Coevolution on tunable fitness landscapes

Benjamin Blonder

15 June 2007

Abstract

I examine the statistical properties of a NKC model of biological evolution in which species interact and compete to maximize their individual fitness. Species are modeled as adaptive walkers moving on deformable fitness landscapes. The nature of these landscapes is controlled by a small set of free parameters; by tuning them, I investigate the different phases of the model and search in particular for a critical phase. The model is also extended to include restricted connectivity between species and also fitness fluctuations dependent on an effective temperature.

1 Introduction

A predictive theory of evolution has yet to be developed; it is unknown how short-time effects at the molecular level translate into long-time evolution at the species and ecosystem level. For example, *Nanoarchaeum equitans* has a the smallest non-viral genome known with length 4.9×10^5; there are thus approximately 10^{300000} possible genotypes for the species. One might wonder how species find the best genotype in such a huge space, what sort of complexity natural selection is capable of generating, by which types of mutations the process is driven, and how the genes of each species are coupled. It is hoped that simple models of evolution on digital organisms with shorter genomes will be capable of giving rough answers to some of these questions, as experiments conducted over millions of years on billions of individuals are clearly impossible.

As a simple approximation, the adaptive evolution of organisms is driven by natural selection at the molecular level: genetic mutations create more and less fit individuals relative to the existing organisms; the fitter ones are better able to reproduce and so transmit their genes to their offspring. Over time, the frequency of genes in the population changes, so that eventually the entire population possesses the new mutation. In this approximation, the process of natural selection causes each species to maximize its fitness by promoting certain beneficial mutations in the genome. Each species locally adapts by optimizing its genome - and so expresses certain traits. This view is now being challenged by some scientists,[16] but remains the dominant model of evolution.

The evolution of a species can then be modeled as an adaptive walk through the space of all possible genotypes. This walk is constrained locally, so that the species can move only to nearby genotypes. Most landscapes are correlated, in that nearby genotypes have

similar fitnesses, while far-away genotypes will not necessarily have similar fitnesses. If every genotype corresponds to a different fitness value, then the adaptive walk locally seeks a genotype with a higher fitness. This is an optimization process in which species climb a fitness landscape to its maxima. However, this adaptive walk may be a complex process, dependent on both the method species use to climb the landscape and on the landscape itself, which may not be constant in time - genotypes which were once very fit may not always remain so. Species which interact with each other may deform the landscape. All these notions will be made more precise later.

The idea of an adaptive landscape was first proposed by Wright[26] in 1932, while that of a genotype/protein space was first considered by Smith[19] in 1970. Adaptive walks on such landscapes were proposed by Levin[15] in 1978 and Agur and Kerszberg[1] in 1987. The NK model itself was invented by Kauffman[14] in 1987, and has since been studied and extended by many others. Other toy models of evolution also exist.[2]

In this report I describe some of the important statistical properties of a model of evolutionary processes. I first reintroduce the NK model of fitness landscapes and survey some natural questions and important results for adaptive walks on different classes of fitness landscapes. I also describe a novel class of adaptive walkers which are capable of exploring a larger fraction of genotype space. Couplings between species are then introduced to extend the $NK(C)$ model to include coevolution. The model is described by four important parameters: S, the number of species, N, the number of genes for each species, K, the number of couplings between each gene and other genes within a species, and C, the number of couplings between each gene and other genes external to the species. By varying the ratio of K and C to N and S, it is possible to vary the form of the fitness landscape and the evolutionary dynamics. Using this model I deduce some general properties of coevolution, and find the model is capable of fixed, periodic, punctuated, and chaotic behavior.

I also discuss the concept of criticality and its possible relevance to scale-free systems and a general theory of evolution. I investigate the capacity of the NKC model to display critical behavior through a variety of methods, and then further extend the model to approximate ecosystems with restricted (nearest-neighbor) interspecies connectivity. I also introduce the idea of a temperature in order to allow for negative mutations. I find that the NKC model, with or without these extensions, does not display critical behavior.

2 Coupled fitness landscapes: the NK model

2.1 Overview

Mutations are the generators of genetic variability and can occur in many different ways. Most common are point mutations which affect a single DNA base pair and can result in the insertion, deletion, or change of base. However many more complex mutations involving longer sections of DNA can occur, where parts of genes or whole genes are created, destroyed, or moved. In sexual organisms recombination is also possible, generating

further diversity. Some mutations change DNA in ways that may have no immediate effect on the organism's function. Mutations can occur for many reasons such as environmental damage, errors made by DNA-copying/repairing proteins, recombination, or viral DNA integration. Additionally, some regions of the genome are more susceptible to mutations than others. Thus, a real organism's genome is capable of complex changes over short time periods and is difficult to model perfectly.

The NK model simplifies this complexity.[12] A species is assumed to have N genes, each with A alleles. Each gene makes an additive fitness contribution to the total (normalized) fitness F of the species; the magnitude of this contribution depends on the gene itself and on the state of $K(< N)$ other genes. Thus, there are A^{K+1} combinations—and thus possible values for the fitness contribution—for each gene. Each combination is assigned a weight value W between 0 and 1. If $\psi_{in}(G)$ determines the allele of nth coupling of the ith gene of a genotype G, then the fitness for a given genotype is

$$F(G) = \frac{1}{N} \sum_{i=0}^{N-1} W_i(\psi_{i,1}(G), \psi_{i,2}(G), \ldots, \psi_{i,K+1}(G)) \tag{1}$$

Without loss of generality, I assume $A = 2$ (larger values of A should be similar to systems with larger values of N and smaller values of A). Genotypes can thus be written simply as binary strings: if $N = 5$, a possible genotype might be 01101. K is intended to capture epistatic couplings between genes: the fitness of one gene may depend on the state of other genes. Epistatic couplings are not symmetric; if gene a depends on gene b, gene b need not necessarily depend on gene a. For simplicity, I assume that the initial genotype, epistatic couplings, and weights should be assigned at random.

2.2 An example

Consider the simple case $A = 2, K = 2, N = 3$ in which every gene is epistatically coupled to every other gene. A random weight table for the three genes may be generated:

0.0	0.2	0.1	W	1.1	1.0	1.2	W	2.2	2.0	2.1	W
0	0	0	0.582230	0	0	0	0.297102	0	0	0	0.906845
1	0	0	0.809567	1	0	0	0.426051	1	0	0	0.294026
0	1	0	0.591919	0	1	0	0.899498	0	1	0	0.936244
1	1	0	0.511713	1	1	0	0.652999	1	1	0	0.414645
0	0	1	0.876634	0	0	1	0.901534	0	0	1	0.308457
1	0	1	0.995085	1	0	1	0.961533	1	0	1	0.514893
0	1	1	0.726212	0	1	1	0.164713	0	1	1	0.395430
1	1	1	0.966611	1	1	1	0.857987	1	1	1	0.789785

The associated fitnesses can then be calculated:

G		F
000	$0.582230 + 0.297102 + 0.906845$	0.595392
100	$0.809567 + 0.899498 + 0.936244$	0.881769
010	$0.876634 + 0.426051 + 0.308457$	0.537047
110	$0.995085 + 0.652999 + 0.395430$	0.681171
001	$0.591919 + 0.901534 + 0.294026$	0.595826
101	$0.511713 + 0.164713 + 0.414645$	0.363690
011	$0.726212 + 0.961533 + 0.514893$	0.734213
111	$0.966611 + 0.857987 + 0.789785$	0.871461

2.3 Fitness landscapes and adaptive walks

It is now possible to define and map the fitness landscape. There are A^N possible genotypes, each of which corresponding to a vertex on a hypersolid of dimension N. Each vertex has $(A-1)N$ nearest neighbors. The fitness of each vertex (genotype) can be calculated, and so local and global fitness maxima can be identified. Suppose the species is capable of jumping a distance D within a certain amount of time. At any instant, the species evaluates its neighborhood and jumps to a genotype within this neighborhood with higher fitness; this is known as an adaptive walk and corresponds to an uphill movement on the fitness landscape. This iterative mutation and selection process is only one of the possible methods for exploring and finding optima in genotype space.

There are many possible types of adaptive walks, capable of taking steps of different sizes and of evaluating fitness gradients with more or less accuracy. At one extreme is the 'greedy' walker with a maximal value of D; it is capable of seeing the entire fitness landscape and always walks up the steepest gradient. In a single step it would obtain the globally maximal fitness value. At the other extreme is a random walker, which is unlikely to obtain any maximal value for any length of time. I have chosen to probe the landscape with a walker that can choose from any nearest-neighbor genotype ($D = 1$). At any instant, the walker selects a random nearest-neighbor. If this neighbor has a higher fitness than the current genotype, a jump is made; if not, nothing happens. This is called a 'fitter $D = 1$' walker. This is an approximation for the common class of organisms that are primarily subject to point mutations. It is unreasonable to expect an organism with $A = 4, N \approx 4.0 \times 10^9$ to be greedy and always choose the best point mutation, since the time required to explore $4^{4\times10^9} \approx 10^{240823997}$ mutations is presumably long relative to the time between mutations. I have also investigated other walkers, the details of which shall be described later.

2.4 Assumptions and limitations

The NK model greatly simplifies biology, but in doing so hopes to capture two essential features of evolution: interactions between genes and genotype optimization by local search. It is helpful to explicitly state the model's assumptions to compare them to reality:

1. A population of organisms (a species) can be described by a single genotype. Mutations occur sufficiently infrequently that they propagate to all members of the species before further mutations occur, so that there is no genetic variability between organisms. This is limiting behavior; in reality, it would be necessary to consider a large number of individuals with a distribution of genotypes about the mean.

2. All species are haploid. There is a single copy of each gene, and no recombination, sex, or conjugation occurs. This eliminates from consideration most higher taxa, but is a good first approximation for bacteria.

3. Mutations occur at a constant rate with equal probability anywhere on the genome. Actually, certain regions of the genome are well-established mutation hot-spots, and some particularly important regions of the genome may be resistant to mutation as any changes would be lethal.

4. The genome length is constant, so that only substitution mutations occur. In reality, different types of organisms have different genome sizes and other mutation types such as insertions and deletions are quite common.

5. The genotype of the organism is the sole determinant of a species' fitness.

6. All genes are equally important and equally coupled to other genes.

2.5 Tunable landscapes

The power of the NK model lies in its ability to characterize a wide range of fitness landscape types. By varying the value of K, it is possible to tune the landscape from singly-peaked and highly correlated to multiply peaked and uncorrelated (a landscape is highly correlated if the fitness at one genotype is similar on average to the fitness at a nearby genotype). Changing K changes the number of interconnections between genes and thus the number of conflicting constraints. It is very unlikely that genes a and b maximize fitness with the same set of inputs; thus, accepting a mutation that is optimal for gene a may be harmful to gene b. Instead of there being a single optimal solution for all genes, there are many compromise solutions which correspond to local fitness maxima. At larger values of K, fitness maximization becomes more difficult, since a small change in one part of the system will have a large effect on the rest of the system. Intuitively, smaller values of K create fewer, higher, fitness peaks, while larger values of K create many small fitness peaks.

Consider the two limiting cases. For $K = 0$, there are no couplings between genes, so every gene is independent. This is a landscape with a single fitness maximum which can be accessed by a 'fitter' walker, since any genotype can be converted to the optimal genotype by a sequence of individual independent mutations. Suppose without any loss of generality that the ideal genotype were 11111; any other genotype (such as 01000) would need only to mutate every 0 to a 1 in order to obtain the optimal genotype. Since

the genes are independent, this process is independent of path. This is also a smooth landscape, since each gene can contribute at most $1/N$ to the total fitness: there are no drastic fitness drops possible. Thus, $K = 0$ landscapes are also highly correlated - since each gene has an independent contribution to the total fitness, nearby genotypes also have close fitness values.The process of reaching the maximum takes approximately $N/2$ adaptive steps, since a random initial genotype is likely to consist equally of 1s and 0s. Thus, global optimization on a $K = 0$ landscape is a rapid and assured process.

At the other limit, $K = N-1$, the fitness landscape is rugged and becomes completely random. Mutating any single gene affects the fitness contribution of that gene and of every other gene. Since the new fitness contributions (weights) are drawn from a flat random distribution, the new fitness value of the mutated species is a random number independent of the previous fitness value. This is a completely uncorrelated landscape. On such a random landscape, there are many local peaks. At any given genotype, there are $N(A-1)$ nearest-neighbors. The probability that the current genotype is a local maximum is then just $1/(N(A-1)+1)$. Since there are A^N possible genotypes, the number of local maxima is $A^N/(N(A-1)+1)$. For $N = 100, A = 2$, there are 1.25×10^{28} local maxima. The probability that an adaptive walker will find the global maximum is very small indeed. This probability is particularly small if the adaptive walker is a 'fitter' one, for as soon as it encounters a local maximum it will stop its walk. Adaptive walks for 'fitter' walkers are very short, on the order of $\log N$.[14] Walkers become quickly trapped at local maxima, unable to reach distant higher peaks. It is difficult to traverse a large range of genotype space on a random landscape.

Numerical simulations support these conclusions. As shown in Figure 1, higher values for K result in shorter adaptive walks. Additionally as shown in Figure 2.5, higher values for K result in lower fitness values.

If the NK model is a good approximation to reality, biological systems must lie somewhere between the $K = 0$ and $K = N - 1$ extremes. Point mutations tend to have small changes on organisms, rather than drastically improving or reducing their fitness, so the landscape is not random. There are many evolutionary solutions and niches, so the landscape is not singly-peaked either. Real fitness landscapes must be locally smooth, but possess high fitness peaks that organisms can stably achieve.

Much work has been done to characterize the behavior of different classes of walkers on NK landscapes, but they are not of immediate interest. A fuller treatment can be found elsewhere.[14] I will restrict myself primarily to analyzing the behavior of 'fitter $D = 1$' walkers on fitness landscapes, for the reasons described previously. Of more interest is an extended version of the model which treats interactions between species in addition to interactions between genes.

