On the interaction of adaptive timescales on networks

Elisabeth zu Erbach-Schoenberg[1,2], Connor McCabe[3] and Seth Bullock[1,3],
[1]Institute for Complex Systems Simulation, [2]School of Management, and
[3]School of Electronics and Computer Science,
University of Southampton, Southampton, UK
e.zuerbach-schoenberg@soton.ac.uk

Abstract

The dynamics of real-world systems often involve multiple processes that influence system state. The timescales that these processes operate on may be separated by orders of magnitude or may coincide closely. Where timescales are not separable, the way that they relate to each other will be important for understanding system dynamics. In this paper, we present a short overview of how modellers have dealt with multiple timescales and introduce a definition to formalise conditions under which timescales are separable.

We investigate timescale separation in a simple model, consisting of a network of nodes on which two processes act. The first process updates the values taken by the network's nodes, tending to move a node's value towards that of its neighbours. The second process influences the topology of the network, by rewiring edges such that they tend to more often lie between similar individuals. We show that the behaviour of the system when timescales are separated is very different from the case where they are mixed. When the timescales of the two processes are mixed, the ratio of the rates of the two processes determines the systems equilibrium state. We go on to explore the impact of heterogeneity in the system's timescales, i.e., where some nodes may update their value and/or neighbourhood faster than others, demonstrating that it can have a significant impact on the equilibrium behaviour of the model.

Introduction

Real-world adaptive systems typically involve many interacting parts and processes operating at multiple timescales. However, models of these systems often proceed by identifying a single substantive timescale. Faster processes are often idealised as essentially instantaneous, while slower processes are often treated as a constant background influence that parametrises the model's dynamics.

For instance, Kauffman's (1993) NK landscape model of adaptation on rugged fitness landscapes has a single substantive timescale. At each step the genotype of a genetically fixated population is updated to one of the fitter adjacent genotypes. In reality, a newly discovered fitter mutant takes time to reach fixation. This process is idealised as instantaneous. During a single run of Kauffman's model, the parameters N and K, which determine the length of a genome and the degree of epistasis within it, are held fixed. They parametrise the system's dynamics. Of course in reality both N and K vary as a consequence of evolutionary change. The security of Kauffman's idealisations hinges on whether these processes are separable: the faster processes are much faster, and the slower processes much slower, than the timescale of the process that he focuses on. There are various interpretations of this kind of concept and in the scope of this paper we define *separation of timescales* as follows. The timescales of two processes are separated if one process leads the system into equilibrium before the other process influences the system. This means when the second process sets in, the system has already reached equilibrium.

Where processes take place over similar timescales and affect each other, i.e., they are coupled, dealing with these interacting timescales becomes an important issue. For real-world systems there are further considerations that may be significant. To what extent is there component-wise *heterogeneity* in the rates at which different components operate? While, on average, a genome's alleles might mutate with probability p, it may be the case that some alleles are more vulnerable to mutation than others. While, on average, the children in a schoolyard might update their social ties at rate r, some might update these ties more often than others. Moreover, a system's timescales might vary with time. The traffic on a network might be diurnal, with higher rates during the day. The rate of plasticity in a neural component might decay with the age of the component.

Here we are interested in exploring these issues in the context of adaptive processes modelled on networks. In the coevolutionary networks literature (Gross and Blasius, 2008), two processes are typically modelled: one governing the tendency for nodes to change their *state*, and one governing *topological change* in the network. These two processes may occur on separable timescales where interactions between the two processes can perhaps be neglected. If the timescales for the two processes are not separable, their interplay will affect the behaviour exhibited by the network.

Here we explore a very simple coevolutionary network in which both state and topology evolve over time. We first vary the rates of change for both processes and demonstrate that their ratio impacts on the equilibrium state of the network. We proceed to explore the impact of heterogeneity in timescale, demonstrating that it can impact on both the distribution of node states and the topology at equilibrium. Before introducing the simple model and its results, we review some literature demonstrating the issue of timescale in modelling adaptive systems. We conclude with discussion of the results presented here and ideas for future work.

Dealing with timescale

Synchrony vs. asynchrony

Several studies have revealed that models of adaptive systems can be sensitive to the updating scheme chosen. Using a synchronous model of the iterated prisoner's dilemma, Nowak and May (1992) found complicated spatial patterns within which co-operation persisted. Using an asynchronous update scheme for the same model, Huberman and Glance (1993) found that spatial patterns disappeared with defection the only strategy adopted. While Kauffman (1993) has shown that synchronous Random Boolean Networks can exhibit many stable cyclic behaviours, glossed as analogous to the multiple cell types that may result from the same genome, Harvey and Bossomaier (1997) showed that the same Random Boolean Networks with asynchronous update would tend to evolve to a fixed point.

