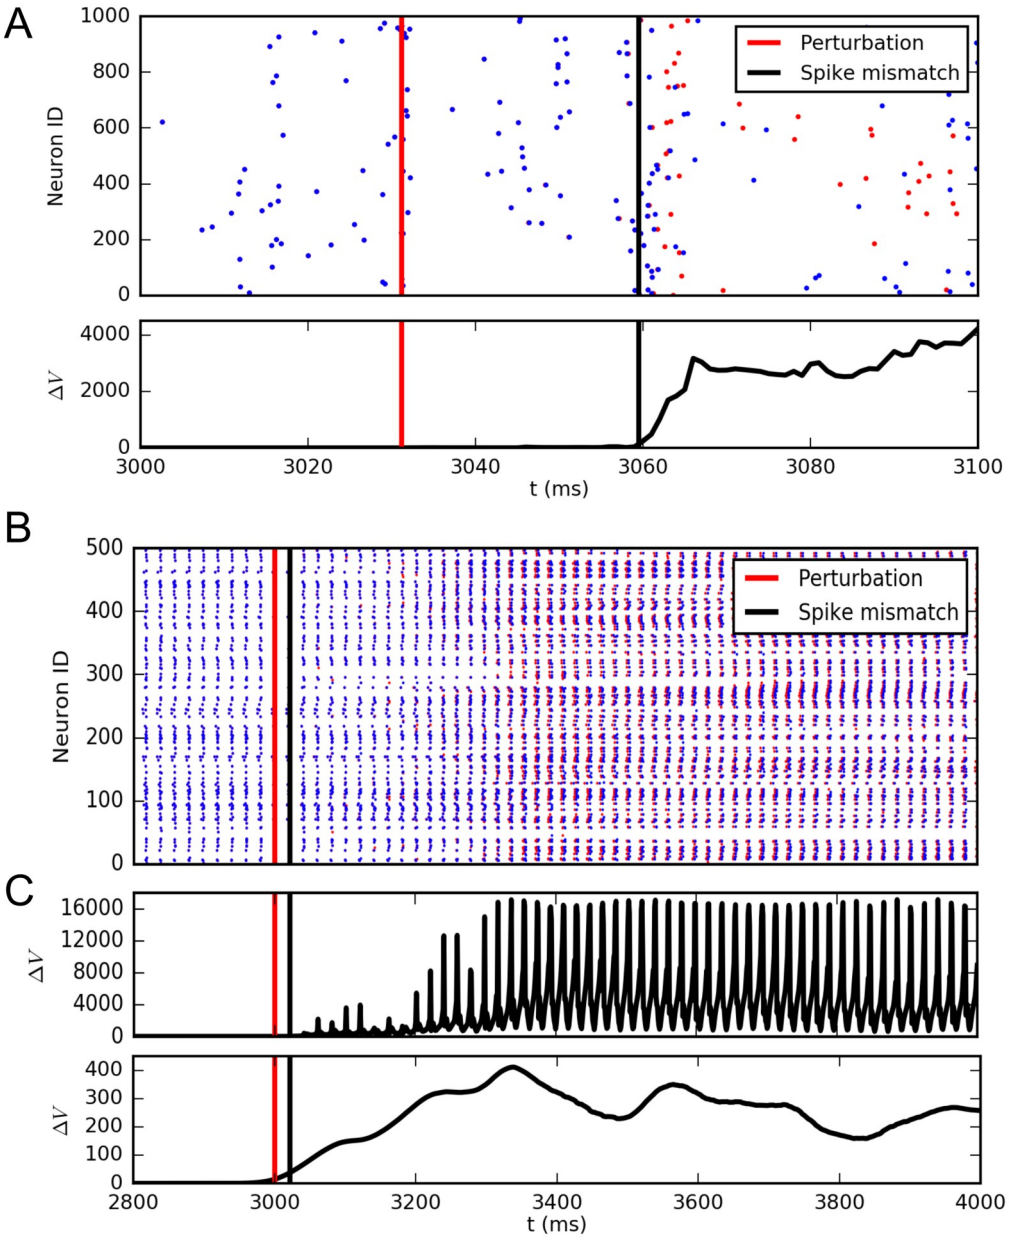

Figure 4.1: Effect of minimal spike time perturbations on spiking activity
(A) Representative divergent activity trajectories of two otherwise identical network simulations in which one is minimally perturbed. Red and black vertical lines indicate times of perturbation and spike mismatch events. (Top) raster plot of excitatory spike activity, in which spikes from the paired networks are plotted as red and blue points overlayed on the same plot. (Bottom) summed difference between membrane potential variable V of excitatory neurons. (B) Perturbation raster plot of network clustered with $p_{in}^{EE}=0.28$ showing synchronous firing.(C) Summed difference of membrane potentials ΔV corresponding to the spike activity shown in (B). (Top) ΔV calculated on raw values of V. (Bottom) ΔV calculated on values of V convolved with a Gaussian kernel with $\sigma=30\,ms$.

