Evolution of Boolean networks under selection for a robust response to external inputs yields an extensive neutral space

Agnes Szejka and Barbara Drossel

Institut für Festkörperphysik, TU Darmstadt, Hochschulstrasse 6, 64289 Darmstadt, Germany

We study the evolution of Boolean networks as model systems for gene regulation. Inspired by biological networks, we select simultaneously for robust attractors and for the ability to respond to external inputs by changing the attractor. Mutations change the connections between the nodes and the update functions. In order to investigate the influence of the type of update functions, we perform our simulations with canalizing as well as with threshold functions. We compare the properties of the fitness landscapes that result for different versions of the selection criterion and the update functions. We find that for all studied cases the fitness landscape has a plateau with maximum fitness resulting in the fact that structurally very different networks are able to fulfill the same task and are connected by neutral paths in network ("genotype") space. We find furthermore a connection between the attractor length and the mutational robustness, and an extremely long memory of the initial evolutionary stage.

PACS number(s): 87.16.Yc, 89.75.Hc

I. INTRODUCTION

Boolean networks, where each node can take on only two possible values, are used to model phenomena as different as social interactions, gene regulation, or communication between neurons [1–3]. In the original "random" form of the model, introduced in 1969 by Stuart Kauffman, each of the N nodes receives an input from K randomly chosen other nodes. The value of each node i is updated at each time step according to

$$\sigma_i(t) = f_i[\sigma_{i_1}(t-1), \sigma_{i_2}(t-1), \dots, \sigma_{i_{K_i}}(t-1)],$$

where f is a randomly assigned Boolean function and σ_{i_1} to $\sigma_{i_{K_i}}$ are the inputs to node i. Starting from an initial state $\vec{\sigma} = \{\sigma_1, \dots, \sigma_N\}$, the dynamics will eventually lead to a periodic attractor. Kauffman intended this model to capture essential features of gene regulation networks, where the two states of a node correspond to a gene being expressed or not expressed. While it has become clear in the meantime that models that are constructed at random cannot fulfill this task, Boolean networks that are constructed based on biological data, can indeed successfully model the dynamical properties of real gene regulatory networks. The segment polarity network of *Drosophila melanogaster* [4] and the yeast cell-cycle network [5], for example, were modeled using Boolean dynamics of the genes, and both models show dynamical attractors that agree with the biological sequence of events.

The dynamical behavior of random Boolean networks differs significantly from that of the mentioned realistic models. Random Boolean networks can be in three different dynamical regimes, depending on the value of K and the statistical weights assigned to the different Boolean functions. These dynamical regimes are usually called frozen, critical or chaotic. The parameter λ, which is K times the probability that the value of a node changes when one of its inputs changes, is used to discriminate between these regimes [6,7]. A network is said to be in the frozen phase if a perturbation at one node propagates during one time step on an average to less

than one other node ($\lambda < 1$). In the chaotic phase, a perturbation propagates on an average to more than one node during one time step ($\lambda > 1$). Networks at the boundary between the two phases are called critical ($\lambda = 1$). In the frozen phase, the length of attractors is either 1 or very small. Most of the nodes are frozen, this means that they do not change their states anymore in the stationary state. In chaotic networks, attractors are very long on average, and a nonvanishing proportion of the nodes change their states on the attractors. Both frozen and chaotic networks typically have a small number of attractors. In critical networks, the number of relevant nodes (i.e., those nodes that determine the attractors) increases with the network size only as $N^{1/3}$, while the average number and length of attractors increases exponentially with the number of relevant nodes.

In contrast, the realistic Boolean models for gene regulation combine features of "chaotic" and "frozen" networks. Like frozen networks, they have a robust dynamics, with attractors being reached after a few update steps. As in chaotic networks, one (nontrivial) attractor dominates the dynamics with its basin of attraction (number of states leading to and lying on the attractor), and a considerable proportion of nodes change their state while the network is going through the biologically relevant sequence of updates.

Model networks with these realistic features can be obtained by starting from random networks and subjecting them to an evolutionary process. In [8], we evolved a random network by selecting for dynamical robustness, i.e., by requiring that a network that has reached an attractor returns to this attractor after a node is perturbed. The network was mutated by changing a function or by rewiring a connection, and after a small number of mutations, a network could be obtained that satisfied fully the selection criterion. When the network was evolved further via neutral mutations, i.e., mutations that do not change the fitness, the basin of attraction of the main attractor became eventually very large. The most striking result of this investigation was that the space of all possible networks contains a percolating subset of networks with 100% fitness, with neighbors in this subset being con-

nected by one mutation. The evolution of an entire population of networks under the mutation and selection rules mentioned above was studied in [9].

Studying the evolution of Boolean networks has a long tradition. Previous publications used only subsets of the above-listed possible mutations on Boolean networks. The mutations chosen by Kauffman and Smith [10] are the rewiring of connections or the changing of bits of the Boolean functions. The fitness measure is the Hamming distance of an attractor to a predefined target state. This simulated evolution is not able to produce networks with 100% fitness. In the articles by Bornholdt and Sneppen [11,12] the possible mutations are adding and/or removing a link, and the selection criterion is mutational robustness, that is the daughter network has to reach the same attractor as the mother network. Long time evolution shows persistence of high K values. The evolved networks have shorter attractors and larger frozen components than random networks. In the publications by Paczuski *et al.* [13] and Bassler *et al.* [14], the nodes of the network represent competing agents and a rule based on game theory determines the least fit node. The update function of this node is changed. The network evolves to a stationary state "at the edge of chaos" and is highly canalized. In the work of Braunewell and Bornholdt [15], the mutations are the rewiring of links, and networks are selected for robustness of attractors against small perturbations in the update times. Only a small number of mutations is required to find a fully stable attractor set. During the evolutionary process, the average number of attractors decreases and the size of the largest basin increases.