Multiple timescales

Artificial life has typically considered multiple adaptive timescales in the context of interactions between learning and evolution (Ackley and Littman, 1992; Belew and Mitchell, 1996), such as the Baldwin effect (Hinton and Nowlan, 1987). Further examples where the separation of timescales is critical to adaptive dynamics include the interaction between processes of neurotransmission and (much slower) neuromodulation (Buckley et al., 2004, 2005; Buckley, 2008; Husbands et al., 2010), and the interaction between the evolution of individual behaviours and ecological relationships (e.g., Powers et al., in press; Watson et al., in press; Van Der Laan and Hogeweg, 1995).

Timescales on networks

Most research involving dynamic networks has focused on addressing either the dynamics 'on' a network, or the dynamics 'of' the network (Gross and Blasius, 2008). The dynamics 'on' a network describe the state transitions of the network's nodes, while the dynamics 'of' a network describe topological changes. Research on so-called coevolutionary networks recognises that these processes are inherently reflexive, with network state influencing topological change (as when edges are formed between similar nodes), and topology constraining state change (as when neighbours exchange information) (Blasius and Gross, 2009; Gross and Blasius, 2008; Gross and Sayama, 2009). Coevolutionary networks have been the subject of recent study in the context of the epidemic spread of diseases (Newman, 2002; Zhong et al., 2010; Funk and Jansen, 2010; Van Segbroeck et al., 2010), cascading network behaviour (Watts, 2002), opinion dynamics (Kozma and Barrat, 2008; Demirel et al., 2011), diffusion of innovations / information (Onnela and Reed-Tsochas, 2009; Ke and Yi, 2008), evolution of social groups (Palla et al., 2007), the growth of social networks (Sun and Wang, 2008), co-operation (Pacheco et al., 2006; Van Segbroeck et al., 2009), community formation (Bryden et al., 2010), synchronisation (Zhu et al., 2010) and global adaptation (Watson et al., in press). The dynamical interplay of state update and rewiring processes are typically central to the evolution of these systems.

Heterogeneous timescales

Typically, models make a simplifying assumption that all components update their state at a shared characteristic rate, while structural relationships change at some other arbitrary rate. However, some models have explored systems with heterogeneous rates. Van Segbroeck et al. (2009), for instance, found that increased diversity in their model accelerates the rate of evolution to an equilibrium state where co-operation is a robust and dominant strategy. Pacheco et al. (2006) employed variable re-wiring rates in a social agent model. Their results suggest that introducing heterogeneity has an effect on the system as a whole which can change the frequency of co-operation observed at equilibrium.

A simple model

To study the influence of timescale separation we introduce an abstract model based on models of opinion dynamics that include adaptive change in network topology as well as the spread of opinions over the network (e.g., Kozma and Barrat, 2008). Here, nodes have an internal value and tend to update this value in the direction of their neighbours' values. The second process changes the network topology by rewiring edges between nodes such that nodes disconnect from dissimilar neighbours and connect to nodes with more similar values.

To illustrate what kind of processes this model could be related to we could assume that each node's value represents the opinion of a different person and that edges represent social interactions between people. In this setup, we can imagine that either rewiring or state update might be the faster process. If we assume a node's value represents something such as the religion a person believes in or a political affiliation, we can assume that this value changes very slowly. We can further assume that therefore a person would more readily change to associate with individuals sharing a similar opinion than change their own opinion to match that of their neighbours. In this case, the rewiring process would be faster than the value update process. At the other end of

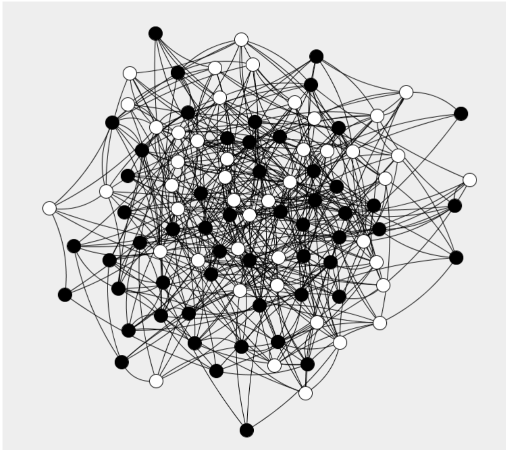

Figure 1: A typical network after initialisation.

the spectrum, we could assume a node's value represented a person's preference for meeting friends at one restaurant rather than another. In this case the individuals would be likely to change their opinion based on the opinions of their friends, rather than changing their friends on the basis of their restaurant preference. In between these two extremes we can think of intermediate cases where individuals have a preference for socialising with individuals that share a similar opinion, but also change their own opinion towards that of their neighbours.