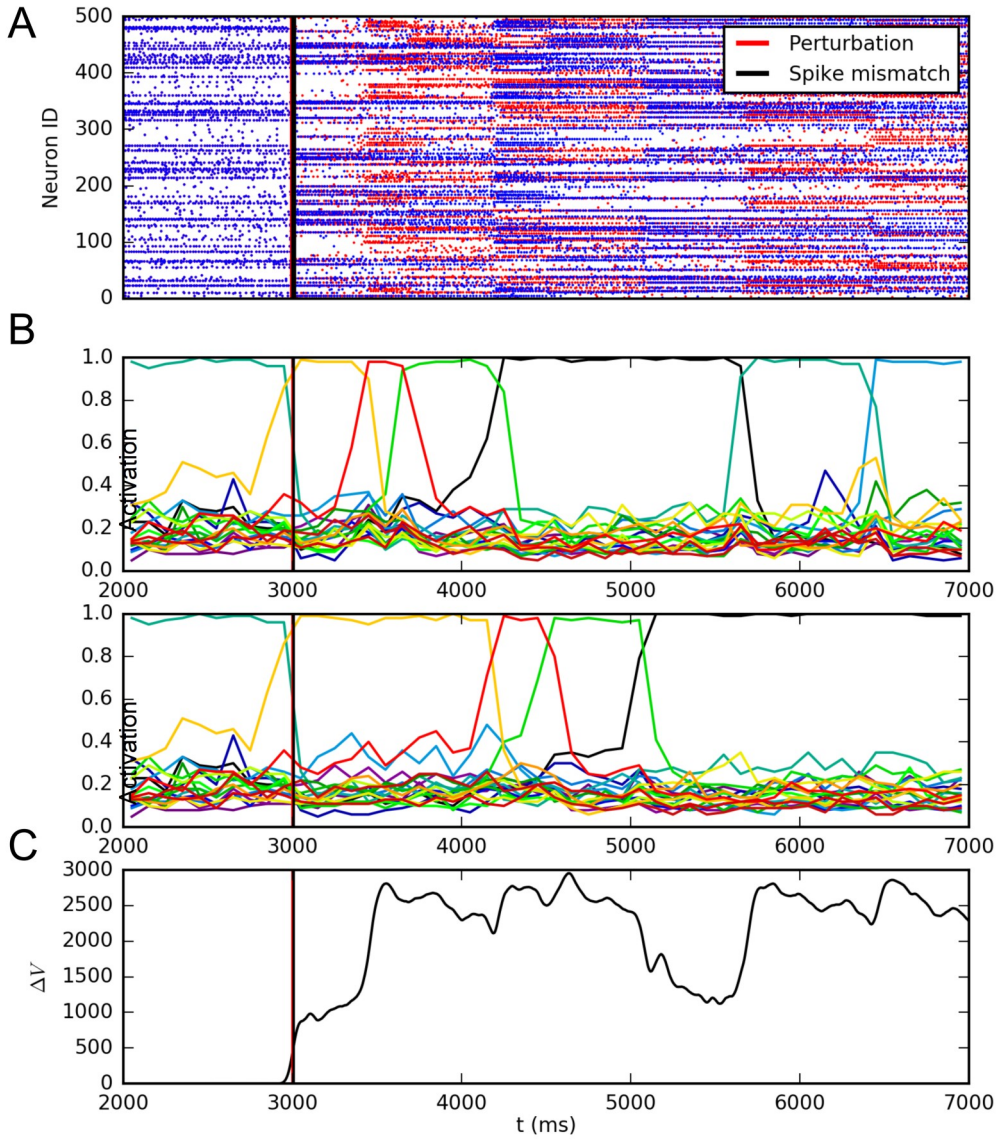

Figure 4.2: Divergence of trajectories of adjacent states in an example clustered network.

All plots correspond to the same simulation data, in which adjacent states in identical networks were created with the perturbation method. (A) Raster plot of excitatory spike activity, in which spikes from unperturbed and perturbed networks are plotted as red and blue points overlayed on the same plot. (B) Assembly activation measure A plots for paired networks. (C) System state difference measure ΔV calculated on V of excitatory neurons convolved with a Gaussian with $\sigma = 30\,ms$. Note the correspondence of ΔV to the agreement or non-agreement of the attractor states of the paired simulations.