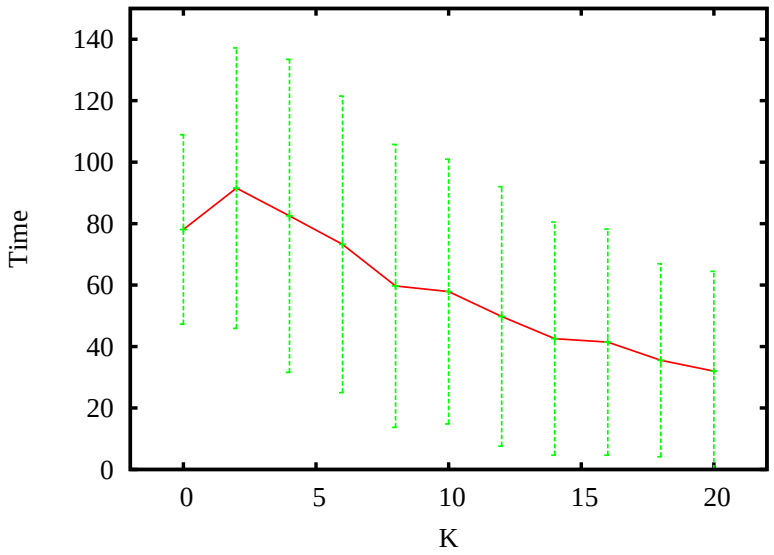

Figure 1: Number of steps required for a 'fitter D=1' walker to reach a local fitness maximum. Results are for $N = 24$, averaged over 500 runs.

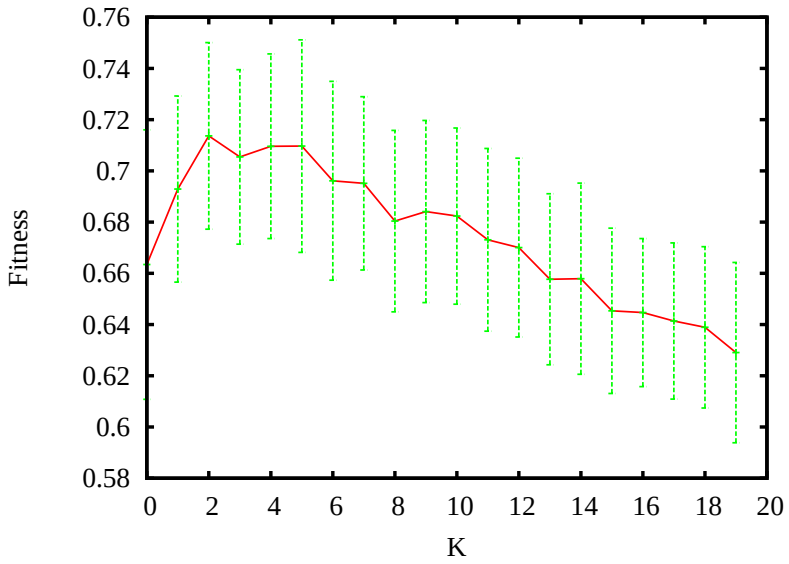

Figure 2: Average fitness achieved by a 'fitter D=1' walker. Results are for $N = 20$, averaged over 100 runs.

3 Coevolutionary dynamics: the NKC model

3.1 Overview

The NK model is able to capture the effects of gene coupling and adaptive walks on fitness landscapes for single species, but has no provision for the interactions between species. Real ecosystems have many interconnections between species, such as those between a predator and its prey, or two species competing for the same resource niche. Thus an advantageous mutation for one species may reduce the fitness of another species - for example, a snake which has evolved a stronger venom reduces the fitness of its prey and also of other snakes competing for the same prey; this might cause the prey to then evolve venom resistance and the other snakes to evolve stronger venom. In the fitness landscape picture, this is equivalent to deformation of the fitness landscape by adaptive walkers. In an ongoing feedback process, local moves by one species change the landscape, affecting the adaptive moves available to other species.

To model these effects, two new parameters are introduced to the model: S, the total number of species, and C, the number of genes in other species to which each gene in a species couples. Thus, each gene is coupled to $1 + K + C$ other genes. The total number of weights W_s for a given gene s in a species n is then given by

$$W_{s,n} = A^{1+K_{s,n}+C_{s,n}} \tag{2}$$

The parameter C controls the ruggedness of the landscape in the same way K does: higher values of C mean an adaptive move by one species is capable of drastically changing the fitness of another species. It is possible to either assign these couplings randomly and uniformly among all genes in all species, though it is also interesting to restrict the subset of genes and species for a given coupling. Ecosystems may be highly heterogeneous with species possessing genomes of different levels and different coupling strengths, or have different degrees of connectivity with some species being more or less connected than others, or highly random with all species being equally likely to be connected to a gene in another species.

The NKC model is highly flexible, allowing the study of how the level and extent of coupling between species affects coevolution. Through varying or restricting the scope of N, K, and C for each species, the model can provide insight into the way in which adaptive walkers traverse different classes of fitness landscapes landscapes and into how couplings richness and walking method affect the stability and optimality of fitness solutions. However, the model is more easily studied with computational rather than analytic techniques - there are very few exact results known for the behavior of the model.[7][6][22] As a result, I will restrict myself to studying trends exhibited by the model and data obtained through computer simulations.

3.2 The simulation

I have implemented the NKC model with the C programming language, using a novel scripting language to control multiple runs of the simulation and a Lehmer random

number generator as a source of randomness.[17] The project contains approximately 2500 lines of code. In order to understand how the model works at a more detailed level it is instructive to see the flow of logic in the simulation.

The simulation uses two different seeds for the random number generator: a 'couplings' seed, used to determine the initial state of a system (genotypes, couplings), and a 'dynamics' seed, used to control the time evolution of the system (mutations). This enables me to separate behavior due to the landscape from behavior due to the dynamics of the adaptive walk.

Because of the randomness inherent in the model, I sometimes average the results of simulations with different sets of seeds. I define a 'run' to be a simulation of a system with a given 'dynamics' and 'couplings' seed. 'Multiple (independent) runs' are several runs of a system with the same (different) 'couplings' seed but different 'dynamics' seeds. In reality, many statistical quantities I choose to measure do not have significant seed dependence, so it becomes not only unnecessary to distinguish between independent and non-independent runs, but also unnecessary to average over multiple runs. Where the quantity I have chosen to measure has a strong seed dependence, I average over multiple independent runs and note error bounds. Where there is weak seed dependence, I present the results of a single run without error bounds and note that the behavior is typical.

For each simulation run, the following process is used:

1. Choose the 'couplings' seed for the random number generator to control the initial state of the system.

2. Initialize and allocate space for S species. For each species s:

   (a) Initialize and allocate space for N_s genes.

   (b) Choose a random initial genotype for the species by constructing a random binary string of length N_s.

   (c) Initialize and allocate space for couplings. For each gene n:

      i. Initialize and allocate space for $1 + K_{s,n} + C_{s,n}$ couplings.

      ii. Set up the self-coupling between the gene and itself.

      iii. Set up a coupling to a gene in species s to which n is not already coupled, $K_{s,n}$ times.

      iv. Set up a coupling to a gene in a different species to which n is not already coupled, $C_{s,n}$ times. The set of species from which n may choose may be limited.

      v. Initialize and allocate space for $W_{s,n}$ weight values. Assign the weights to the couplings. (The details of this process are not important, since the weights are random.)

      vi. For each weight, choose a random weight value.

3. Choose a separate 'dynamics' seed for the random number generator to control the evolutionary dynamics.

4. Until a fixed number of steps or an equilibrium is reached, iterate/step the system forward in time:
   (a) Calculate the fitness of every species, given its genotype and the genotype of other species.
   (b) Sample the system by storing these fitness values for later analysis.
   (c) Choose a species $S_{choice} := (S_{choice} + 1) \bmod S$.
   (d) Choose a random number R_{choice} between 1 and D_s.
   (e) Flip the value of R_{choice} bits of the genotype of S_{choice}.
   (f) Calculate the change in fitness for S_{choice} due to this mutation.
   (g) If the fitness change is positive or neutral, accept the mutation. If the fitness change is negative, discard the mutation; the system does not change during this iterative step.

After a set of simulations has been completed, further analysis and statistics are performed on the sampled fitness data.

3.3 Assumptions

Like the NK model, the NKC model also oversimplifies reality. All the limitations of the NK model are still present, along with several others, which it is helpful to explicitly note.

1. Extinction and speciation events never occur. The model deals only with interactions between continuously evolving species. There is no chance that a species can become two separate species, or cease to exist due to competition. This is reasonable when making comparisons to real systems over short enough time scales.
2. The model is typically run with all species having the same value of K and C; this implies a degree of homogeneity which it is unreasonable to expect in a real ecosystem. This assumption is relaxed later.
3. Only the 'fitter' adaptive walker is used, typically with $D = 1$.
   (a) Negative mutations are simply discarded and never accumulate.
   (b) Species trapped at local maxima stay at local maxima, unless the fitness landscape changes under them.
4. No two species can evolve simultaneously. Evolution is a linear process. This follows from the previous assumption that a species can be described by a single genotype - mutations must propagate quickly relative to the iteration speed.

4 Simple results from the NKC model

4.1 Methods and definitions

All experiments described below were conducted on my personal computer; simulation size and duration were limited by available memory. Simulations with approximately $K > 16$, $C > 8$ were intractable, requiring the allocation of gigabytes of memory to store weight values. Storing fitness values for the system for later analysis also required large memory allocations. Larger and longer simulations are certainly possible with better hardware, but I hope to have captured the essential features of coevolution in this region of the parameter space. Processor speed also became a limiting resource when checking for equilibria, which will be discussed in more depth later. In general, species with large values of N are better approximation to real species; $N \approx 24$ begins to require very large amounts of memory and so I have chosen to run most simulations at this value.

It is important to distinguish between transient and characteristic behavior of the system. After the system has been initialized, very large fitness jumps may occur as species quickly move from their random locations on the fitness landscape to higher points. This is transient behavior, and is unrelated to how the system behaves once it is evolving in a steady set of states. It is possible to either iterate the system over very long time periods so this transient behavior becomes unimportant, or to truncate transient data before calculating the distribution. I have chosen to do a mix of both, iterating for many (typically 10^6) steps and ignoring the first 10% of time-domain data.

I will use the term 'global fitness' to mean the average of the fitnesses of every species at a given time; it is a good indicator of the overall behavior of a system. 'Mean global fitness' is the average of the global fitness over a time period; this time period is taken to be the duration of the simulation, less the time assumed for transients to disappear.

4.2 Mean fitness and the mutation distance

Typically I will perform experiments with $D = 1$ so that species are constrained to only jump to nearest-neighbor genotypes. All other types of mutations for larger D can be built from $D = 1$ mutations, provided there is a positive fitness gradient along the entire path. Thus, increasing D is advantageous if the fitness landscape is sufficiently rugged that there are nearby higher fitness maxima. Large D allows a species to escape a low fitness peak. But a very large value of D is disadvantageous - mutations that take a species far away from its starting point are unlikely to be or be near fitness maxima. Exceeding the correlation length of the system results in a species becoming lost on the fitness landscape. As shown in Figure 3, regardless of the value of K, fitness is maximized for small values of D. When C is small, an intermediate value of D is optimal, because the landscape is slightly rugged - but increasing C further drives the optimal D to be as small as possible, because more couplings greatly decrease the correlation length of the fitness landscape. This suggests that real organisms are optimally fit when their mutation rate is relatively low. Because the actual optimal value for D varies with K

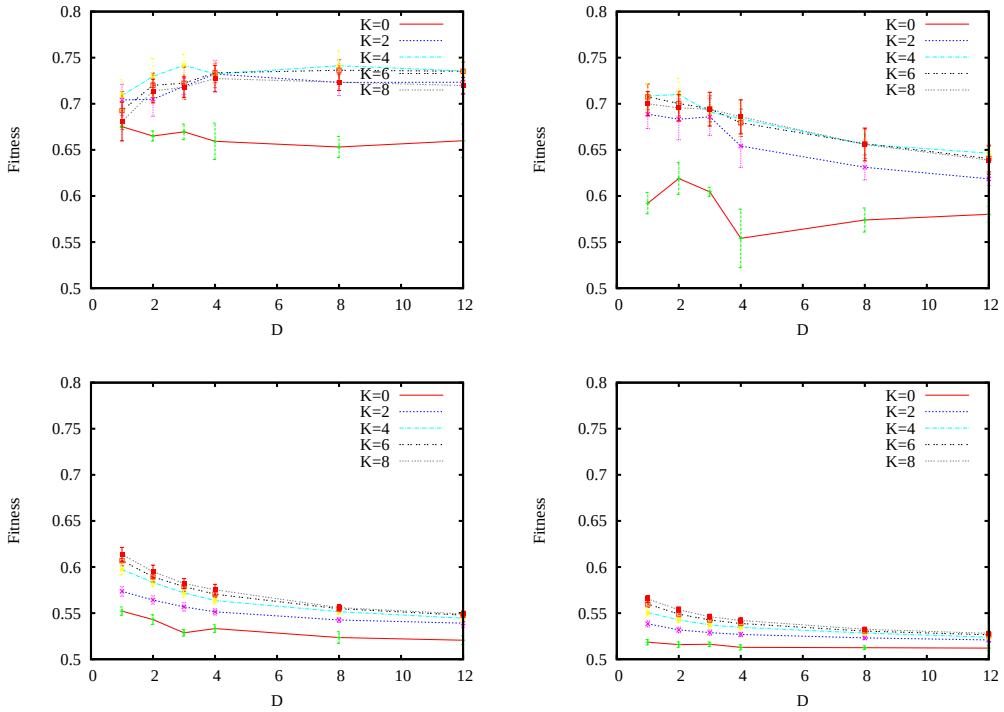

Figure 3: Dependence of the mean global fitness of a system on the mutation length, D. From top left, $C = 0$, $C = 1$, $C = 4$, $C = 8$. In general, smaller values of D are associated with a higher mean global fitness. Simulations were carried out for $S = 4$, $N = 16$ and averaged over 50 runs.

and C, $D = 1$ is a good approximation of the optimum and of biological systems that undergo only single mutations at any given time.

4.3 Interactions between species with different K and C values

From a game-theory perspective, it is interesting to consider the coevolution of multiple pairs of species with different values of K and C. There are be particular combinations of these parameters which are particularly advantageous or disadvantageous for the species pairs. I have not analyzed these interactions in this work, but a full treatment can be found elsewhere.[13]

4.4 Tuning K and C to optimize mean global fitness

On a homogeneous landscape where all species have the same values for K and C and where any gene is allowed to couple to any other gene, it is interesting to consider how fitness depends on the number of couplings between genes and species. As in the simpler NK model, it is expected that the optimal value for K will be small but non-zero. By

the same logic, it is expected that values of C that are too large will have a negative effect on global fitness.

These hypotheses are supported by experiment; as shown in Figure 4, mean global fitness is low at both low and very high values of K, while it peaks at an intermediate value. Lower values of C are uniformly associated with lower global fitness values; this is because any inter-species coupling results in even small fitness changes by one species translating into large fitness changes for others; nonzero C make it more difficult for a species to stay at a high fitness peak. Additionally, the variance in fitness of species with lower values for K and C is higher, supporting the idea that fitness peaks are more varied and higher on less rugged landscapes. When C is too large, species are forced to walk on a landscape with many very small peaks and valleys.