The model

The model consists of a network of N interconnected nodes (here $N = 100$). Each node has a single value in the interval $[0.0, 1.0]$. Even though a node's value can be any value between 0 and 1, each starts with the value 0.0 or 1.0, with equal probability. Nodes are connected by undirected, unweighted edges, meaning an edge is either present or absent and if node a cuts a tie to b, b also loses it's connection to a. Self-connections are not allowed. To initialise the network between the nodes, we specify an average degree d and generate a random network by making an edge between each possible pair of unique nodes with probability $\frac{d}{N-1}$. In the examples presented here, we use an average degree of $d = 10$. A visualisation of a typical network after initialisation is given in Figure 1.

Value Update: When a node i updates its value, it chooses a random individual n from the set of its neighbours. It then discovers the value of its neighbour $v(n)$ and calculates the difference $v(n) - v(i)$ between the neighbour's value and its own. The node then updates its state towards the state of its neighbour, proportional to the difference in values: $v(i)_{t+1} = v(i) + m(v(n) - v(i))$. The factor m determines the maximal change that can occur in one step. Here we choose $m = 0.01$, to ensure that it takes several updates for two nodes to reach the same value. If the updating node and its chosen neighbour have the same state, i.e., $v(i) = v(n)$, the update results in no change to $v(i)$.

Rewiring: When node i rewires, it compares its own value to the values of its neighbours, identifying the neighbour with which it is most dissimilar, n. The node i then generates a list of all neighbours of all of its neighbours, comprising all nodes that are two edges away. Members of this list that are already neighbours of i are discarded. If the list is non-empty, i drops the connection to n and rewires this edge to a randomly chosen member of the list of neighbours' neighbours. This implies that, if an individual is already connected to all neighbours of its direct neighbours, an attempt to rewire will result in no topological change.

Timescales: In each step of the algorithm, a list is generated containing all nodes that are ready to update their state in the current time step. These nodes are then updated in a random order, one at a time. After this, the same procedure is repeated for all nodes ready to rewire. Whether a node is ready to update or rewire depends on the timescales of the two processes. The relation of the timescales is incorporated in the model as follows. Each node is assigned two values, V_i and R_i, specifying the number of time steps in the interval between two consecutive value updates for i and two consecutive rewiring events for i, respectively. In the case of homogeneous timescales, all nodes have identical values for V and identical values for R, i.e., $\forall i V_i = V$ and $\forall i R_i = R$. In the case of heterogeneous timescales this constraint does not hold and values for the two rates may differ from node to node. The algorithm stops when neither the state update nor rewiring process effects any change in the network. We will consider this stopping criterion in more detail next.

Equilibriation: Both the value update and the rewiring process can only change the system's state if there is a local difference between two nodes. A local difference is present if two nodes that are connected by an edge have non-identical values. This difference can be reduced by updating the value of one or both nodes or by deleting the edge between the two nodes and rewiring it to a node with a more similar value. Once there there are no local differences in the system anymore, neither the value update process nor the rewiring process change the system's state when invoked. Therefore both processes need a value difference between connected nodes to operate. Thus, we can see the difference in values between connected nodes as some kind of energy available to the two processes to use for changing the system's state. Both processes can only operate if there is energy left in the system and both processes reduce the energy, at least locally. One way of formally defining this energy is as the sum of absolute value differences between all pairs of connected nodes, $e = \sum_{i,j \text{ connected}} |v(i) - v(j)|$.

The energy specified in this way reduces over time and once it has reached zero, the system's state cannot change any more. Therefore, we can use reaching zero energy as a formal stopping criterion and terminate the algorithm when the energy has reached zero. Note that, from the initial conditions considered here, each process is capable of reducing energy to zero in the absence of the other.

In the case of homogeneous timescales, whether they are separated depends on the ratio of the values V and R. For $R >> V$ the timescales of the two processes are separated, with only the value update process influencing the dynamics. We also have separation of timescales in the opposite case, $V >> R$, where the rewiring process dominates the dynamics. Let us now specify further when exactly the timescales are separated to find values for the parameters V and R for which we can be certain the timescales are separate. Based on the definition presented in the introduction, the timescales of the two processes are separate if one process acts after the other process has reached equilibrium. Based on the equilibrium definition as a zero energy state, we define the equilibrium points t_R and t_V as the number of steps the rewiring or value update process takes in isolation to reduce the energy of the system to zero and therefore reach equilibrium. We measure these two points for a particular set of initial conditions by running the algorithm with only one of the two processes operating. Measuring the time the system takes to reach zero energy when only one process acts on it is the equilibrium time for that process, t_V or t_R. If the second process acts only after the system has reached equilibrium, it is unable to change the system state as there is no energy for it to exploit (i.e., no value difference between connected nodes). This means that the timescales of the two processes are separated in two cases. The first case is when $V > t_R$, meaning that the value update only happens after the rewiring process has brought the system to equilibrium. In the second case, for $R > t_V$, the rewiring process happens after the value update process has already reduced the system's energy to zero. In any other case the timescale are mixed to some degree.