The random Boolean models and selection criteria chosen so far do not take into account a central feature of real gene regulation networks: real regulatory networks respond to external inputs. Therefore, we want to investigate in this paper the evolution of Boolean networks under selection for a robust response to changes in an external input. Depending on the considered application, two versions of the selection criterion are relevant: cells need to be able to respond reversibly to changes in the environment, such as a change in the available food source or a temporary stress. In this case, a reversible switching between attractors has to occur. On the other hand, cells perform an irreversible switching between attractors during embryonic development, where cells change their gene expression pattern when the organism goes to a new developmental stage. In addition to following one of these two switching patterns, we require both attractors to be stable under perturbations of one node (if this node is not the external input node).

In order to check if our results are generic, we use two different sets of update functions, both of which are considered realistic for gene regulation networks. Our first choice of functions is motivated by [5], where threshold functions are used to model the yeast cell-cycle network. Each connection is thus either activating or repressing, and the incoming signals are added up to determine if the total input is above or below the threshold for switching a node on. The second set of functions used in this paper is the set of canalizing functions, which are found to be abundant in eukarotic gene regulatory networks [16].

We perform the evolution of a single network by means of an adaptive walk, where every mutation that does not lower

the fitness is accepted. Adaptive walk simulations were also performed in [8]. They are a valuable tool for obtaining insights into the properties of the fitness landscape. A fitness landscape represents the fitness as function of the genotype. This can be done either by a direct mapping, or via the phenotype. The "genotype" of a network is its structure and the "phenotype" is given by its dynamical behavior.

In the next two sections, we describe the details of our dynamical update rules and of our adaptive walk simulations. Then, in Sec. IV the properties of the evolved networks and of the fitness landscape obtained after short and after very long times are presented and discussed. In the last section, we summarize and evaluate our findings.

II. UPDATE RULES

The first type of update functions used in our investigation is threshold functions. The value of a node in the next time step is determined in the following way:

$$\sigma_i(t+1) = \begin{cases} 1, & \sum_j c_{ij}\sigma_j(t) - h > 0, \\ 0, & \sum_j c_{ij}\sigma_j(t) - h < 0, \\ \sigma_i(t), & \sum_j c_{ij}\sigma_j(t) - h = 0, \end{cases} \quad (1)$$

where h is a threshold that is the same for every node [17]. The couplings c_{ij} are ± 1 with equal probability, and $c_{ij}=0$ if node i receives no input from node j. So the input $s_j=c_{ij}\sigma_j$ from node j to node i can take three different values: $0, +1$ or -1. A node is activated when the sum of its inputs exceeds the threshold value and is inactive when the sum of its inputs is below the threshold. When the sum of the inputs gives exactly the threshold value, the node does not change its state in the next time step. The phase diagram of random Boolean networks with this set of update functions was explored in [17]. The practical advantage of using this set of functions lies in the fact that nodes with many inputs can be implemented with little numerical effort and with little memory usage.

The second type of update functions used in this study are canalizing functions, as implemented in [18] and used in [8]. These update functions fall into the following four subclasses

$$f(\sigma_1, \sigma_2, \ldots) = \sigma_1 \text{ OR } g(\sigma_2, \ldots), \quad (2)$$

$$f(\sigma_1, \sigma_2, \ldots) = (\text{NOT } \sigma_1) \text{ AND } g(\sigma_2, \ldots), \quad (3)$$

$$f(\sigma_1, \sigma_2, \ldots) = (\text{NOT } \sigma_1) \text{ OR } g(\sigma_2, \ldots), \quad (4)$$

$$f(\sigma_1, \sigma_2, \ldots) = \sigma_1 \text{ AND } g(\sigma_2, \ldots), \quad (5)$$

where the four subclasses are chosen with equal probability for each node. Here, σ_1 is the canalizing input that determines the output of the function independently of the other nodes when it takes its canalizing value. A random Boolean function g that depends on the remaining variables deter-

mines the output when σ_1 is not on its canalizing value. The function g is generated by choosing with the same probability the output value 0 or 1 for every input combination.