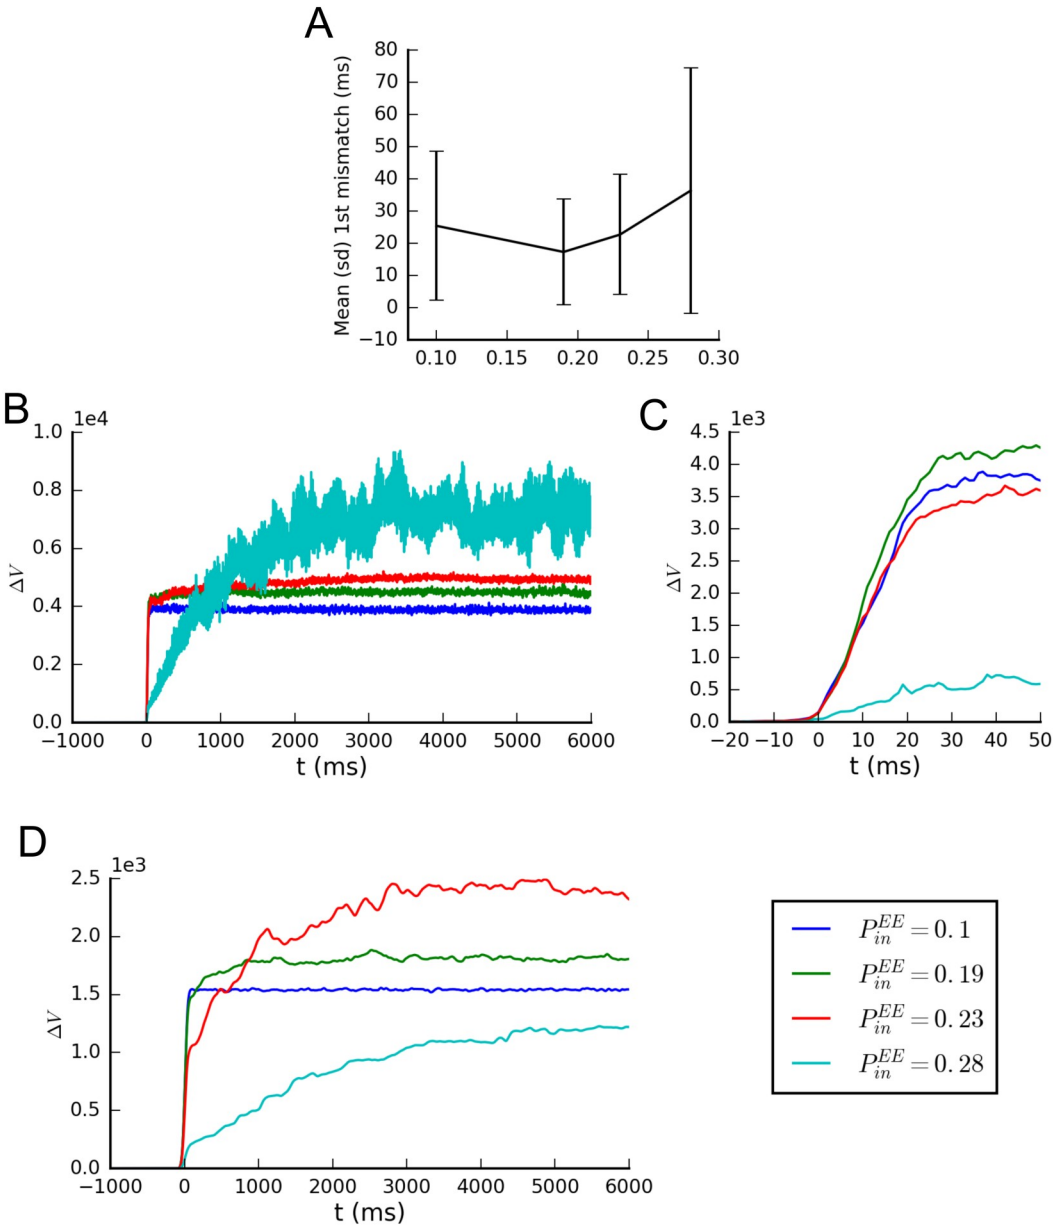

Figure 4.3: Divergence of adjacent networks states tested with perturbation.
(A) Mean (SD) time between induction of minimally different states via perturbation and occurrence of the first "spike mismatch", as a function of network clustering. See text for details. (B) Mean state space difference measure ΔV *over 50 randomly generated iterations for different levels of clustering parameter* p_{in}^{EE}*, with t=0 set as the time of the first spike mismatch. (C) Close-up of the data in (B) around the spike mismatch event. (D) Mean state space difference measure* ΔV *over 50 randomly generated iterations, when* V *series of individual neurons is convolved with a Gaussian filter prior to calculation of* ΔV *.*

5 Conclusion and future outlook

5.1 Recapitulation of main results

We have described here a series of simulation studies which synthesize and extend previous theoretical work regarding spiking neurons, network attractor dynamics and neural computation. We have confirmed previous findings that the addition of clustered groups induces switching activation dynamics in spiking balanced networks. Total within-cluster recurrent connectivity is a critical property in determining dynamics, and setting this property in the case of correlated (i.e. overlapping) clusters requires special measures. The emergence of clustering sufficient to induce switching activation is associated with the emergence of a distinct group in the eigenvalue spectrum of the connection matrix, with the mean real value of of this group having a close relationship with network dynamics. Spiking activity in clustered networks reproduces key features of biological neural networks, particularly firing rate variability at the single neuron level and oscillatory activity in the sub-4KHz delta range. The sequence of spontaneous activations is probabilistic and can be highly state-dependent, and the overall degree of state-dependence and the probability of individual transitions is dependent on structural features of the network, particularly the degree of clustering and the degree of correlation between clusters. The state sequence dynamics and their generation mechanism are comparable to those in previous models with simplified non-spiking units. We demonstrate for the first time the production of these dynamics in a network of spiking units, thereby supporting the hypothesis that similar mechanisms could generate sequences in the brain. Finally, we replicate previous findings of chaos in spiking balanced networks, and find that divergence of minimally separated states in spiking networks goes through two distinct stages with differing timescales, to which switching activation dynamics add a third stage in which the assembly-level activation states of the systems diverge.

5.2 Correspondence to anatomy

This study has been primarily a theoretical simulation study with the aim of revealing general neural computational principles rather than replicating specific systems.

Various simplifications were made for the sake of minimizing computational load and design complexity, particularly in terms of homogeneity of neurons within the excitatory and inhibitory classes, and ranges and distributions of cellular and connection parameter values. Heterogeneity of neuron types and parameter values has been found to be increase statistical realism of dynamics in more anatomically realistic models of the cortex (Hass et al. 2016). It is plausible that the crucial elements of the networks in this study, in particular the clustered assemblies, could be implemented in networks with a more complex structure and range of elements, but this would necessarily be more challenging to design. One possible way of testing hypotheses regarding the in-vivo existence of structures and dynamics similar to those show here is to evaluate experimental data in terms of their fit to potential structure and dynamics in clustered networks (Klinshov et al. 2014).

Despite the simplifications, the clustered network model here contains 30+ parameters, of which only the clustering parameters p_{in}^{EE} and R^{EE} were systematically explored. Exploration of the full parameter space would be a challenging task that could benefit from the extension of previous mean-field analysis of working memory models (Renart, Brunel, and Wang 2004). A potentially interesting form of parameter space analysis would be to test the effect of parameter modifications that mimic experimentally measured effects of gene alleles or neuromodulatory transmitters such as dopamine. Previous simulation studies have used this approach to investigate the effect of neurotransmitters on spiking working memory models, in the context of behavioral effects and functions of neurotransmitters (Durstewitz 2007; Brunel and Wang 2001). The Markov transition description of network dynamics used for our network can also be interpreted in terms of regularity or diversity of dynamics, and it would be relatively simple to perform an equivalent simulation study of quantitative neurotransmitter effects on these dynamics.

We are essentially agnostic about where in the brain dynamics similar to those here might be found, but comparisons can be made with existing experimental neural and behavioral data, such as the HVC cortex in songbirds (Fujimoto, Hasegawa, and Watanabe 2011), PFC activity during delayed localization tasks (Seidemann et al. 1996) and gustatory cortex activity during taste stimuli (Jones et al. 2007)(see also section 3.4).

5.3 Network size and scaling

The question of where in the brain might use mechanisms similar to those shown here is related to the issue of how and to what degree clustered networks could be scaled. The used network size of 1250 neurons has a similar order of magnitude as a single mammalian cortical column (Helmstaedter et al. 2007) but many times smaller than a cortical area defined by other criteria, such as a Brodmann area. This network size was chosen to be large enough to approximate physiology while computationally undemanding enough to allow extended duration simulations to be completed in a moderate amount of time using available hardware.

Changing the absolute size of a neural simulation while maintaining quantitative and qualitative activity typically requires adjustment of other parameters in order to compensate for the increased number of inputs to each unit. Connection weights J in balanced networks are typically scaled by $1/\sqrt{K}$, where K is the number of inputs of the neuron population concerned, in order to maintain the synaptic input variance $Var(I_{synapse}) = J^2 Kr$, upon which the firing rate depends, assuming the balance condition is maintained (Renart, Brunel, and Wang 2004). Litwin-Kumar (2013) noted that scaling a clustered network in such a way that maintains dynamics is complicated by the requirement of maintaining the ratio of within-assembly and without-assembly input, and that the parameter sensitivity of the existence of attractor or quasi-attractor states increases with absolute assembly size, which makes parameter-tuning more difficult.