Results

We now observe the system behaviour for varying ratios $\frac{R}{V}$, first for homogeneous timescales and then for varying degrees of heterogeneity.

Homogeneous timescales

If only the rewiring process is active and its rate is the same for all nodes, the system reaches equilibrium after $t_R \approx 10.0$ steps. If only the state update process is active, it takes longer for the system to reach equilibrium, $t_V \approx 6500$. Having measured these values, we can assign values to the parameters V and R for which the timescales are separated and one of the two processes dominates the dynamics.

Setting $V = 100$ and $R = 1$[1] the timescales are separated as $V > t_R = 10$, with only the rewiring process influencing the dynamics as it reaches equilibrium before the state change process has time to affect the network.

The equilibrium state of the system under these parameters is shown in Figure 2a[2]. Since the network is initially populated by equal numbers of nodes with value 0.0 and value 1.0, the rewiring process removes edges between dissimilar nodes and replaces them with edges linking nodes with identical value, forming two homogeneous components, one containing all the nodes with value 0.0 and the other containing the nodes initialised with value 1.0. At the other extreme, $V = 1$ and $R = 10000 > t_v$, only the state update process shapes the network. Figure 2j depicts the equilibrium state under these conditions. Node values have gradually changed towards the average value of the initial population until all nodes have exactly this value. Since all the nodes have identical values no rewiring can take place and the network topology does not change at all.

Intermediate cases where the timescales are mixed are shown in Figures 2b–2i. Where the rewiring process is fast relative to the state update process, the network breaks up into several components, each eventually consisting of nodes with the same value, but with values differing significantly between the components (see, e.g., Figure 2b). Where the system's dynamics are more influenced by the state update process (see, e.g., Figure 2e) the values adopted by different components tend to be less diverse and closer to the system mean. Eventually, the state update dynamic is fast enough to equilibrate the network before the rewiring process can cause it to fragment (see, Figures 2h–2j).

These results show that the system reaches the predicted equilibrium when the timescales are separated. For the intermediate cases with mixed timescales however, the ratio between the two timescales determines which equilibrium the system ends up in and the character of this equilibrium, in terms of the node values and the network topology.

Figure 3 depicts how the distribution of node values at equilibrium varies with $\frac{R}{V}$. It shows that for very low values of $\frac{R}{V}$, the rewiring process dominates the system dynamics and only the initial values (0.0 and 1.0) are present. As $\frac{R}{V}$ increases, we observe more and more intermediate values, converging to the average value in the system. For high values of $\frac{R}{V}$, there is only one value present in the system, corresponding to the mean of the system's initial values.

A similar transition can be observed for the topology of the network. Figure 4 depicts how the distribution of component sizes at equilibrium varies with $\frac{R}{V}$. Here we observe that when rewiring dominates, the two network components have nearly the same size, consisting of roughly half of the nodes each (one is larger as a consequence of the ini-

[1]Note that the equilibrium times measured above assume that the process considered happens each time step ($V = 1$ or $R = 1$). We therefore set the frequency of the faster process to 1.

[2]In the examples presented here, the same initial network shown in Figure 1 is used.

tial random allocation of value to the population of nodes). For mixed timescales, components are smaller and isolated nodes (with component size 1) exist. As we move towards

