

Criticality of spreading dynamics in hierarchical cluster networks without inhibition

M Kaiser^{1,2,5}, M Görner³ and C C Hilgetag⁴

¹ School of Computing Science, Newcastle University, Newcastle upon Tyne NE1 7RU, UK

² Institute of Neuroscience, Newcastle University, Newcastle upon Tyne NE2 4HH, UK

³ Department of Mathematics, UC Berkeley, Berkeley, CA 94720, USA

⁴ School of Engineering and Science, Jacobs University Bremen, Campus Ring 6, 28759 Bremen, Germany

E-mail: M.Kaiser@ncl.ac.uk, matthias@math.berkeley.edu and c.hilgetag@iu-bremen.de

New Journal of Physics **9** (2007) 110

Received 21 February 2007

Published 2 May 2007

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/9/5/110

Abstract. An essential requirement for the representation of functional patterns in complex neural networks, such as the mammalian cerebral cortex, is the existence of stable network activations within a limited critical range. In this range, the activity of neural populations in the network persists between the extremes of quickly dying out, or activating the whole network. The nerve fibre network of the mammalian cerebral cortex possesses a modular organization extending across several levels of organization. Using a basic spreading model without inhibition, we investigated how functional activations of nodes propagate through such a hierarchically clustered network. The simulations demonstrated that persistent and scalable activations could be produced in clustered networks, but not in random networks of the same size. Moreover, the parameter range yielding critical activations was substantially larger in hierarchical cluster networks than in small-world networks of the same size. These findings indicate that a hierarchical cluster architecture may provide the structural basis for the stable and diverse functional patterns observed in cortical networks.

⁵ Author to whom any correspondence should be addressed.

Contents

1. Introduction	2
2. Materials and methods	3
3. Results	5
3.1. Critical range of nonlocalized activation	5
3.2. Influence of localization	7
3.3. Delay of activation spreading	8
3.4. Parameter range for persistent contained activity	8
3.5. Exhaustion	10
4. Discussion	10
Acknowledgments	12
Appendix A. Spreading of activity depending on the relative density of connections within sub-clusters, clusters and at the network level in the hierarchical cluster network	12
References	13

1. Introduction

Natural systems operate within a critical functional range, sustaining diverse dynamical states [1, 2] on the basis of their intricate system architecture. For instance, in neural systems, such as the cerebral cortical networks of the mammalian brain, this critical range is indicated by the fact that initial activations result in various neuronal activity patterns that neither die out too quickly, nor spread across the entire network. What are the essential structural and functional parameters that allow complex neural networks to maintain such a dynamic balance? In particular, which factors limit the spreading of persistent neural activity to the whole brain, thus preventing a pathological state resembling epilepsy?

Most current models of neural network dynamics focus on maintaining the right balance of network activation and rest through functional interactions among populations of inhibitory and excitatory nodes [3, 4]. However, the topology of the networks may also make a significant contribution towards critical network dynamics, even in the absence of inhibitory nodes.

It is known from the anatomy of the brain that cortical architecture and connections are organized in a hierarchical and modular way, from cellular microcircuits in cortical columns [5] at the lowest level, via cortical areas at the intermediate level, to clusters of highly connected brain regions at the global systems level [6]–[8]. At each level clusters arise, with denser connectivity within than between the modules. This means that neurons within a column, area or area cluster are more frequently linked with each other than with neurons in the rest of the network.

We used a basic spreading model to explore the functional potential of different network topologies, particularly those possessing a clustered organization, for producing persistent yet contained node activations. The approach was inspired by classical ‘netlet’ studies of random neural networks [9]. More recently, spreading analysis has also been applied to cortical networks [10] and to other complex networks with a non-random organization [11]–[13].

The present model operates without explicit inhibition, as we were specifically interested in the contributions of different network topologies to maintaining critical activations. This feature

also reflects a structural attribute of cortical systems networks [14] and other complex networks, for instance, social networks [11], that do not possess explicit inhibition, at least at some level of organization. Thus, while the current model is inspired by structural and functional properties of neural networks, it may also be relevant for understanding spreading phenomena in other complex systems, such as disease or virus spreading in social or computer networks [12, 15, 16].

2. Materials and methods

All simulations were implemented in Matlab 7 (Mathworks Inc, Natick, USA). To investigate the effect of clustered network organization on the spread of activity, simulations were run on three different types of networks: random networks, small-world networks with a fixed fraction $p \in [0, 1]$ of random connections, and hierarchically clustered networks. All networks were undirected graphs with $N = 1000$ vertices and $E = 12\,000$ edges. To create the (non-hierarchical) small-world networks, we used a procedure similar to the algorithm presented in [17].⁶

To create the hierarchical cluster network, the 1000 vertices were parcellated into 10 disjoint sets ('clusters'), each consisting of 100 vertices. Then each cluster was divided into 10 'sub-clusters', containing 10 vertices each. The network was wired randomly, such that 4000 edges (one-third of the total 12 000 connections) connected vertices within the same sub-clusters, a further 4000 edges connected vertices within the same clusters, and the last 4000 were randomly distributed over all nodes of the network (figure 1).

