

Structure and Accessibility of Fitness Landscapes

Inauguraldissertation

zur

Erlangung des Doktorgrades der
Mathematisch-Naturwissenschaftlichen Fakultät
der Universität zu Köln

vorgelegt von

Benjamin Schmiegelt

aus

Bergisch Gladbach

Berichterstatter:
(Gutachter)

Prof. Dr. Joachim Krug
Prof. Dr. Anton Bovier

Tag der mündlichen Prüfung: 23.6.2022

Abstract

In this thesis I take a look at stochastic models for fitness landscapes, specifically the House-of-Cards (HoC) model^[26] of completely uncorrelated fitness values and the NK-type models which are built through combination of HoC landscapes as building blocks.^[24] These models are parameterized by a parameter k , which is considered to determine the “ruggedness” of the fitness landscape, with a maximal value of k corresponding to the HoC model and a value of $k = 1$ corresponding to non-epistatic landscapes. I consider the behavior of two properties related to ruggedness on these landscapes, namely the number of local fitness maxima and the accessibility of genotypes via paths of monotonic fitness increase. Although high ruggedness is connected to a higher number of local maxima and therefore intuitively also a lower probability of distant genotypes being accessible from one another, this turns out to not hold generally. Contrary to assumptions made when the NK model was first introduced,^[61] it can be shown that asymptotic different quantitative results for the number of local maxima can be found for different choices of interaction structures between loci of the genotype.

These models have analogous interpretations in solid state physics as the random energy model^[12] and spin glass models.^[54,61]

Kurzzusammenfassung

In dieser Arbeit betrachte ich stochastische Modelle für Fitnesslandschaften, im Speziellen das House-of-Cards (HoC) Modell^[26] vollständig unkorrelierter Fitnesswerte und NK-artige Modelle welche durch Kombination von HoC Landschaften als Baustein gebildet sind.^[24] Diese Modelle sind von einem Parameter k parametrisiert, welcher die “Rauheit” der Landschaft festlegt, mit einem Maximalwert von k im HoC-Modell und einem Wert $k = 1$ auf nicht epistatischen Landschaften. Ich betrachte das Verhalten von zwei Eigenschaften welche im Zusammenhang mit Rauheit auf diesen Landschaften stehen, namentlich die Anzahl der lokalen Fitnessmaxima und die Zugänglichkeit von Genotypen via Pfaden von monoton steigender Fitness. Obwohl hohe Rauheit mit einer erhöhten Anzahl lokaler Maxima und damit intuitiv auch einer niedrigeren Wahrscheinlichkeit, dass entfernte Genotypen zugänglich sind, zusammenhängt, stellt sich dies im Allgemeinen als nicht wahr heraus. Entgegen Annahmen welche bei Einführung des NK-Modells gemacht wurden,^[61] kann gezeigt werden, dass asymptotisch quantitative unterschiedliche Ergebnisse für die Anzahl der lokalen Maxima für verschiedene Wahlen des Interaktionsnetzwerks zwischen Loci gefunden werden kann.

Diese Modelle haben analoge Interpretationen in der Festkörperphysik als das “random energy model”^[12] und Spin-Glass-Modelle.^[61]

Contents

1	Introduction	5
1.1	Biological context	5
1.2	Other contexts	6
2	Mathematical setting	8
2.1	Genotype spaces	8
2.1.1	Genotypes	8
2.1.2	Mutation graphs	8
2.1.3	Loci and allele graphs	9
2.2	Fitness landscapes	11
2.2.1	Fitness graphs	11
2.2.2	Epistasis	12
2.3	Building block models	14
2.3.1	House-of-Cards (HoC) model	14
2.3.2	Additive model	14
2.4	NK models	14
3	Number of local maxima	18
3.1	Definitions and general bounds	18
3.1.1	Minimal number of local maxima	18
3.1.2	Maximal number of local maxima	19
3.2	Number of local maxima in NK models	20
3.2.1	General behavior	20
3.2.2	Known results for fixed k	21
3.2.3	Known results for large k	22
3.2.4	General bounds	23
3.2.5	Regular- and uniform structures	25
3.2.6	New results for specific NK structures	27

4	Accessibility	30
4.1	Definitions	30
4.2	Accessibility in the HoC model	31
4.3	Accessibility in the NK model	34
4.3.1	Accessibility in the SN structure	35
4.3.2	Accessibility in the BN structure	36
4.3.3	Accessibility in other regular and uniform structures	37
5	Publication 1: Universality classes of interaction structures for NK landscapes.	38
6	Publication 2: Accessibility percolation on cartesian power graphs.	39
7	Discussion	75
A	Proofs and other calculations	77
A.1	Expected number of local optima for normal-distributed uniform and regular NK landscapes	77
A.2	Mean number of local optima in regular and uniform NK models with sum-dominating distribution	80
B	Bibliography	85
C	Teilpublikationen	90
D	Erklärung nach Promotionsordnung	91

Chapter 1

Introduction

1.1 Biological context

The concept of a fitness landscape originates with Sewall Wright, as an assignment of fitness values to all possible combinations of possible alleles in a given evolutionary environment, coupled with a notion of distance in accordance with the number of mutations separating these combinations. He observed that due to the exponential number of combinations, in spite of the high dimensionality, such landscapes ought to have many local fitness maxima, allowing for a view of evolution in which populations cluster around and move from peak to peak.^[65]

Underlying the static view of fitness landscapes is the assumption that a single fitness value determines the adaptiveness or expected success in producing offspring of an individual in the evolutionary dynamics. Other effects, such as phenotypic plasticity,^[52] time-dependent effects of environment changes,^[39] intra-population dependence^[60] and development plasticity^[58] are neglected in the static picture, but extensions to dynamical fitness landscapes are used to incorporate such effects.^[39,64] Popular mathematical models used to simulate dynamics of populations such as the Wright-Fisher^[18,66,67] and Moran^[38] models also take this point-of-view.

The fitness value assigned to an offspring under these conditions is expected to be similar to that of the parent. This similarity of fitness is explained by inheritance of fitness-defining characteristics from the parent. The primary biological mechanism for this inheritance is the passing on of genomic material, in particular the parent's DNA, although other important mechanisms, collectively referred to as *epigenetics* have been identified.^[6,9,30] This genetic material encodes information, used as blueprint for the construction and implementation of biological processes in an organism, in the form of a sequence of discrete values identified by different nucleolus.^[31] The duplication and passing on of genetic material is not perfect and free of error. Mutations can occur via multiple mechanisms, result-

ing for example in substitutions of single nucleotides or deletions and duplications of subsequences of nucleotides.^[31] These mutations provide for the possibility of change of fitness in the dynamical picture described above and combined with natural selection on the fitness landscape form the basis for the modern synthesis of evolution theory.^[22]

Therefore, because the genome is a sequence of individually small information units, the space of genotypes may be considered as a high-dimensional product space of small discrete spaces. While more general forms of mutations could be considered as well, this model effectively considers only point mutations on individual nucleotides or equivalently point mutations on higher level abstractions of the DNA, such as amino acids encoded by codons or alleles of genes.

The high dimensionality of this product space makes it difficult to visualize its geometry. Wright originally provided a picture of fitness landscapes as continuous maps from low dimensional spaces to real values, allowing a visual representation of the meaning of peaks on the landscape and the behavior of populations as clouds on the landscape, although he noted that this reduction of dimensionality is not representative of the actual high-dimensional product space.^[65]

Instead of this intuition, we consider the fitness landscape and its properties rigorously as a function on a high-dimensional discrete space. Populations of individuals can be viewed as a cloud of particles on the fitness landscape. Reproduction and natural selection result in these particles moving along on the fitness landscape and in particular if natural selection is sufficiently strong and the rate of mutation rate small, it is possible to approximate the whole population by its majority genotype, resulting in dynamics described by an adaptive walk of a single particle on the fitness landscape.

In this picture structural properties of the fitness landscape become relevant. Because mutants with lower fitness are less likely to reproduce, movement of the population on the fitness landscape is biased towards increasing fitness. A peak on the landscape implies a hindrance for the population to obtain higher fitness values and it would be expected that a fitness landscape with fewer peaks gives the population more available escape routes to higher fitness, resulting in faster adaptation. Alternatively, the existence of such escape routes, or in general accessible paths meaning monotonically increasing paths, on the landscape can be considered directly. In this thesis I will consider both of these properties.

1.2 Other contexts

In the context of solid state physics and other areas of physics, the notion of fitness landscapes may be recast into a notion of energy landscapes. Instead of genotypes, in physics we often consider systems with states composed of many discrete units.

As a simple example these units could be spin variables in a solid, each of which can be found in one of two states. This again forms a high-dimensional product space as the overall system state, with each system state having an assigned energy, resulting in the equivalent notion of energy landscape as assignment of energy values to system states. In this setting time evolution is biased towards minimization of the energy and in particular at low temperatures the structure of the energy landscape becomes relevant. For example local minima on the energy landscape correspond to meta-stable states, which may prevent the system from reaching its ground state or delaying its relaxation.

In the context of optimization in computer science, a common problem is to optimize a function which takes a binary string as input. If not much about the function is known, one may employ generic heuristic search algorithms to find as good an optimum of the function as possible in reasonable time. Such algorithms depend on functions in practice having the property that small changes to the input do tend to result in small changes in the output overall, equivalently to the biological picture of mutations and fitness effects. These functions can then be viewed as fitness landscapes with the input sequence as genotype and the output as fitness. Generic search algorithms are then often expressible as movement on this fitness landscape with a tendency towards higher fitness values. This analogy is made proper in particular for the so-called genetic search algorithms, which explicitly use the analogy to biological evolution by using models of evolutionary dynamics on a population of searching agents under natural selection with the given fitness landscape.

Chapter 2

Mathematical setting

2.1 Genotype spaces

2.1.1 Genotypes

The evolutionary dynamics of a population on the fitness landscape are modeled as a stochastic process in time. The state of the population is given by the combined state of all individuals in the population and in the picture of a static fitness landscape the state of individuals is characterized by a single discrete type, referred to as *genotype*. Although the term indicates a relation to inheritable genetic information, it can also subsume non-genetic inheritable discrete properties of the individual. Ultimately fitness is determined by *phenotypes*, which are often continuous in nature, but under the constraints of a static fitness landscape, phenotypes are a function of the discrete genotypes, so that they may be subsumed into the characterization of genotypes.

Generally genotypes in a given evolutionary system therefore form some finite set, or more generally some countable set, which I will denote \mathbb{G} . A more general approach to a fitness landscape framework can be found e.g. in.^[53]

In the context of physics, the stochastic process in time describes the evolution of a discrete set of system states, for example the combined states of a number of spin objects. These system states corresponds to genotypes in the biological setting.

2.1.2 Mutation graphs

Modes of reproduction can largely be separated by whether they involve recombination of genotypes of multiple parents or whether only the genotype of a single parent is inherited. For the purpose of the description in this thesis I focus on the non-recombining case. In the process of reproduction, the parent passes on its

genotype to its offspring, but the possibility of mutations allows for changes in the genotype of the offspring.

It is however expected that there is a notion of similarity between genotypes, so that genotypes of the offspring are relatively similar to the genotype(s) of their parent(s). Although the results of mutation are probabilistic in nature, one approximate approach is to separate parent-offspring genotype transitions by allowed and forbidden transitions, corresponding to mutations that are reasonably possible and those that are practically not likely to occur. For example, single-nucleotide substitutions may happen on multiple locations of the genome, but under the reasonable assumption of them occurring independently of one another, the probability that many such substitutions happen in a single reproduction cycle are low, given that individual mutation rates are not large. In this picture, transitions involving small number of nucleotide substitutions are allowed, while those involving many are forbidden.

Based on this idea one can introduce a *mutation graph*, which is a directed graph on the genotype set, such that arrows indicate allowed parent-to-child transitions via mutations in a single reproductive cycle. With this graph all non-recombining reproduction dynamics can be described as long as the likelihood of different mutational transitions occurring can be neglected. If these likelihoods are relevant, the graph's arrows maybe equipped with a weight indicating them. However, I will not consider this generalization in this thesis. The properties under consideration here are static-structural in nature and not dependent on the specific mutation probabilities, which are however relevant in the behavior of specific dynamics on the fitness landscapes.

2.1.3 Loci and allele graphs

Typically the mutation graph exhibits certain structural properties due do the encoding scheme of genetic information. In particular, because the genome encodes data as a sequence of small information units, each of which as a first-order approximation is able to mutate stochastically independently, it is useful to consider the mutation graph as a product graph over smaller graphs describing mutations on individual elements of the genomic sequence.

Under the assumption that mutation is weak enough that each parent-child transition can mutate the genotype only at one element of the encoding sequence at most, the resulting mutation graph can be written as a Cartesian graph product of factor mutation graphs over the states of individual sequence elements.

The Cartesian product of two mutation graphs \mathcal{G} and \mathcal{H} is the mutation graph $\mathcal{G} \square \mathcal{H}$ over the Cartesian product set of the corresponding genotype sets with an arrow from (g, h) to (g', h') if and only if (g, g') is an arrow in \mathcal{G} or, exclusive, (h, h') is an arrow in \mathcal{H} .

I will refer to the individual factor graphs forming the overall mutation graph as *allele graphs* and its vertices as *alleles*. The allele graphs are usually naturally indexed, in which case the index of an allele graph in the mutation graph's factorization is referred to as *locus* or *site* of the genotype. Although the factorization was motivated by the encoding of genetic information in a sequence of nucleobases, alternative interpretations on a higher level are also possible, one of which motivates the terminology chosen, as described below.

If the mutation graph is factorized into individual base pairs on the genome, then the allele graph at each locus is best described as the complete graph on four alleles \mathcal{K}_4 . The four alleles account for the four possible nucleotide configurations for each base pair and mutations are possible from any configuration to any other. With L loci of this kind, the resulting mutation graph is the Hamming graph $H(L, 4) = \mathcal{K}_4^{\square L}$.

In protein-coding regions of the genome each triple of subsequent base pairs encodes an amino acid. These triples are referred to as *codons* and the mapping of triples to amino acids as *genetic code*. The factorization of the mutation graph may also be viewed as a factorization into codons, in which case the allele graph can be chosen as the Hamming graph $H(3, 4) = \mathcal{K}_4^{\square 3}$ to still allow only mutations resulting from single nucleotide substitutions. The complete mutation graph will then coincide with the mutation graph in the single-nucleotide factorization, except that the number of loci is reduced by a factor three.

If however codons are mapped via the genetic code to the amino acid they encode, a structurally different mutation graph emerges. The $4^3 = 64$ possible codon states encode only 22 amino acids plus stop conditions (which I will consider a fictional amino acid for simplicity). Therefore multiple codons are mapped to the same amino acids. This allows for nucleotide substitutions to leave the encoded amino acid unchanged, in which case the mutation is said to be *synonymous*. Other single nucleotide substitutions may change the encoded amino acid and an allele graph can be constructed by considering the 22+1 amino acids as alleles and transitions between amino acids as allowed if there exists a single nucleotide substitution resulting in the codon encoding changing from the first to the second amino acid. In this interpretation the allele graph has a more complex structure, but is still symmetric.

Moving up to an even higher-level abstraction, loci may be considered to be e.g. individual protein-coding genes and the alleles to be sets of mutations of interest on the individual genes. In this interpretation there are more possible choices of allele graphs. For example, mutations are often strongly deleterious producing non-viable offspring, in which case it is useful to exclude these mutations directly from the allele and mutation graphs. As a result an allele graph may for example be a subgraph of a Hamming graph, one example being a symmetric path graph

describing a number of mutations which are required to occur in sequence to be viable.

To avoid some uninteresting edge cases I will assume that the allele graphs have at least two alleles and that the allele graph is at least weakly connected. Furthermore I will assume that the allele graph is the same on all loci. This is done to avoid overly complex formulas, but most results are extendable to the more general case.

The simplest non-trivial allele graphs to consider are then the allele graphs on two loci, of which there are two up to isomorphism. The first one is \mathcal{K}_2 , the complete graph on two vertices, on which it is possible to mutate from the wild-type allele to the mutant allele, as well as back again. The second one is asymmetric and allows only for mutations from the wild-type but not back again. I refer to this as mutation graph without back-mutation and write it \mathcal{K}'_2 . A similar extension of the allele graph without back-mutation for larger number of alleles is the complete graph on A alleles with all arrows to the wild-type allele removed, which I write as \mathcal{K}'_A .

In this thesis, if not stated otherwise in context, I assume that the allele graph is \mathcal{K}_2 .

2.2 Fitness landscapes

In the picture of static fitness landscapes each genotype g is assigned a real value F_g representing its fitness. The full mapping F of genotypes to fitness values is the *fitness landscape*, which determines the natural selection aspect of evolutionary dynamics in this setting.

Instead of analysing individual empirically obtained fitness landscapes, the goal of probabilistic modelling of fitness landscapes is to find a probability distribution or ensemble over the space of all fitness landscapes with a given underlying mutation graph, such that empirical fitness landscapes are typical for this ensemble. This allows calculation of typical values for quantities of interest in the model ensemble, which are then expected to be similar to these quantities observed in empirical landscapes.

2.2.1 Fitness graphs

The properties which I will consider in this thesis are of a specific form, allowing us to ignore detailed quantitative interpretations of fitness. Both the number of local optima, as well as the accessibility properties, the two properties I will discuss, depend only on the ordering of fitness values relative to one another.

Instead of analyzing fitness landscapes themselves it is therefore helpful to consider a modified fitness landscape, which I will call the *ranked fitness landscape*. Instead of the original fitness values it assigns each genotype its numeric rank in the ordering of genotypes by their fitness values. In order to avoid specification of the ranking behavior under fitness value ties, I will assume that fitness values are never exactly equal. This assumption is trivially satisfied in the fitness landscape models which I will use in this thesis almost surely, in the probabilistic sense.

The ranked fitness landscape reduces the space of landscapes to consider from a continuous set to a discrete set, finite even, assuming the number of genotypes is finite as well.

However, even further reduction of information is possible when talking about the properties of interest in this thesis. All of them depend only on *local* orderings of fitness values, meaning they depend only on the ordering of fitness value pairs of genotypes adjacent on the mutation graph. The ordering of pairs of adjacent genotypes can be encoded as an directed graph obtained from the mutation graph by removal of arrows which do not point from a genotype of lower fitness to a genotype of higher fitness. The resulting directed graph is known as *fitness graph* and it encodes all information about the landscape relevant for processes that are restricted to local movement on the mutation graph and are indifferent to magnitudes of fitness differences.^[8] In terms of evolutionary dynamics, this is just the information relevant in the so-called *strong selection weak mutation regime (SSWM)*.^[20,42] Under the SSWM regime mutations are rare enough that the whole population can be modeled as a single particle moving on the fitness landscape with the majority genotype as its genotype, resulting in random walk dynamics. Strong selection further implies that mutations reducing fitness are unlikely to fixate and therefore evolutionary dynamics can be reduced to an adaptive walk along arrows in the fitness graph.

2.2.2 Epistasis

One important notion for the description of fitness landscapes is that of *epistasis*, describing the interaction between mutational effect on different loci. Many definitions for the exact meaning of the notion of epistasis exist,^[10,44] however I will give a common definition of pairwise (in contrast to higher-order^[47,57,62]) epistasis in the context of the study of fitness landscapes in the following.

At a given initial genotype g , I use $\Delta_l^{a \rightarrow b} g$ to denote the genotype obtained from g by replacing the allele at locus l from a to b . Typically (a, b) is assumed to be an arrow in the allele graph and so $(g, \Delta_l^{a \rightarrow b} g)$ is an arrow in the mutation graph. Point mutations are therefore identified by the operators $\Delta_l^{a \rightarrow b}$ and can be applied to any genotype currently containing allele a at locus l . On the binary allele graph I will omit the upper index $a \rightarrow b$ given that at every given genotype

there is only one valid choice.