The stochastic switching activation dynamics described in this study are also partly dependent on finite size effects that necessarily reduce as network size becomes larger. First, some of the stochasticity in the timing and identity of state transitions is driven by fluctuations in synaptic drive within and between assemblies due to finite size effects in the summation of the spiking of single neurons. Cumulative firing rate fluctuations should approximately obey $1/\sqrt{K_{in}}$ and so decrease with network size, which should lead to more predictable state transitions, assuming all other sources of variation remain equal (Deco and Rolls 2006).

A second issue is that, if assembly sizes remain a fixed proportion of the total network size and overlaps between assemblies are random, as they are in the

overlapping assemblies described here, the standard deviation of relative differences in overlap between assemblies, which obeys $\sqrt{N^{E}_{assembly}\,p(1-p)}$ where p is the size of each assembly as a proportion of the total neuron population, will increase with network size in absolute terms but decrease as a proportion of the mean overlap, which is determined by $N^{E}_{assembly}\,p$. Assuming that differences in transition probabilities are determined more by proportional rather than absolute differences in between-assembly connectivity, increasing network size would also lead to a reduction in differences in transition probabilities and an increase in transition entropy, unless greater differential connectivity was set intentionally. It is notable that the signs of the predicted finite size effects of connectivity and fluctuations in synaptic drive on transition predictability are opposite. To what extent they would counteract each other in practice remains to be determined.

5.4 Final remarks

Despite extensive investigation, understanding the human brain remains one of the most difficult challenges in contemporary science. The essential problem is one of complexity, as the brain contains vast numbers of neurons in a structure which is highly detailed at every spatial and functional level (Abbott 2013). This study has considered dynamics at the level of local neural networks, where one of the greatest puzzles is understanding how large numbers of comparatively simple units collectively produce highly complex but useful dynamics. The sheer number of constituent parts makes this particularly challenging, to the extent that even entirely artificial systems that use simplified neuron models require inference from simplified descriptions of dynamics to aid understanding (Sussillo and Barak 2013). Understanding the principles which produce cognition as an emergent property of neural networks is likely to require a long process of proposing, testing, extending and combining theoretical building blocks in order to make a greater whole. We hope that the work presented in this study will become one of these blocks.

Acknowledgements

I am indebted to the many people who have helped me in various ways and made the completion of this thesis possible. First thanks have to go to my supervisor Daniel Durstewitz for guiding the course of the research at its beginning and at important points along the way, and for being an endless repository of apposite analysis methods. I am also very appreciative of the other past and current members of the computational neuroscience group, who provided illuminating discussions, laughs and coffee and cake at critical moments: Emili, Charmaine, Sebas, Joachim, Loreen, Olga, Eleonora, Thomas, Sven, Florian, Georgia, Hazem, Carla, Sadjad and Dominic. Of these, Eleonora, whose own research was an important inspiration for this thesis, deserves special mention.

Additionally, and too numerous to mention, I would like to thank the other members of the Bernstein Center for Computational Neuroscience Heidelberg-Mannheim for providing a stimulating and memorable social and scientific environment. Various people at the Zentralinstitut have also provided some form of professional assistance, particularly Luanna, Hengyi, Anais, Christian and Heike.

In addition to those above, various friends have also provided social and practical support during my time in Mannheim, as well as my family, and most of all Astrid, who has supported me in countless ways and will always have my gratitude.

6 Appendix

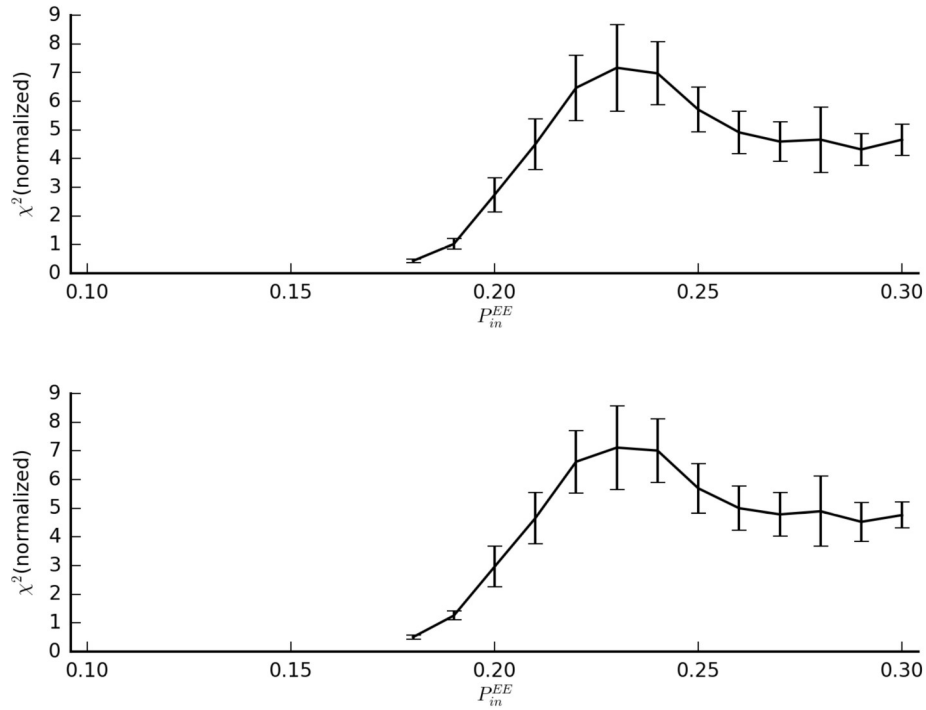

Figure 6.1: Control results for the effect of sample sizes in calculating the normalized chi-squared statistic for state-dependent transition probabilities.
(Top) Mean (SD) when statistic is calculated using all transitions occurring within 30 minutes of simulated time. (Bottom) Mean (SD) when statistic is calculated using only the first 900 transitions that occur during the simulation.