III. RULES FOR THE SIMULATION OF THE ADAPTIVE WALK

The adaptive walk is a hill climbing process that leads to a local fitness maximum. If repeated often enough, it yields valuable insights into the fitness landscape of a system. We start our simulations by creating a random network with K inputs per node and with update function chosen at random from the considered set. One of the nodes is chosen to be the external input. It therefore receives no input from the other nodes. In order to determine the fitness of the initial network (and later that of the mutated networks), first the network is updated according to the Eqs. (1) or (2)–(5), until it reaches an attractor. Then the value of each node (except the one of the external input) is flipped one after the other, and it is counted for how many of these flips the network returns to the same attractor. This can happen at most $N-1$ times. Then the external input is flipped and the network is again updated until it reaches an attractor. If this attractor is the same as before, the fitness is set to 0. If a different attractor is reached the nodes are again flipped one after the other, and the robustness of the second attractor is evaluated. The maximum fitness value, which can be obtained based on the robustness of the two attractors, is $2(N-1)$. The next step depends on the chosen fitness criterion. The external input is switched back to its initial value. The "stay criterion" demands that the network stays on the second attractor, the "return criterion" demands that the network returns to the first attractor. Not fulfilling this criterion leads to a fitness decrease. We performed our simulations with two different fitness functions. In a first version, the fitness was lowered by 1 if the criterion was not fulfilled. In a second version, the fitness was set to zero. While both versions assign to the same networks a fitness value of 100%, the path taken to the fitness maximum differs. For the first version, networks tend to first improve the robustness of their attractors and then to satisfy the "return" or "stay" criterion, while the order is reversed for the second version. After determining the fitness of a network, a mutation is attempted and is accepted if it does not lower the fitness. Neutral mutations are those that do not change the fitness value. This procedure of attempting a mutation and evaluating the fitness is iterated until a certain stopping condition is satisfied. In the cases considered in this paper this stopping condition is a predetermined number of accepted mutations. We performed the following four types of mutations, all of which occur with the same probability: the addition, the deletion or the redirection of an input, or the change of an update function. An update function is changed by choosing a new random function in case of a canalizing network and by changing randomly the signs of the interactions in case of a threshold network. After choosing the type of mutation, the node at which this mutation shall be performed is selected at random. For canalizing networks, we limited the maximum number of inputs to $K_{max}=10$, since simulations become too slow if we go beyond this value. For
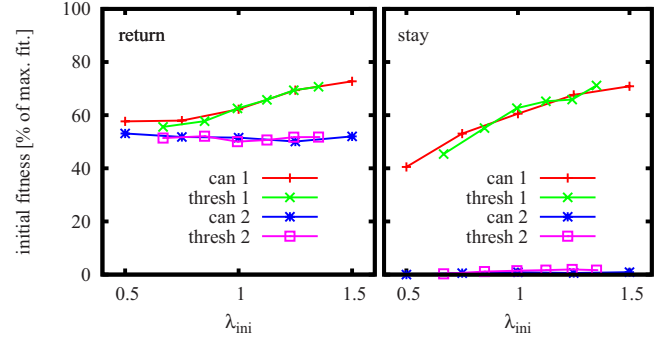

FIG. 1. (Color online) Fitness of the initial randomly generated networks for different initial connectivities K (initial sensitivities λ). Except for the second version of the "stay" criterion, initial random networks already have a quite high fitness.

threshold networks, $K_{max}=N$. The minimum number of inputs is $K_{min}=1$ in both cases. If the chosen mutation could not be performed at the selected node because of these constraints, a different node was chosen.

IV. RESULTS

The simulations were run for network sizes $N=50$ and initial connectivities ranging from $K_{ini}=1$ to 5 for canalizing networks and from 2 to 7 for threshold networks. To be able to compare the properties of the adaptive walk until the networks reach maximum fitness and the evolved properties of the two network types, our data are not plotted against K_{ini} but against λ_{ini} which is calculated according to Eq. (9) in [18] for canalizing networks and according to Eqs. (3) and (5) in [17] for threshold networks. For threshold networks, the parameters were chosen as $K_{ini}=2$ to 7 and $h=-0.5$, so that λ_{ini} lies between 0.5 and 1.5, just as for canalizing networks. λ thus covers the frozen, the critical and the chaotic regime. In the next subsection, we will investigate the adaptive walk up to the moment where maximum fitness is reached. The properties of the path to 100% fitness and the properties of the evolved networks are evaluated. In the subsequent subsection, we investigate the neutral evolution of the networks after they have reached maximum fitness. Finally, in the last subsection we will examine some networks with maximum fitness in detail and describe how they fulfill the selection criteria.

A. Path to maximum fitness

Every point in the plots of this section is an average over at least 600 network realizations. On the left hand side of each figure the results for the "return" criterion are plotted for both versions of defining the fitness function and for canalizing and threshold networks. On the right hand side the same is done for the "stay" criterion.

First, we evaluated the fitness of the initial networks (Fig. 1). Except for the second version of the "stay" criterion (where the fitness is set to zero if the network does not stay on the second attractor after switching again the external input), the initial networks have an average initial fitness of more than 40% of the maximum fitness, which is due to the
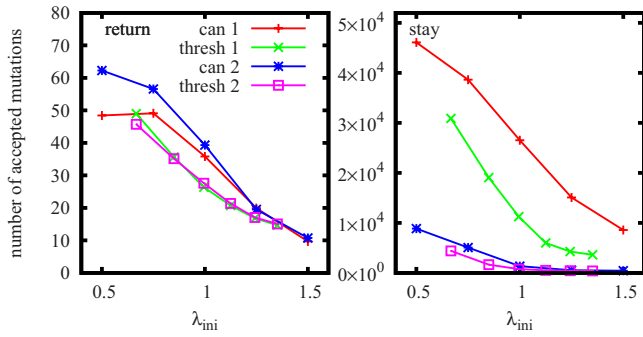
FIG. 2. (Color online) Number of accepted mutations until the fitness maximum is reached. The two versions of the "stay" criterion are much more difficult to fulfill, many more mutations are needed to reach 100% fitness. Independently of the selection criterion or update rule, networks that started with higher connectivity needed less mutations to reach 100% fitness.

fact that the networks have a considerable probability of returning to the same attractor after perturbing one node. For frozen and critical networks, this is due to the large number of frozen nodes. Chaotic networks also have a frozen core, but its size decreases with increasing λ. In our investigation in [8], where dynamical robustness was the only selection criterion, the initial fitness was found to decrease with increasing λ and with increasing N. Figure 1 shows no such decrease with increasing λ, apparently because the smaller dynamical robustness of chaotic networks is compensated for by the larger flexibility, which makes it easier to switch between two attractors upon changing the external input. The probability of returning to the first attractor after switching back the value of the input node, is so high that the average initial fitness decreases by less than 1/3 between the first and second version of the "return" criterion and decreases significantly only for the second version of the "stay" criterion. Since we have plotted fitness as function of the sensitivity λ, the average fitness does not depend on the set of update functions used. For the first version of both selection criteria, initial fitness grows with increasing initial connectivity or initial λ. It is easier for networks with more chaotic dynamics to fulfill our fitness criteria, because there are less (but longer) attractors.