A basic spreading model [2] was modified to simulate the spread of activity through the network. Note that the labels of activated/inactivated neural nodes, as used here, could be easily replaced by labels for infected/susceptible nodes in an analogous model of disease spreading.

The simulation operated in discrete time steps, with nodes being in one of the two states: active or inactive. To advance from one time step to the next:

1. an inactive node became activated in the next time step, if it was connected to at least k (default $k = 6$) active nodes;
2. an active node became inactivated with probability ν (default $\nu = 0.3$) in the next time step, or otherwise stayed active.

For initialization, i ($i \leq i_0$) nodes among the nodes 1 to i_0 were selected randomly and activated in the first time step. The network nodes were numbered consecutively. Hence, by setting i_0 to, for example, 10, 20 or 100, only nodes in the first sub-cluster, the first two sub-clusters, or the first cluster, respectively, were activated during initialization. Thus, i determined the number of initially active nodes while i_0 controlled the localization of initial activations, with smaller values resulting in more localized initial activity.

For a systematic comparison of the criticality behaviour of different network topologies, small-world networks were tested for different fractions of randomly wired links, p . When

⁶ The algorithm started with a regular ring lattice where the vertices were placed on a circle and each vertex was connected to its nearest neighbors, its second nearest neighbours, ... and its d th nearest neighbour, where d was obtained by rounding E/N to the next integer. Then $(1 - p)E$ edges were randomly picked from the regular ring lattice and pE random edges were added to them to obtain a small-world network with E edges.

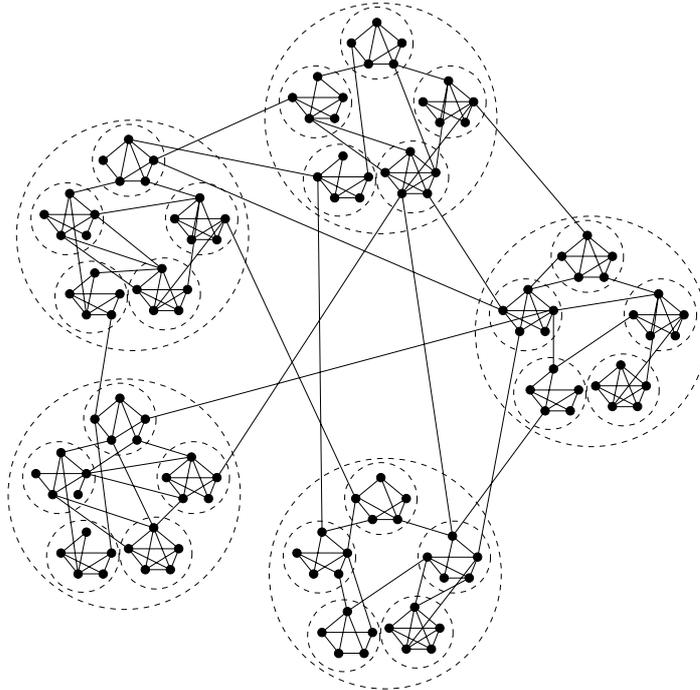


Figure 1. Schematic view of a hierarchical cluster network with five clusters containing five sub-clusters each.

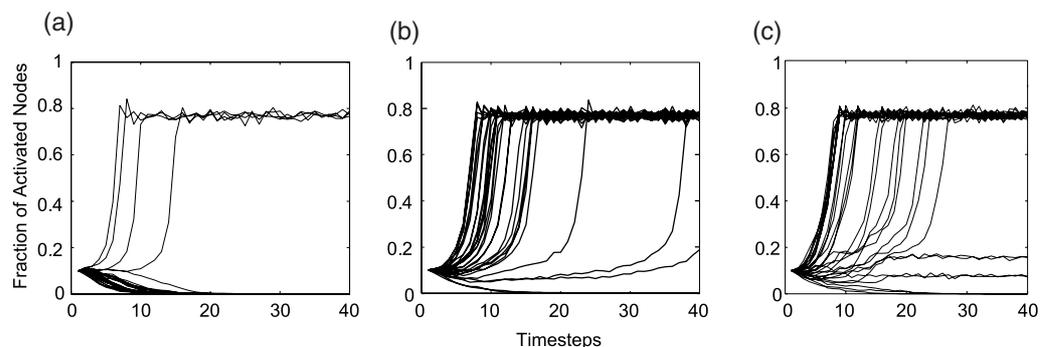


Figure 2. Spread of initially unlocalized activity in (a) random, (b) small-world and (c) hierarchical cluster networks ($i = 100$, $i_0 = 1000$), based on 30 simulations for each network.

plotting figure 2 for various values of p , the simulation yielded sustained intermediate activity (neither dying out nor spreading through the whole network) only if p was between 0.4 and 0.6. Consequently we chose $p = 0.5$ for all simulations. Note that for this p , the small-world and hierarchical network had similar clustering coefficients and characteristic path lengths (see table 1) This means that the hierarchical cluster network also possessed small-world characteristics [17]. Given the high probability of rewiring in the default small-world network, one may wonder if this network still possessed strong *local* clustering comparable to the hierarchical cluster network. To test this aspect, we measured the edge density within each set of 10 consecutive network nodes in the small-world network, that is for 1000 sets in total.