7 Bibliography

Abbott, B Y Alison. 2013. “Neuroscience: Solving the Brain.” *Nature* 499 (7458): 272–74. doi:10.1038/499272a.

Abeles, Moshe. 2009. “Synfire Chains.” *Scholarpedia* 4 (7): 1441. doi:10.4249/scholarpedia.1441.

Balaguer-Ballester, Emili, Christopher C Lapish, Jeremy K Seamans, and Daniel Durstewitz. 2011. “Attracting Dynamics of Frontal Cortex Ensembles during Memory-Guided Decision-Making.” *PLoS Computational Biology* 7 (5): e1002057. doi:10.1371/journal.pcbi.1002057.

Branco, Tiago, and Kevin Staras. 2009. “The Probability of Neurotransmitter Release: Variability and Feedback Control at Single Synapses.” *Nature Reviews. Neuroscience* 10 (5): 373–83. doi:10.1038/nrn2634.

Braun, Jochen, and Maurizio Mattia. 2010. “Attractors and Noise: Twin Drivers of Decisions and Multistability.” *NeuroImage* 52 (3). Elsevier Inc.: 740–51. doi:10.1016/j.neuroimage.2009.12.126.

Brette, Romain, and Wulfram Gerstner. 2005. “Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity.” *Journal of Neurophysiology* 94 (5): 3637–42. doi:10.1152/jn.00686.2005.

Broyd, Samantha J, Charmaine Demanuele, Stefan Debener, Suzannah K Helps, Christopher J James, and Edmund J S Sonuga-Barke. 2009. “Default-Mode Brain Dysfunction in Mental Disorders: A Systematic Review.” *Neuroscience and Biobehavioral Reviews* 33 (3): 279–96. doi:10.1016/j.neubiorev.2008.09.002.

Brunel, Nicolas. 2000. “Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons.” *Journal of Computational Neuroscience* 8 (3): 183–208.

Brunel, Nicolas, and Xiao-Jing Wang. 2003. “What Determines the Frequency of Fast Network Oscillations with Irregular Neural Discharges? I. Synaptic Dynamics and Excitation-Inhibition Balance.” *Journal of Neurophysiology* 90 (1): 415–30. doi:10.1152/jn.01095.2002.

Brunel, Nicolas, and XJ Xiao-jing Wang. 2001. “Effects of Neuromodulation in a Cortical Network Model of Object Working Memory Dominated by Recurrent Inhibition.” *Journal of Computational Neuroscience*, 63–85.

Büsing, L, Benjamin Schrauwen, and Robert Legenstein. 2010. “Connectivity, Dynamics, and Memory in Reservoir Computing with Binary and Analog

Neurons.” *Neural Computation* 1311: 1272–1311.

Carrillo-Reid, Luis, Jae-Eun Kang Miller, Jordan P Hamm, Jesse Jackson, and Rafael Yuste. 2015. “Endogenous Sequential Cortical Activity Evoked by Visual Stimuli.” *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 35 (23): 8813–28. doi:10.1523/JNEUROSCI.5214-14.2015.

Churchland, Anne K., R. Kiani, R. Chaudhuri, Xiao Jing Wang, Alexandre Pouget, and M. N. Shadlen. 2011. “Variance as a Signature of Neural Computations during Decision Making.” *Neuron* 69 (4). Elsevier Inc.: 818–31. doi:10.1016/j.neuron.2010.12.037.

Compte, Albert, Christos Constantinidis, Jesper Tegner, Sridhar Raghavachari, Matthew V Chafee, Patricia S Goldman-Rakic, and Xiao-Jing Wang. 2003. “Temporally Irregular Mnemonic Persistent Activity in Prefrontal Neurons of Monkeys during a Delayed Response Task.” *Journal of Neurophysiology* 90 (5): 3441–54. doi:10.1152/jn.00949.2002.

Compte, Albert, Maria V Sanchez-Vives, David a McCormick, and Xiao-Jing Wang. 2003. “Cellular and Network Mechanisms of Slow Oscillatory Activity (<1 Hz) and Wave Propagations in a Cortical Network Model.” *Journal of Neurophysiology* 89 (5): 2707–25. doi:10.1152/jn.00845.2002.

Constantinidis, C., and M. A. Steinmetz. 1996. “Neuronal Activity in Posterior Parietal Area 7a during the Delay Periods of a Spatial Memory Task.” *J Neurophysiol* 76 (2): 1352–55.

Cossart, Rosa, Dmitriy Aronov, and Rafael Yuste. 2003. “Attractor Dynamics of Network UP States in the Neocortex.” *Nature* 423 (6937). Nature Publishing Group: 283–88. doi:10.1038/nature01614.

Cressie, Noel, and Timothy R C Read. 1989. “Pearson’s X and the Loglikelihood Ratio Statistic G : A Comparative Review.” *International Statistical Review* 57 (1): 19–43.

D’Esposito, Mark, and Bradley R Postle. 2015. “The Cognitive Neuroscience of Working Memory.” *Annual Review of Psychology*, no. 66(September). Annual Reviews: 115–42. doi:10.1146/annurev-psych-010814-015031.

Dayan, Peter, and L F Abbott. 2001. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Edited by Peter Dayan and L F Abbott. Computational Neuroscience. The MIT Press. doi:10.1016/j.neuron.2008.10.019.

Deco, Gustavo, and Edmund T. Rolls. 2006. “Decision-Making and Weber’s Law: A Neurophysiological Model.” *European Journal of Neuroscience* 24 (3): 901–16. doi:10.1111/j.1460-9568.2006.04940.x.

Durstewitz, Daniel. 2007. “Dopaminergic Modulation of Prefrontal Cortex Network Dynamics.” *Monoaminergic Modulation of Cortical Excitability*.

———. 2009. “Implications of Synaptic Biophysics for Recurrent Network Dynamics and Active Memory.” *Neural Networks : The Official Journal of the International Neural Network Society* 22 (8). Elsevier Ltd: 1189–1200. doi:10.1016/j.neunet.2009.07.016.

Durstewitz, Daniel, and Gustavo Deco. 2008. “Computational Significance of Transient Dynamics in Cortical Networks.” *The European Journal of Neuroscience* 27 (1): 217–27. doi:10.1111/j.1460-9568.2007.05976.x.