Figure 2 shows the length of the path, measured in the number of accepted mutations, until reaching maximum fitness. There is a huge difference between the two criteria. Networks evolved with the "stay" criterion need far more accepted mutations until they reach maximum fitness, which shows that the "stay" criterion is much more difficult to satisfy than the "return" criterion. For both criteria and both update rules, less mutations are needed to reach the global fitness maximum when λ_{ini} is larger, even when networks start with similar initial fitness. More chaotic networks have less frozen nodes and can be changed more drastically with one mutation. Dynamical robustness appears to evolve faster with threshold functions than with canalizing functions, since the number of mutations required to reach maximum fitness is lower in the first case.

Figure 3 shows the percentage of accepted mutations among the attempted mutations. For the "return" criterion
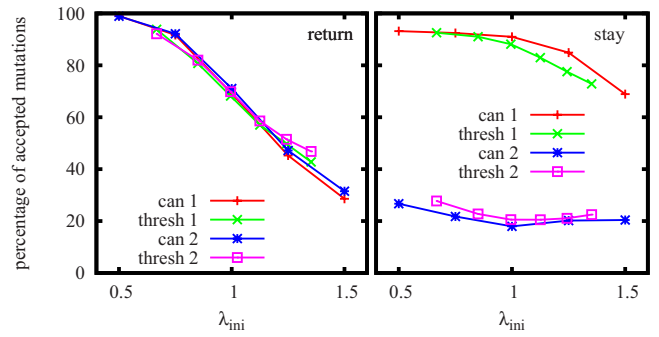

FIG. 3. (Color online) Percentage of accepted mutations until fitness maximum is reached. The percentage of accepted mutations until the fitness maximum is reached differs between the two criteria. Networks with different update rules behave similarly, but there is a big difference between the two versions of the stay criterion.

(Fig. 3 left), it is independent of the version or of the update rules used. The percentage of accepted mutations decreases with increasing λ_{ini} or K_{ini}. Taking together Figs. 2 left and 3 left, one can see that although chaotic networks are closer to maximum fitness, more attempted mutations are needed in order to find the successful mutations. For the "stay" criterion, the decrease in the percentage of accepted mutations with initial λ is not as pronounced as for the "return" criterion. For the second version of the "stay" criterion it even seems to be always 20% independent of λ_{ini}. Apparently, almost all mutations are performed during the search for a network that satisfies the "stay" criterion. Once it is satisfied, high robustness is obtained very quickly.

Next we look at the ratio between the length of the attractors at the end and at the beginning of the adaptive walk to the fitness maximum (Fig. 4). For the "return" criterion the change in attractor length is similar for both versions of the criterion, for both sets of update rules and for both attractors. Attractors of networks that start in the frozen phase become somewhat shorter during the adaptive walk, attractors of networks that start with a $\lambda \geq 1$ become slightly longer on the
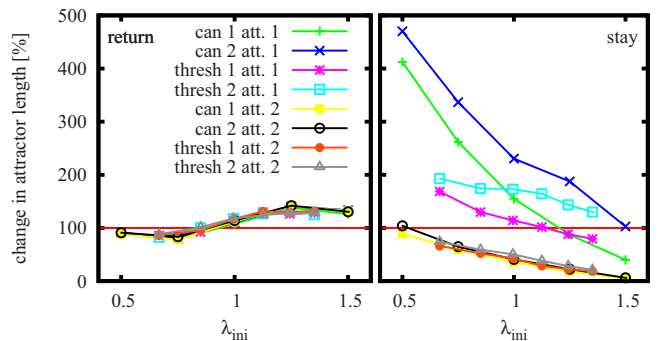

FIG. 4. (Color online) Change in attractor length between initial networks and 100% fitness networks. For the "return" criterion the change in attractor length during evolution is not as pronounced as for the "stay" criterion. For the "return" criterion the lengths of both attractors change in a similar way. For the second "stay" criterion the first attractor always grows during evolution, for the first "stay" criterion it grows for small λ_{ini} and shrinks for large λ_{ini}; the second attractor always shrinks.
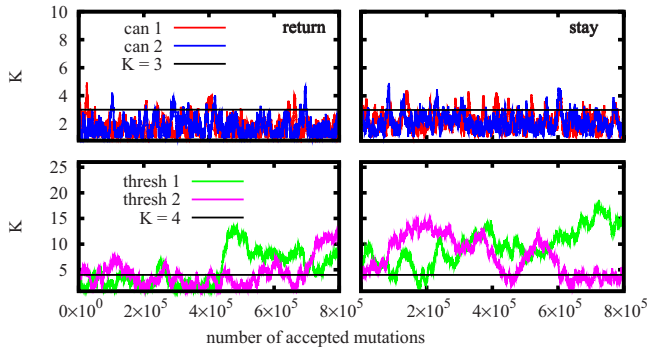
FIG. 5. (Color online) Average connectivity as function of evolutionary time for networks with 100% fitness. Evolution finds many networks with 100% fitness, but the average connectivity is never high.