Table 1. Topological properties of the networks.

Measure	Random network	Small-world network	Clustered network
Clustering coefficient	0.025	0.11	0.15
Characteristic path length	2.5	2.6	2.6

The average local edge density in the small-world network was 0.51 ± 0.08 , and thus comparable to the average local edge density of 0.60 ± 0.15 measured in the same way in the hierarchical cluster network.

3. Results

Across different simulation conditions, the hierarchical cluster network showed a larger variety of behaviours than the random or small-world networks, and it produced persistent yet contained network activity for a wider range of initial activation conditions. Moreover, in contrast to random networks, the activation behaviour of the hierarchical cluster network was also influenced by the localization of the initial activity.

Examples exhibiting the behaviours of the different networks are given in figure 2. The figure shows the result of 30 simulations in all three networks when 10% of the nodes were randomly selected for initial activation. In the random network, activity either quickly spread through the network, or, more frequently, died out. In the small-world network, spread of activity was slower than in the random network, but eventually also resulted in all-or-nothing activation of the network. Compared to the random network, activation of the whole network was more frequent. In contrast to the previous two types of networks, the hierarchical cluster network showed a large variety of delays for the spread of activity and also produced cases in which spreading was limited to one or two clusters with sustained activity. Such persistent activation could be sustained with different patterns and varying extent of involved nodes (figure 3).

3.1. Critical range of nonlocalized activation

To quantify the critical range of network function, we considered the range of parameters for which initial activations could result in intermediate levels of network activity and not just all-or-none network activation.

The critical range for each network was estimated by simulating the spreading of initially unlocalized activity ($i_0 = 1000$) over 80 time steps for varying values of $i = 40, 50, \dots, 120$, while counting how many of the 100 simulations for a specific network and values of i ended in one of the following scenarios: the number of activated nodes was zero (activity died out), between zero and 100 (activity stayed in one cluster), between 100 and 200 (activity stayed in two clusters), or larger than 200 (activity spread over the whole network); see figure 4. The critical range for initially unlocalized activity extended from around $i = 65$ to $i = 105$ in the case of the small-world and clustered hierarchical networks, while the random network only showed all-or-none activation.

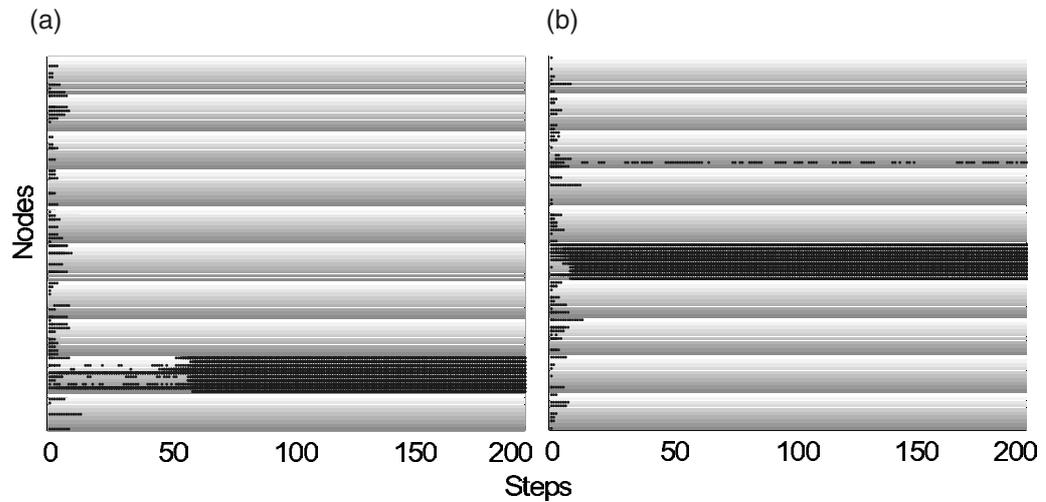


Figure 3. Examples for different sustained activity patterns in hierarchical cluster networks ($i = 90$, $i_0 = 1000$). Graded grey background shading indicates the 10 sub-clusters within each of the 10 clusters. Black dots represent nodes active at the respective time step. (a) One cluster showing sustained activity. (b) One cluster remaining active with frequent co-activation of one external sub-cluster.