Commonly I will also use the short-hand notation

$$\Delta_l^{a \rightarrow b} F(g) = F(\Delta_l^{a \rightarrow b} g) - F(g) \quad (2.1)$$

describing the fitness effect of a mutation $\Delta_l^{a \rightarrow b}$ at a background genotype g .

A mutation $\Delta_l^{a \rightarrow b}$ is then said to be *epistatically dependent* on another mutation $\Delta_{l'}^{a' \rightarrow b'}$ with $l \neq l'$ at background g if

$$\Delta_l^{a \rightarrow b} F(\Delta_{l'}^{a' \rightarrow b'} g) \neq \Delta_l^{a \rightarrow b} F(g) \quad (2.2)$$

The difference between the two sides of this equality is a common measure for the magnitude of epistasis and also differentiates between *positive* and *negative* epistasis, which results in the approach to fitness landscape analysis by *shapes*.^[3]

For the purposes of this thesis magnitudes of fitness differences are not a focus. Instead I am interested in the signs of fitness differences. This leads to the notion of *sign-epistatic* dependence^[63] obtained by replacing the real-valued fitness effects with their signs:

$$\text{sgn} \left[\Delta_l^{a \rightarrow b} F(\Delta_{l'}^{a' \rightarrow b'} g) \right] \neq \text{sgn} \left[\Delta_l^{a \rightarrow b} F(g) \right] \quad (2.3)$$

In other words, a mutation is not sign epistatically dependent on another mutation at a given background if in the square formed by the four genotypes g , $g' = \Delta_l^{a \rightarrow b} g$, $h = \Delta_{l'}^{a' \rightarrow b'} g$ and $h' = \Delta_l^{a \rightarrow b} \Delta_{l'}^{a' \rightarrow b'} g$ and the arrows between them in the mutation graph, either both (g, g') and (h, h') remain (parallel) arrows when transitioning to the fitness graph or both (g', g) and (h', h) remain (parallel) arrows.

The property of sign epistasis is not symmetric, meaning that a mutation X may be sign-epistatically dependent on a mutation Y , but Y not sign epistatically dependent on X . However, if both of these properties hold at the same time at a given background, the two mutations are said to be *reciprocal sign-epistatic* at the given background,^[45] which is known to be a necessary condition for the existence of multiple fitness peaks on the fitness landscape^[46] and demonstrates mutations that are strongly mutually antagonistic.

Generally all of the concepts introduced so far are dependent on the specific background at which epistasis is investigated. Generally it is not to be expected that there are strong correlations between these properties when compared across large distances on the fitness landscape. However, if reciprocal sign-epistasis between two mutations extends to all background genotypes, I speak of *global reciprocal sign epistasis* (GRSE) between the two mutations. As it turns out the specific structure of NK fitness landscapes often can produce GRSE with drastic consequences for the structure of the landscape.^[23,48]

2.3 Building block models

2.3.1 House-of-Cards (HoC) model

In search for a suitable model ensemble of fitness landscapes, one possible approach is to start with a toy model with low complexity in terms of the mathematical description, which nonetheless has interesting non-trivial behavior. A straightforward approach to this which applies to any mutation graph, is the House-of-Cards model.^[24,26,34] In this model a probability distribution over the reals is fixed and the fitness value of each genotype is drawn from it independently.

Not only has this model a low complexity description, but it also has the nice property that, with the exception of non-continuous distributions, the particular choice of the probability distribution does not influence the distribution of the fitness graph, because only the order of the individual fitness values affects the ranked landscape and consequently the fitness graph.

The analogue of the House-of-Cards model in statistical physics is known as the *random energy model*.^[12]

2.3.2 Additive model

As an opposite to the House-of-Cards model, one can consider the so-called *additive model*, *linear model* or *non-epistatic model*.^[17]

As the name suggests this model describes fitness values that are purely additive components per locus, such that no epistatic interactions between loci emerge.

Specifically, each allele on each locus is assigned a fixed fitness contribution. The total fitness of a genotype is then the sum of these contributions for which it contains the corresponding allele at its locus.

Under this construction, the fitness effect of a mutation on one locus and the fitness effect on another locus can always be summed to yield the combined 2-mutant fitness effect, implying that the landscape is completely non-epistatic.

As shown in the next section, the additive model can be viewed as a special case of the NK model and at the same time the NK model can be viewed as an extension of the additive model combined with the HoC model.

2.4 NK models

While the HoC model and additive model are simple and interesting as a toy model, it is unlikely to be a good explanation of biological fitness landscapes. In particular the HoC model ignores the underlying mutational structure of the genotype space completely. Selection coefficients do not correlate with distance of genotypes in

the HoC model. However, in actual organisms, we expect that small modifications of genotypes tend to result in small effects on phenotype, expected survivability and expected reproduction, so that fitness values should be correlated, at least for small distances on the mutation graph.

Furthermore, at least locally, individual loci tend to have somewhat consistent effects on fitness in organisms, often with epistatic interaction between limited sets of loci.

One approach to incorporating these notions is to start from a number N of individual building block landscapes defined on subsets of loci \mathbb{B}_i , each of which produces for every projection of a genotype to the corresponding block a fitness contribution f_i , which can then be combined to an overall fitness by some function. These individual building blocks may also for example be considered individual phenotypes.

This construction causes correlation between fitness values of close genotypes by virtue of mutations on few loci modifying only a small subset of the building block fitness values. It also imposes a structure on the set of loci, which can be described as a hypergraph through the sets \mathbb{B}_i as hyperedges. We refer to this structure as the *neighborhood structure* and the building block sets \mathbb{B}_i as *neighborhoods*.

While it may also be of interest to consider other case, the combining function is typically chosen to be addition, compatible with the additive choice of definition for pairwise epistasis, and the individual building block landscapes are chosen to be HoC landscapes with some fitness distribution. With this choice the resulting model is known as the *NK model*. Originally this model was introduced by Kauffman^[25] in a restricted form, specifically he required the number of building blocks to be identified one-to-one with the loci, in such a way that the locus corresponding to a neighborhood is contained in it, and that each neighborhood ought to share the same size k . I refer to such neighborhood structures as *classical*.

In the description of neighborhood structures I make use of usual terminology from the theory of hypergraphs. In particular a k -uniform neighborhood structure is one in which all neighborhoods have size k and a r -regular neighborhood structure is one in which all loci are contained in exactly r neighborhoods. Even if a structure is not k -uniform or r -regular, I will refer to the average size of neighborhoods as k and the average number of occurrences of loci in neighborhoods as r . As a consequence, by an analogue of the handshake lemma, the general equality $Nk = Lr$ holds. This allows reduction of the base parameters describing neighborhood structures to r and k , both of which are intrinsic quantities with respect to the number of loci. For the classical neighborhood structures $r = k$, but smaller as well as larger r may be considered. In order to avoid neutral mutations I require that each locus appears at least once in a neighborhood, resulting in a value $r \geq 1$.

A few common choices for neighborhood structures are:

- Block structure (BN):^[41,43] The set of loci is divided into disjoint sets of equal size k , each of which is chosen as a neighborhood. The structure is regular and uniform. Choosing each such set k times results in a classical NK structure.
- Generalized block structure:^[50] The requirement of equal size may be lifted, in which case uniformity is lost.
- Adjacent structure (AN):^[25] Loci are given a ring structure and for each locus its $k - 1$ nearest neighbors on the ring are chosen as part of its (classical) neighborhoods. The structure is regular, uniform and classical.
- (Classical) Random structure (RN):^[25] Any classical neighborhood structure with specified k on the locus set is chosen with uniform probability. Here we also consider the special cases of the rRN structure which restricts the choice to regular structures. The RN structure, together with the AN structure, are the specific choices first considered by Kauffman.
- Pure random structure (PR):^[61] For specified values of k and r any, not necessarily classical, neighborhood structure is chosen with uniform probability. Here additional requirements of uniformity and/or regularity may be imposed, resulting in the uPR, rPR and urPR models.
- Mean field (MF):^[23] All k -subsets are chosen as neighborhoods.
- Star structure (SN):^[23,48] In this classical structure for each neighborhood associated to a locus the remaining $k - 1$ loci are chosen as the first $k - 1$ remaining loci by their index.

The RN, PR and uPR structures are not regular, but many statements about NK models with regular structures can be extended to them, since they are quasi-regular in the sense that the distribution of degrees in the hypergraph converges in all moments as L grows to infinity at constant k and r .

In contrast to this, the SN structure is specifically chosen, such that all moments of the degree distribution quickly diverge to infinity, exemplifying a strongly non-regular neighborhood structure.

Aside from the structures mentioned above, some models for spin-glasses also can be directly interpreted as (generalized) NK models. Specifically, in either case without external fields, the Sherrington-Kirkpatrick model^[51] corresponds to an NK model in which all two-element sets are chosen as neighborhoods, resulting in the MF structure with $k = 2$. The Edwards-Anderson model^[14] is obtained

from a d -dimensional lattice on loci, choosing all locus pairs corresponding to edges on the lattice as neighborhoods. Assuming circular boundary conditions on the lattice, this results in a uniform and regular NK structure with $k = 2$ and r the coordination number of the lattice. Specifically, the one-dimensional case is identical to the AN structure with $k = 2$ and therefore classical.

Spin-glass models including external fields may also be represented to some degree by NK models as defined here, assuming Gaussian distributions for interaction terms. External fields in this context correspond to additional neighborhoods, each containing only a single locus. However, this always results in external fields of equal energy as single pair interaction contributions. To achieve different relative strength of the external field, it is possible to add multiplicities to the neighborhoods. Addition of external fields in this manner make the structure non-uniform, but doesn't affect regularity if the field is applied uniformly across sites.

Chapter 3

Number of local maxima

3.1 Definitions and general bounds

A local maximum of a fitness landscape is a genotype g such that $\Delta_l^{a \rightarrow b} F(g) < 0$ for all possible point mutations $\Delta_l^{a \rightarrow b}$ on g .

Note that I consider the orientation of arrows in the mutation graph as relevant in this definition since $\Delta_l^{a \rightarrow b}$ is a point mutation only if there is an arrow from a to b in the allele graph. Although a genotype may be a local fitness maximum by this definition, there may be genotypes h such that there exists a point mutation $\Delta_l^{a \rightarrow b}$ such that $g = \Delta_l^{a \rightarrow b} h$ and $\Delta_l^{a \rightarrow b} F(h) < 0$.

Equivalently a local fitness maximum is a *leaf* in the fitness graph, i.e. a genotype without any outgoing arrows (which point to higher fitness).

Under adaptive walk dynamics a population is unable to escape such a local maximum once it is reached and therefore the number of local maxima is considered a measure of “ruggedness” of the landscape and is therefore often estimated for fitness landscapes.^[11,19,55]

In this chapter I will denote the number of local maxima of the fitness landscape under consideration as N_{\max} and provide a short review of known results for the number of local maxima in HoC and NK models, as well as new results. But first, I will provide some general bounds on fitness landscapes on different allele graphs.

3.1.1 Minimal number of local maxima

In general, a fitness landscape must have at least one local maximum, its global maximum.

With the complete graph as allele graph it is easy to see that it is always possible to produce a fitness landscape with exactly one local maximum by choice of an additive fitness landscape.

Such a fitness graph can however not be produced by all allele graphs. Consider $L = 1$. Then the allele graph and the mutation graph coincide. It may or may not be possible to construct a fitness landscape with only one leaf in the fitness graph. For example the allele graph over three alleles a , b and c with arrows (a, b) and (a, c) already has two leafs and will always have at least two leafs on the fitness graph, since the fitness graph is obtained by removal of arrows from the mutation graph. Under such conditions the minimum number of local maxima may be exponentially growing in L , e.g. with the previous example all combinations of genotypes formed from only the alleles b and c must be local maxima, giving a lower bound of 2^L .

However, if the allele graph is symmetric and connected, then the allele graph can be oriented into an acyclic directed graph with exactly one leaf. This orientation can be interpreted as a fitness graph for $L = 1$ and its L -times product will produce a fitness graph for the corresponding mutation graph with L alleles.

3.1.2 Maximal number of local maxima

A more difficult problem in general is the maximum number of local maxima that a fitness landscape can have, even in the symmetric allele graph case. Certainly there can never be more than A^L local maxima on an allele graph with A alleles, however this is a very trivial bound.

For symmetric allele graphs another point-of-view to take on local maxima of a fitness landscape is that they form an *independent set* of genotypes. A set of genotypes (or generally vertices in an undirected graph) is called an *independent set* if there is no edge between any pair of genotypes in the set.

Any independent set can therefore be chosen to form a subset of the local maxima of the fitness landscape, by assigning high fitness values to the elements of the set and smaller fitness values to all other genotypes. Conversely any set of local maxima must form an independent set, as two local maxima cannot be adjacent on a symmetric mutation graph.

A *maximal independent set* is an independent set to which no further genotype may be added without the independent set property being lost. Equivalently an independent set is maximal if we can't find any remaining genotype which is not adjacent to one of the elements of the set. This coincides with a maximal choice of local fitness maxima on the landscape.

The *independence number* is the size of the largest maximal independent set. The problem of finding the maximal number of local maxima on a fitness landscape corresponding to a given symmetric allele graph therefore reduces to a search for the independence number of the mutation graph.

With allele graph \mathcal{K}_2 a landscape with the maximal number of local maxima is known as an egg-box landscape^[16] and in it local maxima and local minima

alternate starting from the wild-type, resulting in a fraction $\frac{1}{2}$ of all 2^L genotypes being local maxima. The same idea shows that on \mathcal{K}_A , a fraction $\frac{1}{A}$ of the A^L genotypes give the largest possible number of local optima on the fitness landscape. This can also be seen as a consequence of Vizing’s bound on the independence number of Cartesian products of graphs.^[24,29,34,59]

Unfortunately, outside of the complete graph, determining the independence number of the Cartesian product graph from the allele graph is not a simple problem and only few general bounds are known.^[29] In any case, removal of arrows from the complete allele graph will only result in removal of arrows from the genotype space and therefore can only increase the independence number. As a result, the maximum number of local maxima is always at least that of the complete graph on A alleles, meaning A^{L-1} as explained above.

Because the number of genotypes is however also limited by A^L for any allele graph, this then implies that the maximum number of local maxima on a given genotype space is always growing exponentially in L with rate $\ln A$ with at most a factor A variation between allele graphs of size A .

3.2 Number of local maxima in NK models

3.2.1 General behavior

The number local maxima on NK fitness landscapes have been investigated for different choices of neighborhood structures, fitness distributions and limiting behaviors of the ruggedness parameter k .^[13,24,33,43]

In general the expected number of local optima can be written as

$$\mathbb{E}[N_{\max}] = A^L e^{-\lambda L} Q(L) \tag{3.1}$$

with $0 \leq \ln A - \lambda \leq \ln A$ the exponential rate at which the expected number of local maxima increases and $Q(L)$ the correction to the exponential terms. Here A^L is the total number of genotypes and the remaining terms then correspond to the expected probability of a uniformly chosen genotype to be a local optimum. Assuming that the expectation doesn’t fluctuate in its leading order term, the correction $Q(L)$ can be chosen to be sub-exponential. Even if that is not the case, we may consider converging subsequences of the sequence of NK landscapes, giving upper and lower bounds on the exponential rate λ between which the model fluctuates.

The following known results typically consider only the complete graph on two alleles \mathcal{K}_2 as allele graph, which is implied if not stated otherwise. The results can further be categorized by the limiting behavior as $L \rightarrow \infty$. In the first variant k is

assumed a constant in the limit of $L \rightarrow \infty$ and in the second variant k is assumed to be linearly increasing with L as $k = \alpha L + \mathcal{O}(1)$ with $0 < \alpha \leq 1$ a constant.

3.2.2 Known results for fixed k

For some choices of neighborhood structure or fitness distribution it is possible to determine λ , as well as the sub-exponential corrections exactly in these limits.

The HoC model may be viewed as a NK model with only one neighborhood containing all loci. For it, the expected number of local maxima can be obtained exactly for all complete allele graphs by a simple combinatorial argument:^[24]

$$\mathbb{E}[N_{\max}] = A^L \frac{1}{1 + L(A - 1)} \quad (3.2)$$

implying $\lambda = 0$ and therefore the maximal possible exponential rate with which the number of local maxima can increase in L . However, the sub-exponential correction is $Q(L) = \frac{1}{1 + L(A - 1)}$, an order L^{-1} reduction from the maximum possible number of local maxima.

Also using combinatorial arguments, the variance of the number of local optima can be calculated to^[34]

$$\text{Var}[N_{\max}] = A^L \frac{(A - 1)(L - 1)}{2(1 + L(A - 1))^2} \quad (3.3)$$

Furthermore, the two events of a genotype g and a genotype h being local maxima is independent if the distance between g and h is at least 2 on the mutation graph and as a consequence a generalized form of the central limit theorem based on dependency graphs^[2] applies, implying that the number of local optima converges in law to a normal distribution with the given mean and variance as $L \rightarrow \infty$.^[34]

For the BN structure the NK model effectively factorizes into a set of independent HoC landscapes, so that the fitness distribution is, analogously to the HoC model, not affecting its fitness ordering either. The number of local maxima is the product of local maxima on each of these partial HoC landscapes with all factors independent and so with \mathcal{K}_A as allele graph one obtains from the HoC result:^[43,50]

$$\mathbb{E}[N_{\max}] = A^L \left(\frac{1}{k(A - 1) + 1} \right)^{\frac{L}{k}} \quad (3.4)$$

implying, at constant k , that $\lambda = \frac{\ln k(A - 1) + 1}{k}$ and $Q(L) = 0$. At linearly increasing k , it implies that $\lambda = 0$ and $Q(L) = (\alpha L)^{-\frac{1}{\alpha}} + o(1)$.

The structure of a product of independent factors also implies that the distribution of the number of local optima in the BN model converges at constant k in

law to a log-normal distribution as $L \rightarrow \infty$ and to a normal distribution if k scales linearly with L .

For other choices of structure and fitness distribution, for some small constant and non-trivial $k > 1$, specific values of λ have been found. These seem to always fall strictly between 0 and $\ln A = \ln 2$, so that the expected number of local maxima seems to always increase exponentially under these conditions, but with suboptimal rate.

- AN, $k = 2$, positive exponential distribution:^[15] $\lambda \approx 0.5627$
- AN, $k = 2$, negative exponential distribution:^[13]

$$\lambda = \ln \left(\frac{18}{5 + \sqrt{29}} \right) \approx 0.5500 \quad (3.5)$$

- AN; $k = 2$, Gamma distribution with shape parameter 2:^[15] $\lambda \approx 0.5717$
- AN, $k = 2$, Gamma distribution with shape parameter $\frac{1}{2}$:^[40]

$$\lambda = \ln \left(\sqrt{10 + 6\sqrt{3}} - \sqrt{3} - 1 \right) \approx 0.5787 \quad (3.6)$$

- AN, $k = 2$, uniform distribution:^[13] $\lambda \geq 0.4954$
- AN, $k = 3$, positive exponential distribution:^[15] $\lambda \approx 0.4920$
- SK/MF with $k = 2$, normal distribution:^[56] $\lambda_{SK} \approx 0.494$, the minimum of the function $f(x) = \frac{1}{2}x^2 - \ln \Phi(x)$.
- EA structure with normal distribution:^[7,56] $\lambda \approx \lambda_{SK} - \frac{0.0656}{r} + O(r^{-2})$

Analogously to the BN model, under weak conditions on the probability distribution in the AN model at constant k the expected number of local maxima is asymptotic to an exponential in L and the number of local maxima converges to a log-normal distribution.^[13,15]

3.2.3 Known results for large k

With linearly increasing k , known results which will be presented in this section indicate that $\lambda = 0$, implying a maximal exponential rate of increase in the number of local maxima. They also indicate that $Q(L)$ is polynomial in L , as already seen above for the BN structure:

$$Q(L) = e^{-\zeta(\alpha) \ln L + o(\ln L)} \quad (3.7)$$

Specifically at $\alpha = 1$ it is expected that $\zeta(1) = 1$ as the NK model with $k = L$ must coincide with the HoC model. As seen above, for the BN model $\zeta(\alpha) = \frac{1}{\alpha}$.