FIG. 6. (Color online) Length of the first attractor as function of evolutionary time. For the "stay" criterion the first attractor is larger than for the "return" criterion.

way to the maximum. Since the number of mutations between the initial and final network is small, the change in attractor length is also small. Attractors in frozen networks can be made more robust by making them even shorter, while attractors in chaotic networks become more robust by making them larger, because this means there are less attractors left and that their basin of attraction is larger.

With the "stay" criterion, the second attractor always shrinks for all networks. The shrinking is more pronounced for larger λ_{ini} because the networks start with larger attractors but end up with very short attractors. It appears that the "stay" criterion is easier fulfilled when the second attractor is a fixed point and when the network is frozen. In contrast, the first attractor must be sufficiently large that it can respond to a change in the external input. This effect is more pronounced for canalizing functions. For the second version of the "stay" criterion, the first attractor always grows independent of the update rules used (at least for the λ values investigated by us). The initial effort to first find a network that satisfies the "stay" criterion, increases the length of the first attractor by a larger amount than when the network is first optimized for dynamical robustness. In all cases the two attractors have equal weight in state space, that is their basins have similar sizes.

B. Neutral evolution on the plateau of maximum fitness

After the networks have reached maximum fitness, they can evolve further via neutral mutations. Apparently, there is a huge plateau of maximum fitness that spans the entire network space. For each version of the model, we continued the evolution for a very long time, until the networks had undergone 800 000 accepted mutations. Figure 5 shows the evolution of the average connectivity (denoted by K). The large fluctuations of K indicate that structurally very different networks are able to fulfill the same task. The horizontal lines indicate the K values for which $\lambda = 1$. They would separate chaotic networks from frozen networks if the networks were random Boolean networks. Although the evolved networks cannot be classified in this way, there is nevertheless some correlation between the attractor length and the average connectivity (see below).
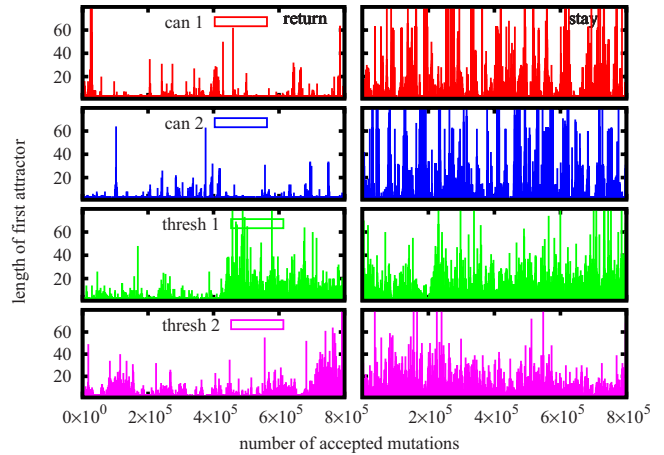
For canalizing networks, the K values behave similarly for the "return" and the "stay" criterion. Most 100% fitness networks have K values between 1 and 3, but solutions with $K > 3$ are also found by the evolutionary process. As mentioned above, we have imposed a cutoff at $K_{max} = 10$, but this is of no relevance since the networks did not come even close to this value. Similarly, the cutoff $K_{max} = 50$ for threshold networks was of no relevance. Compared to canalizing networks, the K value shows correlations over longer time intervals, and they spend more time above the horizontal line. The canalizing networks appear to experience a "force" which tends to decrease the K value, while the K value of the threshold networks is less restricted. This may be due to the fact that the addition of a link in a threshold network makes the response of a node to its inputs only slightly more complex, while it makes the response far more complex in a canalizing network.

For threshold networks, there is a difference between the two fitness criteria. For the "stay" criterion, the typical K values appear to be larger, which means that the above-mentioned "force" is weaker in this case.

The difference between version 1 (where the fitness is lowered only by a small amount when the "stay" or the "return" criterion is not fulfilled) and version 2 (where the fitness is set to zero in this case) has no meaning after 100% fitness have been reached. Nevertheless, the K values appear to be smaller for the first version. This may simply give an impression of the variations between separate simulation runs, or it may indicate an extremely long memory of certain features of the initially evolved networks. Below, we will present evidence for the second interpretation.

In Fig. 6, the lengths of the first attractors are plotted over evolutionary time. In all cases the "stay" criterion produces networks with larger attractors than the "return" criterion. While for canalizing networks, this cannot be explained by larger K values, this might well be the predominating cause for threshold networks. The investigation below of the mechanisms by which networks manage to satisfy the "stay" criterion shows that these networks must contain feedback loops, which can propagate a signal from the first node af-
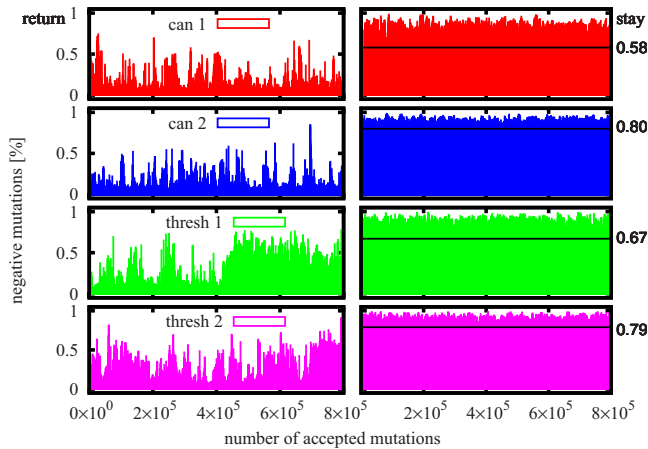
FIG. 7. (Color online) Percentage of mutations that lower the fitness as function of evolutionary time.

fected by the external input back to this node. This means that the first attractor must be far from frozen. As shown in [19] feedback structures lead to more complex dynamics with longer attractors. Threshold networks seem to be able to buffer an emerging feedback structure more efficiently than canalizing networks.