Interestingly, the optimal value of K appears to vary with C, perhaps linearly. More simulations (beyond the resources available to me) would be necessary to determine the precise relationship. The trend itself is reasonable—when a species' fitness is subject to changes by many external couplings, increasing the number of internal couplings reduces the importance of these external effects.

The highest global fitness values are achieved when $C = 0$, $K \approx 6$. This is difficult to achieve in reality, but corresponds to a situation where each ecological niche is filled; there is no competition or feedback due to interactions between species. Most ecosystems do not have $C = 0$, and so the fitness of each species is reduced by competition. It is possible that species are able to adjust their values for K and C; this has been studied in some detail before.[13]

4.5 Nash equilibrium

In general, there are two possible outcomes for a coevolutionary process involving 'fitter' adaptive walkers. Either the system can achieve an equilibrium state, in which every species is at a local fitness maximum, or the system can fail to achieve an equilibrium, in which species are always trying to climb a constantly deforming fitness landscape. From a game-theoretic perspective, the second type of state can be thought of as a Nash equilibrium - each agent (species) is unable to make any further gain (a positive fitness change) through a change in strategy (a mutation) given that the strategies of the other agents remain unchanged. No species chooses to mutate if offered the possibility. A system that has reached Nash equilibrium is not necessarily in the optimal configuration, but simply in one (of many) stable configuration. For multiple runs of a system that begins in the same configuration but has different dynamics, it is unlikely that the same Nash equilibrium will be reached multiple times.

The waiting time to encounter a Nash equilibrium varies non-trivially with all the parameters S, N, K and C. For systems with S or N large, genotype space becomes very large, and it is increasingly difficult for adaptive walkers to find local maxima. As C increases, the landscape becomes increasingly volatile, and finding a local maximum for even one species becomes a matter of chance. As the correlation length of a fitness landscape decreases, so does the probability of achieving equilibrium. This implies that large ecosystems are never stable; species constantly are interacting and competing with

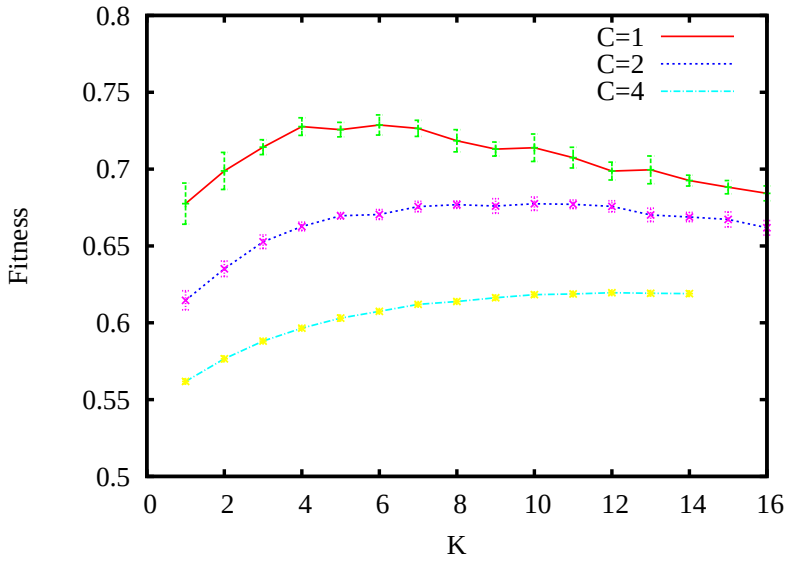

Figure 4: Mean global fitness. There is an optimal value for K given C. Global fitness is higher and also more variable when C is lower. $S = 8$, $N = 24$, averaged over 10 runs per data point. There are no data for $C = 4$, $K > 14$ due to computational limitations.

each other.

However, it is possible for relatively small systems to achieve Nash equilibrium; a sample system is shown in Figure 5. Biologically, this corresponds to a set of interacting species that have succeeded in adapting to the behavior of the rest; stable interactions and an acceptable level of competition have been obtained. In particular, I have studied systems in which a pair of species interact with each other. Figure 6 shows that low values of C, species with higher values for K take more time to reach equilibrium, while for high values of K, species with higher values of K take less time to reach equilibrium. Perhaps at $K \approx C$ there is a crossover between these two regimes. At high C, the landscape is fairly random and small moves by one species deform the landscape significantly. Lower values of K reduce the total number of constraints and make it possible for the system to (randomly) settle on a mutually acceptable equilibrium, though it may be of low global fitness. On the other hand, for low C, the fitness landscape is difficult for one species to deform, and it is possible to apply results from the NK model, where it was found that larger values of K result in shorter adaptive walks—each species relatively independently finds a local fitness maximum more quickly when K is high.

It is computationally expensive to determine if a system is in a Nash equilibrium. One must evaluate the fitness of all genotypes within the neighborhood of each species to ensure they are not more fit. For large N and and $D > 1$, this prohibitively difficult. As a result, I have only studied Nash equilibria in systems with $D = 1$ walkers.

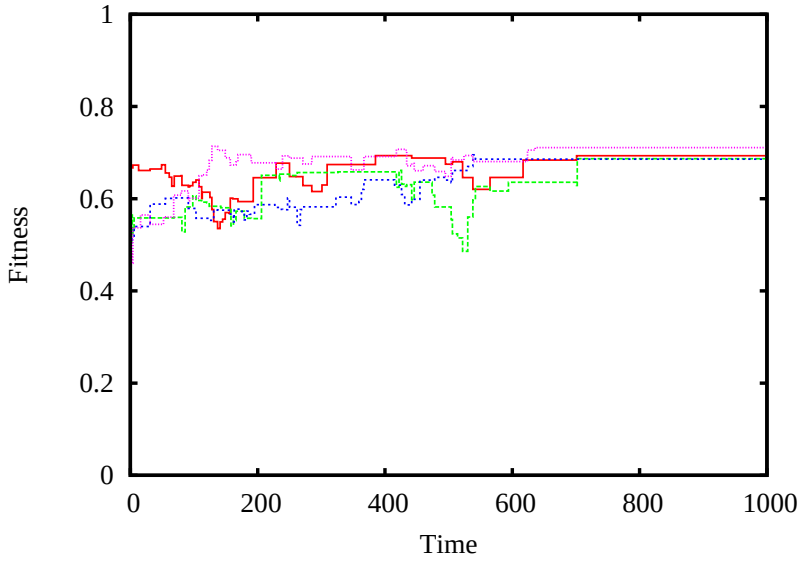

Figure 5: An example of coevolution between four species where Nash equilibrium is reached. $N = 24$, $K = 8$, $C = 1$.

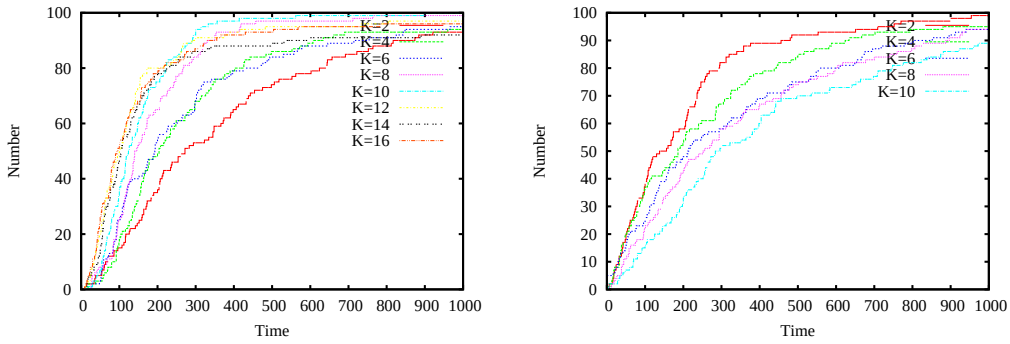

Figure 6: Number of runs (out of 100 total) in which equilibrium has been reached after a given number of steps. Larger values of C generally increase the equilibrium waiting time. For small C, high values of K decrease the waiting times; the opposite is true for large C. $S = 2$, $N = 24$. Left, $C = 1$; right, $C = 8$.

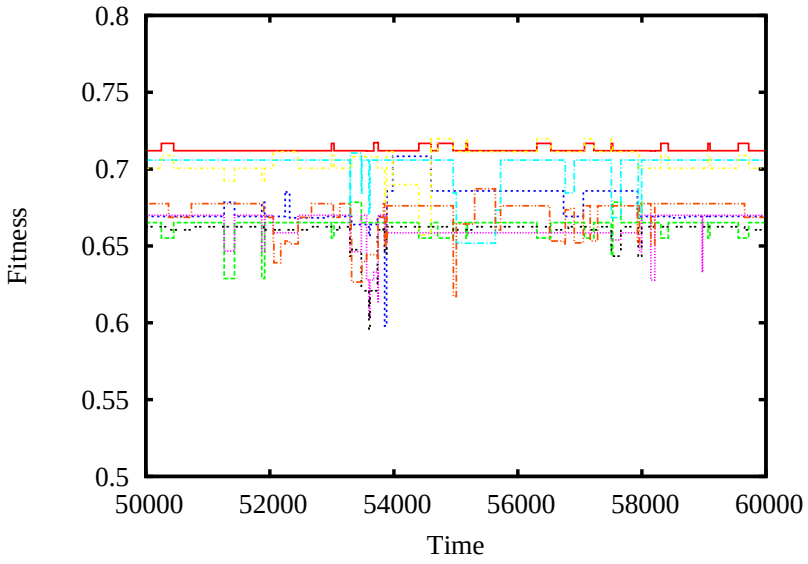

Figure 7: Long periods of stability followed by rapid periods of change are typical of large systems. $S = 8$, $N = 24$, $K = 16$, $C = 1$.

4.6 Punctuated meta-equilibria

Most systems do not quickly reach equilibrium, but continue to evolve. It is difficult to *a priori* determine whether or not a system will ever reach equilibrium—iterating the system is the only way to be certain. I have found many simulations which do not reach equilibrium in under 10^8 steps. For many systems, there is a probability that by chance they will find themselves at Nash equilibrium—but as size of genotype space increases, this probability becomes negligible.

It is common for a system to appear to be in an equilibrium state, but actually be in a meta-equilibrium. That is, all the species in a system may keep the same genotype for an extended period of time, but the state is not stable to mutations. An example of this is shown in Figure 7: the fitness of the top species does not change over thousands of iterations of the system, but is not invariant.

Because of the nature of 'fitter $D = 1$' walkers, the probability that a given gene in a given species will be selected to mutate is only $1/SN_s$. Additionally, there is a longer waiting time between mutations as species approach fitness maxima. Therefore as S and N_s and ftness increase, the average time between beneficial mutations also increases.

However, if K and C are relatively large, a single mutation in one species following a long period of meta-equilibrium may cause an avalanche of further mutations. Once the fitness landscape has been deformed by one species, the remaining species may no longer be no longer as fit as they once were, and so become free to mutate. After enough mutations have occurred, it becomes increasingly difficult for any species to find a beneficial mutation, and the system regains a similar frozen state as it began in. I

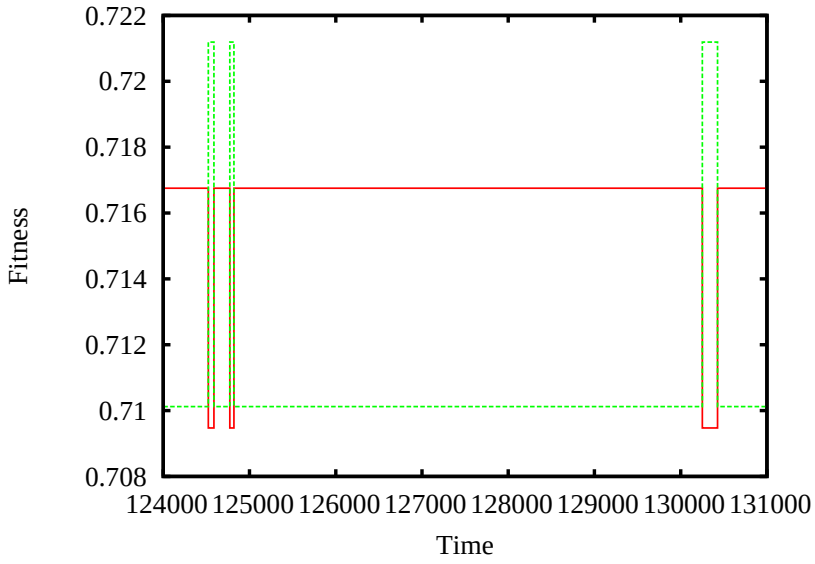

Figure 8: Fitness oscillations can occasionally occur when species repeatedly flip the same subset of coupled genes. Here a pair of coupled species is shown for a simulation with $S = 8$, $N = 24$, $K = 16$, $C = 1$.

shall treat these avalanches and frozen landscapes in more depth later.

These long periods of stability interspersed with short periods of rapid change that I have observed may be analogous to the punctuated equilibrium model of evolution put forward by Gould and others.[9]

4.7 Stable oscillations

Occasionally, it is possible for a system to evolve into a stable non-equilibrium state with periodic behavior. Oscillations are possible when a set of species are coupled to each other in such a way that the same set of mutations keep occurring. In the two-species case, a mutation in the first species will make a mutation in the second species advantageous, which then makes the mutation in the first species disadvantageous; once the first species has mutated to its original form the second species also mutates to its original form, and the process repeats. I have only observed this to occur for pairs of coevolving species (an example is shown in Figure 8), but there is no reason why it should not occur in larger groups.