Durstewitz, Daniel, and Jeremy K Seamans. 2002. “The Computational Role of Dopamine D1 Receptors in Working Memory.” *Neural Networks : The Official Journal of the International Neural Network Society* 15 (4–6): 561–72.

———. 2008. “The Dual-State Theory of Prefrontal Cortex Dopamine Function with Relevance to Catechol-O-Methyltransferase Genotypes and Schizophrenia.” *Biological Psychiatry* 64 (9): 739–49. doi:10.1016/j.biopsych.2008.05.015.

Durstewitz, Daniel, Jeremy K Seamans, and Terrence J Sejnowski. 2000a. “Dopamine-Mediated Stabilization of Delay-Period Activity in a Network Model of Prefrontal Cortex.” *Journal of …*, 1733–50.

Durstewitz, Daniel, JK Seamans, and TJ Sejnowski. 2000b. “Neurocomputational Models of Working Memory.” *Nature Neuroscience* 3 (november).

Durstewitz, Daniel, Nicole M. Vittoz, Stan B. Floresco, and Jeremy K. Seamans. 2010. “Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning.” *Neuron* 66 (3). Elsevier Ltd: 438–48. doi:10.1016/j.neuron.2010.03.029.

Fagiolo, Giorgio. 2007. “Clustering in Complex Directed Networks.” *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 76 (2): 1–8. doi:10.1103/PhysRevE.76.026107.

Fino, Elodie, and Rafael Yuste. 2011. “Dense Inhibitory Connectivity in Neocortex.” *Neuron* 69 (6). Elsevier Inc.: 1188–1203. doi:10.1016/j.neuron.2011.02.025.

Ford, Judith M., Daniel H. Mathalon, Susan Whitfield, William O. Faustman, and Walton T. Roth. 2002. “Reduced Communication between Frontal and Temporal Lobes during Talking in Schizophrenia.” *Biological Psychiatry* 51 (6): 485–92. doi:10.1016/S0006-3223(01)01335-X.

Friston, Karl, Michael Breakspear, and Gustavo Deco. 2012. “Perception and Self-Organized Instability.” *Frontiers in Computational Neuroscience* 6 (January). Frontiers: 44. doi:10.3389/fncom.2012.00044.

Fujimoto, Hisataka, Taku Hasegawa, and Dai Watanabe. 2011. “Neural Coding of

Syntactic Structure in Learned Vocalizations in the Songbird." *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 31 (27): 10023–33. doi:10.1523/JNEUROSCI.1606-11.2011.

Fujisawa, Shigeyoshi, Asohan Amarasingham, Matthew T Harrison, and György Buzsáki. 2008. "Behavior-Dependent Short-Term Assembly Dynamics in the Medial Prefrontal Cortex." *Nature Neuroscience* 11 (7): 823–33. doi:10.1038/nn.2134.

Fuster, J. M., and G. E. Alexander. 1971. "Neuron Activity Related to Short-Term Memory." *Science* 173 (3997). American Association for the Advancement of Science: 652–54. doi:10.1126/science.173.3997.652.

Gewaltig, Marc-Oliver, and Markus Diesmann. 2007. "NEST (NEural Simulation Tool)." *Scholarpedia* 2 (4): 1430. doi:10.4249/scholarpedia.1430.

Grinstead, Charles M., and J. Laurie Snell. 2010. "Markov Chains." *Introduction to Probability*, 1–66.

Gros, Claudius. 2007. "Neural Networks with Transient State Dynamics." *New Journal of Physics* 9. doi:10.1088/1367-2630/9/4/109.

Guntekin, Bahar, and Erol Basar. 2015. "Review of Evoked and Event-Related Delta Responses in the Human Brain." *International Journal of Psychophysiology*, no. February. doi:10.1016/j.ijpsycho.2015.02.001.

Hanuschkin, Alexander, Markus Diesmann, and Abigail Morrison. 2011. "A Reafferent and Feed-Forward Model of Song Syntax Generation in the Bengalese Finch." *Journal of Computational Neuroscience* 31 (3). Springer US: 509–32. doi:10.1007/s10827-011-0318-z.

Harmony, Thalía. 2013. "The Functional Significance of Delta Oscillations in Cognitive Processing." *Frontiers in Integrative Neuroscience* 7 (January). Frontiers: 83. doi:10.3389/fnint.2013.00083.

Harvey, Christopher D, Philip Coen, and David W Tank. 2012. "Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task." *Nature* 484 (7392). Nature Publishing Group: 62–68. doi:10.1038/nature10918.

Hass, Joachim, Loreen Hertäg, Daniel Durstewitz, JM Fuster, JM Fuster, D Durstewitz, NM Vittoz, et al. 2016. "A Detailed Data-Driven Network Model of Prefrontal Cortex Reproduces Key Features of In Vivo Activity." *PLOS Comput Biol* 12 (5). Public Library of Science: 319–33. doi:10.1371/journal.pcbi.1004930.

Helmstaedter, M, C P J de Kock, D Feldmeyer, R M Bruno, and B Sakmann. 2007. "Reconstruction of an Average Cortical Column in Silico." *Brain Research Reviews* 55 (2): 193–203. doi:10.1016/j.brainresrev.2007.07.011.

Hertäg, Loreen, Joachim Hass, Tatiana Golovko, and Daniel Durstewitz. 2012. “An Approximation to the Adaptive Exponential Integrate-and-Fire Neuron Model Allows Fast and Predictive Fitting to Physiological Data.” *Frontiers in Computational Neuroscience* 6 (January). Frontiers: 62. doi:10.3389/fncom.2012.00062.

Hills, Thomas T., Peter M. Todd, David Lazer, a. David Redish, and Iain D. Couzin. 2014. “Exploration versus Exploitation in Space, Mind, and Society.” *Trends in Cognitive Sciences* 19 (1). doi:10.1016/j.tics.2014.10.004.

Hofer, Sonja B, Ho Ko, Bruno Pichler, Joshua Vogelstein, Hana Ros, Hongkui Zeng, Ed Lein, Nicholas a Lesica, and Thomas D Mrsic-Flogel. 2011. “Differential Connectivity and Response Dynamics of Excitatory and Inhibitory Neurons in Visual Cortex.” *Nature Neuroscience* 14 (8): 1045–52. doi:10.1038/nn.2876.