In Fig. 7, the percentage of negative mutations among the attempted mutations is plotted for every 100 accepted mutations. Negative mutations are those that lower the fitness and that are not accepted. The horizontal lines in Fig. 7 indicate the average values. The percentage of negative mutations is generally lower for the "return" criterion. In this case one can see no significant difference between the two versions of the criterion or between networks with different update rules. Comparison with Fig. 5 shows a correlation between the value of K and the percentage of negative mutations for the "return" criterion but not for the "stay" criterion. This means that for the "return" criterion networks with larger K (that are therefore probably also more chaotic) are more sensitive to mutations. Their mutational robustness is lower for larger K. The loss of fitness can be attributed to attractors that are less robust against perturbations. Such a correlation between robustness to noise and robustness to mutations was previously found in evolutionary models of gene regulatory networks [20–22]. A correlation between robustness to noise or temperature and robustness to mutations can also be found in micro-RNA [23]. For the "stay" criterion the percentage of negative mutations fluctuates around an average value that is indicated on the right hand side of the figure. The second version of the criterion shows an increase in negative mutations both for canalizing and for threshold networks. Thus for networks evolved with the "stay" criterion the loss of fitness—as it is equally probable for all values of K—must be due to the incapability of the networks to stay at the second attractor. But it is surprising that there are differences between the first and the second version of the criterion, since this difference does not matter once the maximum fitness has been reached. The initial evolution seems to lead to different regions in networks space depending on the version of the "stay" criterion.

To assess the "phenotypic" similarity between two neighboring networks on the plateau of 100% fitness we computed the difference between the attractors of two subsequent networks of an adaptive walk for 100 steps. The difference between two attractors was measured by counting how many nodes are frozen on one attractor and blinking on the other attractor in each of the two networks, and how many of the nodes that are equally frozen on both attractors have different states. While a thorough quantitative evaluation extends the limits of our study, we were able to see the following trends: the similarity is larger for threshold than for canalizing networks and for networks evolved with the "stay" criterion than for those evolved with the "return" criterion. For threshold networks evolved with the "stay" criterion, for growing connectivity a growing number of mutations is phenotypically neutral, with the second attractor changing more slowly than the first one.

C. Properties of networks with maximum fitness

As we have seen in Sec. IV A, the "stay" criterion is much more difficult to fulfill than the "return" criterion. In Sec. IV B we have further seen that the "stay" property is also easier to destroy.

For each of the eight studied cases (two update rules, two fitness criteria, two ways of assigning the fitness), we picked the first network that had reached 100% fitness and eight additional networks at time intervals of 100 000 mutations, and we examined how exactly these networks fulfill the imposed tasks. The solutions found by the evolutionary process differ obviously between the two selection criteria. However, we found no obvious difference between the different versions of assigning the fitness. For canalizing and threshold functions, the solutions are implemented in ways appropriate to the update rules, but otherwise there are no qualitative differences between the types of solutions found by the evolutionary process for the two sets of update functions.

For the "stay" criterion, the signal from the external input usually propagates only through one of the nodes connected to this input, even when there are many such nodes. This is because there has to exist a feedback through the network to the nodes that directly react to the external input to keep them from switching to their former behavior when the external input is switched back. The external input itself must not canalize this node. In comparison, the "return" criterion imposes far less constraints on the networks. The external signal may propagate through more than one output of the external node without detrimental effects, and the external node may be canalizing or not for these nodes.

Switching between the two attractors generally happens either through a cascade of switching events where frozen nodes change their value, or by freezing blinking nodes or unfreezing frozen nodes. Usually, just one of these possibilities occurs in one network. The switching cascade happens when both attractors are fixed points. Mixed forms, where switching cascades of frozen nodes are combined with the freezing or unfreezing of nodes, also exist but are frequent only for threshold networks evolved with the "stay" criterion. Rarely it also occurs that the same nodes are blinking but in a different fashion.

In Fig. 8, the network parts relevant for the switching between the attractors are plotted for canalizing networks
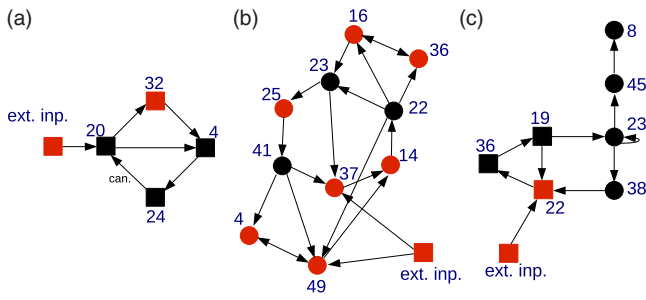
FIG. 8. (Color online) Relevant parts of three canalizing networks (a, b, c) with 100% fitness evolved with the "stay" criterion. Nodes that are depicted as circles blink at least on one of the attractors. Nodes depicted as squares are frozen on both attractors. Gray (red) nodes have the value 1, black nodes have the value 0.
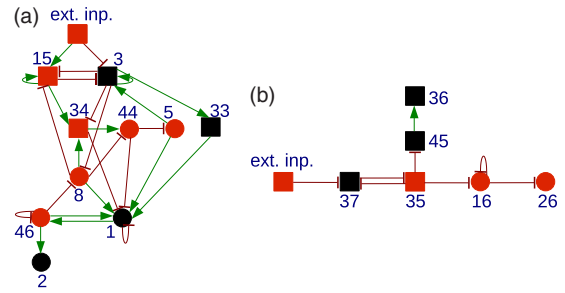