4.8 Chaos

By increasing K and C, is possible to make the landscape sufficiently rugged and deformable that at any step, there are several positive mutations available to each species. Thus, the system never reaches equilibrium, because each species is always able to find

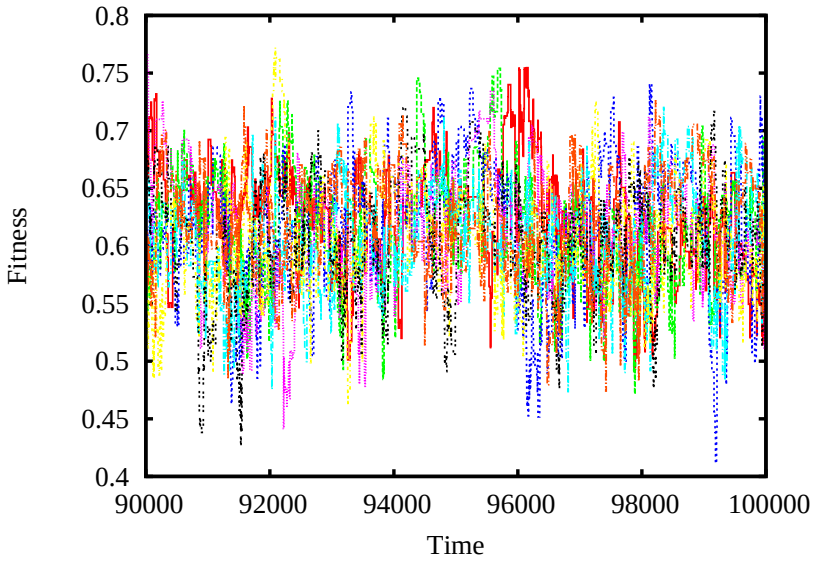

Figure 9: Chaotic dynamics can occur on very rugged and deformable landscapes. $S = 8$, $N = 24$, $K = 8$, $C = 4$.

a beneficial mutation and change the fitness landscape. This results in a larger subset of the genotype space being explored, but has deleterious consequences for the fitness of each species and the mean global fitness. When changes occur so often on such a rugged landscape, the evolutionary path taken by each species is closer to a random walk than an adaptive walk. Species never stay at local fitness maxima, and tend to rapidly jump between a wide range of fitness values. This is analogous to an ecosystem where too many species have been introduced. The system attempts to balance all the conflicting couplings between genes, without success. In a real ecosystem, extinctions or spontaneous changes to the number of couplings would result - but these are impossible in the NKC model. Chaotic behavior is the result. Examples of this are shown in Figure 9 and Figure 10.

5 Criticality and the edge of chaos

5.1 Motivation

I have shown that the NKC model is capable of displaying several types of behavior: equilibrium, oscillation, punctuated equilibrium, and chaos. The coevolutionary dynamics can range from being completely ordered (where every species finds a local fitness maximum) to completely disordered (where no species ever finds a stable local fitness maximum). Clearly, real evolution cannot occur at either extreme: if the system is too ordered, species stagnate and evolution ceases to occur, while if the system is too disordered, no adaptation is possible and evolution ceases to be a productive process. The

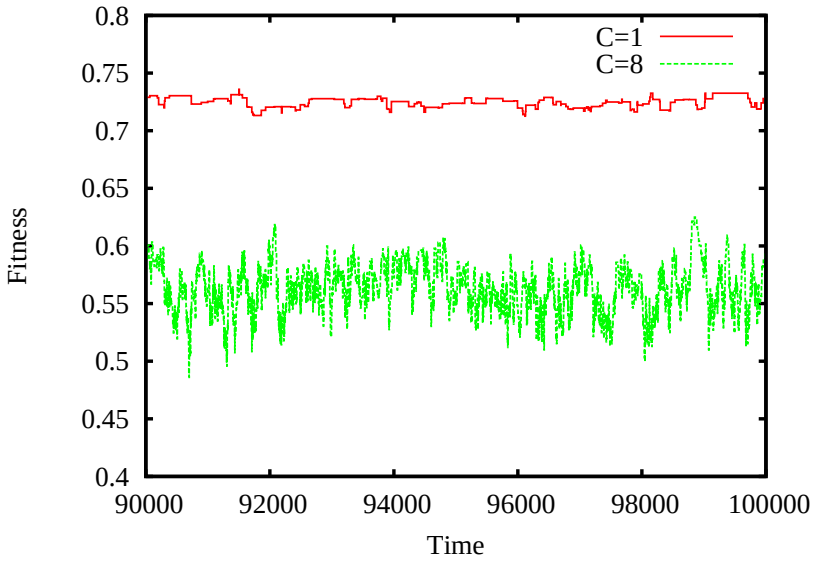

Figure 10: Mean global fitness can behave chaotically. On a highly coupled landscape, the fitness is very low, highly variable, and time-uncorrelated. $S = 8$, $N = 24$, $K = 8$.

intermediate state must have properties that correspond to the richness of biological evolution.

Stephen Wolfram, in his studies of computation in cellular automata, has hypothesized that any particular cellular automata can be assigned to one of four different phases: fixed, ordered, critical, and chaotic.[25] Fixed systems are non-functional (constant), while ordered systems are governed by a set of rules and switch between well-defined states; these correspond to many simple sorts of calculations and to algorithms that halt. Chaotic systems may also be governed by simple rules but are capable of exploring many states, and correspond to non-halting algorithms. In between is the critical phase, which corresponds to cellular automata capable of complex and deep behavior. Importantly, this depth can be—and often is—generated by simple rules. Complex behavior does not presuppose complex rules.

Criticality is not a loose generic term, but has a precise meaning in the context of statistical physics and phase transitions. In the next section I will be more precise; for the present it suffices to note that critical systems are governed by the same simple rules as non-critical systems, but exhibit radically different behavior. In the critical regime, the system becomes scale-free and show fractal (power law) behavior. Complex patterns and trends are capable of emerging from simple rules.

Many biological systems have been shown to have fractal properties or exhibit scale-free behavior. For example, it has been shown that some rain forests are self-similar, with fractal properties.[20] Yeast colonies, resource networks, heartbeats, and many other systems also have fractal, scale-free properties.[10] Additionally, metabolism has

been shown to be related to body mass by a power law.[24] Evolution also appears to be a scale-free process; for example, the extinction record seems to obey a power law.[3].

It is also hoped that these scale-free systems exist in a critical phase—and so, perhaps, their complexity can be explained by a small set of rules. If real evolution is a critical process, then it is possible that there is no need for a general theory of macroevolution, because patterns of extinction and speciation are entirely (with the exception of certain global catastrophes) the result of simple interactions between individual species. There is already evidence for this hypothesis in many realms of biology.[8] Real biological evolution also exhibits many features of criticality.[21] These have been studied for macroscopic trends, due to the difficulty of getting data for individual mutations and tracking them over long times.

It is hoped that the NKC model of evolution is also capable of existing in the critical phase and displaying complex behavior. It would be a major success if such a simple model, in which each species seeks only to maximize its own fitness, were capable of showing trends similar to those found in macroevolution.

5.2 Critical phenomena

In statistical physics, a phase transition occurs when a singularity occurs in the derivatives of a thermodynamic potential; when this singularity is present in the second (or higher) derivative, the transition is called continuous or critical. The potential and the approach to this singularity is governed by at least one parameter; suppose that there is only one, T, and that the critical point occurs at T_c. Let $t = (T - T_c)/T_c$ be a measure of the deviation from the critical point. Then for a function $F(t)$ associated with the potential, $F(t) = A|t|^\lambda(1 + bt^{\lambda_1} + \ldots)$. For $t \ll 1$, this can be approximated to first order as $F(t) \sim |t|^\lambda$, where λ is known as a critical exponent. Typically, $\lambda < 0$. Therefore near critical points, functions describing the system are described by power laws which diverge at the critical point. In particular, the correlation length of the system becomes infinite at the critical point; this implies that a small perturbation of a critical system can have an arbitrary result—large changes in the behavior of the system are possible. Additionally, the system becomes scale-free, since $F(kt)/F(t) = k^\lambda$ has no t-dependence. For some systems it is possible to calculate these critical exponents with mean-field or renormalization group theory.[27][5] Here, it is important only to understand the general characteristics of criticality; in particular, the presence of power laws.

5.3 Self-organized criticality

Most thermodynamic systems require tuning of at least one free parameter in order to achieve criticality; the critical point is not necessarily an attractor for all parameters. However, many natural systems self-organize, developing patterns and structures without any external control (for example, zebra stripes, beehive structure, or clathrin lattices). It has been suggested that many of these structures are scale-free and may naturally exist in a critical state without tuning any parameters. This self-organized criticality is then dependent not on the exact values of parameters, but instead on the local dynamics

of the system. Small simple interactions between parts of the system spontaneously generate large complex orderings of the system as a whole.

This idea was first proposed by Bak, Tang, and Wiesenfeld, who tried to explain the prevalance of time- and space-fractal (scale-free) structures in nature.[4] Such structures are not the lowest possible energy configuration possible, so some sort of dynamical selection must occur. They imagined that these sorts of systems are governed by two types of processes: internal relaxation, and external drive. Such a network becomes correlated by the interactions of the internal relaxation processes counteracting the external driving force. Suppose such a system fully relaxes before any external drive. It is then possible for such a system to continually exist in a meta-stable states near some threshold for the system. Any external perturbation induces a relaxation process which brings the system to a nearby stable state, until another perturbation brings the system to another nearby stable state. It was hypothesized that dynamic coupling occurs between many different parts of the system, possibly leading to fractal structure. If such a system has no characteristic scale, then it becomes possible for relaxation events of all scales to occur.

The canonical example of such a system is the sandpile model. Consider a two-dimensional grid, each point of which can contain a column of sand grains of height $h_{i,j}$. Very occasionally, a grain of sand is added to a random point on the grid, so that $h_{i,j} \mapsto h_{i,j} + 1$; this is the external drive. The system is relaxed by assuming there is a critical height h_c: if ever $h_{i,j} > h_c$, then the sand at such a point is moved to its nearest neighbors: $h_{i,j} \mapsto h_{i,j} - 4$ and $h_{i,j+1} \mapsto h_{i,j+1} + 1$, $h_{i,j-1} \mapsto h_{i,j-1} + 1$, $h_{i+1,j} \mapsto h_{i+1,j} + 1$, $h_{i-1,j} \mapsto h_{i-1,j} + 1$. This process continues until no column of sand is above the critical height. One may either assume the grid is infinite, or that sand is removed from the edges of a finite size grid. In this model, it is useful to define the size of an 'avalanche' as the number of grains of sand that move after a grain of sand is added to the pile. Remarkably, one finds through numerical simulation that the number and size of avalanches are related by a power law, and that the critical exponent does not depend on h_c. The model displays criticality without external tuning.

This sort of self-organized critical system shows very long time correlations, since small fluctuations in one part of the system can result in large fluctuations elsewhere in the system—avalanches can become arbitrarily large. The temporal correlation function for a time signal $N(t)$ is defined as

$$G(\tau) = \langle N(\tau_0)N(\tau_0 + \tau)\rangle - \langle N(\tau_0)\rangle_{\tau_0}^2 \tag{3}$$

If there is no correlation between the signal at τ and τ_0, then $G(\tau) = 0$. The speed at which $G(\tau)$ decreases from its instantaneous value $G(0)$ to zero measures the duration of correlations in the system. Often these long time correlations can be detected in the power spectral density

$$S(f) = \lim_{T\to\infty} \frac{1}{2T} \left| \int_{-T}^{T} N(\tau) e^{2\pi i f \tau} d\tau \right|^2 \tag{4}$$

which can also be related to the temporal correlation function,

$$S(f) = 2\int_0^\infty G(\tau)\cos(2\pi f\tau)d\tau \tag{5}$$

Suppose $S(f) \sim 1/f^\beta$ and $G(\tau) \sim 1/\tau^\alpha$. Then from the previous equation, $1/f^\beta \sim 1/f^{1-\alpha}$. As $\beta \to 1$, $\alpha \to 0$. But if $\alpha = 0$ the temporal correlation becomes infinite. Thus, power spectral densities of the form $1/f$ can be associated with very long time correlations.[11]

5.4 Quantification and measurement techniques

Before measuring the properties of the NKC model in a critical state, it is important to specify how such a state might be obtained. I hypothesize that natural selection promotes the survival of systems that have a higher (mean global) fitness than others, and that criticality is present in nature; therefore, I expect that systems which display critical behavior are more fit than those that do not. Ideally the model would realistically enable species to self-tune their values of K and C in order to produce higher mean global fitness, but this is difficult. I instead have chosen to artificially vary these parameters, measuring the mean global fitness and signs of criticality for each combination.

There are several methods available to quantify and distinguish between critical and non-critical states. Ideally, I would measure the correlation length of the fitness landscape, which is expected to diverge when the system becomes critical. However for multiple species with large numbers of genes, this quickly becomes computationally impossible. There are several less resource-intensive methods which measure a property or characteristic implied by the system being in a critical state. However the causation is unidirectional—if the property exists, the system is not necessarily in a critical state. Therefore to separate critical states from non-critical states in the absence of a correlation length measurement, it is useful to verify that multiple quantities associated with the property are or are not present.

5.5 Power spectral density

The first and simplest method is to quantify the frequency of fitness changes. This is done by computing the power spectral density of the time-ordered fitness data for each species. Changes in fitness, whether due to mutations or deformations in the fitness landscape, can be quantified by their magnitude and effective length; these are equivalent to the amplitude and frequency of different Fourier components. A good indicator of criticality is whether or not the power spectral density is best fitted by a $1/f$ curve—that is, whether or not the time-ordered fitness data fluctuate in such a way that the system has very long time correlations.

In practice, the power spectral density is expensive to compute, especially for large data sets (long time series). To ameliorate this problem, the raw time-ordered fitness data is sampled at a uniform frequency. A real discrete Fourier transform is then performed on the reduced data and the absolute value of each resulting complex value is

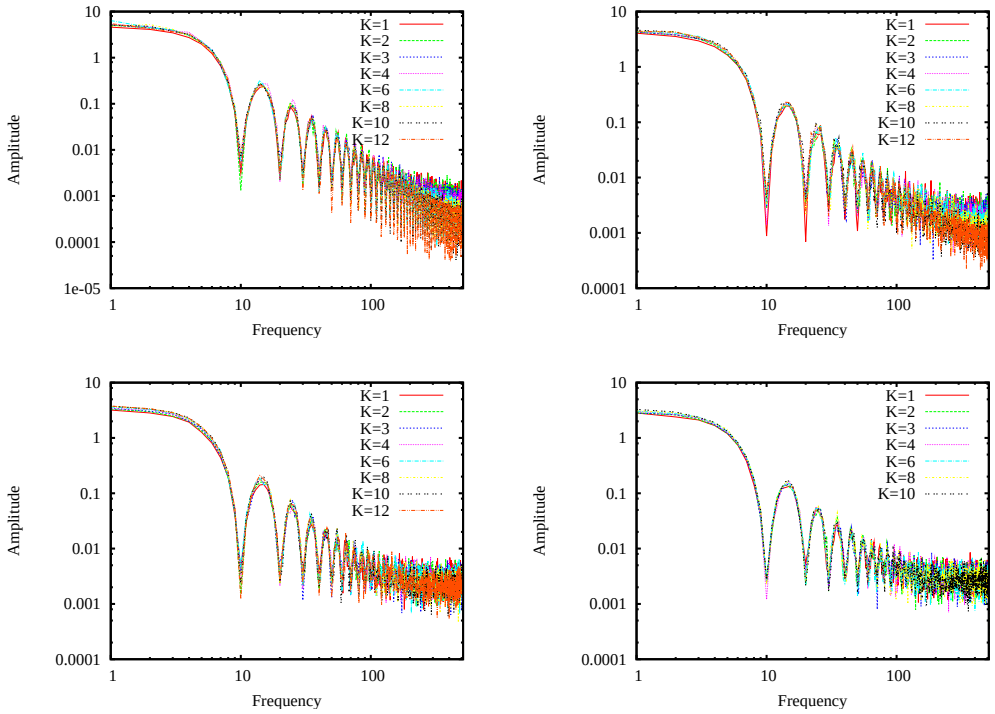

Figure 11: Typical mean power spectral density for all species in a $S = 8$, $N = 24$ system. From top left, $C = 1$, $C = 2$, $C = 4$, $C = 8$.

then computed in order to obtain the power spectral density for a given frequency. The sampling of the raw data affects the power spectral density only in that it ignores high frequency components of the signal. But this is not problematic because only the low frequency components are use for distinguishing between the agreement to a $1/f$ fit. I typically take 1024 samples from a system iterated 10^6 times.