For both the AN and uRN structures with normal distributed fitness Weinberger^[61] gives $\zeta(\alpha) = \frac{1}{\alpha}$, coinciding with the BN structure. This seems to indicate a universal behavior, at least for classical NK structures. However, in his derivation he applies some mean approximations without rigorous justification. The result turns out to be correct for the AN structure,^[33] but the assumptions made by Weinberger apply also e.g. to the rRN model, for which more specific calculations yield different results, as will be discussed in the next section. This issue has also been noted in.^[15] It seems therefore unclear whether Weinberger's result for the uRN structure is asymptotically exact.

Limic and Pemantle^[33] show that in the AN structure $\zeta(\alpha)$ depends on the choice of fitness distribution, but in such a way that always $\frac{1}{\alpha} \leq \zeta(\alpha) \leq \frac{3}{\alpha}$. They also show that with a standard normal distribution $\zeta(\alpha) = \frac{1}{\alpha}$ and conjecture that this holds for all distributions.

If k is constant in the increase of L , but taken to infinity after L , it is expected from the above, that λ will coincide with the limit obtained by taking $\alpha \rightarrow 0$. In all of the cases explained above asymptotically in that limit $\zeta(\alpha) \sim \frac{\zeta'}{\alpha}$ for some constant ζ' and so the expected asymptotic behavior is

$$\lambda \sim \zeta' \frac{\ln k}{k} \quad (3.8)$$

which is indeed found in all the cases mentioned above.

- BN, any distribution:^[43,50] $\zeta' = 1$
- AN, normal distribution:^[33,61] $\zeta' = 1$
- AN, any distribution:^[13,33] $1 \leq \zeta' \leq 3$

Based on these results it seems that the expected number of local maxima in the NK model with regular structure is always of the forms eq. (3.7) and eq. (3.8) in the respective limits. In the following section I investigate this apparent general form, also for arbitrary (finite) allele graphs.

3.2.4 General bounds

I will start by considering some general bounds on the number of local maxima in the NK model.

In the House-of-Cards model fitness values of different genotypes are mutually independent and therefore the probability of a genotype being a local maximum is simply $\frac{1}{n+1}$ with n the number of out-degrees of the genotype. For r -regular

allele graphs the expected number of local maxima is therefore exactly $\frac{1}{Lr+1}A^L$. If the allele graph is not regular, $\frac{1}{n+1}$ will vary between genotypes, however as $L \rightarrow \infty$, except for an exponentially small fraction of genotypes and an arbitrarily small fraction of loci, all alleles will appear equally often on genotypes. Therefore asymptotically the expected number of local maxima is still $\frac{1}{Lr+1}A^L$ with r the average (out-)degree on the allele graph, which has to lie between $1 - \frac{1}{A}$ and $A - 1$.

The number of local optima is therefore always growing with the maximal exponential rate $\ln A$. This makes the House-of-Cards fitness landscapes very rough. This is not surprising since the House-of-Cards model does not take into account the graph structure at all. In terms of an optimization problem it is the worst possible case since no amount of knowledge about the fitness of a set of genotypes can help in further optimization.

However, the expected number of local maxima in the House-of-Cards model is still asymptotically negligible by a factor of order $\frac{1}{L}$ relative to the maximum possible number of local maxima, which always is up to a constant of order A^L as seen before.

In the NK model the choice of neighborhood enforces strict limitations on the fitness landscapes which can be generated, because mutations on loci which are independent in the neighborhood structure, that is which have no mutually shared NK neighborhoods, are non-epistatic. This can be seen directly from the fitness effect $\Delta_l^{a \rightarrow b} F(g)$ decomposing into a sum of fitness effects on the individual partial HoC landscapes associated with the neighborhoods. In particular this also restricts the possible rank orderings that can be generated by an NK model.

Suppose there is a subset \mathcal{M} of loci, such that all loci in the set are mutually independent with respect to the neighborhood structure. Then for any given state of all the loci outside \mathcal{M} , there can be only one local maximum:

$$N_{\max} \leq A^{L-|\mathcal{M}|} \quad (3.9)$$

If the NK neighborhoods are sufficiently small and few, then this results in an exponentially smaller number of local optima than in the HoC model. In particular, if all neighborhoods are at most of size k and the degree of loci in the NK structure is at most r , then we can always choose $\frac{L}{kr}$ loci into any maximal independent set and therefore

$$N_{\max} \leq A^{L(1-\frac{1}{kr})} \quad (3.10)$$

In general for classical NK models $r = k$ and this bound only give a weak bound of $\lambda \geq \frac{1}{k^2} \ln A$. However in many classical NK neighborhood structure, such as the block model and the AN model, the strong overlap of neighborhoods allows one to draw independent sets of order $\frac{L}{k}$, resulting in a bound that is asymptotically of the form $\lambda \geq \frac{1}{k} \ln A$.

While the NK neighborhood imposes a strict constraint on the maximum number of local maxima, it doesn't do so for the lower bound on the number of local maxima. It is always possible to construct for each NK structure a realization which has the same ranked fitness landscape as the additive model on the same landscape and therefore the minimal possible number of local maxima. For this consider the ranked fitness landscape induced by the additive model on a partial landscape corresponding to a single neighborhood. Because the partial landscape is of HoC type, it generates this particular ordering with a finite probability. If then all partial landscapes generate this ordering, the overall landscape has the same ranked landscape as the additive model and the probability of this realization occurring is non-zero.

A simple lower bound on the expected number of local maxima can however be obtained if the number of NK neighborhoods is sufficiently small. If a given genotype is a local maximum on the projections onto each partial landscape, then it is one with respect to the full landscape as well. Since the i -th partial landscape is of HoC type where this probability is $\frac{1}{(A-1)k_i+1}$, this implies the bound

$$\mathbb{E}[N_{\max}] \geq A^L \exp\left(-\sum_i \ln((A-1)k_i+1)\right) \quad (3.11)$$

where i sums over all NK neighborhoods and k_i is the size of the i -th neighborhood. By Jensen's inequality

$$\mathbb{E}[N_{\max}] \geq A^L \exp\left(-N \sum_i \ln((A-1)k+1)\right) \quad (3.12)$$

where $N = \frac{Lr}{k}$ is the number of NK neighborhoods, r is the mean degree of loci and k is the mean size of neighborhoods. For r and k constant with $r < \frac{k \ln A}{\ln((A-1)k+1)}$, this bound then shows that the expected number of local maxima must be growing exponentially in L . For classical NK models with $k = r$ or even more than L neighborhoods this bound is however not useful.

In fact it can be seen that there are classical NK neighborhoods, for which the number of local optima is constant as $L \rightarrow \infty$ at all constant k . For example in the star neighborhood with $A = 2$ the number of local maxima is at most (and asymptotically almost surely) 2^{k-1} , see^[23] attached as chapter 5.

3.2.5 Regular- and uniform structures

While the above does not provide any useful lower bounds on the number of local maxima in general, the situation changes for uniform and regular NK structures.

Uniformity and regularity constrain the structure of the neighborhood hypergraph and these restrictions may be used to derive lower bounds on the number of

local maxima, in particular through the existence of global reciprocal sign epistasis and similar motifs in the fitness graph.

The following is a slight generalization and summary of the reasoning I presented in^[48] and in^[23] which is attached as chapter 5.

For any given locus I define the induced distance- n subgraph around this locus as the sub-hypergraph obtained by keeping only loci in distance n or less to the focal locus, as well as hyperedges which are fully contained in this set of loci. In regular and uniform hypergraphs the possible number of induced distance- n subgraphs is then finite for all fixed n , up to isomorphism. Essentially, there are only a finite number of ways that the regular uniform structure can look locally.

Given two loci l and m which share a neighborhood, we can consider the induced distance-2 subgraph of the neighborhood structure around one of the loci. All loci not present in this subgraph will not share a neighborhood with either l or m and therefore their state cannot influence the effect of mutations on l and m .

It is always possible to find an assignment of fitness contributions on the partial landscapes of the subgraph so that the sign of effects of mutations on l and m are completely independent of the state of the remaining loci. To see this consider first a fixed background genotype and the partial landscape that would be obtained by only mutations on loci l and m starting from g . This partial landscape can generate all possible rank orderings of the corresponding NK model with $L = 2$ on the same allele graph. We can choose any such rank ordering of interest and the HoC model will generate it with some non-zero probability. We can then consider all other backgrounds formed by loci in the induced subgraph and since the graph is uniform and regular, there is an upper bound on the number of these backgrounds in terms of r and k . The probability that the rank ordering of interest is generated on all possible backgrounds is therefore a non-zero probability for fixed k and r . Similarly, by expanding the induced subgraph to distance-3 we can further obtain a non-zero probability that the chosen rank ordering on l and m is not only persistent independent of background but also that mutations on l and m do not affect the mutation effect of mutations on other loci.

Because the NK structure is uniform and regular at fixed k and r as L goes to infinity, it is always possible to find a linearly increasing number of adjacent locus pairs so that the distance-3 induced subgraphs mutually do not overlap. For each pair then the probability of the ordering of interest being observed is independent and as a consequence as $L \rightarrow \infty$, the number of locus pairs which show the ordering of interest will be stochastically dominated by a normal distribution with linearly increasing mean.

As seen before, we can always find a fitness rank ordering on the mutation graph with $L = 2$ and the given allele graph which has at least 2 local maxima. For each pair of loci having this ordering globally as described above, there will

then be 2 local maxima for any given local maximum and as a consequence the number of local maxima will be at least 2^Z where Z is the number of pairs with this property. By Jensen's inequality the expected number of local maxima is then growing exponentially in L .

This extends the previously known result of exponentially increasing expected number of local maxima in the AN structure to any uniform and regular structure or with slight modification any NK structure with bounded k_i and r_i .

The same approach as above can be used to show that GRSE will be almost surely present as $L \rightarrow \infty$ and that this pair of globally reciprocal loci will satisfy the even stronger variant of GRSE which I call *separable global reciprocal sign epistasis* under which mutations of the two loci do not affect the sign of mutations on other loci in addition to satisfying the GRSE property.

For more detailed argumentation see,^[23,48] which is specifically discussing the biallelic case, but also extends to the multi-allelic case in a straight-forward manner. Obtaining explicit bounds on the exponential rate via this approach is however somewhat convoluted and the resulting bounds are very weak for even moderate k , so that I will not present them here. I gave some results in this direction in.^[48]

3.2.6 New results for specific NK structures

I will now give some further new results presented in the publication^[23] attached as chapter 5 and in the appendices of this thesis.

For r -regular structures, one can consider the negative Gamma distribution with shape parameter $\frac{1}{r}$, meaning the distribution of a random variable X for which $-X$ is Gamma distributed with the given shape parameter. With this choice, the expected number of local maxima can be calculated exactly to^[23]

$$\mathbb{E}[N_{\max}] = 2^L \left(\frac{1}{k+1} \right)^{\frac{L}{k}} \quad (3.13)$$

coinciding with the results for the BN model.

While the BN model offers a choice of neighborhood structure under which the expected number of local optima is invariant to the fitness distribution, the negative Gamma distribution with the r -dependent shape parameter offers a fitness distribution for which the expected number of local maxima is invariant to the choice of neighborhood structure, assuming regularity.

For structures which are not sufficiently regular, the above does not apply. As seen before strongly irregular landscapes can behave in strongly different ways. Specifically the star neighborhood is here again a counter example since its number of local optima is always strictly limited by a constant at constant k , also for the negative Gamma distribution.^[23]

For the MF structure and the rRN structure with normal distributed fitness one can find by asymptotically exact methods $\zeta(\alpha) = \frac{2}{\alpha} - 1$ and $\zeta' = 2$.^[23] Weinberger obtains $\zeta(\alpha) = \frac{1}{\alpha}$ for the uRN model, but uses a derivation that relies only on approximations which also apply to the rRN structure. This result is therefore an indication that Weinberger's approximations do not result in the asymptotically fully exact expression.

The normal distribution's nice mathematical properties allow us to make more general statements about larger classes of NK structures. In,^[23] attached as chapter 5, I show that for all k -uniform and r -regular neighborhood structures with normal distributed fitness $\frac{1}{\alpha} \leq \zeta(\alpha) \leq \frac{2}{\alpha}$. A more careful derivation gives

$$\frac{1 + \frac{s}{r}}{\alpha} \leq \zeta(\alpha) \leq \frac{2}{\alpha} \quad (3.14)$$

$$\frac{s}{r} \leq \zeta' \leq 2 \quad (3.15)$$

where s is the smallest singular value of the incidence matrix of the neighborhood structure. The proof is attached as appendix A.1. The fraction $\frac{s}{r}$ can be seen to lie between 0 and 1 and for the upper bound it results in tight bounds for ζ' . A possible realization of such a NK structure with $\frac{s}{r} = 1$ is the mean field structure, for which we had already obtained $\zeta' = 2$ by other means.^[23]

Universality for uniform and regular structures

Based on the previous results, $\zeta' = 1$ seems to pose a general upper bound on the number of local maxima and interestingly it is always possible to reach this upper bound by adjustment of either structure or fitness distribution individually with the k -dependent negative Gamma distribution always satisfying it for any uniform and regular structure, while the BN structure generates it without any dependence on the choice of distribution. Both structure and distribution do seem to impose varying limits on the lower bounds though, but always staying within the expected asymptotic form of eq. (3.8) with ζ' bounded above by some constant.

The general result on uniform and regular NK models with normal distributed fitness is also expected to be relatively stable with respect other sufficiently regular fitness distributions.

For extremely heavy-tailed distributions however, the behavior seems to change and this is also where Limic and Pemantle find the upper bound $\zeta' \leq 3$ in the AN model.^[33] In particular they consider a distribution so heavy-tailed that sums of many i.i.d. random variables are asymptotically almost surely dominated by the addend with the largest absolute value.

I consider a similar distribution with the same property in order to investigate whether the apparently universal behavior of (3.8) can be violated. Indeed

I find interesting indications that there exist specific uniform- and regular NK structures, such that the expected number of local maxima with this distribution is asymptotically significantly smaller than in the usual cases discussed before, namely that

$$\lambda = \omega\left(\frac{\ln k}{k}\right) \quad (3.16)$$

A reasoning is provided in appendix A.2. As shown in the same appendix, there is a general bound applying to the sum-dominating distribution of

$$\lambda = \mathcal{O}\left(\frac{(\ln k)^2}{k}\right) \quad (3.17)$$

which I expect to be actually achievable by some uniform and regular NK structure.

The considerations above show in any case, that the possible behavior of expected number of local maxima is quite limited at larger k . Only for strongly irregular structures or strongly heavy-tailed fitness distributions can the expected number of local maxima vary strongly. For strongly irregular structures it is easy to reduce the expected number of local maxima drastically, as in the SN structure, but heavy-tailed fitness distributions are also able to reduce the expected number of local optima asymptotically significantly in the exponential rate.

Interestingly, either of these irregularities only decreases the expected number of local optima and it does in fact seem that the BN structure has the highest achievable value among all NK structures.

Chapter 4

Accessibility

4.1 Definitions

Local maxima determine the possible final states of populations under adaptive walk dynamics. They do not however alone determine the likelihood of populations reaching specific maxima. The structure of non-maximal genotypes may skew the probability of specific maxima being reached in such a way that high fitness values among the local maxima are more or less likely to be reached by adaptive walk dynamics.

In particular, the structure of the landscape can, if it is sufficiently rough, strongly constrain the number of local maxima and genotypes in general that are reachable from any given starting genotype by an adaptive walk. We say that one genotype h is *accessible* from another genotype g , if there exists a (directed) path on the mutation graph from g to h in the usual graph-theoretic sense, along which fitness values are strictly increasing on the fitness landscape. Such paths are called *accessible paths*. If a genotype is accessible from another, then the exact probability of it being reached will depend on the specifics of the population dynamics model chosen. Wanting to abstract from these details, we may therefore first consider the question of the probability that a given pair of genotypes is accessible in a stochastic fitness landscape model. Furthermore, we may consider the distribution of the number of accessible paths between the genotypes in the model. This number does generally not correspond to the likelihood of the given genotype being reached in adaptive dynamics, but may serve as a general indicator to differentiate genotypes which are highly likely to be reached and genotypes which are unlikely to be reached.

For the random variable representing the number of accessible paths between genotypes g and h I will write $N_{g,h}$ in this chapter.

In many stochastic fitness landscape models, such as the HoC or the NK model

with RN neighborhood, the model is invariant under relabelling of loci and other models can be made invariant by imposing the model to be applied to a random permutation of loci. Under this symmetry $N_{g,h}$'s distribution is then invariant under relabelling of loci and consequently only the equivalence classes of pairs (g, h) under locus relabelling need to be considered. These classes can be identified by a matrix $M \in \mathbb{N}_0^{A \times A}$, such that $M_{a,b}$ determines the number of loci l on which $g_l = a$ and $h_l = b$. With this $\sum_{a,b \in A} M_{a,b} = L$. In the limit $L \rightarrow \infty$, this matrix determines the direction and relative distance in which the two points g and h are considered. In particular, we are interested in the limiting behavior where $\frac{M_{a,b}}{L}$ converges for all a and b , in which case this limit is denoted $p_{a,b}$. The error term $R_{a,b} = M_{a,b} - Lp_{a,b}$ determines the rate of convergence to this limit and may be relevant for (higher-order) behavior of the probabilities involved. Finally, the null matrix, which corresponds to sub-linear distances between the genotypes under consideration will be excluded in the following. The probability of paths typically has non-trivial behavior only in the linear-distance case as a result of the linear increase in dimension of the mutation graph.

Specifically for the biallelic case, the notion of the matrix M can be simplified. There are only two non-trivial allele graphs on two alleles, a directed one and a symmetric (undirected) one. In the directed case, loci with $a_l = b_l$ are effectively irrelevant to the accessibility of paths, since there is no walk on the mutation graph from a to b mutating locus l . The result is simply the behavior of the model for a reduced value of L . Therefore, with the directed biallelic case it is only useful to consider accessibility from the wild-type to its single antipodal genotype in maximal distance. For the symmetric biallelic case $a_l = b_l$ does result in walks mutating locus l , but the allele graph is invariant under exchange of the two alleles and this symmetry is preserved by the fitness landscape models which we consider here analogously to the symmetry of locus exchange. Therefore only a single quantity, the number of loci with $a_l \neq b_l$, or equivalently the distance between a and b on the mutation graph is relevant. In this case the distance is denoted d and the relative distance $\delta = \frac{d}{L}$, which in accordance with the requirements on M is assumed to converge as $L \rightarrow \infty$.

4.2 Accessibility in the HoC model

Accessibility between pairs of genotypes has been previously considered in the HoC model for the biallelic case.^[4,5,21,32,35,37] Typically in this setting a HoC model with fitness values distributed according to a standard-uniform distribution is considered. The number of local maxima is not affected by the particular choice of distribution.

Commonly in this setting the model is modified, imposing that the fitness at

the destination d is the global maximum of the landscape, meaning that it has rank 1 on the ranked fitness landscape. This point of view is useful to demonstrate the maximum possible accessibility of a local maximum. At the same time it turns out that additional fixation of the fitness at a , results in a sharp transition of accessibility in these models depending on this fixation. Instead of fixation of the rank of the fitness at a directly, fixation is typically done in terms of the uniform fitness. As a result the actual rank of the fitness value at a varies, but for large L the relative ranking quickly converges to the equivalent fixed fitness.