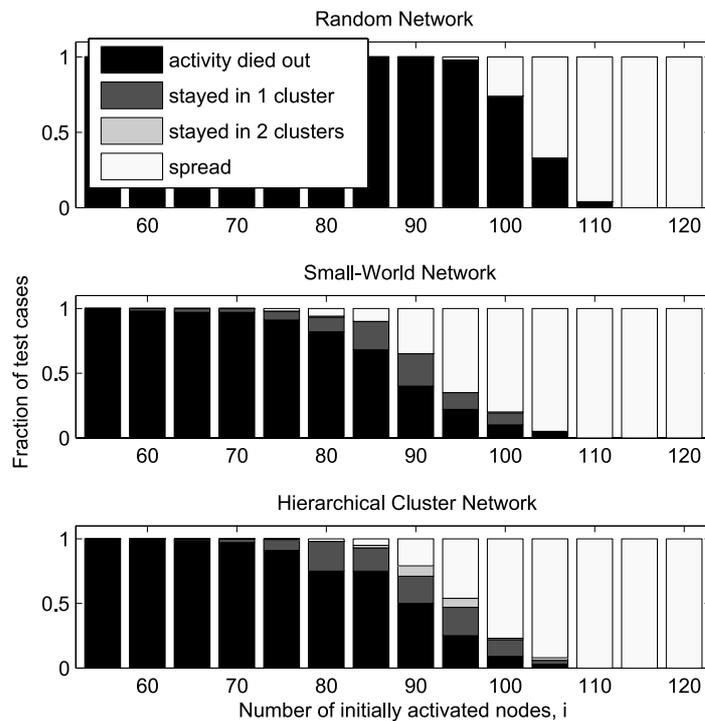


Figure 4. Critical range of unlocalized node activations: fraction of test cases resulting in different activation scenarios after 80 time steps versus number (i) of initially activated, nonlocalized ($i_0 = 1000$) nodes. For small-world networks, which did not possess an explicit multi-cluster organization, intermediate grey levels represent persistent contained activity.

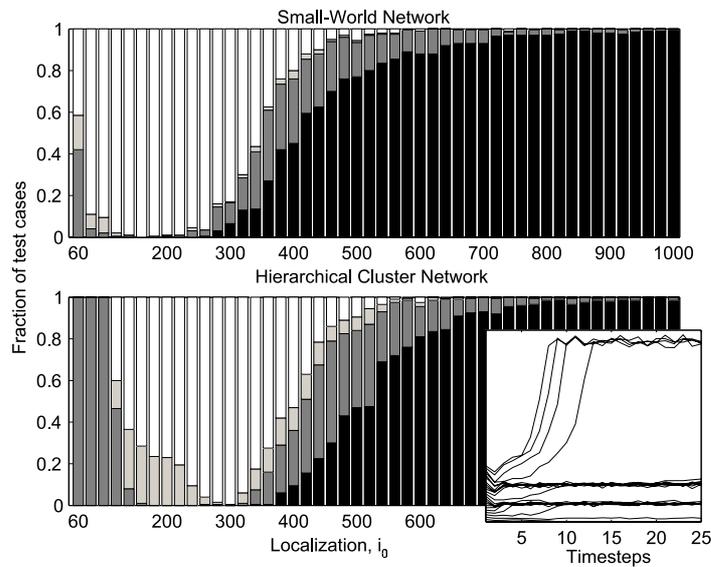


Figure 5. Critical range of initially localized node activations: fraction of test cases ending in different activation scenarios after 80 time steps (see figure 4 for legend) versus localization ($i_0 = 60, 80, \dots, 1000$) of initially activated nodes. Number of initially activated nodes was fixed at $i = 60$. Data not shown for random networks, in which localization had no influence on the spread of activity and activity always died out for $i = 60$. Sustained activity mostly remaining in clusters (inset): for each value of $i = 0, 10, \dots, 220$, one simulation of localized activity ($i_0 = i$) for 25 time steps is plotted for the hierarchical cluster network.

3.2. Influence of localization

Initially activated nodes may be localized in only a few clusters or distributed over the whole network. To investigate the impact of initial localization on the spread of activity, the simulations from the previous section were repeated with fixed $i = 60$ and for values of i_0 ranging from 60 to 1000 in steps of 20 (figure 5).

For the random network, activity always died out for $i = 60$. Thus, localization had no impact on spreading in this type of network. The spread of activity in small-world and hierarchical cluster networks, however, depended on where the initial activation was localized. When initially active nodes were widely dispersed, activity in the network would eventually die out. More focused initial activation, on the other hand, could activate the whole network, or lead to persistent activation of network subregions. This scaled partial activation was particularly pronounced in the hierarchical cluster network (example given in inset of figure 5).

The simulations also indicated a localization effect for the small-world network. The plots in figure 5 for the small-world and clustered hierarchical networks appear similar, but shifted, with a larger range of intermediate activations in the hierarchical cluster network. Notice, however, that the appearance of intermediate activation in small-world networks may have been partly due to the very slow spreading of activity in this type of network. Thus, many settings eventually leading to complete network activation had not yet activated the whole network at the end of the observation period and were therefore classified as partial activation.