As shown in Figure 11, for every combination of K and C tested, the power spectral density is governed by a power law of the form $1/f^{\alpha}$. The oscillation superimposed on the signal is presumably an artifact of the sampling and is not important. Qualitatively, larger values of C decrease α while varying K has a minimal effect.

To obtain more quantitative results, I averaged the power spectral densities for different initial configurations with the same values of K and C and found the best value of α for each combination of the parameters. These results are shown in Figure 12. For high C, $\alpha < 1$, while for low C, $\alpha < 1$ for sufficiently large K. However, it is possible to choose values of K and C such that $\alpha \approx 1$; for example, $K = 2, 3, 4$, $C = 2$ or $K = 12$, $C = 4$. But as shown in Figure 4, these configurations have a mean global fitness of approximately 0.65, while a value of 0.75 is not impossible for other configurations of the system. This suggests that $1/f$ noise, typical of systems in a critical state, can be obtained in the NKC model. However, the parameters for which this noise appears are

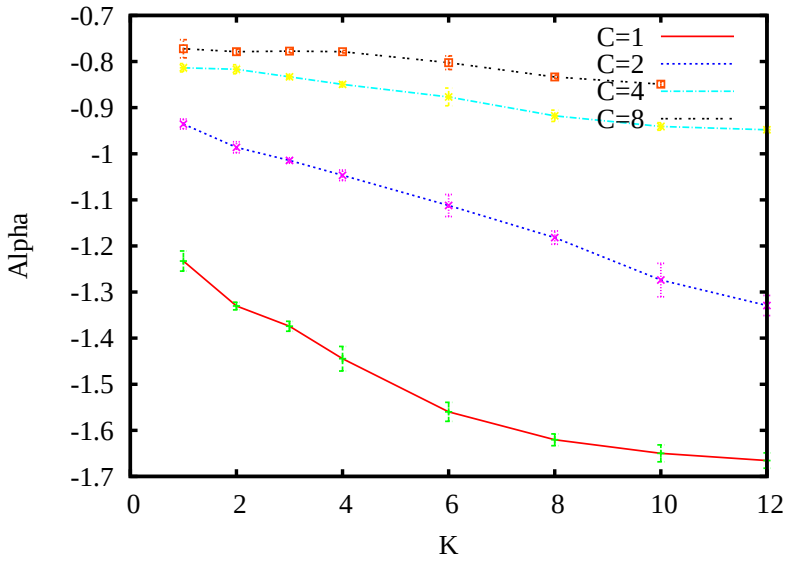

Figure 12: The scaling exponent α of the power spectral density decreases as C increases. Data were obtained by averaging over three runs of a $S = 8$, $N = 24$ system.

not ones to which the system would naturally gravitate.

5.6 Mutation avalanche distribution

The second method of quantifying criticality involves the distribution of mutation avalanches in the system. If the system is critical, we expect the distribution of avalanches to follow a power law distribution - larger avalanches should occur less frequently than smaller avalanches, but arbitrarily large and small avalanches should be possible. This is of course limited by the finite size of the system. It is difficult to choose a single good definition for the size of an avalanche, because there are many possibilities. I have chosen to define the size of a mutation avalanche as the number of iterations of the system after an initial mutation in one species before there is an iteration in which no mutation in any species occurs. This captures the following sort of behavior: if a beneficial mutation occurs in one species, it enables a second beneficial mutation to occur in another (or the same) species. This in turn allows a third beneficial mutation to occur.... Eventually this process will halt if the last mutation fails to create a beneficial environment for future mutations. The avalanche is then over; its size is equivalent to the number of mutations following the first mutation. In practice, the size of the avalanche is determined by the number of times the global fitness changes before reaching a stable value. If a mutation in one species occurs, it is very likely to change the fitness of either itself or another species, and it is exceedingly unlikely that the net change in fitness across the system will be zero. Thus, a change in the global fitness almost certainly signifies a mutation. This reduces the memory usage of the simulation, because it becomes unnecessary to

store the genotype of each species at every iteration of the system.

This avalanche definition differs from Kauffman's, which is stated as "the number of species which have changed at each ecosystem generation from the start of the avalanche until the avalanche stops... [it] includes both the number of species which are affected, and the number of ecosystem generations in which each is affected." Avalanche sizes increases in my definition when a mutation occurs, whereas for Kauffman, any change in fitness in any species increases the avalanche size. I think it is more realistic to only deal with mutations, since fitness values themselves are artificial constructs of the model and have no direct consequence except to control when mutations occur.

However, these definitions are very sensitive to the type of adaptive walk used by each species, and may not necessarily reflect the number of mutations caused by an initial mutation. Avalanches are defined to stop after a single step of the system in which there are no mutations by any species. However, 'fitter $D = 1$' walkers do not always mutate when mutations with positive fitness changes are possible. This is because of the randomness inherent in the process: sometimes the walker tries and rejects a non-beneficial mutations. Moreover, the system only offers a single species the chance to mutate during a given time step. If a mutation in one species enables a beneficial mutation in only one other species which is not offered the possibility to mutate in the next time step, the avalanche is considered to be over—even if the beneficial mutation is actualized one more time step afterwards.

Or instead, suppose a mutation which does not change the fitness landscape occurs in one species. The next time step, a mutation which also does not change the fitness landscape occurs in a separate species, and the next time step, no mutations occur. Two independent events have occurred, yet an avalanche of size two has been recorded! I have not found an easy and computationally feasible way to determine causation in the model. To decide whether a mutation was caused by a previous mutation (and should therefore be counted in the size of the resulting avalanche) is impossible without creating two instances of the system at every mutation: one in which the mutation occurs, and one in which it does not. My definition of an avalanche captures the general character of such an event; it is hoped that miscounting due to recording independent events and miscounting due to ignoring dependent events cancel each other.

It is instead possible to allow avalanches to continue after multiple steps of inactivity. Since N iterations take place before each species has been given the chance to mutate, it is reasonable to relax the avalanche end condition from one step without a mutation to at least N steps without a mutation. This better captures the interactions of weak couplings, where a mutation in one species may only affect a single gene in another species. This process results in more larger and fewer smaller avalanches being recorded, since the total number of mutations does not depend on how avalanches are measured. As shown in Figure 13, the avalanche distribution for a system is qualitatively invariant regardless of the number of allowed inactive steps: a linear distribution does not become exponential, nor does an exponential distribution become a power law. As a result, there is no benefit to increasing the allowed period of inactivity in the avalanche definition. For all further experiments, avalanches are defined to end after a single step in which no

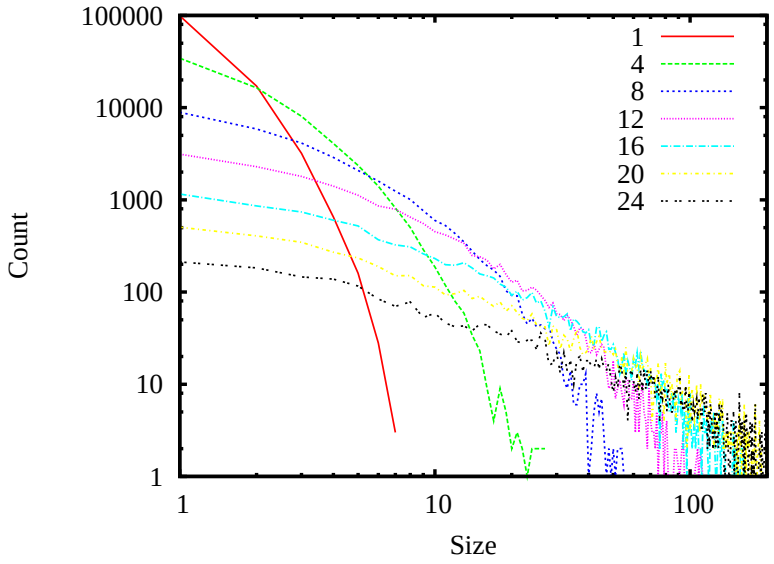

Figure 13: Mutation avalanche size distributions calculated for various allowed periods of inactivity (1 is the default behavior). Systems with longer allowed periods of inactivity are associated with larger avalanches. All distributions were calculated from a single run of 10^6 iterations of a $S = 8$, $N = 24$, $K = 8$, $C = 4$ system.

mutation occurs.

There is little evidence for critical behavior provided by the distributions of mutation avalanches, since power laws do not appear to be present in the mutation avalanche distribution. As shown in Figure 14, for almost all values of K and C, the avalanche count decays exponentially with size. I have fitted each distribution to both a power law and an exponential decay, and only in the single case of $K = 10$ for $C = 8$ was a power law a very slightly more appropriate choice. Moreover, I have previously shown that systems with lower values of K and C have higher mean global fitness, suggesting that the highly coupled system would not be found in a real ecosystem. This exponential behavior suggests that the probability of a mutation occurring at a given time is independent of the previous state of the system. It is likely that large avalanches are not the byproduct of a particularly important single mutation but instead are the result of nothing more than chance.

Moreover, there are no very large avalanches—the largest avalanche found in 10^6 iterations of the system only involved 19 mutations (and only 31 in 10^7 iterations) although arbitrarily large avalanches are permitted. This suggests that whatever the underlying nature of the distribution, it is very difficult for a single mutation to indirectly cause a very large series of mutations. No individual species can cause a catastrophic change in the rest of the ecosystem. Very large avalanches are found in the transient data, but these correspond to the system moving from its random initial state to a

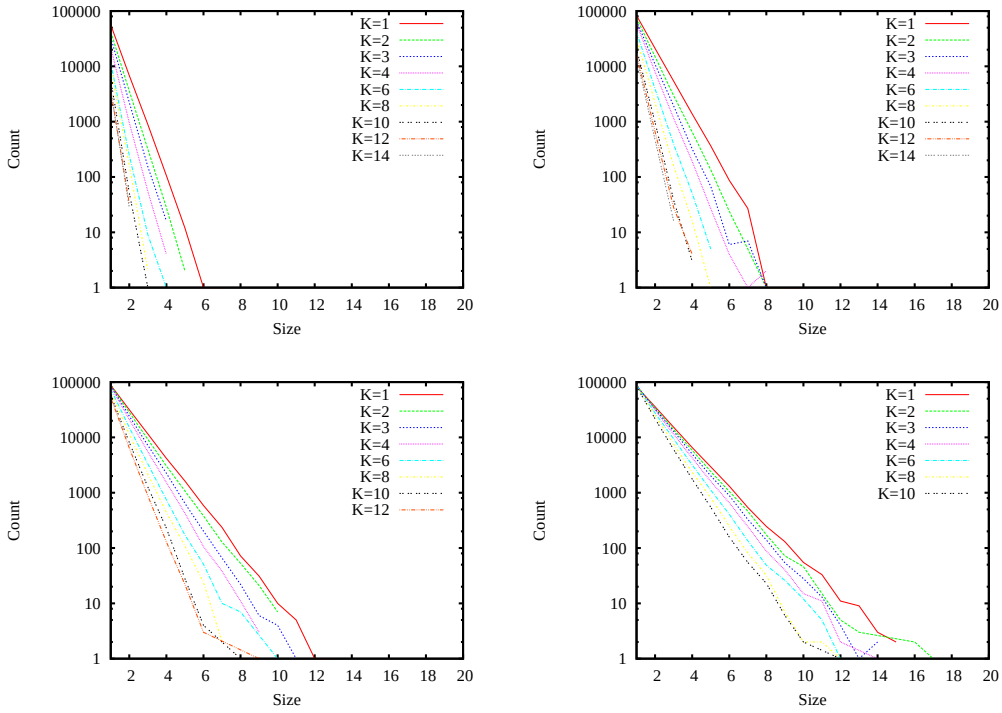

Figure 14: Number of fitness avalanches of a given size in 10^6 iterations of a single run of a $S = 8$, $N = 24$ system. From top left, $C = 1$, $C = 2$, $C = 4$, $C = 8$.

typical and steady state. Without a large enough sample size, these transient effects can give the distribution a longer tail, in effect making an exponential distribution appear to obey a power law.

5.7 Delta-fitness avalanche distribution

Another type of event that can be thought of as an avalanche is a fitness change due to a mutation or landscape deformation (delta-fitness). Here, an avalanche's size is defined to be the absolute value of the difference between a species' fitness at a given time and its fitness after one time step. An avalanche of a given size is sorted into a fitness bin with a specified width; the count for avalanches in the specified range is incremented when an avalanche of the proper size occurs. Fitness changes of zero (iterations during which no landscape deformation or mutation occurs) are ignored.

This may seem to be equivalent to the power spectral density, but there is one important difference: when calculating a power spectral density, all phase information present in the Fourier components is destroyed; one does not know how different frequency components of the signal overlap in the time domain. But for the delta-fitness calculation the relative phase information is retained, because the size of any fitness jump depends on how each Fourier component is aligned relative to all other components.

What this type of avalanche represents is not necessarily obvious, since fitness values are all normalized and calculated relative to other fitness values and weights, the components of fitness, are drawn from a random distribution. It could be argued that only the sign (not the magnitude) of the fitness change is important, since the adaptive walkers' decision to accept or reject mutations is based entirely on the sign of the fitness change. Nevertheless, it is hoped that species in critical systems undergo fitness changes of all sizes. However there is an upper bound to the magnitude of fitness changes. Certainly no fitness change can be larger than 1, and in an uncoupled landscape, no larger than $1/N$, since each gene independently contributes to the total fitness of the species. In a coupled landscape, no fitness change can be larger than $K(1 + C)/N$, since each gene can affect K genes in the same species and C genes in another species, each of which may affect K of the other species' genes. Since $K(1 + C)/N$ can be made arbitrarily large, this is not a sharp bound but only functions as a heuristic: larger values of K and C increase the size of the largest possible delta-fitness value.