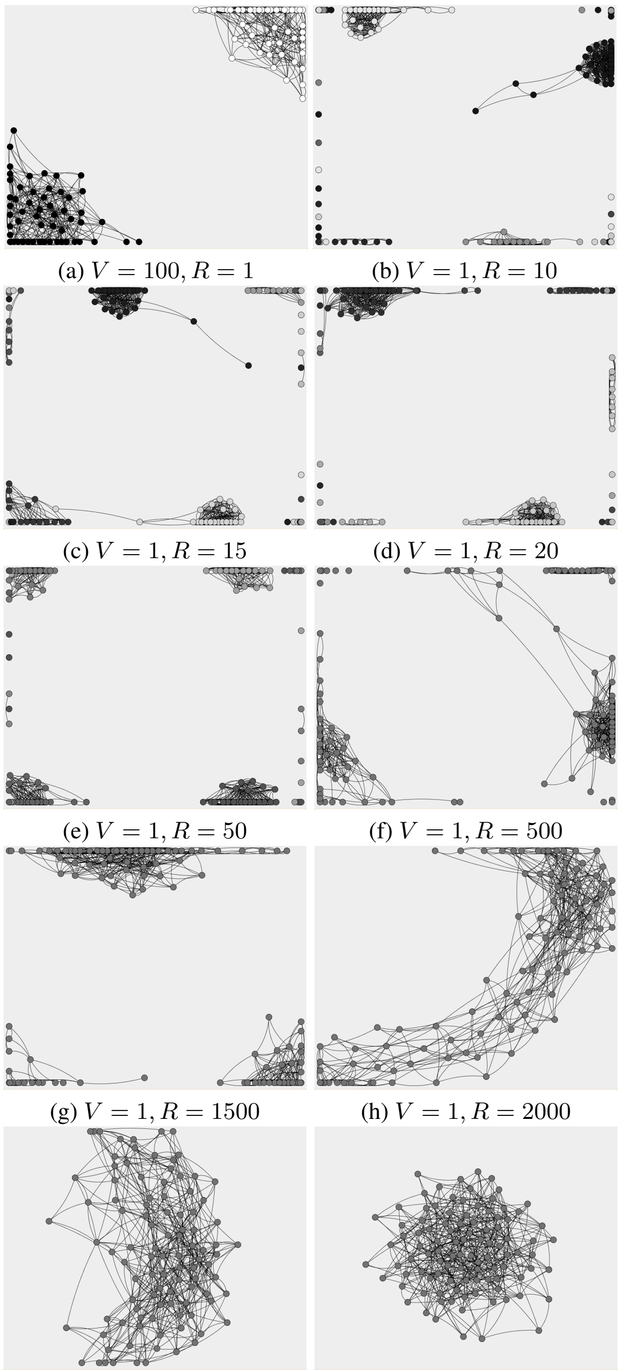

(a) $V = 100, R = 1$ (b) $V = 1, R = 10$

(c) $V = 1, R = 15$ (d) $V = 1, R = 20$

(e) $V = 1, R = 50$ (f) $V = 1, R = 500$

(g) $V = 1, R = 1500$ (h) $V = 1, R = 2000$

(i) $V = 1, R = 3000$ (j) $V = 1, R = 10000$

Figure 2: Networks at equilibrium for different values of R and V. Node shading indicates the nodes states, with the heaviest shading indicating 0.0 and no shading indicating 1.0.

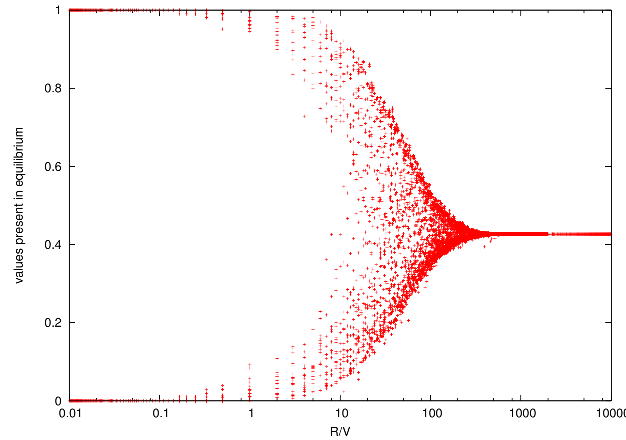

Figure 3: Values present in the equilibrium state for different ratios $\frac{R}{V}$. For each value the system is started with the same initial conditions.

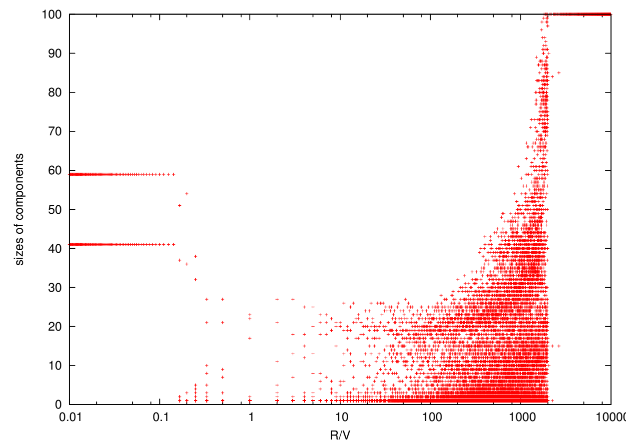

Figure 4: Component sizes present in the equilibrium state for different ratios $\frac{R}{V}$. For each value the system is started with the same initial conditions.

the regime where the state update process dominates the systems dynamics, larger components exist at equilibrium. Once state update is the only active process, only one connected component is present at equilibrium.

Comparing these two graphs, we observe that the apparent thresholds in system behaviour exhibited by node values and network topology are different. From the perspective of node values, we can see three regimes separated by two threshold values of $\frac{R}{V}$. First, a transition occurs around $\frac{R}{V} = 0.5$, with a second qualitative change in the equilibrium behaviour at around $\frac{R}{V} = 250$. However, when we consider the network's equilibrium topology, the equivalent transitions seem to occur around $\frac{R}{V} = 0.1$ and $\frac{R}{V} = 2500$.