Hopfield, J. J. 1982. “Neural Networks and Physical Systems with Emergent Collective Computational Abilities.” *Proceedings of the National Academy of Sciences* 79 (8): 2554–58. doi:10.1073/pnas.79.8.2554.

Ikegaya, Yuji, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste. 2004. “Synfire Chains and Cortical Songs: Temporal Modules of Cortical Activity.” *Science (New York, N.Y.)* 304 (5670): 559–64. doi:10.1126/science.1093173.

Jahnke, Sven, Raoul-Martin Memmesheimer, and Marc Timme. 2008. “Stable Irregular Dynamics in Complex Neural Networks.” *Physical Review Letters* 100 (4): 2–5. doi:10.1103/PhysRevLett.100.048102.

———. 2009. “How Chaotic Is the Balanced State?” *Frontiers in Computational Neuroscience* 3 (November): 13. doi:10.3389/neuro.10.013.2009.

Jones, Lauren M, Alfredo Fontanini, Brian F Sadacca, Paul Miller, and Donald B Katz. 2007. “Natural Stimuli Evoke Dynamic Sequences of States in Sensory Cortical Ensembles.” *Proceedings of the National Academy of Sciences of the United States of America* 104 (47): 18772–77. doi:10.1073/pnas.0705546104.

Katahira, Kentaro, Kenta Suzuki, Kazuo Okanoya, and Masato Okada. 2011. “Complex Sequencing Rules of Birdsong Can Be Explained by Simple Hidden Markov Processes.” Edited by Gonzalo G. de Polavieja. *PloS One* 6 (9). Public Library of Science: e24516. doi:10.1371/journal.pone.0024516.

Klinshov, Vladimir V, Jun-nosuke Teramae, Vladimir I Nekorkin, and Tomoki Fukai. 2014. “Dense Neuron Clustering Explains Connectivity Statistics in Cortical Microcircuits.” Edited by Bard Ermentrout. *PloS One* 9 (4). Public Library of Science: e94292. doi:10.1371/journal.pone.0094292.

Knyazev, Gennady G. 2012. “EEG Delta Oscillations as a Correlate of Basic Homeostatic and Motivational Processes.” *Neuroscience and Biobehavioral*

Reviews 36 (1). Elsevier Ltd: 677–95. doi:10.1016/j.neubiorev.2011.10.002.

Lakatos, Peter, George Karmos, Ashesh D Mehta, Istvan Ulbert, and Charles E Schroeder. 2008. “Entrainment of Neuronal Oscillations as a Mechanism of Attentional Selection.” *Science (New York, N.Y.)* 320 (5872): 110–13. doi:10.1126/science.1154735.

Latham, Peter E, and Sheila Nirenberg. 2004. “Computing and Stability in Cortical Networks.” *Neural Computation* 16 (7): 1385–1412. doi:10.1162/089976604323057434.

Legenstein, Robert, and Wolfgang Maass. 2007. “Edge of Chaos and Prediction of Computational Performance for Neural Circuit Models.” *Neural Networks : The Official Journal of the International Neural Network Society* 20 (3): 323–34. doi:10.1016/j.neunet.2007.04.017.

Linkerhand, Mathias, and Claudius Gros. 2013. “Generating Functionals for Autonomous Latching Dynamics in Attractor Relict Networks.” *Scientific Reports* 3 (June). The Author(s): 2042.

Litwin-kumar, Ashok. 2013. “Relationship between Neuronal Architecture and Variability in Cortical Circuits.”

Litwin-Kumar, Ashok, and Brent Doiron. 2012. “Slow Dynamics and High Variability in Balanced Cortical Networks with Clustered Connections.” *Nature Neuroscience* 15 (11). Nature Publishing Group: 1498–1505. doi:10.1038/nn.3220.

London, Michael, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E Latham. 2010. “Sensitivity to Perturbations in Vivo Implies High Noise and Suggests Rate Coding in Cortex.” *Nature* 466 (7302): 123–27. doi:10.1038/nature09086.

Long, Michael A, Dezhe Z Jin, and Michale S Fee. 2010. “Support for a Synaptic Chain Model of Neuronal Sequence Generation.” *Nature* 468 (7322). Nature Publishing Group: 394–99. doi:10.1038/nature09514.

Luczak, Artur, Peter Barthó, Stephan L Marguet, György Buzsáki, and Kenneth D Harris. 2007. “Sequential Structure of Neocortical Spontaneous Activity in Vivo.” *Proceedings of the National Academy of Sciences of the United States of America* 104 (1): 347–52. doi:10.1073/pnas.0605643104.

Lukoševičius, Mantas, and Herbert Jaeger. 2009. “Reservoir Computing Approaches to Recurrent Neural Network Training.” *Computer Science Review* 3: 127–49. doi:10.1016/j.cosrev.2009.03.005.

Mante, Valerio, David Sussillo, Krishna V Shenoy, and William T Newsome. 2013. “Context-Dependent Computation by Recurrent Dynamics in Prefrontal Cortex.” *Nature* 503 (V). Nature Publishing Group: 78–84. doi:10.1038/nature12742.

McDonnell, Mark D, and Lawrence M Ward. 2014. “Small Modifications to Network Topology Can Induce Stochastic Bistable Spiking Dynamics in a Balanced Cortical Model.” *PloS One* 9 (4): e88254. doi:10.1371/journal.pone.0088254.

Mel, B. W. 1993. “Synaptic Integration in an Excitable Dendritic Tree.” *J Neurophysiol* 70 (3): 1086–1101.

Monteforte, Michael, and Fred Wolf. 2010. “Dynamical Entropy Production in Spiking Neuron Networks in the Balanced State.” *Physical Review Letters*, 1–4.

———. 2012. “Dynamic Flux Tubes Form Reservoirs of Stability in Neuronal Circuits.” *Physical Review X* 2 (4): 1–12. doi:10.1103/PhysRevX.2.041007.

Moreno-Bote, Rubén, John Rinzel, and Nava Rubin. 2007. “Noise-Induced Alternations in an Attractor Network Model of Perceptual Bistability.” *Journal of Neurophysiology* 98 (3): 1125–39. doi:10.1152/jn.00116.2007.