Given fixation of fitness values at a and b with $\beta = F(b) - F(a)$, assumed positive, the accessibility question can always be mapped back to the HoC model with fitness 0 at a and fitness value 1 at b with an additional Bernoulli percolation imposed on all other genotypes, removing each independently with probability $1 - \beta$.^[49]

Therefore, only the difference β , rather than the choice of global maximum, is relevant to the following results. In the following I write $N'_{a,b}(\beta)$ to refer to the number of accessible paths in this model restricted by β . Because both the fitness value at a and at b are independently uniformly distributed, the accessibility in the original model can then be written as

$$\mathbb{P}[N_{a,b} \geq 1] = \int_0^1 d\beta (1 - \beta) \mathbb{P}[N'_{a,b}(\beta) \geq 1] \quad (4.1)$$

In the directed biallelic case, there is a threshold function at $\beta = 1 - \frac{\ln L}{L}$ of width at most $\frac{\sqrt{\ln L}}{L}$ below which accessibility converges to 0, but above which it converges to 1 as shown by Hegarty and Martinsson.^[21] In particular, if the fitness difference is fixed in L , it must be chosen as 1 in order to achieve accessibility. For the accessibility in the original model, it follows that

$$\mathbb{P}[N_{a,b} \geq 1] \sim \frac{1}{2} \left(\frac{\ln L}{L} \right)^2 \quad (4.2)$$

On the symmetric biallelic allele graph the number of walks between the initial and final points are much larger, resulting in a critical value at $\beta = \beta^*$ with β^* the unique positive solution to^[4,32,35,37]

$$\cosh(x)^{1-\delta} \sinh(x)^\delta = 1 \quad (4.3)$$

In particular at $\delta = 1$ the value of β^* becomes $\operatorname{asinh}(1) \approx 0.881$ and it falls with decreasing δ to zero. As a consequence, in the unrestricted model, any two genotypes are reachable from one another asymptotically with non-zero probability.

Moreover, it turns out that the critical point β^* is also exactly the fitness difference at which the expected number of accessible paths changes convergence behavior from convergence to zero to exponential divergence. Application of Markov's bound yields that convergence of the expected number of accessible paths to zero implies convergence of the probability of accessibility to zero. Therefore the expected number of paths is sufficient to determine a lower bound β^* . However, since β^* turns out to be exactly the value of this bound we can say that the expectation value of the number of accessible paths “tells the truth”^[4] about the critical point.

Martinsson noticed that the accessibility question in the restricted and non-restricted model can be mapped to a first-passage percolation problem on the same graph with uniformly distributed weights on vertices.^[35]

He also considered the equivalent first passage percolation problem with weights on edges rather than vertices, which doesn't directly map back to the accessibility percolation problem. However he considered this first passage percolation problem on arbitrary Cartesian graph products in analogy to allele graphs in the accessibility percolation setting.^[36]

Results equivalent to those by Martinsson can be obtained for the multi-allelic House-of-Cards model accessibility percolation problem as I demonstrate in the attached preprint^[49] attached as chapter 6.

In the following I summarize the result: Assuming \mathbf{A} is the adjacency matrix of the allele graph we can define β^* as the unique positive solution to

$$\Gamma(x) = \sum_{i,j \in \mathcal{A}} p_{i,j} \ln \left(e^{x\mathbf{A}} \right)_{i,j} = 0 \quad (4.4)$$

The term on the left-hand side is the limit of $\frac{\mathbb{E}[N_{a,b}]}{L}$ as $L \rightarrow \infty$. In other words $\Gamma(x)$ is the exponential growth rate of the expected number of accessible paths. For small β this term will have a negative exponential rate, implying that the expected number of accessible paths decreases asymptotically to zero. At β^* the exponential rate switches to an increasing exponential rate, implying that the number of accessible paths above β^* is exponentially large in L . β^* is similarly a lower bound on the percolation time in any of the first-passage percolation problems mentioned above. As before this behavior of the expectation value gives a lower bound on the critical point β^* , and the obvious question is whether this expectation also “tells the truth”.

Martinsson shows^[36] that this is the case in the edge-weighted first-passage percolation problem if and only if a certain function of the adjacency matrix fulfills certain properties. As I show in the attached publication, the same function determines whether or not β^* is the critical value in the accessibility problem. Martinsson considers only the case where the pairs (a_l, b_l) are equal on all loci, but in the following I give the condition in its straight-forward generalization to

arbitrary a and b with the conditions laid out at the beginning. Specifically, the condition is that

$$\mathfrak{M}(s, t) = \sum_{i,j \in \mathcal{A}} p_{i,j} \frac{\sum_{k,l \in \mathcal{A}} \left(e^{\beta^*(1-s)t\mathbf{A}} \right)_{i,k} \left(e^{\beta^*s\mathbf{A}} \right)_{k,l} \ln \left(e^{\beta^*s\mathbf{A}} \right)_{k,l} \left(e^{\beta^*(1-s)(1-t)\mathbf{A}} \right)_{l,j}}{\left(e^{x\mathbf{A}} \right)_{i,j}} \quad (4.5)$$

is negative for all $0 \leq t \leq 1$ and $0 < s < 1$ and that it satisfies certain additional conditions at the boundaries which vary between the results obtain by Martinsson and those I obtain in chapter 6. If these conditions are not satisfied, then the critical value is strictly larger than β^* . The expectation value therefore doesn't *always* tell the truth. but while such cases do exist, as shown by Martinsson,^[36] I am not aware that any critical values have been obtained for these unusual cases.

The conditions above seem to only rarely be false, so most examples of allele graphs can be applied to obtain the critical fitness difference from the expectation value's behavior. In particular for the complete graph on multiple alleles, numerical derivations of the critical values were carried out by Zagorski et al.^[68] with results matching the analytical ones.

In addition to the critical fitness difference I also show in the derivation in the attached preprint^[49] in chapter 6 that at the critical fitness difference the length of accessible walks can be determined to be up to lower order corrections $\Gamma'(\beta^*)L$. The equivalent result for the first passage percolation problem with weights on edges as well as more detailed analysis of path properties in this setting have been given independently by Kistler et al.^[27,28]

4.3 Accessibility in the NK model

In this section I will consider only \mathcal{K}_2 as allele graph.

While the expected number of accessible paths in the HoC model is always growing at most exponentially with L and has a critical behavior for the probability of accessibility of two genotypes in linear distance, the same does not generally hold for NK model landscapes.

As a simple example, in the non-epistatic model the accessibility of two genotypes in distance d depends only on their relative position to the global optimum. If they are accessible, then it is possible to choose any ordering of the d mutations required to reach the destination genotype in any order, resulting in an expected number of paths growing factorially in d . On the other hand the probability for the two genotypes to be ordered such that they are accessible is only 2^{-d} , implying an unconditional exponential decrease in accessibility at linear distances in L . However, in particular, choosing either the destination as global maximum or the source as global minimum of the landscape, results in guaranteed accessibility.

4.3.1 Accessibility in the SN structure

First I will consider the star neighborhood structure as an example of an NK landscape which is easy to handle combinatorially, although it is very unusual in its strong irregularity.

For each choice of alleles for the center loci of the star neighborhood, the sub-landscape formed by the remaining ray loci is completely non-epistatic. This has some immediate consequences for the accessibility of genotypes. Two genotypes g and h may differ in some or all of the $k - 1$ center loci, as well as at least $d - (k - 1)$ of the ray loci. We can first consider only paths which are crossing the distance on the center loci first. Since the partial landscape formed by the center loci is a HoC landscape with the state of the remaining loci fixed, there is an expected number of paths crossing this distance which is non-zero and either increasing or decreasing exponentially in k . At the same time, once the center loci are crossed, there is a probability of $\frac{1}{2}$ for each of the remaining loci to be individually and independently accessible. But since each ordering of the remaining loci in a path implies the same accessibility, it follows that the expected number of accessible paths is at least of order $d!$ up to exponential factors.

Therefore, as in the non-epistatic model, the SN model always has a significantly larger expected number of accessible paths between genotypes in linear distance than the HoC model does.

It is however less clear with what probability the two genotypes are actually accessible, i.e. whether the large expectation value “tells the truth” about accessibility, as it does for the non-epistatic model only if either the source or destination genotype’s fitness is conditioned to be small/large.

For each state of the center loci, the ray loci form a non-epistatic sublandscape and for each ray locus on which g and h differ, there is a probability $\frac{1}{2}$ that the mutation towards h increases fitness. This condition is satisfied independently for all 2^{k-1} states of the center loci, since any state change on the center loci changes the state of all partial landscape. There is then a chance of $2^{-2^{k-1}}$ that a locus is mutable towards h on any of the sub-landscapes. Because there are at least $d - (k - 1)$ loci which could independently satisfy this condition, there is a probability of at least $\Omega\left(\exp\left(-\left(d - (k - 1)\right)2^{-2^{k-1}}\right)\right)$ that one of the relevant loci satisfies the condition. But if any locus satisfies this condition, then there is no way to construct an accessible path from g to h . As a consequence, at fixed k , the probability that two randomly chosen loci are accessible from one another is decreasing exponentially to zero in the SN model.

If the source and destination genotypes are constrained to be at least global minimum and maximum on their respective sub-landscapes formed by the ray loci, the situation changes. It is then always possible to mutate the loci on which the two loci differ either at the beginning or the end of the path. At the same

time, there is a non-zero probability that it is possible to find an accessible path across the HoC sub-landscape formed by the center loci directly at the beginning of the path and this probability is also converging to a non-zero value as seen in the previous section. Consequently under these constraints there is always an asymptotically non-zero probability that any two genotypes are accessible.

It can therefore be said that the expectation value “tells the truth” in the SN model only under constraint of the fitness values of the source and destination genotype.

4.3.2 Accessibility in the BN structure

As example of a combinatorial tractable NK structure with the regularity property, the block neighborhood may be considered.

It is easy to see that due to the modular structure, a path is accessible if and only if it is accessible with respect to the sub-landscape formed by the NK neighborhoods of a given block of the BN model. And since each of these block sub-landscapes are effectively of HoC type, the overall accessibility can be derived from the results for the HoC model.

In particular, at constant k , the size of blocks and the accessibility of individual blocks remains unchanged, while the number of blocks increases linearly with L . As a consequence the probability that two genotypes in linear distance are accessible decreases exponentially in L . At the same time, if there is an accessible path, it doesn't matter in which order the individual sub-paths on blocks are ordered into a path on the whole landscape. Therefore the expected number of accessible paths will actually increase factorially with L , as in the SN model.

For k increasing linearly with L , the individual block sizes also increase linearly, resulting for the complete graph on two vertices as allele graph in an asymptotically non-zero probability of accessibility on each block. At the same time, the number of blocks will be constant, implying that the BN model accessibility in this limit also has a critical fitness difference threshold as the HoC model does. The expected number of accessible paths under this limit also follows an exponential behavior as in the HoC model.

In the BN model it can therefore be said that the expected number of accessible paths “tells the truth” only if k is growing sufficiently fast with L .

Contrary to the irregular SN model, in the BN structure at constant k even constraining the source and destination fitness values does not yield asymptotic accessibility.

4.3.3 Accessibility in other regular and uniform structures

In other regular- and uniform NK structures, it can be seen that the behavior at constant k and r always qualitatively follows that of the BN structure.

As seen in the case of the number of local maxima, regularity and uniformity of the NK structure implies that any given rank ordering of a two-locus landscape will eventually be generated a linearly growing number of times between pairs of adjacent loci in the neighborhood structure given that k and r are constant and that they can also be generated in such a way that all mutations on the pair of loci do not affect the signs of mutation effects on other loci.

In particular *separable global sign epistasis* will occur as one special case of such a rank ordering. It is impossible to find an accessible path crossing both of the two loci with this property. As a consequence the presence of a single such pair is sufficient to make it impossible to find accessible paths of between genotypes in distance L on the landscape at all. Even more so, since the set of pairs with this property will linearly increase in L , eventually the probability that one such pair is contained in the set of loci on which the source genotype g and the destination genotype h differ, will converge to one and accessibility between any pair of accessible loci in linear distance will decrease to zero.

Instead of separable global sign epistasis, simply all orderings of the two-locus landscape can be considered, while still keeping the restriction on not influencing signs on mutations of other loci. Such motifs will also appear in linear growing number and therefore effectively there will be a sublandscape over a subset of loci which behaves like a BN landscape with $k = 2$.

As seen before on the BN landscape accessibility at constant k converges to zero, while the expected number of accessible paths increases factorially. While the fraction of the overall landscape contained in this BN-like landscape may be small, it is still sufficient to determine the qualitative behavior of both the expected number of accessible paths and accessibility of uniform- and regular NK models.

More details on this approach are provided in the attached^[23] in chapter 5 as well as in my previous work.^[48]

Chapter 5

Publication 1: Universality classes of interaction structures for NK landscapes.

This chapter contains the first publication incorporated into this thesis. It provides a more detailed review and new results for the behavior of local maxima in NK models, as well as some other properties of NK landscapes, such as the existence of accessible paths and properties of adaptive walks.

The following is the published online version of

S. Hwang et al., “Universality classes of interaction structures for nk fitness landscapes,” *Journal of Statistical Physics* **172**, 226–278 (2018).

The results and wordings of sections “Mathematical Background and Definitions”, “Accessible Pathways” and appendix C are primarily contributions by myself with supportive and editorial contributions by the co-authors. My contributions to other sections of the article were to varying degrees editorial and supportive on obtaining of results.

In accordance with the doctoral regulations, this article, already published in a peer-reviewed scientific journal, is not attached in the published version of this thesis.

Chapter 6

Publication 2: Accessibility percolation on cartesian power graphs.

The following is a preprint reproduction of

B. Schmiegelt and J. Krug, “Accessibility percolation on cartesian power graphs,” 2021

Results were obtained primarily by myself under guidance by the co-author. With the exception of parts of the introductory section and editorial influence, the wording of the article is my own.

manuscript No.
(will be inserted by the editor)

Accessibility Percolation on Cartesian Power Graphs

Benjamin Schmiegelt · Joachim Krug

Abstract A fitness landscape is a mapping from a space of discrete genotypes to the real numbers. A path in a fitness landscape is a sequence of genotypes connected by single mutational steps. Such a path is said to be accessible if the fitness values of the genotypes encountered along the path increase monotonically. We study accessible paths on random fitness landscapes of the House-of-Cards type, on which fitness values are independent, identically and continuously distributed random variables. The genotype space is taken to be a Cartesian power graph \mathcal{A}^L , where L is the number of genetic loci and the allele graph \mathcal{A} encodes the possible allelic states and mutational transitions on one locus. The probability of existence of accessible paths between two genotypes at a distance linear in L displays a transition from 0 to a positive value at a threshold β_c for the fitness difference between the initial and final genotype. We derive a lower bound on β_c for general \mathcal{A} and show that this bound is tight for a large class of allele graphs. Our results generalize previous results for accessibility percolation on the biallelic hypercube, and compare favorably to published numerical results for multiallelic Hamming graphs.

1 Introduction

In the strong-selection weak-mutation (SSWM) regime evolutionary dynamics reduces to an adaptive walk on what is known as a fitness landscape, the map from genotypes to fitness values [20]. For low mutation rates the nearly monomorphic population can be represented by a single majority genotype moving through the space of genotypes by individual mutations that fix with

bschmie1@uni-koeln.de

jkrug@uni-koeln.de

B. Schmiegelt E-mail: schmiegb@thp.uni-koeln.de · J. Krug E-mail: jkrug@uni-koeln.de
Institute for Biological Physics, University of Cologne

arXiv:1912.07925v2 [q-bio.PE] 20 Apr 2021

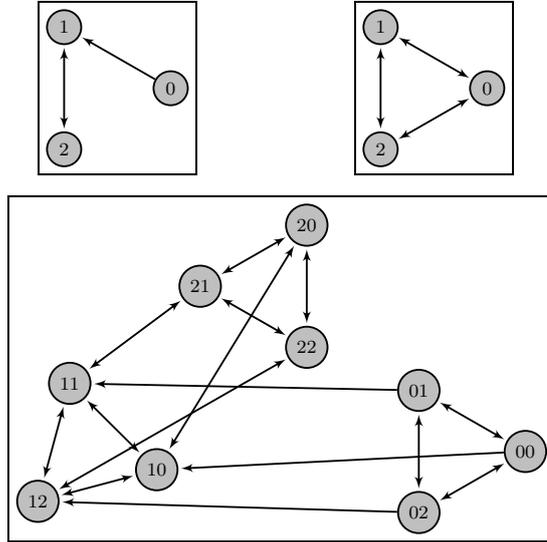


Fig. 1 Example of a genotype space as the Cartesian graph product of two allele graphs. The two allele graphs are shown on top with the genotype space below. While the second factor graph represents a locus with three possible alleles, all of which may mutate freely from one to another, the first factor graph represents a locus on which not all mutations between the alleles are considered possible. Specifically mutations between 0 and 2 must take an intermediate mutation through 1 and additionally the mutation from 0 to 1 is considered irreversible and does not allow backstepping to 0. Although in this work we define the genotype graph as the direct Cartesian power of a single allele graph, different allele graphs as shown here can still be modeled without loss of generality by assuming that A is the distinct sum of the individual graphs. Since the individual constituent graphs are not connected in this sum, this does not increase the accessibility.

a probability depending on the fitness of the mutant relative to the parental genotype [8, 19]. Under strong selection, the movement of such a walker is additionally constrained towards increasing fitness values, making it an adaptive walk [10]. This limits the number of selectively *accessible paths* a population can take through the genotype space [5, 7, 21].

Here we investigate the impact that the mutational structure of the genotype space has on the number of evolutionary paths available to SSWM dynamics. We use a simple stochastic model for fitness landscapes known as the House-of-Cards (HoC) model, in which each genotype \mathbf{g} is assigned an i.i.d. continuous random fitness value $F_{\mathbf{g}}$ [10, 11]. As the property of a path to be accessible only depends on the rank ordering of fitness values the particular distribution chosen will not be of relevance. For simplicity we will assume the standard uniform distribution.

A genotype is made up of many individual sites or loci, which can be found in some given number of states called *alleles* and can be mutated individu-

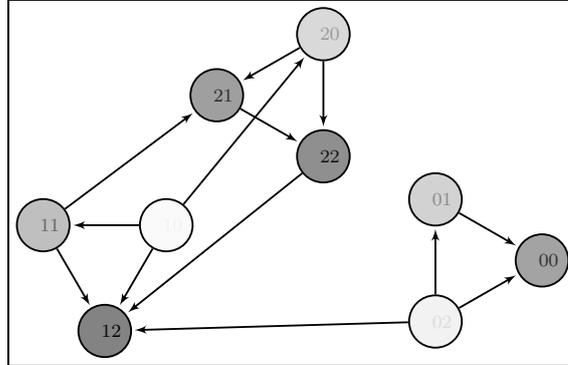


Fig. 2 Example of a fitness graph generated from the genotype space in Figure 1 according to the HoC model. The opacity of nodes indicates the randomly chosen fitness value. Only arrows representing mutations that were originally allowed in the genotype graph and also point towards increasing fitness remain, resulting in an acyclic directed graph on genotypes. The global minimum and maximum in this realization are (10) and (12) respectively and the latter is accessible from the former by multiple accessible paths, for example the direct one (10) \rightarrow (12), but also (10) \rightarrow (20) \rightarrow (22) \rightarrow (12). As a counter-example to accessibility consider (01) and (12). Although (12) has higher fitness than (01), it is not accessible from (01).

ally. For simplicity, but without loss of generality, we will assume that all loci have the same set of possible states. Therefore genotypes are sequences $\mathbf{g} = (g_1, \dots, g_L)$, with L determining the number of loci. Individual (point) mutations, which are the only ones to be considered here, mutate only one of the loci. The mutational structure of the system determines whether every state of one locus is able to mutate to any other or whether some restrictions apply. For example, whereas point mutations in the DNA sequence can mutate any nucleotide base into any other, the genetic code constrains the possible one-step transitions between amino acids. To accommodate general mutational structures we describe the loci by a simple directed *allele graph* \mathcal{A} . The vertex set \mathbb{A} of this graph is the set of all alleles, and its arrows indicate possible one-step mutations between alleles. Again, for simplicity but without loss of generality, we assume the allele graph to be the same on all loci and we will identify the vertices by natural numbers.