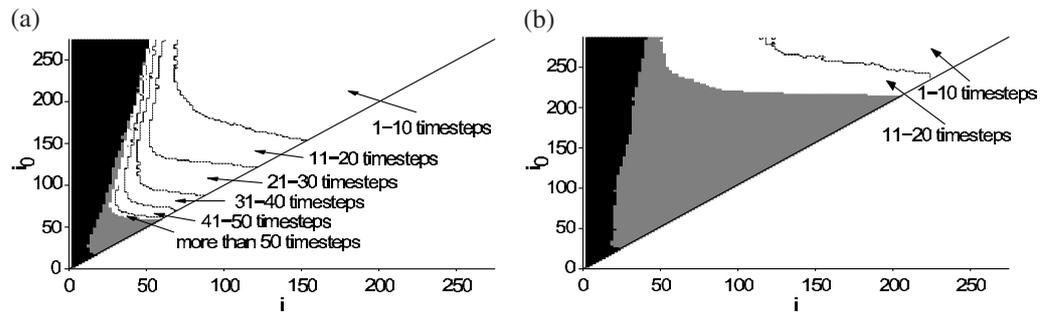


Figure 6. Parameter space exploration of the critical range for all combinations of initial activation parameter i and localization parameter i_0 , based on 1000 test cases. Simulation outcomes are indicated by grey level (black: activity died out; grey: limited spreading; white: complete spreading). (a) Small-world network. (b) Hierarchical cluster network.

3.3. Delay of activation spreading

We also investigated the time needed to activate 50% of the nodes in a network, considering only those test cases in which activity did not die out.

In the random network, if activity spread at all, it did so quickly, typically in less than 10 time steps. Even if the initial activity was in the borderline range for all-or-none network activation, not more than 15 time steps were required in any of the cases. This was in contrast to the small-world and hierarchically clustered networks, for which a wide range of delay times was observed. For the small-world network, delayed spreading could be produced by selecting the initial activity in the critical range, as shown in figure 4. However it was also important that the initial activity was strictly localized ($i_0 = i$). Setting $i_0 = i = 90$ typically resulted in about 40 time steps for spreading, whereas for $i_0 = 190, i = 90$, spreading in the small-world network appeared similar to that in the random network.

By contrast, for the hierarchically clustered network, spreading to the global level did not arise when the initial activation was too strictly localized (figure 5). A maximum delay for spreading was achieved by localizing the initial activity within two or three clusters (e.g. delay around 40 steps for $i_0 = 200, i = 90$).

3.4. Parameter range for persistent contained activity

We systematically explored the network activation behaviours resulting from different settings of the initial node activation and localization parameters. Since initial activation is typically small in neural systems, we limited the maximum number of initially active nodes to 250, that is, one quarter of all network nodes. Persistent contained activity in hierarchical networks was robust for a wide range of initial localization and activation parameters (indicated by grey parameter domain in figure 6). For small-world networks, however, parameters needed to be finely tuned in order to yield sustained activity. Thus, hierarchical networks showed sustained activity for a wider range of initial conditions.

We also explored if these results were robust for variations in the spreading parameters k and ν . For these control studies, we used a Monte Carlo approach that, for each pair of k and ν , generated 20 small-world and 20 hierarchical networks. For each network, the dynamics for

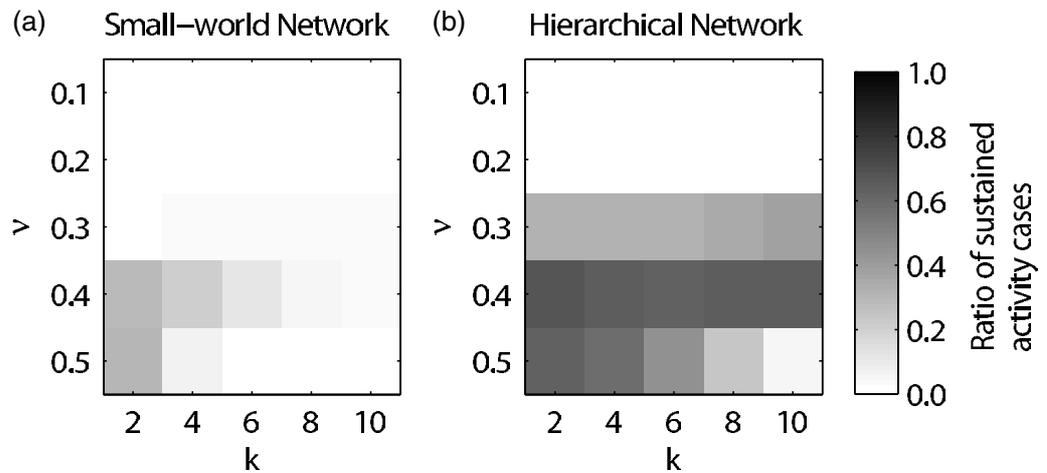


Figure 7. Ratio of sustained activity cases depending on the spreading parameters k (activation threshold) and ν (deactivation probability) for (a) small-world and (b) hierarchical cluster networks.