Natschläger, Thomas, Nils Bertschinger, and Robert Legenstein. 2005. “At the Edge of Chaos: Real-Time Computations and Self-Organized Criticality in Recurrent Neural Networks.” *Advances in Neural …*.

Pehlevan, Cengiz, and Haim Sompolinsky. 2014. “Selectivity and Sparseness in Randomly Connected Balanced Networks.” Edited by Olaf Sporns. *PloS One* 9 (2). Public Library of Science: e89992. doi:10.1371/journal.pone.0089992.

Perin, Rodrigo, Thomas K Berger, and Henry Markram. 2011. “A Synaptic Organizing Principle for Cortical Neuronal Groups.” *Proceedings of the National Academy of Sciences of the United States of America* 108 (13): 5419–24. doi:10.1073/pnas.1016051108.

Pezzulo, Giovanni, Matthijs A A van der Meer, Carien S Lansink, and Cyriel M A Pennartz. 2014. “Internally Generated Sequences in Learning and Executing Goal-Directed Behavior.” *Trends in Cognitive Sciences* 18 (12). Elsevier: 647–57. doi:10.1016/j.tics.2014.06.011.

Rabinovich, M., a. Volkovskii, P. Lecanda, R. Huerta, H. D. I. Abarbanel, and G. Laurent. 2001. “Dynamical Encoding by Networks of Competing Neuron Groups: Winnerless Competition.” *Physical Review Letters* 87 (6): 68102. doi:10.1103/PhysRevLett.87.068102.

Rajan, Kanaka, and L. F. Abbott. 2006. “Eigenvalue Spectra of Random Matrices for Neural Networks.” *Physical Review Letters* 97 (November): 2–5. doi:10.1103/PhysRevLett.97.188104.

Renart, Alfonso, Nicolas Brunel, and XJ Wang. 2004. “Mean-Field Theory of Irregularly Spiking Neuronal Populations and Working Memory in Recurrent Cortical Networks.” In *Computational Neuroscience: A Comprehensive Approach*.

Rubinov, Mikail, and Olaf Sporns. 2010. “Complex Network Measures of Brain Connectivity: Uses and Interpretations.” *NeuroImage* 52 (3): 1059–69. doi:10.1016/j.neuroimage.2009.10.003.

Russo, Eleonora, Vijay M K Namboodiri, Alessandro Treves, and Emilio Kropff. 2008. “Free Association Transitions in Models of Cortical Latching Dynamics.” *New Journal of Physics* 10 (1). IOP Publishing: 15008. doi:10.1088/1367-2630/10/1/015008.

Russo, Eleonora, and Alessandro Treves. 2011. “An Uncouth Approach to Language Recursivity.” *Biolinguistics* 5 (1–2): 133–50.

Saramäki, Jari, Mikko Kivelä, Jukka Pekka Onnela, Kimmo Kaski, and János Kertész. 2007. “Generalizations of the Clustering Coefficient to Weighted Complex Networks.” *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 75 (2): 2–5. doi:10.1103/PhysRevE.75.027105.

Schaub, Michael T, Yazan N Billeh, Costas a Anastassiou, Christof Koch, and Mauricio Barahona. 2015. “Emergence of Slow-Switching Assemblies in Structured Neuronal Networks.” *PLoS Computational Biology* 11 (7). Public Library of Science: 1–18. doi:10.1371/journal.pcbi.1004196.

Seidemann, Eyal, Isaac Meilijson, Moshe Abeles, Hagai Bergman, and Eilon Vaadia. 1996. “Simultaneously Recorded Single Units in the Frontal Cortex Go through Sequences of Discrete and Stable States in Monkeys Performing a Delayed Localization Task.” *Journal of Neuroscience* 16 (2). [Baltimore, MD: The Society, c1981-: 752–68.

Strogatz, SH. 2014. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*.

Sussillo, David, and Omri Barak. 2013. “Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks.” *Neural Computation* 25 (3). MIT Press 55 Hayward Street, Cambridge, MA 02142-1315 email: journals-info@mit.edu: 626–49. doi:10.1162/NECO_a_00409.

Treves, Alessandro. 2005. “Frontal Latching Networks: A Possible Neural Basis for Infinite Recursion.” *Cognitive Neuropsychology* 22 (3): 276–91. doi:10.1080/02643290442000329.

Tsodyks, MV, and H Markram. 1997. “The Neural Code between Neocortical Pyramidal Neurons Depends on Neurotransmitter Release Probability.” *… of the National Academy of …* 94 (January): 719–23.

Tsuda, Ichiro. 2001. “Toward an Interpretation of Dynamic Neural Activity in Terms of Chaotic Dynamical Systems.” *The Behavioral and Brain Sciences* 24 (5): 793-810-48.

van Vreeswijk, C, and H Sompolinsky. 1996. “Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity.” *Science (New York, N.Y.)* 274 (5293): 1724–26.

Wallace, Edward, Hamid Reza Maei, and Peter E Latham. 2013. “Randomly Connected Networks Have Short Temporal Memory.” *Neural Computation* 25 (6): 1408–39. doi:10.1162/NECO_a_00449.

Wang, X J. 1999. “Synaptic Basis of Cortical Persistent Activity: The Importance of NMDA Receptors to Working Memory.” *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 19 (21): 9587–9603.

Welch, PD. 1967. “The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging over Short, Modified Periodograms.” *IEEE Transactions on Audio and Electroacoustics*.

Wennekers, Thomas, and Günther Palm. 2009. “Syntactic Sequencing in Hebbian Cell Assemblies.” *Cognitive Neurodynamics* 3 (4): 429–41. doi:10.1007/s11571-009-9095-z.

Yuste, Rafael. 2015. “From the Neuron Doctrine to Neural Networks.” *Nature Reviews Neuroscience*, no. July. Nature Publishing Group. doi:10.1038/nrn3962.

Zheng, Pengsheng, and Jochen Triesch. 2014. “Robust Development of Synfire Chains from Multiple Plasticity Mechanisms.” *Frontiers in Computational Neuroscience* 8 (June). Frontiers: 66. doi:10.3389/fncom.2014.00066.