The genotype space can then be described as the Cartesian graph product \mathcal{A}^L [17], a directed graph whose vertices form the genotypes and with arrows between genotypes that can be reached via one-step mutations (Figure 1). The fitness landscape constrains which of these arrows may be taken by an adaptive walker and we call the directed sub-graph of the genotype graph obtained by removing arrows which do not point towards increasing fitness the fitness graph [6] (Figure 2).

Our goal is to determine, for a given pair of genotypes \mathbf{a} and \mathbf{b} , how likely it is possible to reach \mathbf{b} from \mathbf{a} on the fitness graph when the fitness landscape is of HoC type. In particular this question is non-trivial if the directed distance $d_{\mathbf{a}\mathbf{b}}$ from \mathbf{a} to \mathbf{b} on the genotype space is of linear order in L . Of special importance for this question is the value β , which we define as the fitness difference $F_{\mathbf{b}} - F_{\mathbf{a}}$. Conditioned on this value it is known from previous work on the case of two alleles, $\mathbb{A} = \{0, 1\}$, and linear distance $d_{\mathbf{a}\mathbf{b}} \sim \delta L$, that there is a critical value β_c , depending on δ , such that for constant choices of β above or below β_c , asymptotically the probability of \mathbf{b} being accessible from \mathbf{a} converges to 1 or 0, respectively, as $L \rightarrow \infty$. [4,3,9,14,15]. The transition occurring at $\beta = \beta_c$ has been referred to as *accessibility percolation* [13,18]. Apart from a computational study [23], so far accessibility percolation has been studied only for the biallelic case for which \mathcal{A}^L space is the L -dimensional (binary) hypercube.

Results related to those presented here have been obtained in the context of first-passage percolation [12,16,17], which is linked to accessibility percolation through a mapping described in [15]. In [17] Martinsson considers a first-passage percolation model that would map to the HoC accessibility percolation problem if fitness values were assigned to edges rather than vertices. We will adapt and extend his methods to directly resolve the specific accessibility problem introduced here without requiring the mapping to first-passage percolation. Instead our results may be mapped to the first-passage percolation problem with uniform weights on vertices as described in [15].

Another way of looking at the HoC accessibility problem conditioned on the value β is to consider it as a Bernoulli percolation on a certain ensemble of orientations of \mathcal{A}^L . Letting without loss of generality $F_{\mathbf{a}} = 0$ and $F_{\mathbf{b}} = \beta$, only genotypes with fitness values below β are relevant in determining whether an accessible path exists. Furthermore, after removal of the ineligible genotypes, accessibility will depend only on the order of the fitness values on the remaining vertices. Because the fitness values are chosen i.i.d. this implies that the problem with some β is equivalent to the same problem with $\beta = 1$ after additional removal of each vertex (that is not \mathbf{a} or \mathbf{b}) with probability $1 - \beta$, i.e. a Bernoulli site percolation with rate β . In the base case with $\beta = 1$, the problem reduces to whether or not \mathbf{b} is reachable from \mathbf{a} in a random orientation of \mathbf{A} obtained by intersection of its edge relation with a uniformly chosen linear order of vertices. This perspective may be more suitable if one is to consider the effect of decreasing β .

2 Results

In this section we present our general results for arbitrary allele graphs \mathcal{A} of which specific applications will be demonstrated in the following section. First we define some attributes of the problem more carefully and state the limitations imposed by our following proofs.

2.1 Prerequisites

We assume that the allele graph \mathcal{A} has a maximal (out-)degree of at most $\Delta < \infty$. In particular Δ can always be chosen as the number of alleles for finite \mathcal{A} . If the degrees are not sufficiently bounded in an infinite allele graph, then the number of walks from \mathbf{a} to \mathbf{b} may become so large that the problem results in trivial accessibility.

In the following we denote the number of accessible paths from \mathbf{a} to \mathbf{b} by $Z_{\mathbf{ab}}$. The *accessibility* of \mathbf{b} from \mathbf{a} is then the probability $\mathbb{P}[Z_{\mathbf{ab}} \geq 1]$, usually conditioned on $\beta = F_{\mathbf{b}} - F_{\mathbf{a}}$ as $\mathbb{P}[Z_{\mathbf{ab}} \geq 1 | \beta]$. We study the behavior of this quantity in the large- L limit under some given sequence of β in L . For convenience the dependence of β on L is taken to be implicit and usually not reflected in notation. As L , and with it the graph \mathcal{A}^L , changes we need to define a sequence of endpoint pairs (\mathbf{a}, \mathbf{b}) in L . For general such sequences calculations may become tediously complex and so we impose a few restrictions described in the following on the sequences for which our results will apply. Because the position of a locus in the sequence of the genotypes does not matter, all relevant properties of the pair (\mathbf{a}, \mathbf{b}) can be expressed as an integer-valued matrix:

$$M_{vw} = |\{a_l = v \wedge b_l = w \mid l = 1 \dots L\}| \quad (1)$$

where $v, w \in \mathbb{A}$ are alleles. This matrix counts for each pair of alleles the number of loci on which the path is required to move from v to w , thereby dividing out the permutation-symmetry of loci. A sequence of such matrices M in L is then equivalent to a sequence of pairs (\mathbf{a}, \mathbf{b}) up to the irrelevant symmetry of \mathcal{A}^L . To avoid trivial special cases, we assume that for each locus l , on the single allele graph, b_l is always reachable by at least one walk of finite non-zero length from a_l and that $\mathbf{a} \neq \mathbf{b}$. The former requirement assures that there are no loci which effectively cannot mutate and therefore either inhibit accessibility a priori or could be removed from the problem without influencing the result. For infinite allele graphs we require that there are only finitely many pairs (v, w) for which M_{vw} is ever non-zero for any L . This assures that distributions of quantities over loci are always well-behaved. We also require that this sequence of matrices converges element-wise in the sense that

$$p_{vw} = \lim_{L \rightarrow \infty} \frac{M_{vw}}{L} \quad (2)$$

converges for all $v, w \in \mathbb{A}^L$. This guarantees that there are no significant oscillations in the sequence, which would be tedious to handle and not informative. Further, we require that at least one of the off-diagonal terms of p_{vw} is non-zero. This assures that the directed distance from \mathbf{a} to \mathbf{b} is increasing linearly in L . While the sub-linear case can be considered with the same methods, it requires some special considerations which are out-of-scope here. We then define

$$R_{vw} = M_{vw} - Lp_{vw} \quad (3)$$

as the remaining corrections to the linear order.

In order to succinctly state our results, we introduce the following quantity for pairs of alleles $v, w \in \mathbb{A}$ and $t > 0$:

$$\Gamma_{vw}(t) = \ln \left((e^{tA})_{vw} \right). \quad (4)$$

Here and in the following A stands for the adjacency matrix of the allele graph and the exponential is a matrix exponential from which the element representing alleles (v, w) is extracted, rather than the exponential of the element of the matrix. In addition to the quantity $\Gamma_{vw}(t)$, also its first two derivatives $\Gamma'_{vw}(t)$ and $\Gamma''_{vw}(t)$ with respect to t will be important. We indicate the derivative with respect to the t argument by backticks as shown.

We define then the same quantity for pairs of genotypes $\mathbf{v}, \mathbf{w} \in \mathbb{A}^L$ as averages over the per-locus quantity:

$$\mathbf{\Gamma}_{\mathbf{vw}}(t) = \langle \Gamma_{v_l w_l}(t) \rangle_l \quad (5)$$

where

$$\langle X_l \rangle_l = \frac{1}{L} \sum_{l=1}^L X_l. \quad (6)$$

As we already did for genotypes, quantities acting on \mathcal{A}^L will be written in boldface, while equivalent quantities acting on a single \mathcal{A} copy will be denoted in normal font-face. The canonical connection between the former and latter is averaging over loci.

$\Gamma_{vw}(t)$ is non-decreasing in t and for $v = w$ it is 1 at $t = 0$. For $v \neq w$, it diverges to $-\infty$ as $t \rightarrow 0$. If w is reachable from v by at least one non-trivial walk, i.e. a walk of finite non-zero length, then $\Gamma_{vw}(t)$ diverges to infinity as $t \rightarrow \infty$. Special care must be taken if w is not reachable from v , as the matrix exponential is then 0 for all t . In this case we formally interpret $e^{\Gamma_{vw}(t)}$ as $(e^{tA})_{vw}$, which would be zero.

We also require the following function with domain $0 \leq r, s \leq 1$, which is a slight generalization of a function introduced by Martinsson in [17]:

$$\mathfrak{M}(s, r, \beta) = \left\langle \langle \Gamma_{x_l y_l}(\beta s) \rangle_{x_l, y_l}^{s, r} \right\rangle_l, \quad (7)$$

where $\bar{r} = 1 - r$ and $\bar{s} = 1 - s$ and $\langle \cdot \rangle_{x_l, y_l}^{s, r}$ is the mean over $x_l, y_l \in \mathbb{A}$ weighted by

$$e^{\Gamma_{a_l x_l}(\beta \bar{s} r) + \Gamma_{x_l y_l}(\beta s) + \Gamma_{y_l b_l}(\beta \bar{s} \bar{r})}, \quad (8)$$

meaning that

$$\mathfrak{M}(s, r, \beta) = \left\langle \frac{\sum_{x_l, y_l \in \mathbb{A}} \Gamma_{x_l y_l}(\beta s) e^{\Gamma_{a_l x_l}(\beta \bar{s} r) + \Gamma_{x_l y_l}(\beta s) + \Gamma_{y_l b_l}(\beta \bar{s} \bar{r})}}{\sum_{x_l, y_l \in \mathbb{A}} e^{\Gamma_{a_l x_l}(\beta \bar{s} r) + \Gamma_{x_l y_l}(\beta s) + \Gamma_{y_l b_l}(\beta \bar{s} \bar{r})}} \right\rangle_l. \quad (9)$$

We refer to this function as Martinsson's function. For the case that y_l is not reachable from x_l or that $s = 0$, the formula yields negative infinities for $\Gamma_{x_l y_l}(\beta s)$. We assume that in this case formally the natural choice

$$\Gamma_{x_l y_l}(\beta s) e^{\Gamma_{x_l y_l}(\beta s)} = 0 \quad (10)$$

holds.

With our assumptions on \mathbf{a} and \mathbf{b} , $\mathbf{\Gamma}_{\mathbf{ab}}(\beta)$ has a single positive root in β , which we call $\hat{\beta}$. This $\hat{\beta}$ is our candidate for the critical point.

The objects introduced so far are dependent on L implicitly through the averaging process over loci. In order to be able to make statements about the limiting behavior, it is useful to consider the limits of these quantities as $L \rightarrow \infty$. We use the non- L -dependent mean

$$\langle X_{vw} \rangle_{a_l, b_l}^p = \sum_{v, w \in \mathbb{A}} p_{vw} X_{vw} \quad (11)$$

as a replacement for $\langle X_{vw} \rangle_L$. Assuming that the X_{vw} converge as $L \rightarrow \infty$, the restrictions we put on M result in convergence of $\langle X_{vw} \rangle_L$ to $\langle X_{vw} \rangle_{a_l, b_l}^p$ as $L \rightarrow \infty$.

In particular we write β^* for the limit of $\hat{\beta}$ as $L \rightarrow \infty$, which is also the unique positive root of $\langle \mathbf{\Gamma}_{a_l b_l}(\beta) \rangle_{a_l, b_l}^p$ in β . Similarly we write for the limit of Martinsson's function

$$\mathfrak{M}^*(s, r, \beta) = \left\langle \left\langle \Gamma_{x_l y_l}(\beta s) \right\rangle_{x_l, y_l}^{s, r} \right\rangle_{a_l, b_l}^p. \quad (12)$$

With the necessary quantities defined we can classify the different choices of allele graphs and sequence pairs (\mathbf{a}, \mathbf{b}) as follows. $\mathfrak{M}(s, r, \beta)$ can be seen to always be zero at $s = 0$ and to be equal to $\mathbf{\Gamma}_{\mathbf{ab}}(\beta)$ at $s = 1$. Both statements can be derived immediately by application of matrix multiplication. Specifically $\mathfrak{M}(s, r, \hat{\beta})$ is 0 at $s = 1$ by definition of $\hat{\beta}$ and equivalently for $\mathfrak{M}^*(s, r, \beta^*)$. If not only at $s = 0$ and $r = 0$, but everywhere in its domain $\mathfrak{M}^*(s, r, \beta^*) \leq 0$, then we say that the problem is of *semi-regular* type. Otherwise we say that it is of *irregular* type. If $\mathfrak{M}^*(s, r, \beta^*) < 0$ holds strictly everywhere except at $s = 0$ and $s = 1$ and if additionally the derivative $\partial_s \mathfrak{M}^*(s, r, \beta^*)$ is not zero at $s = 1$, then we say that the problem is of *regular* type.

For our statements and in the following proofs we make use of Landau notation with the usual meanings of $\mathcal{O}(\cdot)$, $o(\cdot)$, $\omega(\cdot)$ and $\Theta(\cdot)$. In our notation of arithmetic terms and equations these symbols are stand-ins for some function in the respective class. The limit variable to which these symbols apply should be evident from context, but is usually $L \rightarrow \infty$. Functions in these classes are not required to be non-negative. In particular e.g. $|\Theta(1)|$ is used to enforce positiveness of a term that is of constant (non-zero) asymptotic order in L . If not stated otherwise, these symbols are assumed to be uniform in the sense that the functions represented depend only on the limit variable and model parameters, but not on other local variables.

2.2 Statements

We state our results in terms of (weak) threshold functions defined as follows. We say that a sequence c_L in L is a $f(L)$ -threshold function for some function $f(L)$ if for all $g(L) = \omega(f(L))$, when the fitness difference between \mathbf{a} and \mathbf{b} is conditioned to be $\beta_L = c_L + |g(L)|$ it is true that

$$\liminf_{L \rightarrow \infty} \mathbb{P}[\mathbf{Z}_{\mathbf{ab}} > 0] > 0, \quad (13)$$

and when it is conditioned to be $\beta_L = c_L - |g(L)|$

$$\limsup_{L \rightarrow \infty} \mathbb{P}[\mathbf{Z}_{\mathbf{ab}} > 0] = 0. \quad (14)$$

In other words, c_L determines the asymptotic transition from zero accessibility to non-zero accessibility if we condition the fitness difference between initial and final genotype, with a window of uncertainty of the same order as $f(L)$. In particular if c_L is a $f(L)$ -threshold for some $f(L) = o(1)$, then the limit of c_L is the critical value β_c . The notion of threshold chosen here is weak in the sense that it doesn't imply a transition from zero to one, but only from zero to some non-zero probability. We do not think that our results are actually restricted to this weak bound and we expect that arguments analogous to those made in [17] may be used to extend our weak threshold result to a strong threshold with $\liminf_{L \rightarrow \infty} \mathbb{P}[\mathbf{Z}_{\mathbf{ab}} > 0] = 1$, but we did not pursue this improvement here.

Our first main statement is that for problems of regular type a $\frac{1}{L}$ -threshold function is given by

$$c_L = \hat{\beta} - \mathbf{\Gamma}'_{\mathbf{ab}}(\hat{\beta})^{-1} \frac{\ln L}{L}, \quad (15)$$

which in particular implies that $\beta_c = \beta^*$ in this case. Additionally the zero-accessibility side of this threshold statement holds irrespectively of (semi-)regularity.

Our second main statement is that for problems of irregular type β_c , if it exists, is strictly larger than β^* .

Lastly we show that for problems of regular type at the critical point $\beta_c = \beta^*$, all accessible walks are asymptotically almost surely of length $\mathbf{\Gamma}'_{\mathbf{ab}}(\beta^*)\beta^*L \pm o(L)$.

The zero-accessibility side of the threshold functions for (semi-)regular types can be derived directly from a consideration of the expected number of (quasi-)accessible walks and an application of Markov's inequality. This approach will be explained in Sect. 4, where we also introduce the notion of quasi-accessibility as a tool to simplify the counting of accessible paths. In addition, the first-moment approach allows us to prove the last statement about accessible walk lengths by consideration of the expected values separated by walk length.

To prove the positive accessibility side of the threshold function, it is necessary to bound a higher moment of the expected number of (quasi-)accessible

walks in relation to the mean. In particular, using a generalized version of the second moment method, it is sufficient to bound moments of the form

$$\mathbb{E} \left[\frac{Z_{ab}}{\mathbb{E}[Z_{ab}]} \ln \frac{Z_{ab}}{\mathbb{E}[Z_{ab}]} \right]$$

to show asymptotic boundedness of the accessibility away from zero. The evaluation of this expected value will follow the general ideas used by Martinsson in [17] to bound for every given (quasi-)accessible focal walk the number of other (quasi-)accessible walks, through the deviating *arcs* on the focal walk that generate all such other walks. In the mean taken over x_l and y_l in Martinsson's function (9), the focal walk is represented by the walk sequence

$$a_l \rightarrow x_l \rightarrow y_l \rightarrow b_l \quad (16)$$

and the corresponding three Γ -terms in the weights, while the deviating arcs are represented by the additional term corresponding to $x_l \rightarrow y_l$ over which the average is performed. Because Martinsson considers a model that corresponds to putting weights on edges rather than nodes, our calculations need to be adjusted accordingly (see Sect. 6).

The lower bound on β_c for the irregular case is again obtained following an approach used by Martinsson, by considering walks through pairs of edges $(\mathbf{x}, \mathbf{x}')$ and $(\mathbf{y}, \mathbf{y}')$, applying Markov's inequality separately, and union bounding the resulting probability to improve on Markov's inequality from the total expected number of (quasi-)accessible walks (see Sect. 5).

2.3 Asymptotic form

The theorems as stated in the previous section are dependent on $\hat{\beta}$ and $\langle \cdot \rangle_l$ averages, which are L -dependent quantities. From the assumptions, we do however know that $\hat{\beta}$ converges to β^* and averages of the form $\langle \cdot \rangle_l$ are asymptotically of the form $\langle \cdot \rangle_{a_l, b_l}^p$, both of which are L -independent quantities. Depending on the specific choice of pairs (\mathbf{a}, \mathbf{b}) , the rate of convergence for these quantities may however differ and add additional significant terms in the threshold function, which we detail in this section.