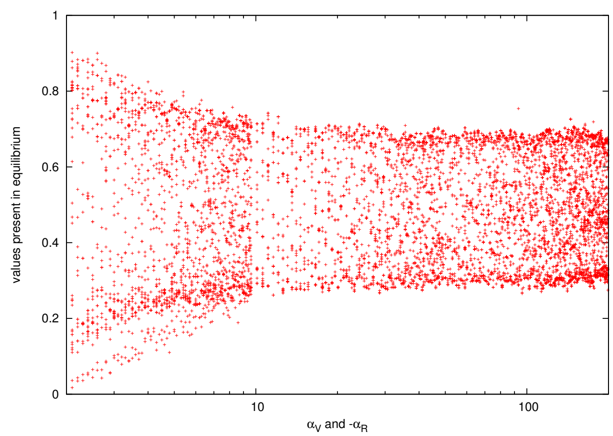

Figure 5: Node values present at equilibrium for different levels of heterogeneity in the value update and rewiring intervals. Only α, which specifies the heterogeneity of timescales, is varied here (with $\alpha = \alpha_V = -\alpha_R$). The runs are otherwise identical and share the same initial conditions. For each value of α all node values, $v(i)$, present at equilibrium are displayed.

Heterogeneity in timescales

We now consider the case where some nodes might update their value or their neighbourhood faster than others. We model this by allocating each node, i, a pair of values, V_i and R_i, governing the individual rates of change for value and neighbourhood, respectively. The V_i and R_i values are Pareto-distributed, meaning that while most of the values are close to the characteristic population mode, V or R, a few are significantly different, due to the long tail of the distribution. Values are generated by transforming a uniform random variable U by the functions $\frac{V}{U^{1/\alpha_V}}$ for V_i and $\frac{R}{U^{1/\alpha_R}}$ for R_i (Newman, 2004). The parameters α_V and α_R determine the spread of values in each distribution. We set α_R to a negative value and α_V to a positive value so that the tails of the distribution point towards each other. Large absolute magnitudes for α (such as $\alpha = 100$) lead to a relatively small average distance between the resulting values and the modal value, V or R, whereas small absolute values for α (e.g., $\alpha = 2.5$) produce a larger spread. This introduction of heterogeneity into the model means that each node has its own internal clocks governing when to update its state and when to rewire.

The effect of heterogeneity on the value process is assessed for the case in which $V = 1$ and $R = 50$, as this is an intermediate case where both processes influence the dynamics. Figure 5 shows that for a low degree of heterogeneity in both processes (higher values of α) the distribution of values present at equilibrium is not very different from the base case without heterogeneity. For higher levels of heterogeneity, however, the diversity of values increases significantly. The effects of heterogeneity on the network topology are illustrated in Figures 6 and 7 for $V = 1$ and $R = 2000$, as this ratio of $\frac{R}{V}$ is the threshold separating single component equilibria from multi-component equilibria. Without heterogeneity, the network forms one component (Figure 7a) with a degree distribution that differs from that of the initial network (compare Figures 6a and 6b). In the presence of heterogeneity however, the network fragments into eleven components (Figure 7b) with a qualitatively different degree distribution (Figure 6c).

Discussion

The results presented here show that in the cases where the timescales are separated, the system behaves as we would expect: if only the value update process is active, there is no topological change and the values of all nodes converge to the average of the initial network. If only the rewiring process acts on the system state, we only observe changes in topology and the network splits into two components, with nodes being sorted according to their initial value. The number of components in that case depends only on the number of initial values present in the system. For example, if we initialise the system with three (e.g. 0.0, 0.5, 1.0) different values instead of two the network fractures into three clusters. To sum up, when timescales are sufficiently separated, the system behaves in the same way as an equivalent system with the slower process ‘switched off’.

The results also show that if the timescales are not separated, the exact ratio between the rates of the two processes influences the system’s equilibrium state. If the rewiring process dominates the dynamics, the values we find in the system in equilibrium differ significantly. As the value update process gains more influence, the values of the components found in the equilibrium state become more and more similar. We can explain this behaviour by observing the system dynamics over time. Starting from a random initial network, the rewiring process stretches the network into a predominantly white and a predominantly black end. In between, there are nodes of intermediate value. At this stage, if the rewiring is fast, the network fractures at several points. In the case where the value process is the main influence on the system, the values of nodes are more similar at the point when the rewiring sets in, as sufficient time has passed for the node values to become more similar. Therefore, the rewiring fractures the network into fewer and larger clusters.

Furthermore, we have shown that heterogeneity changes the state the system reaches in equilibrium. Although the influence of heterogeneity is clearly visible, it is not as strong as we had anticipated. The heterogeneous case needs to be investigated further as we do not fully understand how heterogeneity in the rates influences the dynamics.