FIG. 9. (Color online) Relevant parts of two threshold networks (a, b) with 100% fitness evolved with the "stay" criterion. Arrows (green) represent activating links, blunt-end lines (red) represent repressing links. Squares depict frozen nodes and circles depict nodes that blink on at least one of the attractors. Gray (red) nodes have the value 1, black nodes have the value 0.

evolved with the "stay" criterion. Nodes that are depicted as circles blink on either one of the attractors or on both. Nodes depicted as squares are frozen on both attractors. The first example in Fig. 8 is the simplest. Here both attractors are fixed points. When the external input is switched on, node 20 is switched off then node 32 is switched on, then node 4 is switched on, and finally node 24 is switched off and canalizes node 20 to the "off" value so that it does not switch back when the external input is switched off again. (The external input has a second output to the network that is not relevant for the attractors.) The second example is also the second most common in networks evolved with the "stay" criterion. If not both attractors are fixed points, the first attractor is usually much bigger than the second, which is a fixed point or has sometimes the length 2. Thus, on the first attractor there are blinking nodes that become frozen after the external input has changed. In Fig. 8(b), all nodes apart from the external input are blinking on the first attractor. On the second attractor all of them are frozen. Here, the external signal propagates through two of the four possible links, and both nodes have to be fixed to their frozen value by feedback loops. The third example shows the rather uncommon case where first the three frozen nodes 22, 36, and 19 are switched, and then node 23 stops blinking, and with it also the three other blinking nodes. These three very different solutions of the "stay" criterion occurred during the same adaptive walk and are therefore connected by neutral mutations.

In Fig. 9, examples for the relevant parts of threshold networks are plotted. Again nodes that are frozen on both attractors are depicted as squares and blinking nodes as circles. Arrows depict activating (+1) links and blunt-end lines repressing links (−1). In both examples, the blinking nodes become frozen on the second attractor. Like in canalizing networks one can see feedback loops that prevent the network from going back to the first attractor. The most obvious difference between canalizing and threshold networks is that the relevant parts of threshold networks appear more complex than the relevant parts of canalizing networks with comparable attractor length. This happens because the connectivity of these networks is much higher since the value of the parameter λ rises much slower with connectivity than for canalizing networks. As the connectivity of threshold net-

works is higher than that of canalizing networks most of the time, the number of outputs of the external input is higher in these networks, but the signal propagates only through one or two of them as in the canalizing counterparts. Also, the probability of finding fixed point attractors in networks with large K seems to be higher in threshold networks than in canalizing networks.

V. CONCLUSIONS

In this paper we investigated the evolution of Boolean networks with two different update schemes under selection for two robust attractors, between which the networks switch reversibly or irreversibly upon action of an external signal. Evolution was simulated by means of an adaptive walk, a hill climbing algorithm where a single network is mutated and the mutation is accepted if it does not lower the fitness.

We found that canalizing and threshold Boolean networks show similar behavior during evolution to maximal fitness but differ at later times. Requiring that the networks stay at the second attractor when the external input returns to its former value turned out to be a much stricter criterion than requiring that the networks return to the first attractor. More mutations are needed to fulfill the first criterion, and maximum fitness is much easier destroyed, that is many more mutations lead away from 100% fitness. It seems that creating and sustaining a feedback that prevents nodes from switching to their former behavior is hard to achieve. In contrast, a network that is able to switch robustly between two different attractors seems to be easy to implement. However, networks always reached maximum fitness independently of the selection criterion and were additionally able to evolve further on huge plateaus of 100% fitness. On this plateau, threshold networks evolve more often to connectivities in the "chaotic" regime. The networks reached even higher connectivities when evolved with the "stay" criterion than when evolved with the "return" criterion. For networks with high connectivity, many mutations are "phenotypically neutral," that is, the attractors do not change. When investigating in more detail the evolved networks, we found that structurally and dynamically very different network components are able to fulfill the same task. We also found that networks with

shorter attractors are more robust against mutations, and that the initial evolutionary phase may bring the networks into different parts of network space, to which they remain confined for very long times.

Although the toy networks employed in our simulations are much simpler than real gene regulation networks, they may teach some lessons that can be relevant for real evolutionary processes:

First, selection is based on dynamical behavior, i.e., on the "phenotype" of the networks, while mutations operate on the network structure, i.e., the "genotype." The fitness landscape that is obtained in this way is very different from the rugged landscapes with many isolated fitness peaks which are often envisaged when modeling evolution by assigning a fitness value directly to a genotype [24]. In the high-dimensional genotype space given by network models, there are many neutral directions and percolating clusters of high fitness. This finding is in agreement with more general arguments proposed by Gavrilets [25]. Inspired by the view that biological evolution does not take place on rugged fitness landscapes, he investigated "holey adaptive landscapes." They consist of nearly neutral networks of well-fit genotypes, with all other genotypes being unviable. In [25], he concluded from some simple models for genotype to fitness mapping that the existence of percolating nearly neutral networks appears to be a general property of fitness landscapes with a large number of dimensions, as is the case for real genotype space.