For finite graphs immediately, and for infinite allele graphs due to the additionally stated restrictions on M_{vw} , the linear order of the (directed) distance from \mathbf{a} to \mathbf{b} is given by the sum of off-diagonal terms of p_{vw} :

$$\mathbf{d}_{ab} = \delta L + o(L) \quad (17)$$

where $\delta = \sum_{v \neq w} p_{vw}$. Because we have $\delta > 0$, the value of β^* will be positive,

i.e. not zero, and then we can expand $\hat{\beta}$ around β^* in L :

$$0 = \Gamma_{ab}(\hat{\beta}) = \langle \Gamma_{a_l b_l}(\beta^*) \rangle_{a_l, b_l}^p \quad (18)$$

$$+ \frac{1}{L} \sum_{v, w \in \mathbb{A}} R_{vw} \Gamma_{vw}(\beta^*) + \langle \Gamma'_{a_l b_l}(\beta^*) \rangle_{a_l, b_l}^p (\hat{\beta} - \beta^*) + \dots \quad (19)$$

The term $\langle \Gamma_{a_l b_l}(\beta^*) \rangle_{a_l, b_l}^p$ is zero by definition of β^* and for $\langle \Gamma'_{a_l b_l}(\beta^*) \rangle_{a_l, b_l}^p$ we introduce the short-hand notation $\mathbf{\Gamma}'^*$. So, to highest order

$$\hat{\beta} = \beta^* - \frac{\mathbf{\Gamma}'^{*-1}}{L} \sum_{v, w \in \mathbb{A}} R_{vw} \Gamma_{vw}(\beta^*) + o(L^{-1}). \quad (20)$$

Inserting this into the candidate threshold function (15) we have

$$c_L = \beta^* - \mathbf{\Gamma}'^{*-1} \frac{\ln L + \sum_{v, w \in \mathbb{A}} R_{vw} \Gamma_{vw}(\beta^*) + o(1)}{L}. \quad (21)$$

In general, the candidate critical value does not depend on the non-linear corrections in the behavior of (\mathbf{a}, \mathbf{b}) , but the leading correction to the critical value is dependent if the non-linear corrections are of order $\ln L$ or higher. In particular, if we are looking in a completely linear direction, such that $\sum_{v, w \in \mathbb{A}} |R_{vw}| = \mathcal{O}(1)$, then the formula reduces to

$$c_L = \beta^* - \mathbf{\Gamma}'^{*-1} \frac{\ln L + \mathcal{O}(1)}{L} \quad (22)$$

where the $\mathcal{O}(1)$ contribution is irrelevant since c_L represents a $\frac{1}{L}$ -threshold function.

3 Applications

3.1 Complete graph

The simplest application is to the complete graph on \mathfrak{A} alleles, which leads to genotype spaces known as Hamming graphs. By symmetry, in this case there are only two choices for the initial and final allele on a locus, either $a_l = b_l$ or $a_l \neq b_l$. Therefore the setup can be fully described by just the relative distance δ , which is then also the relative Hamming distance. As shown in [17], the problem is always of regular type for the complete graph. One obtains

$$\mathbf{\Gamma}_{ab}(\beta) = -\ln \mathfrak{A} - \beta + \delta \ln(-1 + e^{\mathfrak{A}\beta}) + \bar{\delta} \ln(\mathfrak{A} - 1 + e^{\mathfrak{A}\beta}), \quad (23)$$

where $\bar{\delta} = 1 - \delta$. In the biallelic case $\mathfrak{A} = 2$ the condition $\mathbf{\Gamma}_{ab}(\hat{\beta}) = 0$ reduces to the relation $\sinh(\hat{\beta})^\delta \cosh(\hat{\beta})^{\bar{\delta}} = 1$ which was first conjectured in [3] and proved in [15, 14]. At full distance $\delta = 1$ without any variation of δ with L , $\hat{\beta} = \beta^*$ and

$$\mathbf{\Gamma}_{ab}(\beta) = -\ln \mathfrak{A} - \beta + \ln(-1 + e^{\mathfrak{A}\beta}). \quad (24)$$

The values of β^* and $\mathbf{\Gamma}'^*$ for small \mathfrak{A} are shown in Table 1.

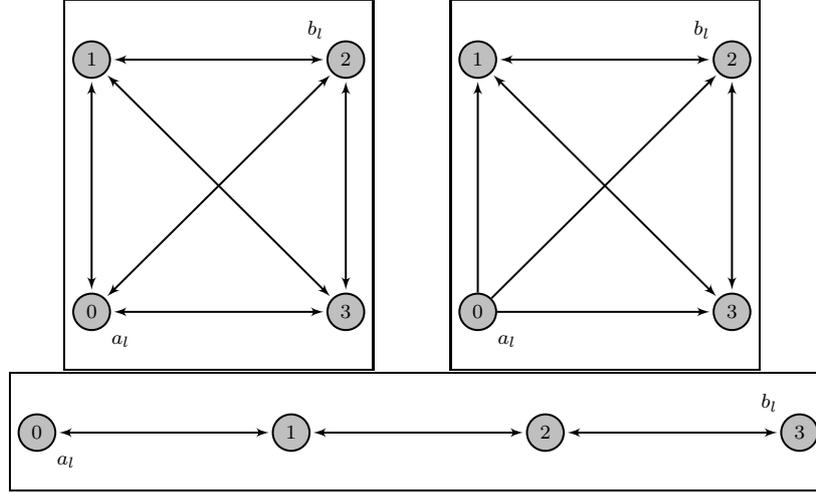


Fig. 3 Allele graph structures described in this section. Top left: Complete graph on four alleles with backmutations to the wild-type. Top right: Complete graph on four alleles with backmutations to the wild-type removed. Bottom: Path graph on four alleles. In each case possible pairs of initial and final alleles as used in this section are indicated by the labels a_l and b_l .

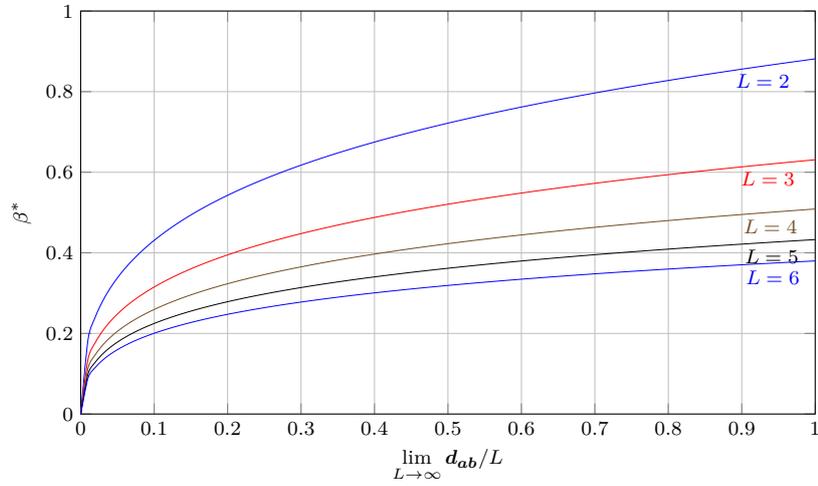


Fig. 4 β^* as a function of the relative distance $\delta = \lim_{L \rightarrow \infty} \frac{d_{ab}}{L}$ for the complete allele graph with 2 – 6 alleles. On the complete graph the distance specifies the choice of \mathbf{a} and \mathbf{b} up to irrelevant symmetries, so that this figure captures the full behavior.

\mathfrak{A}	β^*	Γ'^*	$\beta^* \Gamma'^*$
2	$\arcsin(1) \approx 0.881$	$\sqrt{2} \approx 1.41$	≈ 1.25
3	$\ln\left(2 \cos \frac{\pi}{9}\right) \approx 0.631$	$1 + 2 \cos\left(\frac{2\pi}{9}\right) \approx 2.53$	≈ 1.82
4	$\ln\left(\frac{1}{\sqrt{2}} + \sqrt{\sqrt{2} - \frac{1}{2}}\right) \approx 0.509$	≈ 3.60	≈ 1.83

Table 1 Results for the complete allele graph with 2-4 loci at full distance $\delta = 1$. The last column shows the prefactor of the asymptotic walk length at the critical point. In the biallelic case $\mathfrak{A} = 2$ the result for the walk length was also obtained in [12].

In general e^{β^*} is the unique positive solution of the polynomial equation

$$\left(e^{\beta^*}\right)^{\mathfrak{A}} - \mathfrak{A}e^{\beta^*} - 1 = 0. \quad (25)$$

For $\mathfrak{A} \geq 5$ the solution of this equation cannot be expressed in closed form, however it can be expanded around $\mathfrak{A} \rightarrow \infty$ as

$$\beta^* = \frac{\ln \mathfrak{A}}{\mathfrak{A}} + \frac{1 + \ln \mathfrak{A}}{\mathfrak{A}^2} + \mathcal{O}\left(\frac{\ln \mathfrak{A}}{\mathfrak{A}^3}\right), \quad (26)$$

$$\Gamma'^* = \mathfrak{A} + \mathcal{O}\left(\frac{1}{\mathfrak{A}}\right), \quad (27)$$

$$\beta^* \Gamma'^* = \ln \mathfrak{A} + \frac{1 + \ln \mathfrak{A}}{\mathfrak{A}} + \mathcal{O}\left(\frac{\ln \mathfrak{A}}{\mathfrak{A}^2}\right). \quad (28)$$

As the number of alleles increases, accessibility increases and the required fitness difference between the start and end point decreases. In fact this quantity vanishes to zero for $\mathfrak{A} \rightarrow \infty$. At the same time the length of accessible walks close to the critical fitness difference increases, but slowly. The minimal length of a path covering the full distance d_{ab} is L , and hence $\beta^* \Gamma'^* - 1$ is the fraction of mutational reversions (where a mutated locus reverts to the allele it carried in the initial genotype a_i) and sideways steps (where a mutation occurs to an allele that is part of neither the initial nor the target genotype) [22]. The fraction of all alleles on a given locus that appear along an accessible path close to the critical point is given by $\mathfrak{A}^{-1} \beta^* \Gamma'^*$ which decreases with increasing \mathfrak{A} . Zargorski, Burda and Waclaw carried out simulations of this model, giving β^* with two digit precision for different values of \mathfrak{A} [23]. Their results match the values derived here up to ± 0.01 .

3.2 Complete graph without return to the wild type allele

We can modify the complete graph slightly to disallow mutations back to the allele that was present in the initial genotype (the wild type allele), while still allowing mutations between all other alleles. In this case the expressions

simplify significantly to

$$\Gamma_{ab}(\beta) = \ln \frac{e^{(\mathfrak{A}-2)\beta} - 1}{\mathfrak{A} - 2}, \quad (29)$$

$$\beta^* = \frac{\ln(\mathfrak{A} - 1)}{\mathfrak{A} - 2}, \quad (30)$$

$$\Gamma'^* = \mathfrak{A} - 1. \quad (31)$$

The asymptotic behavior for large \mathfrak{A} is the same as for the complete graph. For $\mathfrak{A} = 2$, the expressions are ill-defined, but the correct expressions coincide with the limits:

$$\Gamma_{ab}(\beta) = \ln \beta, \quad (32)$$

$$\beta^* = 1, \quad (33)$$

$$\Gamma'^* = 1. \quad (34)$$

In the diallelic case $\mathfrak{A} = 2$ this describes accessibility percolation on the directed hypercube, which was considered by Hegarty and Martinsson in [9]. In this case $\beta^* = 1$, which implies that the directed hypercube is marginally accessible under the HoC model [7]. For the diallelic case not only the critical value, but also the leading order corrections in the threshold function are known [9] and coincide with the order $\frac{\ln L}{L}$ contribution in our candidate threshold function and the value of Γ'^* given above.

3.3 Path graph

The complete graph is in some sense the best-case scenario for accessibility. On the opposite side of the spectrum of possible (undirected) allele graphs one can choose the path graph on \mathfrak{A} vertices. In this case the distance between the two end points increases linearly with the number of alleles and there is a unique order in which mutations on a locus must be applied. This causes accessibility to become very low. For $\mathfrak{A} = 2$ the path graph is identical to the complete graph. However, already for $\mathfrak{A} = 3$ we find

$$\beta^* = \frac{\ln(3 + 2\sqrt{2})}{\sqrt{2}} \approx 1.25. \quad (35)$$

Since β^* represents a fitness quantile which must lie between 0 and 1, this value implies that the path graph on three vertices can never be accessible for any fitness difference if (almost) all loci need to mutate from one end of the graph to the other. For higher \mathfrak{A} this effect becomes more pronounced. As a possible biological application of the path graph the description of copy-number variants of genes can be mentioned [2].

Since the complete graph on two vertices without reversions has $\beta^* = 1$ as shown before in (30) and adding edges can only decrease β^* , it is actually required that the distance between a_i and b_i on the allele graph is at least 2 in order for $\beta^* > 1$ to be possible.

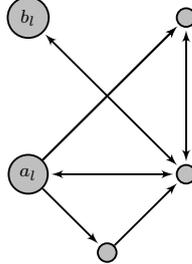


Fig. 5 Example of an allele graph that leads to a problem setup of irregular type with $\beta^* < 1$.

3.4 Example of non-trivial irregular type

Many graphs seem to be of completely regular type in the sense that no matter which sequence of pairs (\mathbf{a}, \mathbf{b}) are chosen, the problem is always of regular type. Martinsson [17] considered different sufficient conditions on graphs to have this property. But he also lists the smallest graph, of order 4, which does not have it. While this example demonstrates that it is possible to have problems of irregular type, it can also be used to generate semi-regular, but not regular, problem types by carefully interpolating the matrix p_{vw} defined in (2) between a regular and irregular type pair of alleles.

However, the example shown by Martinsson turns out to have $\beta^* > 1$, which automatically implies asymptotic inaccessibility in the accessibility percolation context due to the defined range of $\beta = F_{\mathbf{b}} - F_{\mathbf{a}}$ as a difference of uniform random variables. We therefore searched for the smallest graph without the regularity property and $\beta^* \leq 1$ numerically and found the example in Fig. 5 which has $\beta^* \approx 0.983$.

3.5 Genetic code

While the complete graph with $\mathfrak{A} = 4$ may serve as a model for the allele graph of single-nucleotide mutations on DNA or RNA, the expected effect of such a substitution depends significantly on whether or not it changes the amino acid that is encoded by the corresponding three-nucleotide codon. Mutations not affecting the encoded amino acids are known as synonymous. To specifically model the fitness effects of non-synonymous point mutations we therefore consider the allele graph of all amino acids with edges representing the mutual reachability by single-nucleotide substitutions (Figure 6).

This graph is considerably less symmetric than the complete graph and in particular the resulting quantities β^* and Γ'^* will depend on the particular choices of the path endpoints \mathbf{a} and \mathbf{b} rather than simply on their distance. We consider here all pairings of amino acids a_l and b_l , assuming them to be

equal for all loci. Other cases may be interpolated from these. The results are shown in Table 2. Whether the given values determine the asymptotic behavior of accessibility exactly depends on whether the regularity criteria relating to Martinsson's function (9) are satisfied. Due to the degree of the graph we limited ourselves to numerical tests, which did not indicate any violation of the criteria, although such violations may be more subtle than our tests could verify.

The critical point β^* and in particular the expected walk length $\beta^* \mathbf{I}^{L^*} L$ are, as one would expect, strongly correlated with the distance between alleles. The only distance-3 pair of amino acids is Tyr/Met which also corresponds to the largest walk length with a value of $\beta^* \mathbf{I}^{L^*} \approx 4.7567$. All other amino acids lie at mutual distance 1 or 2. Nonetheless, the critical point $\beta^* \approx 0.4527$ for the distance-3 pair lies slightly below that of the distance-2 pair Asp/Met with $\beta^* \approx 0.4570$, demonstrating that the overall structure of the allele graph can have a significant impact on accessibility beyond distance. For comparison, the accessibility of paths between any pair of codons can be obtained from the values for the complete graph with 4 alleles (Table 1). This gives $\beta^* \approx 0.51$, while accounting for the multiplication of three bases per codon yields the expected mean critical walk length per codon as $\beta^* \mathbf{I}^{L^*} \approx 5.5$.

4 First moment bound and walk length

We start with an upper bound for accessibility based on the mean number of accessible paths, or rather the mean number of *quasi-accessible* walks. We define the term *quasi-accessible* as a generalization of the notion of accessibility used up to now as described in the following.

4.1 Quasi-accessibility

In the original definition of accessibility, a non-self-avoiding walk is never accessible, because it would have to visit the same fitness value twice, which makes it impossible for the walk to have increasing fitness. Handling self-avoidance is non-trivial. To remedy this in a simpler manner, instead of considering self-avoiding paths on \mathcal{A}^L , we consider an extension of \mathcal{A}^L to $\mathcal{A}^{L'}$ as follows. The vertex set of $\mathcal{A}^{L'}$ is the set $\mathbb{A}^{L'} = \mathbb{A}^L \times \mathbb{N}$ and there is an arrow from (\mathbf{v}, n) to (\mathbf{w}, m) iff there is an arrow from \mathbf{v} to \mathbf{w} in \mathcal{A}^L . In other words we duplicate every genotype a countable infinite number of times in such a way that traversal of one of its copies can always be replaced by traversal of another copy. The 1-section containing vertices $(\mathbf{v}, 1)$ can be identified with the vertices on \mathcal{A}^L . We then assign each of the vertices in $\mathcal{A}^{L'}$ i.i.d. fitness values. The mentioned 1-section then corresponds to the original HoC problem. All other fitness values do not affect this underlying model. However, it is convenient to introduce these additional fitness values for the following reasons.

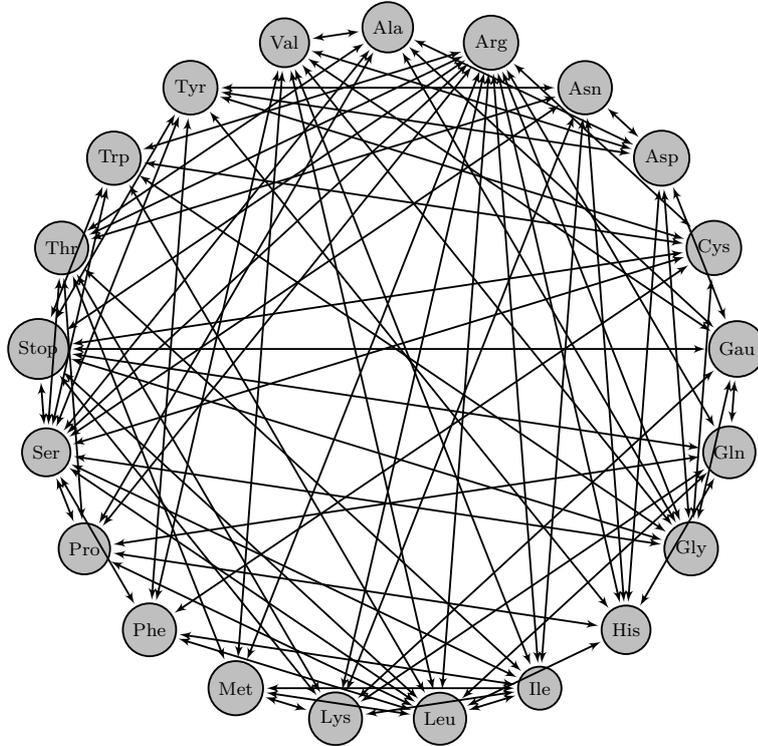


Fig. 6 Allele graph constructed from possible point-mutations on codons. Two amino acids are connected by an arrow iff there is a possible point mutation on a single nucleotide that changes one into the other.