1000 randomly chosen parameters i and i_0 were tested (figure 7). A trial was considered to show sustained activity if $\leq 50\%$ of the nodes were activated at the end of the simulation. For each pair of spreading parameters k and ν , the average proportion of cases for which sustained activity occurred, which is related to the ratio of the grey space in figure 6, was larger for the hierarchical cluster networks than for small-world networks. The maximum ratio for small-world networks was 30%, whereas it was 67% for hierarchical cluster networks.

Finally, spreading of activity in the hierarchical cluster network also depended on the relative distribution of connections over sub-clusters, clusters and the whole network. To explore this effect, 1000 trials were run for each of the connection distributions (appendix, table A.1), with random values of i and i_0 , and $k = 6$ and $\nu = 0.3$. We measured the proportion of trials showing sustained activity (this approach was analogous to the one used for figures 7 and 8). Reducing the number of connections within clusters or sub-clusters blurred the boundaries of local network modules and reduced the proportion of cases with sustained activity, but the proportion was still larger than that for small-world networks. On the other hand, increasing the density of connections within sub-clusters and clusters led to a higher proportion of cases with sustained activity.

For the default arrangement of the hierarchical network (with 4000 connections distributed at each of the three levels of organization), an activation of more than two clusters always resulted in spreading of activity through the whole network, due to the relatively high density of connections between clusters. Conversely, a reduction in the density of inter-cluster connections allowed persistent contained activity in more than two clusters (data not shown). In these cases, the limited number of inter-cluster connections formed a bottleneck for activation of the remaining clusters.

In summary, the prevalence of persistent yet contained activity in hierarchical cluster networks was robust over a large range of model parameters and initial conditions, whereas in small-world networks, sustained activity occurred only for a relatively small parameter range.

3.5. Exhaustion

Energy resources for sustaining neural network activations may be limited. For instance, exhaustion occurs during epileptic seizures, reducing the duration of large-scale cortical activation. Similar phenomena arise in networks where sustaining an active/infected state or preserving the existence of nodes relies on limited resources. Therefore, we also tested the effect of restricting the number of time steps that nodes could be consecutively active.

While in the previous analyses nodes could be active for an unlimited time, we here adjusted the number of simulation steps for which a node could be active without interruption, with upper limits from 7 to 1 steps (figure 8). The diagrams show that persistent yet contained network activation could still occur in the hierarchical cluster network, despite different degrees of limiting node exhaustion. As the number of steps was reduced, a lower value of the deactivation parameter ν was sufficient to keep activation in balance and allow sustained activity. Sustained activity was, however, largely independent of the exhaustion threshold parameter. The range of parameters for which sustained activity occurred remained similar to that in the previous analyses, with no clear correlation with the number of steps (average ratio of sustained activity cases over all pairs of the spreading parameters was 0.272 ± 0.068).

4. Discussion

Our simulations demonstrated a strong impact of network topology on spreading behaviour. Clustered networks were more easily activated than random networks of the same size. This was due to the higher density of connections within the clusters which facilitated local activation. At the same time, the sparser connectivity between clusters prevented the spreading of activity across the whole network. Thus, in contrast to random networks, clustered networks possessed an expanded critical functional range for which initial activations resulted in persistent but non-global network activity.

In addition, there were also differences in the dynamical behaviour of the two clustered network types. Compared to the small-world network, the hierarchical cluster network possessed a substantially larger critical parameter space, defined by the extent of localizations and initial network activations that could result in limited yet sustained network activity patterns (figure 6).

Simulations with similar results as described here were also performed for networks in which all nodes were represented by integrate-and-fire (IF) neurons [18]. These simulations also showed easier activation of the hierarchical cluster network compared to the random network, as well as the existence of intermediate states of activation in the hierarchical cluster, but not the random network [19]. Thus, our results do not appear to depend on the specific neuronal model, but are generally based on the topology of the studied networks. This conclusion is also supported by the systematic variation of model parameters presented here, which did not affect the principal model behaviour.

The present hierarchical cluster model, which reflects the distributed multi-level modularity found in biological neural networks, is different from ‘centralistic’ hierarchical modular networks (with most nodes linked to network hubs) that have been described previously [20]. While developmental algorithms have been suggested for the latter type of network, there are currently no algorithms for producing the hierarchical cluster networks presented here. However, single-level clustered network architectures can be produced by models for developmental spatial growth [21] or dynamic self-organization of neural networks [22]; and such models may serve