Second, if a task can be solved by evolution, it may well be that it cannot only be solved in one way but in many different ways. This behavior of our toy model is in agreement with the phenomenon of convergent evolution [26] observed for the adaptation to niches and the evolution of the various sensory organs. In all these cases, the same challenge was met independently several times by finding solutions that agree in their general features, but that are very different at the genetic level and in the phenotypic details.

Third, solutions of high fitness can be found without going through fitness valleys and without attempting to make "hopeful monsters," i.e., to perform many mutations at the same time that carry the genotype to a distant point in the fitness landscape. If this was not the case, the evolutionary process would not have been possible. Neutral mutations are an important means to navigate through genotype space and to explore new possibilities. Studies of RNA secondary structures show not only that there are many neutral paths that percolate through the entire sequence space but also that they are interwoven. In a small neighborhood of every random sequence, one can find every statistically relevant secondary structure [27,28]. It was shown recently that biological RNA structures have a large number of sequences that fold into these relevant structures, much larger than for random phenotypes [29]. Additionally, the authors of [30] were able to design neutral paths for two very different ribozymes that ended in structures that were only four (not neutral) point mutations apart. That is, four point mutations were enough for a change between one functionality and the other. The ribozymes in between were able to catalyze both reactions but with a lower activity. In [31] it was shown that populations evolving on networks with more neutral neighbors per node are more widely extended in genotype space and can thus access a larger number of different phenotypes by just one or two mutations.

To conclude, the fitness landscape derived from evolutionary simulations based on network models appears to reflect many features of biological evolution. Exploring more complex network models and more complex selection criteria will be a next step at deepening our understanding of biological evolution.

---

[1] J. Klüver and J. Schmidt, Complexity **5**, 45 (1999).

[2] S. Bornholdt, Science **310**, 449 (2005).

[3] K. E. Kürten, Phys. Lett. A **129**, 157 (1988).

[4] R. Albert and H. G. Othmer, J. Theor. Biol. **223**, 1 (2003).

[5] F. Li, T. Long, Y. Lu, Q. Ouyang, and C. Tang, Proc. Natl. Acad. Sci. U.S.A. **101**, 4781 (2004).

[6] B. Luque and R. V. Solè, Phys. Rev. E **55**, 257 (1997).

[7] I. Shmulevich and S. A. Kauffman, Phys. Rev. Lett. **93**, 048701 (2004).

[8] A. Szejka and B. Drossel, Eur. Phys. J. B **56**, 373 (2007).

[9] T. Mihaljev and B. Drossel, Eur. Phys. J. B **67**, 259 (2009).

[10] S. A. Kauffman and R. G. Smith, Physica D **22**, 68 (1986).

[11] S. Bornholdt and K. Sneppen, Phys. Rev. Lett. **81**, 236 (1998).

[12] S. Bornholdt and K. Sneppen, Proc. R. Soc. London, Ser. B **267**, 2281 (2000).

[13] M. Paczuski, K. E. Bassler, and A. Corral, Phys. Rev. Lett. **84**, 3185 (2000).

[14] K. E. Bassler, C. Lee, and Y. Lee, Phys. Rev. Lett. **93**, 038101 (2004).

[15] S. Braunewell and S. Bornholdt, Phys. Rev. E **77**, 060902(R) (2008).

[16] S. E. Harris, B. K. Sawhill, A. Wuensche, and S. Kauffman, Complexity **7**, 23 (2002).

[17] A. Szejka, T. Mihaljev, and B. Drossel, New J. Phys. **10**, 063009 (2008).

[18] A. A. Moreira and Luis A. Nunes Amaral, Phys. Rev. Lett. **94**, 218702 (2005).

[19] L. Ciandrini, C. Maffi, A. Motta, B. Bassetti, and M. Cosentino Lagomarsino, Phys. Rev. E **80**, 026122 (2009).

[20] S. Ciliberti, O. C. Martin, and A. Wagner, PLOS Comput. Biol. **3**, e15 (2007).

[21] K. Kaneko, PLoS ONE **2**, e434 (2007).

[22] V. Sevim and P. A. Rikvold, J. Theor. Biol. **253**, 323 (2008).

[23] G. J. Szöllősi and I. Drényi, Mol. Biol. Evol. **26**, 867 (2009).

[24] S. A. Kauffman and S. Levin, J. Theor. Biol. **128**, 11 (1987).

[25] S. Gavrilets, in *Towards a Comprehensive Dynamics of Evolution—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, edited by J. Crutchfield and P. Schuster (Oxford University Press, New York, 2003).

[26] S. Conway Morris, *Life's Solution: Inevitable Humans in a*

Lonely Universe (Cambridge University Press, Cambridge, England, 2003).

[27] P. Schuster, W. Fontana, P. F. Stadler, and I. L. Hofacker, Proc. R. Soc. London, Ser. B **255**, 279 (1994).

[28] W. Fontana, BioEssays **24**, 1164 (2002).

[29] T. Jörg, O. C. Martin, and A. Wagner, BMC Bioinf. **9**, 464 (2008).

[30] E. A. Schultes, Science **289**, 448 (2000).

[31] A. Wagner, Proc. R. Soc. London, Ser. B **275**, 91 (2008).