We define the following map of walks on \mathcal{A}^L to $\mathcal{A}^{L'}$. Each self-avoiding walk is mapped to the corresponding walk on the 1-section of $\mathcal{A}^{L'}$. But instead of mapping non-selfavoiding walks from \mathcal{A}^L to the 1-section of $\mathcal{A}^{L'}$, we can make use of the additional vertex copies to replace all vertices that are visited multiple times in \mathcal{A}^L with distinct copies in $\mathcal{A}^{L'}$. To make this unique, we assume that the n -th visit of vertex v in \mathcal{A}^L is mapped to the vertex (v, n) in $\mathcal{A}^{L'}$, except if v is the final vertex of the walk, in which case we map the n -th visit in reverse order to (v, n) in $\mathcal{A}^{L'}$. The resulting walk is always selfavoiding in $\mathcal{A}^{L'}$ and the special case assures that every walk in \mathcal{A}^L is mapped to a walk with endpoints on the 1-section in $\mathcal{A}^{L'}$. A walk on \mathcal{A}^L is then said to be *quasi-accessible* if its mapped walk on $\mathcal{A}^{L'}$ is accessible. We say that a walk on $\mathcal{A}^{L'}$ is *valid* if there is a walk in \mathcal{A}^L mapped to it and define *valid-accessibility* as accessibility via valid walks. This guarantees that *all* walks of equal length in

	Phe	Leu	Ile	Met	Val	Ser	Pro	Thr	Ala	Tyr	Stop	His	Gln	Asn	Lys	Asp	Gau	Cys	Trp	Arg	Gly		
Leu	0.34																						
Ile	0.35	0.30																					
Met	0.42	0.33	0.33																				
Val	0.37	0.33	0.34	0.37																			
Ser	0.32	0.28	0.30	0.36	0.35																		
Pro	0.42	0.32	0.37	0.41	0.41	0.31																	
Thr	0.41	0.35	0.32	0.35	0.40	0.30	0.34																
Ala	0.43	0.38	0.39	0.43	0.36	0.33	0.37	0.36															
Tyr	0.37	0.37	0.39	0.45	0.42	0.32	0.42	0.41	0.43														
Stop	0.38	0.30	0.35	0.39	0.38	0.29	0.37	0.37	0.38	0.34													
His	0.43	0.33	0.39	0.43	0.42	0.35	0.35	0.40	0.43	0.37	0.38												
Gln	0.44	0.32	0.38	0.41	0.42	0.35	0.35	0.39	0.42	0.42	0.33	0.36											
Asn	0.42	0.38	0.34	0.42	0.42	0.32	0.41	0.35	0.42	0.36	0.39	0.37	0.42										
Lys	0.43	0.35	0.33	0.35	0.40	0.34	0.39	0.33	0.41	0.42	0.33	0.40	0.35	0.36									
Asp	0.44	0.40	0.41	0.46	0.36	0.37	0.43	0.42	0.36	0.38	0.40	0.38	0.43	0.38	0.43								
Gau	0.45	0.38	0.41	0.43	0.36	0.37	0.42	0.42	0.36	0.43	0.34	0.43	0.37	0.43	0.37	0.37							
Cys	0.36	0.35	0.38	0.43	0.41	0.30	0.41	0.40	0.42	0.35	0.31	0.42	0.41	0.42	0.41	0.43	0.42						
Trp	0.41	0.33	0.39	0.42	0.41	0.31	0.40	0.40	0.42	0.41	0.32	0.42	0.40	0.43	0.41	0.44	0.42	0.34					
Arg	0.36	0.27	0.29	0.32	0.36	0.27	0.30	0.30	0.36	0.36	0.28	0.32	0.31	0.35	0.30	0.39	0.36	0.31	0.31				
Gly	0.40	0.35	0.37	0.41	0.34	0.30	0.39	0.38	0.34	0.39	0.31	0.40	0.39	0.40	0.38	0.35	0.34	0.33	0.33	0.30			

	Phe	Leu	Ile	Met	Val	Ser	Pro	Thr	Ala	Tyr	Stop	His	Gln	Asn	Lys	Asp	Gau	Cys	Trp	Arg	Gly		
Leu	2.68																						
Ile	2.65	2.60																					
Met	3.91	2.68	2.48																				
Val	2.70	2.71	2.67	2.71																			
Set	2.62	2.61	2.63	3.66	3.54																		
Pro	4.13	2.63	3.66	3.85	4.08	2.66																	
Thr	4.07	3.52	2.53	2.62	3.84	2.61	2.66																
Ala	4.14	3.85	3.86	4.13	2.58	2.72	2.74	2.74															
Tyr	2.68	3.70	3.87	3.76	4.10	2.62	4.10	4.12	4.10														
Stop	3.69	2.68	3.68	3.89	3.84	2.59	3.67	3.84	3.85	2.76													
His	4.13	2.69	3.85	4.13	4.13	3.54	2.62	3.85	4.09	2.69	3.68												
Gln	4.43	2.62	3.85	3.85	4.11	3.67	2.64	3.84	4.11	4.13	2.61	2.61											
Asn	3.89	3.83	2.65	3.86	4.06	2.71	3.85	2.62	3.88	2.58	3.88	2.69	4.09										
Lys	4.34	3.52	2.60	2.61	3.86	3.55	3.81	2.58	4.07	4.15	2.71	3.85	2.64	2.72									
Asp	4.14	4.08	4.11	4.53	2.56	3.68	4.13	4.13	2.53	2.65	3.86	2.68	4.15	2.66	4.14								
Gau	4.47	3.87	4.10	4.15	2.58	3.83	4.11	4.11	2.56	4.14	2.67	4.10	2.71	4.14	2.73	2.52							
Cys	2.71	3.53	3.87	4.42	4.10	2.53	4.08	4.11	4.10	2.61	2.54	4.10	4.09	4.11	4.12	4.13	4.11						
Trp	3.85	2.71	3.88	4.16	4.11	2.58	3.87	4.09	4.12	3.84	2.54	4.12	3.86	4.41	4.08	4.44	4.10	2.54					
Arg	3.67	2.51	2.63	2.64	3.68	2.54	2.59	2.61	3.69	3.68	2.54	2.72	2.58	3.55	2.60	4.05	3.68	2.67	2.59				
Gly	3.84	3.55	3.86	4.12	2.65	2.62	3.86	3.88	2.55	3.66	2.58	4.07	3.85	4.08	3.87	2.61	2.54	2.59	2.58	2.70			

Table 2 Top: Critical fitness difference β^* for accessible paths between homopolymer amino acid sequences consisting of the indicated pairs. Bottom: Value of $\Gamma^{i*} \beta^*$ representing the expected per-locus length of accessible walks at the critical point. All values are obtained numerically and rounded to two digits. Colors represent the magnitude of the displayed values from low (yellow) to high (red).

\mathcal{A}^L are equally likely to be quasi-accessible and that valid-accessibility on $\mathcal{A}^{L'}$ coincides with quasi-accessibility on \mathcal{A}^L . Additionally, while quasi-accessibility is different from accessibility for individual walks on \mathcal{A}^L , accessibility and quasi-accessibility of one genotype from another on \mathcal{A}^L coincide, because non-selfavoiding walks are never accessible, but for each quasi-accessible walk there is a selfavoiding accessible walk obtainable by removal of all cycles from the walk or equivalently by restriction of the walk to the 1-section of $\mathcal{A}^{L'}$. This implies that we can restrict our investigation to quasi-accessible walks. We denote the number of valid-accessible walks from $(v, 1) \in \mathbb{A}^{L'}$ to $(w, 1) \in \mathbb{A}^{L'}$, or equivalently quasi-accessible walks from v to w with $v, w \in \mathbb{A}^L$, by \tilde{Z}_{vw} .

4.2 Upper bound

In order to give an upper bound on quasi-accessibility, we will consider the mean number of quasi-accessible walks from \mathbf{a} to \mathbf{b} . We condition here on the difference $F_{\mathbf{b}} - F_{\mathbf{a}}$ being β . Each walk of length N from \mathbf{a} to \mathbf{b} on \mathcal{A}^L is then quasi-accessible with probability

$$\frac{\beta^{N-1}}{(N-1)!}$$

where the numerator accounts for the probability that all inner vertices of the walk are found inside the range of fitness values $F_{\mathbf{a}}$ to $F_{\mathbf{b}}$ and the denominator accounts for the increasing order required on these values. The number of walks taking n steps from a_l to b_l on one locus l is given by $(A^n)_{a_l b_l}$. A walk of length N could take each step on any of the loci, so that the total number of walks of length N can be written as

$$\sum_{n_1+\dots+n_L=N} \binom{N}{n_1, \dots, n_L} \prod_{l=1}^L (A^{n_l})_{a_l b_l} \quad (36)$$

where $\binom{N}{n_1, \dots, n_L}$ is the multinomial coefficient accounting for the different orderings of steps on individual loci. Multiplication of this expression with the probability of quasi-accessibility of each such walk gives the mean number of quasi-accessible paths

$$\mathbb{E}[\tilde{\mathbf{Z}}_{\mathbf{ab}}] = \sum_{N=0}^{\infty} \sum_{n_1+\dots+n_L=N} \binom{N}{n_1, \dots, n_L} \frac{\beta^{N-1}}{(N-1)!} \prod_{l=1}^L (A^{n_l})_{a_l b_l}. \quad (37)$$

The term $(N-1)!$ can be reduced to $N!$ by introduction of a derivative

$$\mathbb{E}[\tilde{\mathbf{Z}}_{\mathbf{ab}}] = \partial_{\beta} \sum_{N=0}^{\infty} \sum_{n_1+\dots+n_L=N} \binom{N}{n_1, \dots, n_L} \frac{\beta^N}{N!} \prod_{l=1}^L (A^{n_l})_{a_l b_l} \quad (38)$$

and redistributing all the factorials and β^N into the product yields

$$\mathbb{E}[\tilde{\mathbf{Z}}_{\mathbf{ab}}] = \partial_{\beta} \sum_{N=0}^{\infty} \sum_{n_1+\dots+n_L=N} \prod_{l=1}^L \frac{\beta^{n_l}}{n_l!} (A^{n_l})_{a_l b_l}. \quad (39)$$

Finally the sums and the product can be interchanged and

$$\mathbb{E}[\tilde{\mathbf{Z}}_{\mathbf{ab}}] = \partial_{\beta} \prod_{l=1}^L \sum_{n=0}^{\infty} \frac{\beta^n}{n!} (A^n)_{a_l b_l} = \partial_{\beta} \prod_{l=1}^L (e^{\beta A})_{a_l b_l} = \partial_{\beta} e^{L\Gamma_{\mathbf{ab}}(\beta)} \quad (40)$$

$$= L\Gamma'_{\mathbf{ab}}(\beta) e^{L\Gamma_{\mathbf{ab}}(\beta)}. \quad (41)$$

This function is monotonically increasing in β and by expansion around some $\beta_0 > 0$ we find

$$\mathbb{E} [\tilde{\mathbf{Z}}_{\mathbf{ab}}] = L(\mathbf{\Gamma}'_{\mathbf{ab}}(\beta_0) + \mathcal{O}(\beta - \beta_0))e^{L(\mathbf{\Gamma}_{\mathbf{ab}}(\beta_0) + \mathbf{\Gamma}'_{\mathbf{ab}}(\beta_0)(\beta - \beta_0) + \mathcal{O}((\beta - \beta_0)^2))}. \quad (42)$$

If $(\beta - \beta_0) = o(1)$ and $\mathbf{\Gamma}_{\mathbf{ab}}(\beta_0) < 0$, the mean falls to zero exponentially quickly. On the other hand for $\beta_0 = \hat{\beta}$, by definition $\mathbf{\Gamma}_{\mathbf{ab}}(\beta_0) = 0$, and the higher-order behavior of $(\beta - \beta_0)$ is relevant. In particular the expected number of quasi-accessible walks converges to any given constant $\eta > 0$ for all sequences

$$\beta = \hat{\beta} - \mathbf{\Gamma}'_{\mathbf{ab}}(\hat{\beta})^{-1} \frac{\ln L}{L} + \mathbf{\Gamma}'_{\mathbf{ab}}(\hat{\beta})^{-1} \frac{\ln \eta - \ln \mathbf{\Gamma}'_{\mathbf{ab}}(\hat{\beta})}{L} + o\left(\frac{1}{L}\right) \quad (43)$$

By subtracting an additional term $\frac{|\omega(1)|}{L}$ from β it decreases to zero and by application of Markov's inequality so does the (quasi-)accessibility of \mathbf{b} from \mathbf{a} , proving the upper bound on accessibility in our first main statement.

4.3 Walk length

The upper bound can be strengthened by considering intervals of walk lengths. Let h_N be the expected number of quasi-accessible walks of length N at some β . This number is an expectation value over realizations of fitness values, but in the following we will consider it as just a number indexed by some number N representing the walk length. Summation of all of these numbers then yields the total expected number of quasi-accessible walks which we calculated already above:

$$\mathbb{E} [\tilde{\mathbf{Z}}_{\mathbf{ab}}] = \partial_{\beta} e^{L\mathbf{\Gamma}_{\mathbf{ab}}(\beta)} = \sum_{N=1}^{\infty} h_N \quad (44)$$

We can interpret this as the value $\phi(1)$ of the function

$$\phi(z) = \partial_{z\beta} e^{L\mathbf{\Gamma}_{\mathbf{ab}}(z\beta)} = \sum_{N=1}^{\infty} h_N z^{N-1}. \quad (45)$$

This function can be viewed as an (ordinary) generating function for the sequence h_N shifted by one. The generating function here is not related to the probability distribution of fitness values, but is rather to be understood as simply a counting tool that separates the total expectation value into slots for different walks lengths using the additivity of the expectation value.

The effect of the derivative $\partial_{z\beta}$ in the generating function can be reversed by integration of each of the monomials, so that

$$\tilde{\phi}(z) = e^{L\mathbf{\Gamma}_{\mathbf{ab}}(z\beta)} = \sum_{N=1}^{\infty} \frac{\beta}{N} h_N z^N = \sum_{N=1}^{\infty} \tilde{h}_N z^N \quad (46)$$

is the generating function of the unshifted h_N multiplied by $\frac{\beta}{N}$, which is another sequence that we define as \tilde{h}_N . Normalizing $\tilde{\phi}(z)$ through division by $\tilde{\phi}(1) = e^{L\Gamma_{\mathbf{ab}}(\beta)}$ turns the generating function into a probability generating function over the parameter N as random variable and this allows us to apply theorems from probability theory. Again, this probability is not related to the distribution of fitness values, but is introduced here artificially as a counting tool. The integrated (probability) generating function factorizes over loci

$$\frac{\tilde{\phi}(z)}{\tilde{\phi}(1)} = \prod_{l=1}^L \frac{e^{\Gamma_{a_l b_l}(z\beta)}}{e^{\Gamma_{a_l b_l}(\beta)}} \quad (47)$$

and therefore the random variable N under the generating function's distribution can be written as a sum $N = \sum_{l=1}^L n_l$, where n_l are random variables with probability generating function

$$e^{\Gamma_{a_l b_l}(z\beta) - \Gamma_{a_l b_l}(\beta)} = \frac{\sum_{n_l=0}^{\infty} \frac{\beta^{n_l}}{n_l!} (A^{n_l})_{a_l b_l} z^{n_l}}{\sum_{n_l=0}^{\infty} \frac{\beta^{n_l}}{n_l!} (A^{n_l})_{a_l b_l}} \quad (48)$$

Because the degree of A is bounded by Δ , $(A^{n_l})_{a_l b_l} \leq \Delta^{n_l}$ and the tail of the distribution is dominated by an exponential. This bound is also independent of the chosen loci a_l and b_l and therefore, with β being bounded in L as well, the central limit theorem applies to the sum N . The mean $L\mu$ and variance $L\sigma^2$ of N under this distribution can be obtained from the first and second derivatives of the probability generating function as

$$\mu = \beta \Gamma'_{\mathbf{ab}}(\beta) \quad (49)$$

$$\sigma^2 = \beta \Gamma'_{\mathbf{ab}}(\beta) + \beta^2 \Gamma''_{\mathbf{ab}}(\beta), \quad (50)$$

and the central limit theorem implies that for constants $c > 0$:

$$\sum_{|N - \mu L| \geq c\sigma\sqrt{L}} \tilde{h}_N = 2e^{L\Gamma_{\mathbf{ab}}(\beta)} \Phi(-c)(1 + o(1)) \quad (51)$$

where Φ is the standard normal CDF. Since the sum's upper bound is asymptotic to μL , for all \tilde{h}_N terms appearing in the sum $\tilde{h}_N \geq \frac{\beta}{\mu L} h_N$, so that

$$\sum_{|N - \mu L| \geq c\sigma\sqrt{L}} h_N \leq 2 \frac{\mu L}{\beta} e^{L\Gamma_{\mathbf{ab}}(\beta)} \Phi(-c)(1 + o(1)). \quad (52)$$

In particular at the threshold function of the form of eq. (43) the bound converges:

$$\sum_{|N - \mu L| \geq c\sigma\sqrt{L}} h_N \leq 2\eta\Phi(-c)(1 + o(1)). \quad (53)$$

This allows one to reduce the mean number of quasi-accessible walks of length outside the interval $\mu L \pm c\sigma\sqrt{L}$ to any arbitrarily small value by choosing c large enough. In other words, if there are quasi-accessible walks at the suggested threshold function, then they are of length μL with fluctuations of at most order of \sqrt{L} . Since the total mean number of quasi-accessible walks at this threshold function is η , this then implies that the mean number of quasi-accessible walks inside the stated interval is converging to $\eta(1 - 2\Phi(-c))$ and for any $c > 0$ and $\eta > 0$ this mean still converges to a non-zero value.

5 Improved upper bound on accessibility

The upper bound on accessibility obtained from the expected value does not take into account any dependence between walks. We can improve the bound by including some of the dependencies. This will make it possible to show our second main statement that $\beta_c > \beta^*$ for irregular types.

Let in this section $0 < r < 1$ and $0 < s < 1$ be constant. The intention is to choose them later such that $\mathfrak{M}^*(s, r, \beta^*) > 0$ as application to the irregular type. This choice is always possible in the irregular case since by continuity $\mathfrak{M}^*(s, r, \beta^*)$ cannot be strictly positive only on the boundaries.

The interpretation of r and s is as determining the fitness spanned by three segments of each walk with s determining the fitness fraction spanned by the middle segment, and r determining the distribution of the remaining fitness span onto the first and last segment. More concretely the intended fitness span of the first segment is $\beta\bar{s}r$, of the second βs and the third $\beta\bar{s}\bar{r}$, adding up to the full fitness span β that needs to be crossed (see Sect. 2.2, eq. (16)).

For each edge on the genotype space we can consider the interval formed by the fitness values of the two nodes incident to it. If a walk from \mathbf{a} to \mathbf{b} is quasi-accessible, then it contains exactly one edge with a fitness interval containing the fitness value $\beta\bar{s}r$. Let this edge be $(\mathbf{x}, \mathbf{x}')$. Similarly there is exactly one edge containing the fitness value $\beta(1 - \bar{s}\bar{r})$. Let this edge be $(\mathbf{y}, \mathbf{y}')$. These two edges segment the walk in the closest possible way according to the intended fitness spans mentioned above. A walk is quasi-accessible only if each of the three segments $\mathbf{a} \rightarrow \mathbf{x}$, $\mathbf{x}' \rightarrow \mathbf{y}$ and $\mathbf{y}' \rightarrow \mathbf{b}$ are quasi-accessible. In the following we refer to these segments as segment 1, 2 and 3 respectively. To obtain an upper bound on the (quasi-)accessibility of \mathbf{b} from \mathbf{a} it is therefore sufficient to form a union bound over these edges:

$$\mathbb{P}[\tilde{\mathcal{Z}}_{ab} \geq 1] \leq \sum_{(\mathbf{x}, \mathbf{x}'), (\mathbf{y}, \mathbf{y}')} \mathbb{P}[\tilde{\mathcal{Z}}_{a\mathbf{x}} \geq 1 \wedge \tilde{\mathcal{Z}}_{\mathbf{x}'\mathbf{y}} \geq 1 \wedge \tilde{\mathcal{Z}}_{\mathbf{y}'\mathbf{b}} \geq 1]. \quad (54)$$

Here the sum is over pairs of edges of \mathcal{A}^L and the probability is assumed to be implicitly conditioned on the edges containing the fitness values as mentioned above. It is not necessary to sum over all allowed edges in $\mathcal{A}^{L'}$, since quasi-accessibility and accessibility coincide. Effectively we ignore edges in $\mathcal{A}^{L'}$ only belonging to non-selfavoiding walks in \mathcal{A}^L when forming the sum.