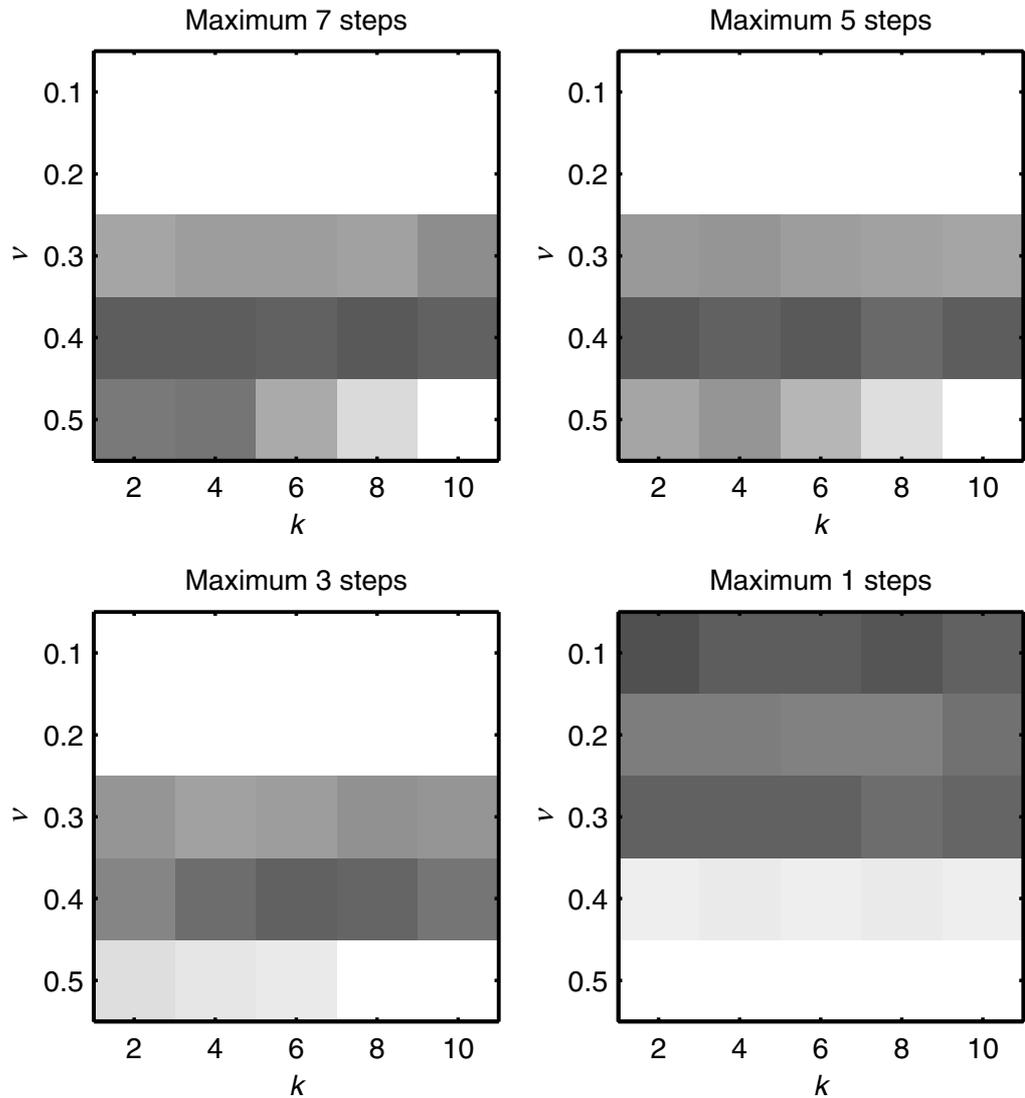


Figure 8. Influence of exhaustion on sustained spreading for hierarchical cluster networks. As the maximum number of steps is reduced from seven down to one, sustained activity only occurs for lower values of v whereas the total parameter range of v remains the same. Grey levels indicate the proportion of cases with sustained activity (cf colourbar in figure 7).

as a starting point for exploring the biological mechanisms for developing multi-level clustered neural architectures.

In conclusion, the present results provide a proof of concept for three points. Firstly, persistent but contained network activation can occur in the absence of inhibitory nodes. This might explain why cortical activity does not normally spread to the whole brain, even though top-level links between cortical areas are exclusively formed by excitatory fibres [14]. While the involvement of inhibitory neurons and other dynamic control mechanisms may further extend the critical range, the present results indicate that the hierarchical cluster architecture of complex neural networks, such as the mammalian cortex, may provide the principal structural basis for their stable and scalable functional patterns. Secondly, in hierarchical clustered networks, activity

Table A.1. Proportion of cases with sustained activity for different ratios of connections within clusters or sub-clusters. Note that the total number of connections within the network remained unchanged.