Experimental results, as shown in Figure 15, reveal that the delta-fitness distribution has a non-trivial dependence on N, K, and C. As expected, large values of C and K increase the number of large delta-fitness events, and also the total number of events (that is, iterations during which a mutation occurred). Interestingly, high values of K indicate fewer small fitness changes, but more large fitness changes. This is because more intra-species couplings make the landscape more random, increasing the probability that a very large event could occur. If it were possible (given computational resources) to set $K = N - 1$, undoubtedly events of size approaching unity would be found.

The delta-fitness distribution never becomes a power law for any combination of K or C tested. On the linear-log plots displayed in the figure, every curve is piecewise convex, whereas a power law would appear to be concave. Therefore, the delta-fitness

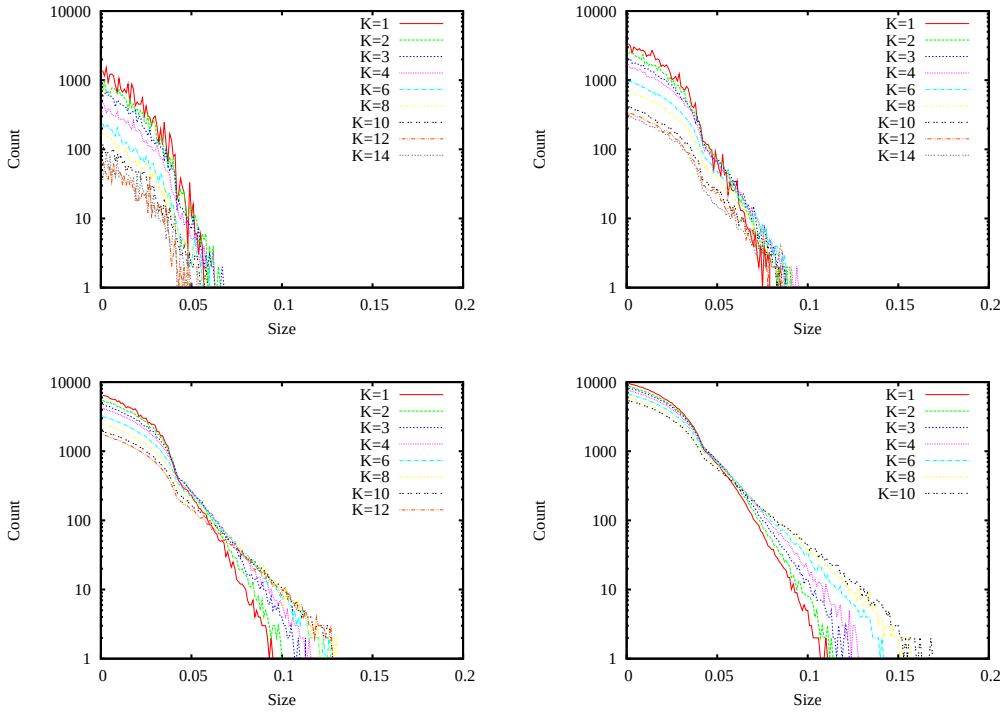

Figure 15: Number of fitness changes of a given size in 10^6 iterations of a single run of a $S = 8$, $N = 24$ system. From top left, $C = 1$, $C = 2$, $C = 4$, $C = 8$.

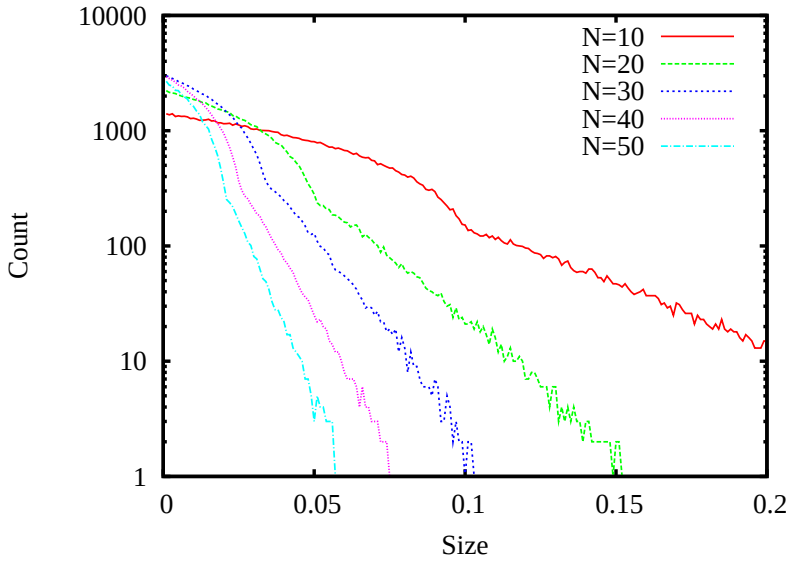

Figure 16: The transition between linear and exponential decay in the delta-fitness distribution occurs at size $1/N$. Results are shown for 10^6 iterations of a single run of a $S = 8$, $K = 8$, $C = 4$ system.

data do not support the claim that the NKC model can exist in a critical state, because of the lack of power law, scale-free behavior over any range of sizes.

Interestingly, the delta-fitness distribution appears to be split into two regimes for any combination of K and C: at small sizes, the distribution is linear, while at large sizes, the distribution decays exponentially. There is a very rapid transition between these two regimes. As shown in Figure 16, I have found that the point at which this transition occurs depends only on N, and in fact occurs at size $1/N$. I have no explanation for this transition yet.

5.8 Discussion

The result of these experiments is not encouraging; the NKC model does not appear to be capable of existing in a critical state. In the case of power spectral densities, I found that for one combination of K and C does the system show characteristic $1/f$ noise, while I found that for a different combination of K and C the mutation avalanche distribution could follow a power law. These two cases were unusual exceptions to the normal non-critical behavior of the system, and neither combination of parameters was associated with a high mean global fitness. Additionally, the delta-fitness distribution showed no evidence of power law behavior. From the negative results of these three methods, I conclude that there is no critical phase in the model. This differs from the power law avalanche distribution previously presented by Kauffman.[13]

6 Extending the model

6.1 Motivation and dangers

The NKC model is not self-tuning; it was necessary to artificially choose parameters and evaluate the mean global fitness of each configuration to determine if the choice was a good one. This artificial tuning affected every species equally—no species became more important or more connected than any other species. Additionally, every species always increased its fitness when given the chance to mutate. The NKC model assumed that any deviation from these assumptions cancels as one averages across enough similar species—however, ignoring these effects is not necessarily reasonable. As previously mentioned, many real biological systems violate these assumptions. Additionally, the absence of critical phenomena in the model suggests that its assumptions are incorrect.

In this section, I relax both of these assumptions (always positive and equally likely mutations, and equally connected species) and study their effect on the system. In particular, I look for signs of criticality and higher mean global fitness. However, it becomes necessary to introduce additional parameters to the model. The dimensions of parameter space grow, not only increasing the difficulty of search, but also increasing the chance that spurious results can be obtained. Stanislaw Ulam has been quoted as saying, "Give me 15 parameters and I can make an elephant; give me 16 and I can make it dance."[23] It is important to ensure that any new free parameters of the system are also associated with relevant physical properties of an evolving ecosystem.

6.2 Temperature

One can hypothesize that negative mutations may occasionally occur and become a part of a species' genotype. This may happen when the mutation is very slightly deleterious, so that it can propagate to all members of the species; if the mutation has only a small contribution to fitness and is not caught the first time it occurs, there is no reason why it should not spread to all members of a species. Evolution does not necessarily select the best genotype for a species, but rather one which is sufficiently good. Additionally, it is possible that mutations which have immediately negative consequences may enhance the evolvability of a species, enabling other mutations which have a greater positive effect on the species. In the language of fitness landscapes, negative mutations momentarily displace species from local fitness maxima, enabling them to follow a different positive fitness gradient. It is also possible that two sequential negative mutations may occur, enabling the species to find a more distant fitness peak. Thus on smooth landscapes (where a negative effect will not immediately randomize a species' fitness), it is possible that species which occasionally accept negative mutations may be more fit than species which do not.

It is useful to now introduce some further concepts from statistical physics. Consider two systems A and A' with energies E and E' which are free to exchange energy with each other. Further, suppose that the systems are in contact with each other, able to exchange energy but that the total energy $E^{(0)}$ is conserved: $E' + E = E^{(0)}$ If $E^{(0)}$ is

fixed, then the energy of system A is the only free parameter of the system. The energy of each system can assume a large range of values, but these values do not occur with equal probability. Using the fundamental postulate of statistical physics, which demands that an equilibrium system is equally likely to be found in any one of its states, it is possible to define the probability of system A obtaining energy E: $P(E) = C\Omega^{(0)}(E)$ where C is a constant of proportionality and Ω are the number of states of the entire system corresponding to energy E. But if system A has energy E, then A' must have energy $E^{(0)} - E$ and can be in any one of its $\Omega'(E') = \Omega'(E^{(0)} - E)$ states. Therefore the number of distinct states accessible to the combined system $A^{(0)}$ is given by the product $\Omega^{(0)}(E) = \Omega(E)\Omega'(E^{(0)} - E)$. The probability of system A having energy E can finally be rewritten as $P(E) = C\Omega(E)\Omega'(E^{(0)} - E)$. The most probable energy for the system can be found by finding the maximum of $P(E)$ or equivalently of its logarithm: $\partial_E \ln P(E) = 0$. Substituting the previous result, this becomes

$$\frac{\partial \ln \Omega(E)}{\partial E} - \frac{\partial \ln \Omega'(E')}{\partial E'} = 0 \tag{6}$$

It now becomes possible to define the temperature T of the system:

$$\frac{1}{kT} = \frac{\partial \ln \Omega}{\partial E} \tag{7}$$

where k is an positive constant with the dimensions of energy. Thus, the temperature quantifies how the number of states available to the system changes with energy. The systems are both in their most probable set of states when their temperatures are equal.[18]

Now suppose that A' is much larger than A, so that the energy of any particular state of A is much less than the the energy of the total system (A' acts as a reservoir or 'heat bath' so that an energy change in A has little effect on its energy). Further, suppose that A is in a single definite state r corresponding to energy E_r. Then, $\Omega(E_r) = 1$ so $P_r = C'\Omega'(E^{(0)} - E_r)$ where C' is a constant of proportionality. Now because $E_r \ll E^{(0)}$ it is possible to expand the logarithm of $\Omega'(E')$ about $E^{(0)}$. Thus $\ln \Omega'(E(0) - E_r) = \ln \Omega'(E(0)) - [\frac{\partial \ln \Omega'}{\partial E'}]_0 E_r + \ldots$ But $\Omega'(E(0))$ is constant with respect to r and the coefficient of the second term is equal to $1/kT$, so $\ln \Omega'(E^{(0)} - E_r) = \ln \Omega'(E(0)) - E_r/kT$ or after integrating, $\Omega'(E^{(0)} - E_r) = \Omega'(E^{(0)} e^{-E_r/kT}$. This implies that

$$P_r = Ce^{-E_r/kT} \tag{8}$$

where C is determined by the condition $C^{-1} = \sum_r e^{-E_r/kT}$. The probability of A obtaining a large energy decreases exponentially with energy—but if the temperature of the combined system A and A' is high, the probability of obtaining a large energy increases.

It is now possible to apply this formalism to the NKC model. Thermal systems can exist in a variety of states corresponding to certain energy values; when placed in contact with a heat bath at a given temperature, the system is perturbed in such a way as to produce a new distribution of states. Similarly, a species in an evolutionary system

can exist in many genotypes which have certain fitness values; when placed in contact with a 'mutation reservoir' which is described by a parameter acting as an effective temperature, the genotypes accessible to the species change in a defined manner. The temperature of the reservoir controls the scale of mutations - at higher temperatures, very large mutations become probable events. This temperature might be biologically analogous to the background rate of mutations, or perhaps another environmental effect. It is easy to modify the algorithm for adaptive walkers to enable the negative mutations to be accepted with probability dependent on the magnitude of the fitness change and the temperature of the system, which becomes a new parameter T of the model. I have chosen to set $k = 1$ for convenience.

This is difficult to implement because it is necessary to calculate the normalization factor C to set the overall scale of mutations. To find the true value of C, it would be necessary to calculate the change in fitness for every possible mutation, weighted by the number of times such a change in fitness occurs. For large number of genes, this rapidly becomes hopeless. One might think that by using a randomly chosen sample of mutations (fitness changes), one could easily approximate the true value of C, but this method fails: the exponential factor varies too quickly with fitness changes, so that very few fitness changes make a significant contribution to the average and the estimate becomes unreliable.

It is instead possible to generate a Markov chain of configurations which generates a representative set of configurations, but it is necessary to impose the requirement of detailed balance on each transition probability. That is, the rate of transitions between states in one direction is equal to the rate of transitions between states in the opposite direction. If the probability of state m becoming state n is P_{mn} and states m and n have M and N occupants respectively, then $MP_{mn} = NP_{nm}$. One possible way to satisfy this requirement is to use the Metropolis algorithm.[27] The probability that a system can move from an initial state i to a final state f is given by

$$P(i \to f) = \begin{cases} e^{-(E_f - E_i)/kT}, & E_f > E_i \\ 1, & E_f \leq E_i \end{cases} \tag{9}$$

This satisfies the requirement of detailed balance. Suppose there are m_r systems in state r and m_t in state t such that $E_t < E_r$. Then the number of transitions between states M is given by $M_{r\to t} = m_r$ and $M_{t\to r} = m_t e^{-(E_r - E_t)/kT}$. Then the net number of transitions in the system is $\Delta M = m_r - m_t e^{-(E_r - E_t)/kT}$. When the system has converged to a steady distribution of states, $\Delta M = 0$ or, $m_r e^{-E_t/kT} = m_t e^{-E_r/kT}$, as desired.

In thermal systems, energy is minimized, while in the NKC model, fitness is maximized. In order to implement the model, it is thus necessary to use a modified version of Equation 9. The new mutation algorithm can now be stated: if the fitness change ΔF is positive or neutral, accept the mutation. If it is negative, calculate $B = e^{\Delta F/kT}$. Choose a random number R between 0 and 1. If $B < R$, accept the mutation; otherwise, discard it. For $T = 0$, this reduces to the previous mutation algorithm.