We have presented a definition for timescale separation in the case of homogeneous and therefore well defined rates, but we need an extended definition for the case of hetero-

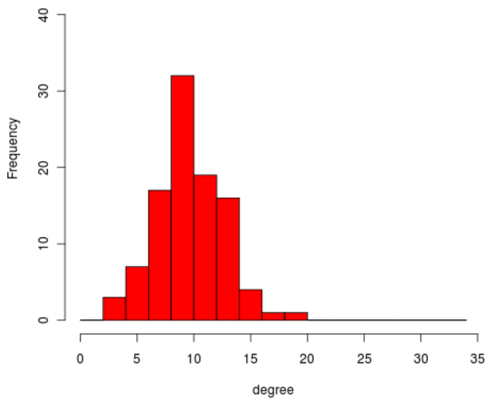

(a) initial distribution

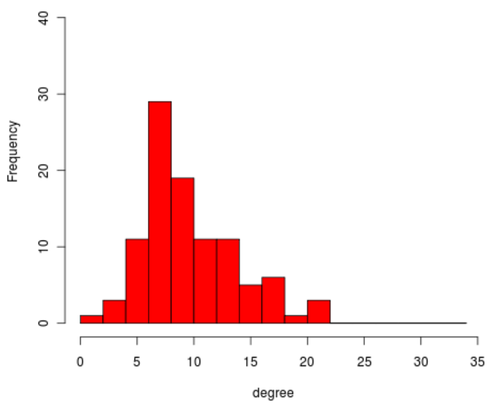

(b) no heterogeneity

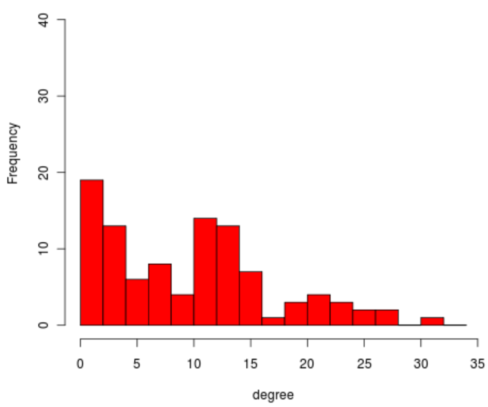

(c) heterogeneity in rates of both processes

Figure 6: Histogram of the degree distribution under different conditions. The rate parameters used are $V = 1$ and $R = 2000$. In the heterogeneous case the α values are $\alpha_V = 2.1$ and $\alpha_R = -4.1$.

geneous rates. There are of course further complications in real-world systems that we have not considered in the model presented here. For example, processes often have dynamic rates, i.e., the change of the rate is a process itself, perhaps influenced by the current state of the system.

(a) no heterogeneity

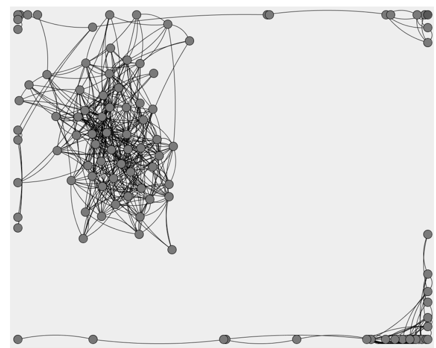

(b) heterogeneity in rates of both processes

Figure 7: The effect of heterogeneity on network topology for $V = 1$, $R = 2000$, $\alpha_V = 2.1$ and $\alpha_R = -4.1$.

Conclusions

In this paper we have presented an initial investigation of timescale separation in adaptive networks, by identifying examples from the literature of different ways of dealing with multiple timescales and proposing a definition of timescale separation, based on the time taken by a system to reach equilibrium under the action of individual processes. Given this definition, we confirmed that, if the timescales of two processes are sufficiently separated, we can ignore their interaction. Where timescales do not separate cleanly, however, the system dynamics exhibit higher variability and hence become more difficult to predict. Heterogeneity complicates matters further as it can result in the system relaxing to different equilibria in comparison to the same system under homogeneous conditions. Where we can not be certain that the timescales are sufficiently separated in a system under consideration, we should expect the dynamics to be sensitive to the interplay between the timescales of the processes present.

Acknowledgements

This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1) and grant (EP/G036926/1.)

References

Ackley, D. H. and Littman, M. L. (1992). Interactions between learning and evolution. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 487–509. Addison-Wesley, Redwood City, CA.

Belew, R. K. and Mitchell, M., editors (1996). *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison-Wesley, Reading, MA.

Blasius, B. and Gross, T. (2009). Dynamic and topological interplay in adaptive networks. In Schuster, H. G., editor, *Annual Reviews of Nonlinear Dynamics and Complexity*, volume 2, chapter 3, pages 63–106. Wiley-VCH, Weinheim.

Bryden, J., Funk, S., Geard, N., Bullock, S., and Jansen, V. A. A. (2010). Stability in flux: community structure in dynamic networks. *Journal of the Royal Society Interface*, 7:1257–1274.