	Network connections	Connections within clusters	Connections within sub-clusters	Proportion of cases with sustained activity
Original calculations	4000	4000	4000	0.436
Increased cluster/network ratio	3600	4400	4000	0.546
Reduced cluster/network ratio	2000	6000	4000	0.851
	5000	3000	4000	0.183
	6000	2000	4000	0.159
Increased sub-cluster/network ratio	7000	1000	4000	0.079
	3600	4000	4400	0.537
Reduced sub-cluster/network ratio	5000	4000	3000	0.203
	6000	4000	2000	0.237
	7000	4000	1000	0.184

can be sustained without the need for random input or noise as an external driving force. Thirdly, multiple clusters in a network influence activity spreading in two ways: bottleneck connections between clusters limit global spreading whereas a higher connection density within clusters sustains recurrent local activity.

These findings may have practical implications and may guide future research. For instance, it might be worthwhile to test whether epileptic patients show a higher degree of connectivity between cortical network clusters or other changes in structural connectivity which would facilitate spreading. Such changes might be reflected in certain aspects of functional connectivity [23] or be demonstrated more directly by diffusion tensor imaging. It would also be interesting to evaluate the impact of inhibition added to the lowest level of the network organization, such as that provided by inhibitory interneurons in the cerebral cortex. However, by abstracting from specific characteristics of biological neural networks, the present study may also help to elucidate general features of spreading behaviour in other hierarchically organized networks without limiting it to neural systems.

Acknowledgments

We thank Tadas Jucikas for performing the simulations with IF networks and Bernhard Kramer for helpful comments on this paper. MK and MG acknowledge financial support from the German National Merit Foundation. MK was supported by EPSRC (EP/E002331/1) and the Royal Society (2006/R2).

Appendix A. Spreading of activity depending on the relative density of connections within sub-clusters, clusters and at the network level in the hierarchical cluster network

Simulating activity spreading for $k = 6$ and $\nu = 0.3$, sustained activity occurred in 18.4% of the cases for 1000 sub-cluster and 7000 global connections and in 7.9% of the cases for 1000 cluster and 7000 global connections (table A.1). These proportions are lower than for the original

network arrangement with 4000 cluster, sub-cluster, and global connections, resulting in 43.6% of cases with sustained activity, but still higher than the respective proportion of 1.96% for small-world networks.

On the other hand, the number of cases with sustained activity was increased if relatively more connections existed within clusters or sub-clusters. For example, increasing edge density within sub-clusters from 90 to 98% led to sustained activity in 53.7% of the cases. Increasing edge density within clusters from 17 to 21% led to sustained activity in 85.1% of the cases (table A.1).

References

- [1] Bak T, Tang C and Wiesenfeld K 1987 *Phys. Rev. Lett.* **59** 381
- [2] Newman M E J 2005 *Contemp. Phys.* **46** 323–51
- [3] Beggs J M and Plenz D 2003 *J. Neurosci.* **23** 11167–77
- [4] Haider B, Duque A, Hasenstaub A R and McCormick D A 2006 *J. Neurosci.* **26** 4535–45
- [5] Binzegger T, Douglas R J and Martin K A C 2004 *J. Neurosci.* **24** 8441–53
- [6] Hilgetag C C, Burns G A P C, O’Neill M A, Scannell J W and Young M P 2000 *Phil. Trans. R. Soc. Lond. B* **355** 91–110
- [7] Hilgetag C C and Kaiser M 2004 *Neuroinformatics* **2** 353–60
- [8] Sporns O, Chialvo D R, Kaiser M and Hilgetag C C 2004 *Trends Cogn. Sci.* **8** 418–25
- [9] Anninos P A, Beek B, Csermely T J, Harth E M and Pertile G 1970 *J. Theor. Biol.* **26** 121–48
- [10] Kötter R and Sommer F T 2000 *Phil. Trans. R. Soc. Lond. B* **355** 127–34
- [11] Pastor-Satorras R and Vespignani A 2001 *Phys. Rev. Lett.* **86** 3200
- [12] Dezsó Z and Barabási A L 2002 *Phys. Rev. E* **65** 055103
- [13] Hufnagel L, Brockmann D and Geisel T 2004 *Proc. Natl Acad. Sci. USA* **101** 15124–9
- [14] Latham P E and Nirenberg S 2004 *Neural Comput.* **16** 1385–412
- [15] Eguiluz V M and Klemm K 2002 *Phys. Rev. Lett.* **89** 108701
- [16] Holme P, Kim B J, Yoon C N and Han S K 2002 *Phys. Rev. E* **65** 056109
- [17] Watts D J and Strogatz S H 1998 *Nature* **393** 440–2
- [18] Koch C 1999 *Biophysics of Computation* (New York: Oxford University Press)
- [19] Jucikas T 2006 private communication
- [20] Ravasz E, Somera A L, Mongru D A, Oltvai Z N and Barabási A L 2002 *Science* **297** 1551–5
- [21] Kaiser M and Hilgetag C C 2004 *Phys. Rev. E* **69** 036103
- [22] Izhikevich E M, Gally J A and Edelman G M 2004 *Cereb. Cortex* **14** 933–44
- [23] Salvador R, Suckling J, Coleman M R, Pickard J D, Menon D and Bullmore E 2005 *Cereb. Cortex* **15** 1332–42