Unfortunately, temperature has a uniformly negative effect on the mean global fitness of the system for any combination of K and C. As shown in Figure 17, the mean global

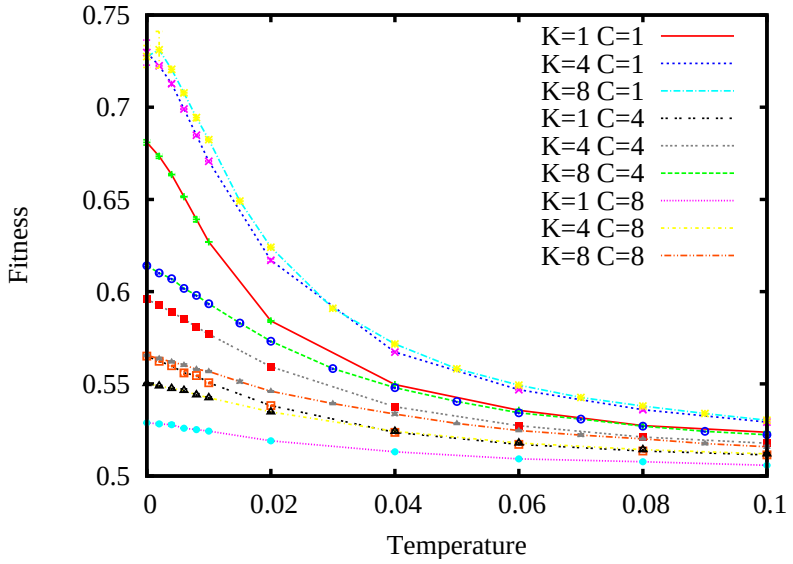

Figure 17: Mean global fitness decreases as temperature increases. Results are averaged over four runs of a $S = 8$, $N = 24$ system.

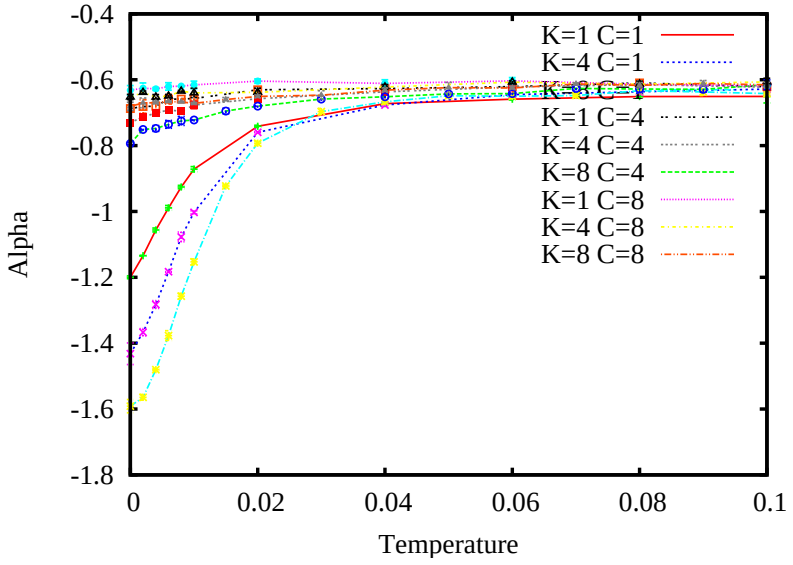

Figure 18: The mean power spectral density obeys a power law $1/f^{\alpha}$. The exponent decreases as temperature increases. Results are averaged over four runs of a $S = 8$, $N = 24$ system.

fitness rapidly decreases as T increases. This is not unreasonable; if T is too high, large negative mutations will often occur, moving species into a random section of the fitness landscape. As expected, this effect is greater when C is large due to the lower correlation length of highly coupled landscapes. For very large T, the mean global fitness approaches 0.5—completely random. It was hoped that for some intermediate value of T, there would in fact be a positive effect on fitness—but this hypothesis was not supported by experiment.

The power spectral density is also affected by the temperature. In every case shown in Figure 18, it remains governed by a power law of the form $1/f^{\alpha}$, but the exponent changes. As temperature gets very large, the exponent decreases, indicating more randomness in the system; the limit case of $\alpha = 0$ is equivalent to white noise.

A nonzero temperature also affects the character of avalanche distributions. Figure 19 shows mutation avalanches for various parameter combinations. Although larger avalanches do occur, they become exponentially unlikely. This is because the temperature introduces an independent probability for a mutation to occur at any iteration of the system, so to a first approximation there is a multiplicative probability that an avalanche will continue for a given number of steps. Finite temperature does not by this method appear to push the system to become critical. The delta-fitness avalanche distribution also remains governed by an exponential decay, as shown in Figure 20. Higher temperatures uniformly increase the number of fitness changes of a given size.

I have shown that there is no benefit to introducing a temperature to the system. Perhaps it is more difficult to escape local fitness maxima than expected—that is, after a single mutation away from a local maximum, there is a guaranteed route to a higher fitness (the previous genotype) and many routes to lower fitnesses (most other nearest-neighbor genotypes) and only a very slim probability of a better route to an ever higher fitness. In this model it has been shown that negative mutations do not improve future fitness values, since even small negative mutations have significant effects. It might be interesting to study temperature in systems with larger values of N, so that each negative mutation has a more negligible effect. When N is large, it is more likely that slightly negative mutations will be propagated and may eventually become useful. $N = 24$ is nowhere near the biologically realistic $N \approx 10^6$ for prokaryotes; it is possible that I have been unable to study the system in the proper regime.

6.3 Ecosystem connectivity

Almost all real ecosystems are not uniformly connected, in that the genes in each species do not all have the same number of couplings, nor are the couplings made to any other species. It is far more realistic to model inter-species interactions locally, by restricting the set of other species to which each species couples. The number of connections each species has may also vary; some species are more important than others. In the language of the model, this corresponds to species with variable values of C, and different ranges of species to which the C couplings can be made.

To study this sort of variability, I consider an ecosystem in which species are arranged on a ring. Each species is allowed to establish a coupling to any species within a certain

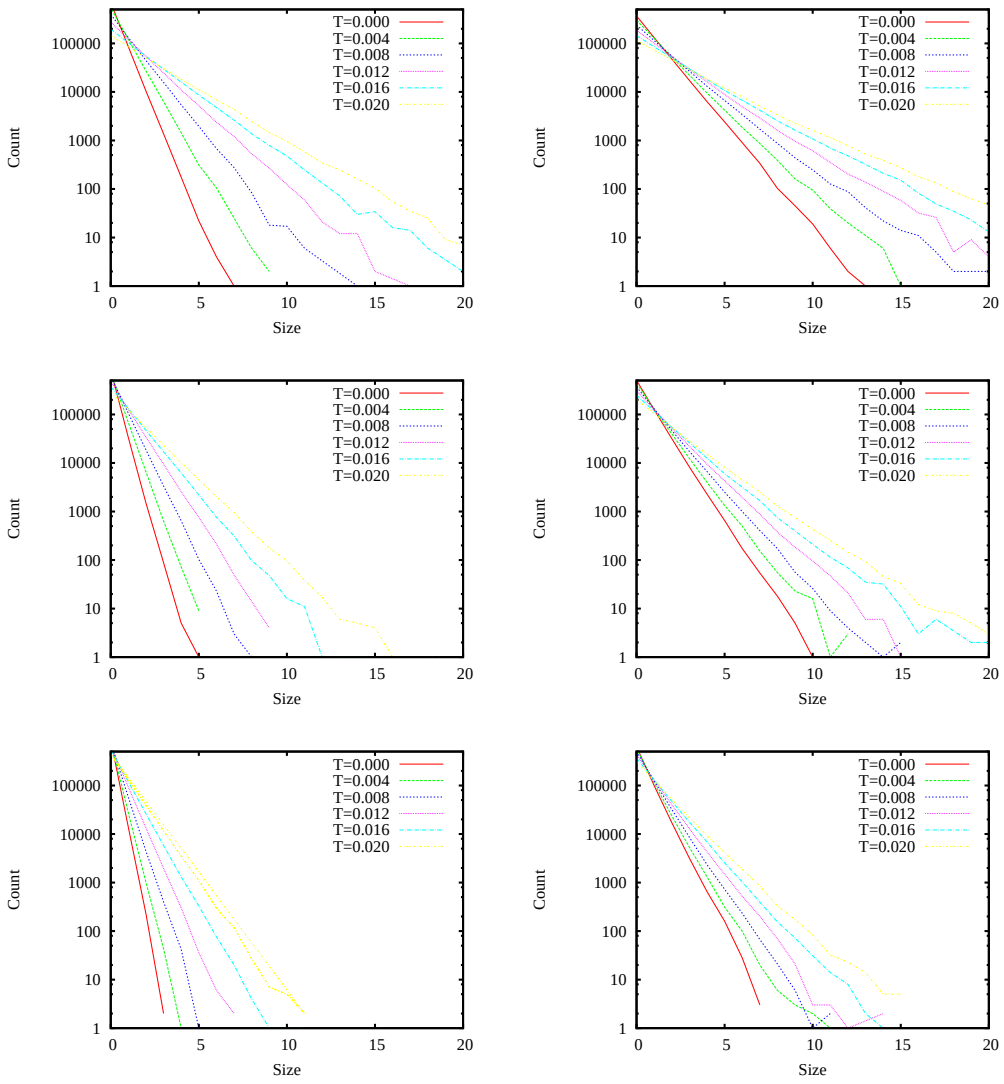

Figure 19: Regardless of temperature, mutation avalanche frequency decays exponentially with size. Typical results are shown for a $S = 8$, $N = 24$ system. Left column, $C = 1$; right, $C = 4$; from top to bottom, $K = 1, 4, 8$.

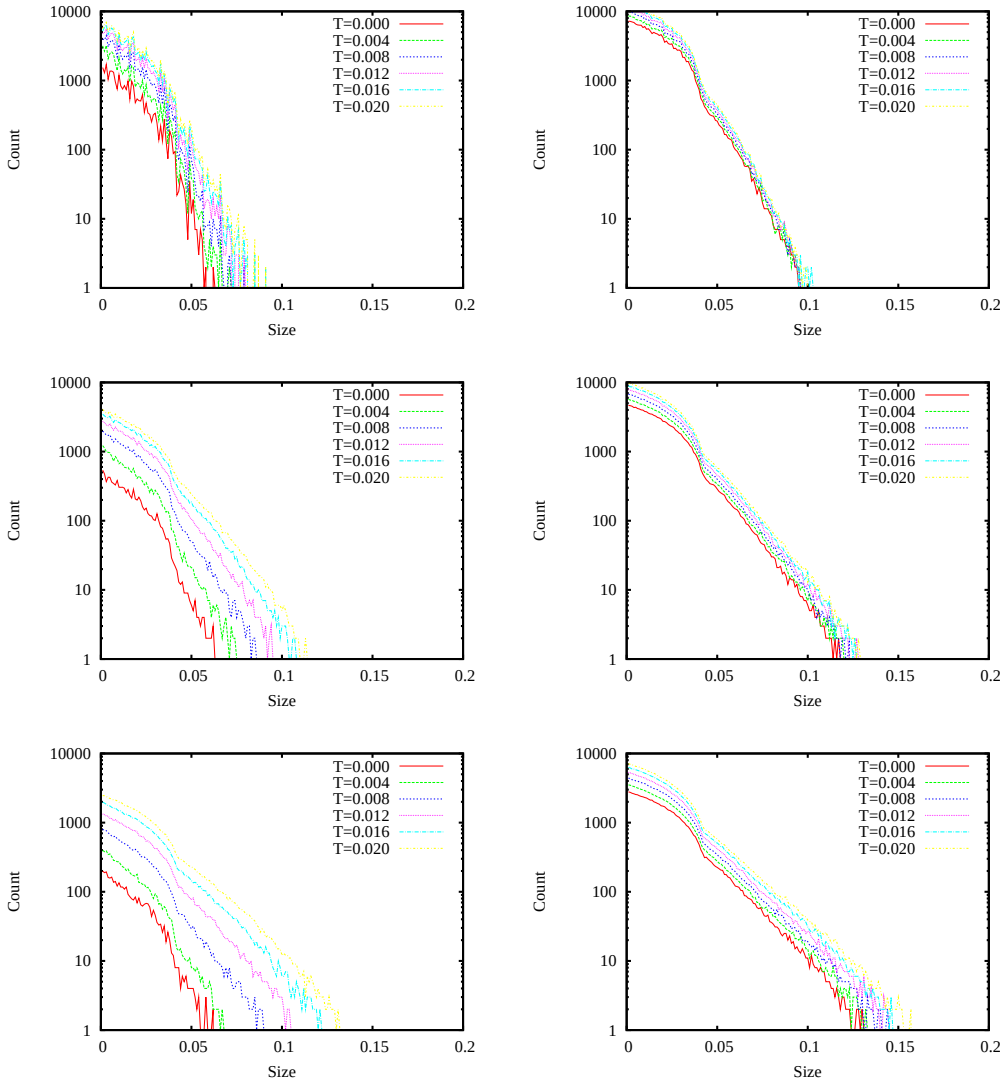

Figure 20: Delta-fitness events uniformly become more frequent at higher temperatures, but the general character of the distribution remains an exponential decay. Typical results are shown for a $S = 8$, $N = 24$ system. Left column, $C = 1$; right, $C = 4$; from top to bottom, $K = 1, 4, 8$.

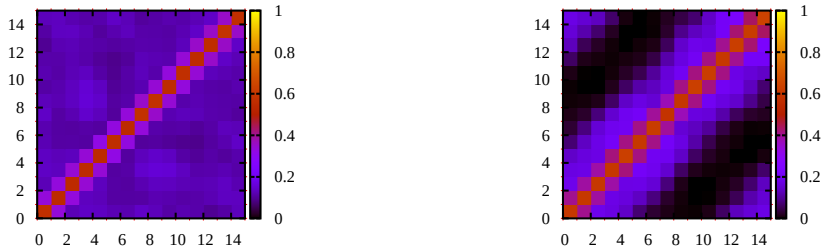

Figure 21: Typical delta-fitness correlation values for a system with external couplings randomly assigned between all species (left) and between those within $d = 4$ on a ring (right). Higher correlations are found for species which are closer to each other. Trivially, a correlation of unity is found between a species and itself. $S = 16$, $N = 24$, $K = 8$, $C = 4$.

distance $d \leq S/2$ (setting $d = S/2$ corresponds to fully random connections) For example, if $S = 8$ and $d = 2$, species 0 would be able to couple to species 1, 2, 6, and 7. and species 4 would be able to couple to species 2, 3, 5, and 6. This not only restricts the range of couplings, but also increases the probability that two species will have reciprocal couplings, hopefully increasing the stability and fitness of the ecosystem. The parameter d controls how possibly connected the ecosystem could be: for very low values, a mutation by one species could affect only a few other species, so avalanches do not propagate quickly. If d is large, only a few iterations of the system might be necessary for an avalanche to propagate to all species. Even if d is large, small values of C may prevent the restriction from having a significant effect on the system—that is, if a small number of couplings makes it unlikely that the full range of species indicated by d will be fully sampled.