Buckley, C., Bullock, S., and Cohen, N. (2004). Toward a dynamical systems analysis of neuromodulation. In Schaal, S., Ijspeert, A. J., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J.-A., editors, *Eighth International Conference on Simulation of Adaptive Behavior*, pages 334–343. MIT Press, Cambridge, MA.

Buckley, C., Bullock, S., and Cohen, N. (2005). Timescale and stability in adaptive behaviour. In Bentley, P., Capcarrere, M., Freitas, A. A., Johnson, C. G., and Timmis, J., editors, *Eighth European Conference on Artificial Life*, pages 292–301. Springer, Berlin.

Buckley, C. L. (2008). *A systemic analysis of the ideas immanent in neuromodulation*. PhD thesis, School of Electronics and Computer Science, University of Southampton, UK.

Demirel, G., Prizak, R., Reddy, P. N., and Gross, T. (2011). Cyclic dominance in adaptive networks. *The European Physical Journal B - Condensed Matter and Complex Systems*, pages 1–8.

Funk, S. and Jansen, V. A. A. (2010). Interacting epidemics on overlay networks. *Physical Review E*, 81(3):036118.

Gross, T. and Blasius, B. (2008). Adaptive coevolutionary networks: a review. *Journal of The Royal Society Interface*, 5(20):259–271.

Gross, T. and Sayama, H., editors (2009). *Adaptive networks: Theory, Models and Applications*. Springer, New York.

Harvey, I. and Bossomaier, T. (1997). Time out of joint: Attractors in asynchronous random Boolean networks. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 67–75. MIT Press, Cambridge MA.

Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.

Huberman, B. A. and Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences USA*, 90(16):7715–7718.

Husbands, P., Philippides, A., Vargas, P., Buckley, C. L., Fine, P., Di Paolo, E., and O'Shea, M. (2010). Spatial, temporal, and modulatory factors affecting gasnet evolvability in a visually guided robotics task. *Complexity*, 16(2):35–44.

Kauffman, S. A. (1993). *The Origins of Order*. Oxford University Press.

Ke, H. and Yi, T. (2008). Information diffusion on adaptive network. *Chinese Physics B*, 17(10):3536.

Kozma, B. and Barrat, A. (2008). Consensus formation on adaptive networks. *Physical Review E*, 77(1):016102.

Newman, M. (2004). Power laws, Pareto distributions and Zipf's law. *Arxiv preprint cond-mat/0412004*.

Newman, M. E. J. (2002). The spread of epidemic disease on networks. *Physical Review E*, 66(1):016128.

Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829.

Onnela, J. P. and Reed-Tsochas, F. (2009). The spontaneous emergence of social influence in online systems. *Proceedings of the National Academy of Sciences*, 107(43):18375–18380.

Pacheco, J. M., Traulsen, A., and Nowak, M. A. (2006). Coevolution of strategy and structure in complex networks with dynamical linking. *Physical Review Letters*, 97(25):258103.

Palla, G., Barabási, A., and Vicsek, T. (2007). Quantifying social group evolution. *Nature*, 446(7136):664–667.

Powers, S. T., Penn, A. S., and Watson, R. A. (in press). The concurrent evolution of cooperation and the population structures that support it. *Evolution*.

Sun, C. and Wang, S. (2008). Modeling adaptive behaviors on growing social networks. In *Fourth International Conference on Natural Computation*, volume 1, pages 465–469, Los Alamitos, CA, USA. IEEE Computer Society.

Van Der Laan, J. and Hogeweg, P. (1995). Predator-prey coevolution: interactions across different timescales. *Proceedings: Biological Sciences*, 259(1354):35–42.

Van Segbroeck, S., Santos, F. C., Lenaerts, T., and Pacheco, J. M. (2009). Reacting differently to adverse ties promotes cooperation in social networks. *Physical Review Letters*, 102(5):058105.

Van Segbroeck, S., Santos, F. C., and Pacheco, J. M. (2010). Adaptive contact networks change effective disease infectiousness and dynamics. *PLoS Computational Biology*, 6(8):e1000895.

Watson, R. A., Mills, R., and Buckley, C. L. (in press). Global adaptation in networks of selfish components: Emergent associative memory at the system scale. *Artificial Life*.

Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences USA*, 99(9):5766–5771.

Zhong, L.-X., Qiu, T., Ren, F., Li, P., and Chen, B. (2010). Time scales of epidemic spread and risk perception on adaptive networks. *Europhysics Letters*, 94(1):18004.

Zhu, J.-F., Zhao, M., Yu, W., Zhou, C., and Wang, B.-H. (2010). Better synchronizability in generalized adaptive networks. *Physical Review E*, 81(2 Pt 2):026201.