In order to verify that nearest-neighbor type couplings restrict the range of fitness change events, it is possible to calculate the correlation between delta-fitness events in different species. If $F_{i,t}$ is the fitness of species i after t iterations of the system, then one can construct a list of values $c_{i,t} = \{F_{i,1} - F_{i,0}, \ldots, F_{i,t} - F_{i,t-1}\}$. The normalized correlation between delta-fitness events in species i and j is then defined as $\chi_{i,j} = \sum_t (c_{i,t} - \langle c_i \rangle) \sum_t (c_{j,t} - \langle c_j \rangle) / \sigma_i \sigma_j$. It is expected that for species pairs with $|i - j| \leq d$ the correlation will be significantly higher than for other pairs. A mutation by one species can cause fitness changes in only those species nearest-neighbor to it. This does indeed occur, as shown in Figure 21.

Varying d has little effect on other quantities that describe the system. As shown in Figure 22, there is no significant change in the mutation avalanche distribution or power spectral density. The mean global fitness does decrease slightly as d increases, but this effect is almost insignificant. There is certainly no optimal value for d, as

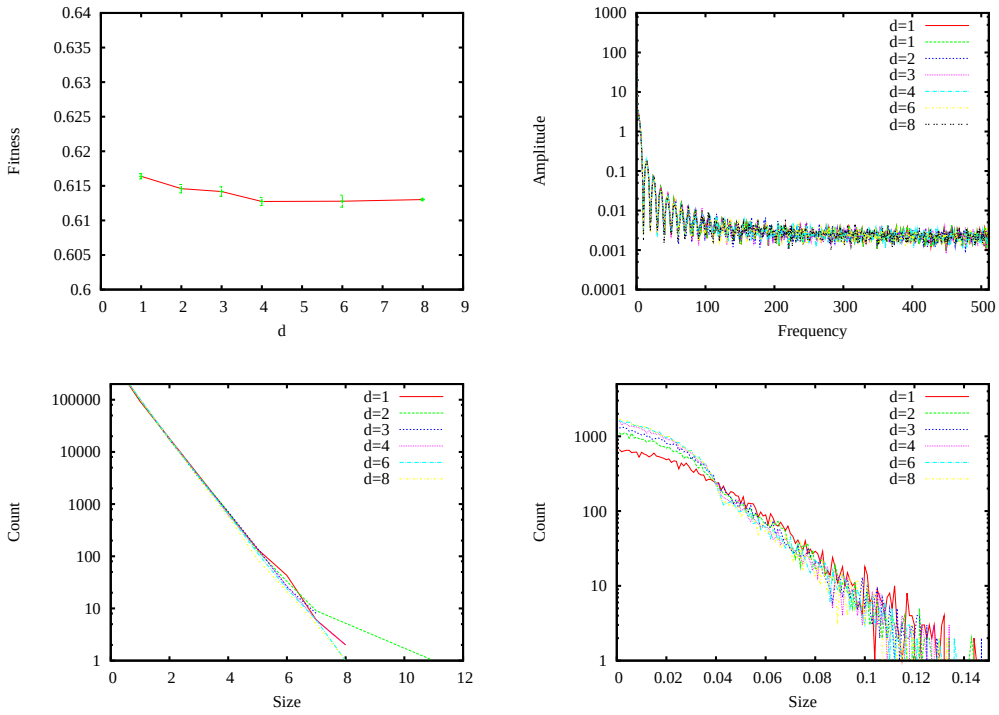

Figure 22: From top left, typical mean global fitness, power spectral density, mutation avalanche distribution, and delta-fitness distribution for a $S = 16$, $N = 24$, $K = 8$, $C = 4$ system. Varying d has little significant effect on these quantities.

hoped. However, smaller values of d result in fewer total fitness-change events, since a mutation in one species does not affect many others. However, the general character of the distribution remains the same. I have tested many combinations of parameters; these general conclusions do not appear to depend on the values of K and C.

Although I have restricted the range of inter-species connections, I have not privileged any species with more connections or a longer range than other; each species is still equally important relative to all others. This may be why d does not significantly affect the behavior of the system beyond correlations in delta-fitness events.

As a result, I have also investigated the effect of species which are more connected than others on the ecosystem. The simplest case occurs when all species are connected with the same values of C and d—except for one 'super species', which has $C = C + sss$ and $d = d + sss$ where sss is the 'strength' of the super species. If sss is large, then the super species is able to couple to more species more frequently than a typical species. For $sss > S/2$, couplings can be to any species and the effective value of C increases for the super species.

The effect on delta-fitness correlations is shown in Figure 23. As expected, for high values of sss, the super species has more of an effect on its nearest-neighbors than other species do. Since $sss = 8$ is the same as $S/2$, the effect is not yet that large. Similarly, there are more delta-fitness events for a super species than for a typical species. Mean global fitness for the super species is also lower than for typical species, because of the increased number of couplings. However, the power spectral density and mutation avalanche distribution do not change when a super-species is introduced. This is shown in Figure 24.

Thus, the limit-case for variable couplings does not show any signs of criticality, nor does it improve fitness. I have also varied the definition of sss such that it only affects either C or d; this also does not produce any interesting results. Certainly I have not had the chance to explore all of parameter space, but it appears as though variable connectivity does not have an important effect on the system.

It might be interesting to instead study this for far larger values of S and N, so that the behavior for low c and d could be better understood. However, this is presently intractable computationally.

7 Conclusions

I have studied the properties of evolutionary dynamics within the context of the NKC model. This model describes coevolving species by their genotype; each gene within each species can be coupled to other genes within the species and within other species, and each gene provides a contribution to the total fitness of the species. This creates a fitness landscape; species can mutate their genotype in order to climb the landscape to fitness maxima, but this process may deform the fitness landscape of other species. By tuning the number of couplings and the method of fitness maximization, I have been able to create many different types of fitness landscapes and evolutionary dynamics.

I have found that the NKC model exhibits many features of real evolution such

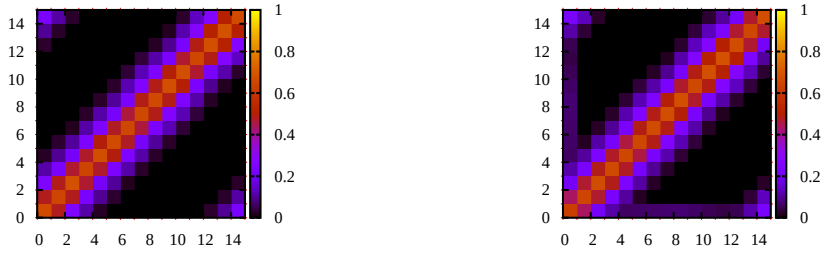

Figure 23: Typical delta-fitness correlation values for a system with left, $sss = 0$ (no super species); right, $sss = 8$. The super species is given label 0. Both results are for $d = 2$ and $S = 16$, $N = 24$, $K = 8$, $C = 4$.

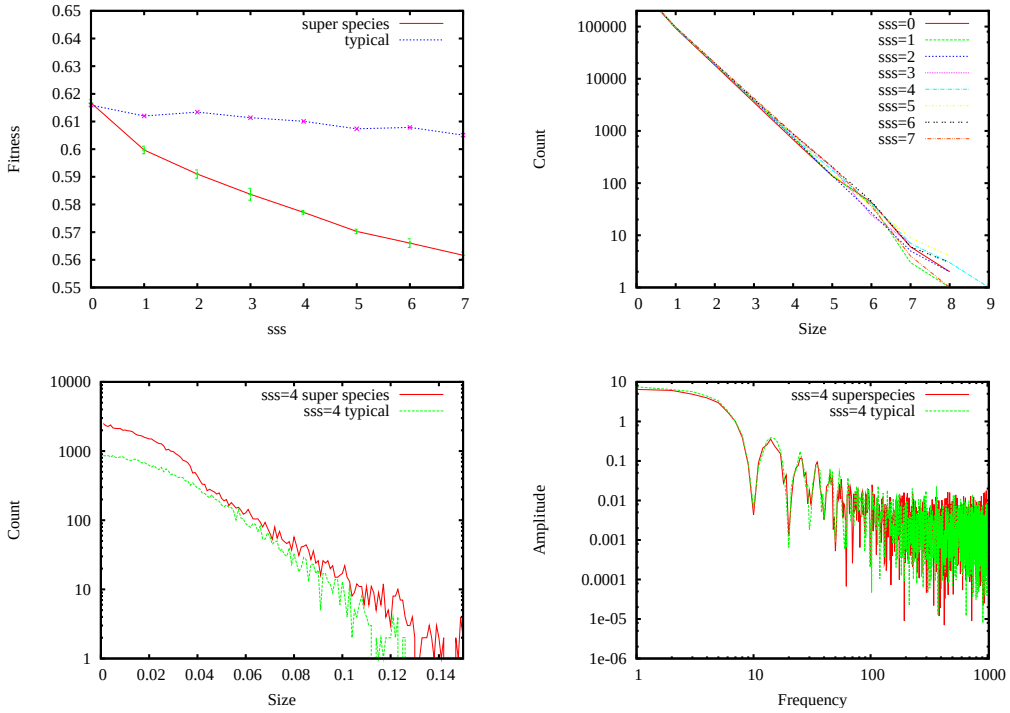

Figure 24: From top left, typical mean global fitness, mutation avalanche distribution, delta-fitness distribution, and power spectral density for a system with a 'superspecies'. The data in the first plot is averaged over four runs; the others are typical. All results are for a $S = 16$, $N = 24$, $K = 8$, $C = 4$ system.

as Nash equilibrium, punctuated equilibrium, oscillations, and chaos. By tuning the parameters of the models I have been able to determine what level of couplings and what fitness values are associated with each class of behavior. In particular, I have found that more couplings (interactions) between species leads to lower fitness and a reduced chance of equilibrium being reached.

Additionally, I have investigated the possibility of criticality in the NKC model, in the hopes of matching and possibly explaining the scale-free properties that have been found in biological evolution. Through the analysis of power spectral densities and avalanche distributions, I have not found any evidence for the existence of a critical phase in the model, though I have not allowed species to individually self-tune their values of coupling parameters.

Finally, I have extended the NKC model to include the possibility of negative mutations by introducing an effective temperature to the system, and have also investigated the effect of variable connectivity between species; in particular, restricting external couplings to nearest-neighbor species. Neither of these extensions produces criticality in the model, nor has a beneficial effect on fitness.

It is possible that more realistic simulations would have far more genes and species than couplings, so that the ratio of couplings to genes and species could be controlled more finely ($K + C \ll SN$). However, this would require more computational resources than were available to me. Nevertheless I hope that my limited studies of the model have provided some insight into the dynamics and principles of coevolution on fitness landscapes.

I would like to thank Adam Gamsa (Rudolf Peierls centre for Theoretical Physics, University of Oxford) for the invaluable support and feedback he gave while supervising this project. I am also indebted to José Luis Machado (Department of Biology, Swarthmore College) for sparking my interest in criticality in biological systems.

References

[1] Agur, Z., Kerszberg, M. *The emergence of phenotypic novelties through progressive genetic change.* American Naturalist (1987) 129(6), 862-875.

[2] Bak, P. *Coevolution in a rugged fitness landscape.* Phys. Rev. A (1992) 46(10), 6724-6730.

[3] Bak, P., Paczuski, M. *Complexity, contingency, and criticality.* PNAS (1995) 92, 6689-6696.

[4] Bak, P., Tang, C., Wiesenfeld, K. *Self-organized criticality: an explanation of* $1/f$ *noise.* Phys. Rev. A (1988) 38(1), 364-374.

[5] Binney, J.J. et al. *The theory of critical phenomena: an introduction to the renormalization group.* Oxford University Press, 1992.

[6] Durret, R., Limic, V. *Rigorous results for the NK model*. Annals of Probability (2003) 31(4), 1713-1753.

[7] Flyvbjerg, H., Lautrup, B. *Evolution in a rugged fitness landscape*. Phys. Rev. A (1992) 46(10) 6714-6724.

[8] Gisiger, T. *Scale invariance in biology: coincidence or footprint of a universal mechanism?* Bio. Rev. (2001) 76(2), 161-209.

[9] Gould, S.J., Eldredge, N. *Punctuated equilibrium comes of age*. Nature (1993) 366, 223-227.

[10] Havlin, S., et al. *Fractals in biology and medicine*. Chaos, Solitons and Fractals (1995) 6, 171-201.

[11] Jensen, H.J. *Self-organized criticality*. Cambridge University Press, 1998.

[12] Kauffman, S. *At home in the universe*. Oxford University Press, 1995.

[13] Kauffman, S., Johnsen, S. *Coevolution to the edge of chaos: coupled fitness landscapes, poised states, and coevolutionary avalanches*. J. Theor. Bio. (1991) 149, 467-505.

[14] Kauffman, S., Levin, S. *Towards a general theory of adaptive walks on rugged landscapes*. J. Theor. Bio. (1987) 128, 11-45.

[15] Levin, S.A. *On the evolution of ecological parameters*, in *Ecological genetics: the interface*. Springer-Verlag, 1978.

[16] Lynch, M. *The frailty of adaptive hypotheses for the origins of organismal complexity*. PNAS (2007) 104, 8597-8604.

[17] Park, S., Miller, K. *Random number generators: good ones are hard to find*. Communications of the ACM. (1988) 10.

[18] Reif, F. *Fundamentals of statistical and thermal physics*. McGraw-Hill, 1965.

[19] Smith, J.M. *Natural selection and the concept of a protein space*. Nature (1970) 225, 563.

[20] Solé, R.V., Manrubia, S.C. *Self-similarity in rainforests: evidence for a critical state*. Phys. Rev. E (1995) 51, 6250-6253.

[21] Solé, R.V., Bascompte, J. *Are critical phenomena relevant to large-scale evolution?* Proc. R. Soc. Lond. B (1996) 263, 161-168.

[22] Weinberger, E. D. *Local properties of Kauffman's NK model*. Phys. Rev. A (1991) 44, 6399-6413.

[23] Weiss, J.N., Qu, Z., Garfinkel, A. *Understanding biological complexity: lessons from the past.* FASEB J. (2003) 17(1), 1-6.

[24] West, G., Brown, J.H., Enquist, B.J. *A general model for the origin of allometric scaling laws in biology.* Science (1997) 276, 122-126.

[25] Wolfram, S. *A new kind of science.* Wolfram Media, 2002.

[26] Wright, S. *The roles of mutation, inbreeding, crossbreeding and selection in evolution.* Proceedings, Sixth International Congress on Genetics (1932) 1, 356.

[27] Yeomans, J.M. *Statistical mechanics of phase transitions.* Oxford University Press, 1992.