In the following we condition the probability on the fitness values $F_{\mathbf{x}}$, $F_{\mathbf{x}'}$, $F_{\mathbf{y}}$ and $F_{\mathbf{y}'}$ to be fixed. The total probability is then the expectation over these conditioned probabilities. As a result of this conditioning, the quasi-accessibilities of the three segments mentioned in the equation are negatively dependent. For example if $\tilde{Z}_{\mathbf{a}\mathbf{x}}$ is at least 1, then $\tilde{Z}_{\mathbf{x}'\mathbf{y}} \geq 1$ becomes less likely since the existence of a quasi-accessible walk from \mathbf{a} to \mathbf{x} implies that some fitness values of other genotypes fall in the range $[F_{\mathbf{a}}, F_{\mathbf{x}}]$, excluding them for consideration in the range $[F_{\mathbf{x}'}, F_{\mathbf{y}}]$ required for them to be part of a quasi-accessible walk from \mathbf{x}' to \mathbf{y} . Consequently:

$$\mathbb{P}[\tilde{Z}_{\mathbf{a}\mathbf{b}} \geq 1] \leq \sum_{(\mathbf{x}, \mathbf{x}'), (\mathbf{y}, \mathbf{y}')} \mathbb{E}[\mathbb{P}[\tilde{Z}_{\mathbf{a}\mathbf{x}} \geq 1 | \cdot] \mathbb{P}[\tilde{Z}_{\mathbf{x}'\mathbf{y}} \geq 1 | \cdot] \mathbb{P}[\tilde{Z}_{\mathbf{y}'\mathbf{b}} \geq 1 | \cdot]] \quad (55)$$

Here the conditioning on the fitness values incident to the chosen edges is implied by \cdot and the outer expectation value is over these values. Since probabilities lie in $[0, 1]$, we can use the upper bound $x \leq x^{1-\alpha}$ for $0 < \alpha < 1$ on the middle factor and afterwards we can apply Markov's inequality to all three terms to obtain

$$\mathbb{P}[\tilde{Z}_{\mathbf{a}\mathbf{b}} \geq 1] \leq \sum_{(\mathbf{x}, \mathbf{x}'), (\mathbf{y}, \mathbf{y}')} \mathbb{E}[\mathbb{E}[\tilde{Z}_{\mathbf{a}\mathbf{x}} \geq 1 | \cdot] \mathbb{E}[\tilde{Z}_{\mathbf{x}'\mathbf{y}} \geq 1 | \cdot]^{1-\alpha} \mathbb{E}[\tilde{Z}_{\mathbf{y}'\mathbf{b}} \geq 1 | \cdot]] \quad (56)$$

The remaining inner expectation values depend only on the differences of the fitness values that they are conditioned on, not the actual placement of that difference. We introduce the following quantities:

$$\epsilon_1 = \beta \bar{s} r - (F_{\mathbf{x}} - F_{\mathbf{a}}) \quad (57)$$

$$\epsilon_2 = \beta s - (F_{\mathbf{y}} - F_{\mathbf{x}'}) \quad (58)$$

$$\epsilon_3 = \beta \bar{s} \bar{r} - (F_{\mathbf{b}} - F_{\mathbf{y}'}) \quad (59)$$

$$(60)$$

These quantities measure how much the fitness difference allocated to one of the three walk segments differs from what it would be assigned if r and s determined it exactly. For example the $(\mathbf{x}, \mathbf{x}')$ edge is required to contain the fitness value $\beta \bar{s} r$. Therefore the first walk segment can span a fitness distance of at most $\beta \bar{s} r$, but this happens exactly only if $F_{\mathbf{x}} - F_{\mathbf{a}} = \beta \bar{s} r$ is chosen. All other valid choices set the fitness value lower than this and ϵ_1 measures the reduction of the segment's length. As it will turn out only the point with all epsilons equal to zero contributes to the expectation value in leading order. Intuitively any constant offset from the intended segment length corresponds to an effective reduction of β by a constant, resulting in an exponentially lower likelihood of walks being quasi-accessible. Nonetheless we will carry the epsilons through the calculation.

The remaining inner expectation values are of the same form as the simple expectation of walks from \mathbf{a} to \mathbf{b} calculated in the previous section:

$$\mathbb{E}[\tilde{Z}_{\mathbf{v}\mathbf{w}} | t] = \partial_t e^{L\Gamma_{\mathbf{v}\mathbf{w}}(t)} = \Gamma'_{\mathbf{v}\mathbf{w}}(t) L e^{L\Gamma_{\mathbf{v}\mathbf{w}}(t)}. \quad (61)$$

The expectation values are dominated by the exponential terms $e^{L\Gamma_{\mathbf{vw}}(t)}$ with an additional linear factor L resulting from the derivative. However, the derivative also adds the term $\Gamma'_{\mathbf{vw}}(t)$. As an average over loci it can be seen that pointwise in t and uniformly over \mathbf{v} and \mathbf{w} , this quantity is bounded by a constant from above. However the bound is not uniform in t . At $t = 0$ it diverges, as can be seen from the expansion

$$L\Gamma'_{\mathbf{vw}}(t) \sim \frac{d_{\mathbf{vw}}}{t}. \quad (62)$$

To avoid this issue we rewrite the expectation value including the sum resulting from application of the product rule of differentiation

$$\mathbb{E}[\tilde{\mathcal{Z}}_{\mathbf{vw}}|t] = \sum_{l'=1}^L (Ae^{tA})_{v_{l'}w_{l'}} \prod_{l \neq l'} e^{\Gamma_{v_l w_l}(t)} \quad (63)$$

In each summand the value is a product over terms, each of which depends only on quantities on a single locus and the bulk of the contributions of loci contribute simply the exponential $e^{\Gamma_{v_l w_l}(t)} = (e^{tA})_{v_l w_l}$. Only the locus l' gives a different contribution, namely the derivative of the exponential term, $(Ae^{tA})_{v_{l'} w_{l'}}$.

Our goal is to bring eq. (56) into the form of a sum over products, such that the product factorizes in the same sense as it does for a single expectation value. In particular the current form is a sum of a product of three expectation values. If we expand each expectation as shown in eq. (63), we obtain three sums, each accounting for one special locus on which the corresponding derivative is taken. We name these special loci l_1 , l_2 and l_3 , corresponding to the means in eq. (56) in the order they appear there. The sum in the middle term can be taken out of the $(\cdot)^{1-\alpha}$ form to give an upper bound, because $1 - \alpha \in [0, 1]$ and therefore the form is subadditive. Having done so, the sum over the pair of edges on the genotype space may similarly be factorized over loci. Each edge on the genotype graph corresponds to a step on one locus. Therefore it is sufficient to sum over individual genotypes together with another special locus, and one edge on the allele graph corresponding to that locus. We denote the sum over loci for these two edges l_{12} and l_{23} respectively. The initial sum then factorizes over loci:

$$\mathbb{P}[\tilde{\mathcal{Z}}_{\mathbf{ab}} \geq 1] \leq \sum_{l_1, l_2, l_3, l_{12}, l_{23}=1}^L \mathbb{E} \left[\prod_{l=1}^L \mathfrak{F}_l \right] \quad (64)$$

Here \mathfrak{F}_l is the resulting factor collecting all sums over quantities on locus l and all factors of the product of the three expectations that are functions of quantities on locus l , as well as potentially e.g. a form $\sum_{x_{l'}} (A)_{x_l x_{l'}}$ if $l = l_{12}$. \mathfrak{F}_l is implicitly dependent on l_1, l_2, l_3 , since these three variables decide whether the contribution resulting from any of the three expectation values has the usual exponential form or that of its derivative. If $l \notin \{l_1, l_2, l_3, l_{12}, l_{23}\}$, then

the contribution of all three expectation values and the edge sum is of the usual form, i.e. the exponential term of the expectation value and no sum over edges and we give it the name \mathfrak{G}_l :

$$\mathfrak{G}_l = \sum_{x_l, y_l} e^{\Gamma_{a_l x_l}(\beta \bar{s} r - \epsilon_1) + (1-\alpha)\Gamma_{x_l y_l}(\beta s - \epsilon_2) + \Gamma_{y_l b_l}(\beta \bar{s} r - \epsilon_3)} \quad (65)$$

$$= \sum_{x_l, y_l} \left(e^{(\beta \bar{s} r - \epsilon_1) A} \right)_{a_l x_l} \left(\left(e^{(\beta s - \epsilon_2) A} \right)_{x_l y_l} \right)^{1-\alpha} \left(e^{(\beta \bar{s} r - \epsilon_3) A} \right)_{y_l b_l}. \quad (66)$$

If l is equal to any of the set of special loci, then some of these exponential terms will be modified and there might be additional sums. For example if l is equal to l_3 and l_{12} , but not equal to any of the other special loci, then

$$\mathfrak{F}_l = \sum_{x_l, y_l, x_l'} \left(e^{(\beta \bar{s} r - \epsilon_1) A} \right)_{a_l x_l} (A)_{x_l x_l'} \left(\left(e^{(\beta s - \epsilon_2) A} \right)_{x_l' y_l} \right)^{1-\alpha} \left(A e^{(\beta \bar{s} r - \epsilon_3) A} \right)_{y_l b_l}. \quad (67)$$

By assumption $s \neq 0$ and also $0 \neq r \neq 1$. Then, with the fixed s and r , \mathfrak{G}_l is bounded away from zero everywhere except at the boundary, because as long as one of the matrix exponentials has non-zero argument, it contributes a finite term to the sum by adequate choice of x_l and y_l so that the indices of the matrix exponential become (a_l, b_l) . More generally \mathfrak{G}_l is also uniformly bounded over s and r , since by definition of s and r at least one of the matrix exponential arguments must be at least $\frac{\beta}{3}$, epsilon shifts notwithstanding. This then allows us to write each \mathfrak{F}_l as a product $\mathfrak{G}_l \mathfrak{H}_l$, with \mathfrak{H}_l bounded away from infinity except at the mentioned boundary. In the next section we will use the same approach with a more detailed handling of \mathfrak{H}_l , but here it is sufficient to apply such a simple uniform bound with a constant.

However first we consider the behavior at the boundary where all epsilon shifts force the matrix exponential arguments to become zero. Due to the bounded degree of the graph, as $t \rightarrow 0$, the diagonal terms of the matrix exponential with argument t drop to 1 and the off-diagonal ones to 0 uniformly. If $a_l \neq b_l$, \mathfrak{F}_l therefore falls to zero as all the epsilons reach their maximum boundary and similarly it falls to 1 for $a_l = b_l$. Since there is by assumption at least a finite fraction of loci with $a_l \neq b_l$, this then implies that eventually, at a finite distance to the boundary

$$\prod_{l=1}^L \mathfrak{F}_l \leq \mathcal{O}(C^L) \quad (68)$$

for some $C < 1$. The special loci on which $\mathfrak{F}_l \neq \mathfrak{G}_l$ are not relevant to this, since there are only finitely many of them and each one is bounded. The contribution to the probability from the boundary is therefore asymptotically zero, since the exponential decay in the integrand cannot be compensated by the additional L^5 factor from the special loci sum.

Returning to the general case away from the boundary, we can bound all \mathfrak{H}_l with $l \in \{l_1, l_2, l_3, l_{12}, l_{13}\}$ by some constant C uniformly, yielding a factor of at most C^5 , while all other \mathfrak{H}_l are 1. This removes the dependence of the product on the particular choice of the special loci:

$$\mathbb{P}[\tilde{\mathcal{Z}}_{ab} \geq 1] \leq L^5 C^5 \mathbb{E} \left[\prod_{l=1}^L \mathfrak{G}_l \right] + o(1) = L^5 C^5 \mathbb{E} [e^{L\mathfrak{X}}] \quad (69)$$

with

$$\mathfrak{X} = \left\langle \ln \sum_{x_l, y_l} e^{\Gamma_{a_l x_l}(\beta \bar{s} r - \epsilon_1) + (1-\alpha) \Gamma_{x_l y_l}(\beta s - \epsilon_2) + \Gamma_{y_l b_l}(\beta \bar{s} r - \epsilon_3)} \right\rangle_l \quad (70)$$

The epsilons are always non-negative in the valid domain and \mathfrak{X} is decreasing in all of them. Therefore we can give an upper bound by setting all of them to 0 and obtain the upper bound on accessibility:

$$\mathbb{P}[\tilde{\mathcal{Z}}_{ab} \geq 1] \leq e^{L(\mathfrak{X}_0 + o(1))} \quad (71)$$

with

$$\mathfrak{X}_0 = \left\langle \ln \sum_{x_l, y_l} e^{\Gamma_{a_l x_l}(\beta \bar{s} r) + (1-\alpha) \Gamma_{x_l y_l}(\beta s) + \Gamma_{y_l b_l}(\beta \bar{s} r)} \right\rangle_{a_l, b_l}^p \quad (72)$$

where it is assumed that β is constant in L . This value is then independent of L and if it is negative, the probability that \mathbf{b} is accessible from \mathbf{a} is asymptotically exponentially falling to zero. We may choose $\alpha \in (0, 1)$ as well as s and r freely except for their boundary values. But specifically for α close to zero, we obtain the following expansion by matrix multiplication:

$$\mathfrak{X}_0 = \langle \Gamma_{a_l b_l}(\beta) \rangle_{a_l, b_l}^p - \alpha \langle \langle \Gamma_{x_l y_l}(\beta s) \rangle_{x_l, y_l}^{s, r} \rangle_{a_l, b_l}^p + \mathcal{O}(\alpha^2) \quad (73)$$

At β^* the zeroth order term is simply zero. The coefficient of the linear order term is exactly $\mathfrak{M}^*(s, r, \beta)$ and in the irregular type problem r and s can be chosen such that it is negative at β^* . With this choice there is then some suitable small $\alpha > 0$, so that \mathfrak{X}_0 is negative at β^* . \mathfrak{X}_0 is continuous as a function of β and therefore we can then also find some $\beta > \beta^*$ such that \mathfrak{X}_0 is still negative at the same choice of α , r and s . This shows that the critical point β_c is strictly larger than β^* in the irregular case, if it exists at all.

6 Lower bound on accessibility

In this section we derive a lower bound on accessibility, allowing us to show that the candidate threshold function found for the regular type from the first moment method is indeed a $\frac{1}{L}$ -threshold function.

6.1 Moment bounds

To prove the lower bound on (quasi-)accessibility, we use a generalization of the second moment method. The idea of the second moment method is to bound the second moment of \tilde{Z}_{ab} from above in order to apply the inequality [1, 9]

$$\mathbb{P}[\tilde{Z}_{ab} > 0] \geq \frac{\mathbb{E}[\tilde{Z}_{ab}]^2}{\mathbb{E}[\tilde{Z}_{ab}^2]}. \quad (74)$$

\tilde{Z}_{ab} is bounded from above through the maximum length of walks and the bounded degree limiting the possible choices in each step and therefore the second moment always exists. In our proof method we do however find that, at least with our non-tight bounds on it, the second moment grows too quickly for some allele graphs to give a non-trivial bound. On the other hand, for some class of allele graphs this bound may be used to obtain a sufficient bound.

To generalize the applicability of the result, we will use a modification of the second moment method which relies on a lower order moment. In particular we know from Hölder's inequality that for all $\xi > 0$:

$$\mathbb{E}[\tilde{Z}_{ab}] = \mathbb{E}[\tilde{Z}_{ab} \mathbb{I}_{\tilde{Z}_{ab} \geq 1}] \quad (75)$$

$$\leq \mathbb{E}[\tilde{Z}_{ab}^{1+\xi}]^{\frac{1}{1+\xi}} \mathbb{E}\left[\mathbb{I}_{\tilde{Z}_{ab} \geq 1}^{\frac{1}{1-\frac{1}{1+\xi}}}\right]^{1-\frac{1}{1+\xi}} \quad (76)$$

$$= \mathbb{E}[\tilde{Z}_{ab}^{1+\xi}]^{\frac{1}{1+\xi}} \mathbb{P}[\tilde{Z}_{ab} \geq 1]^{\frac{\xi}{1+\xi}} \quad (77)$$

and therefore

$$\mathbb{P}[\tilde{Z}_{ab} \geq 1] \geq \left(\frac{\mathbb{E}[\tilde{Z}_{ab}]^{1+\xi}}{\mathbb{E}[\tilde{Z}_{ab}^{1+\xi}]}\right)^{\frac{1}{\xi}}. \quad (78)$$

Because the number of walks of length N is at most exponential due to the bounded degree of \mathcal{A} , while the probability of a walk to be quasi-accessible falls as fast as $\frac{1}{N!}$, the tail of \tilde{Z}_{ab} is dominated by an exponential decay. In particular all moments of \tilde{Z}_{ab} exist. This allows us to take the limit $\xi \rightarrow 0$, dropping all higher order terms:

$$\mathbb{P}[\tilde{Z}_{ab} \geq 1] \geq \mathbb{E}[\tilde{Z}_{ab}] e^{-\mathfrak{R}} \quad (79)$$

where

$$\mathfrak{R} = \mathbb{E}\left[\frac{\tilde{Z}_{ab}}{\mathbb{E}[\tilde{Z}_{ab}]} \ln \tilde{Z}_{ab}\right]. \quad (80)$$

Our goal here is to show that $\liminf \mathbb{P}[\tilde{Z}_{ab} \geq 1] > 0$, i.e. that there is at least a non-vanishing probability of \mathbf{b} being accessible from \mathbf{a} asymptotically.

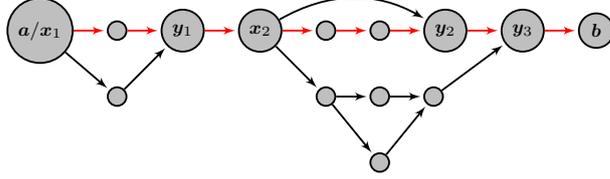


Fig. 7 Illustration of the estimation of \tilde{Z}_{ab} . The focal walk π from \mathbf{a} to \mathbf{b} is marked in red. Relative to π there is one non-trivial arc from \mathbf{x}_1 to \mathbf{y}_1 , one non-trivial arc from \mathbf{x}_2 to \mathbf{y}_2 and two non-trivial arcs from \mathbf{x}_2 to \mathbf{y}_3 . Between any two adjacent vertices on π there is furthermore one trivial arc. All other non-empty path subgraphs are not arcs, since they intersect π more than twice. Each walk is fully specified by choice of one of the arcs for each pair of loci. In fact it would be sufficient to choose one out-going arc per site to account for all walks. By the construction of quasi-accessibility the graph is guaranteed to be free cycles as in the example.

This is enough to show that the candidate threshold function is indeed a threshold function at which there is a transition from zero accessibility to non-zero accessibility, but it remains to be shown that the accessibility above the threshold is one. In any case we assume for this section that β is chosen as the threshold function c_L from our main statement, so that in particular $\mathbb{E}[\tilde{Z}_{ab}]$ converges to a non-zero value. This assures that it is sufficient to show that \mathfrak{R} does not diverge. The following method of bounding \mathfrak{R} adapts the idea used in [17] to account for the correlations of accessible walks using the notion of shortcuts or arcs to obtain alternative walks from a focal one.

Let X_π be the indicator variable that the walk π is quasi-accessible, then

$$\mathfrak{R} = \sum_{\pi} \mathbb{E} \left[\frac{X_\pi}{\mathbb{E}[\tilde{Z}_{ab}]} \ln \tilde{Z}_{ab} \right] \quad (81)$$

where the sum is over all walks from \mathbf{a} to \mathbf{b} on \mathcal{A}^L or equivalently all valid walks on $\mathcal{A}^{L'}$.

Similarly we can expand the right-hand \tilde{Z}_{ab} over individual walks. We will however intentionally over-count these in the following way: Each valid walk π' from $(\mathbf{a}, 1)$ to $(\mathbf{b}, 1)$ in $\mathcal{A}^{L'}$ trivially crosses π in at least two vertices, namely $(\mathbf{a}, 1)$ and $(\mathbf{b}, 1)$. Furthermore if we list out for each valid walk π' the vertices it shares with π in $\mathcal{A}^{L'}$, then the segment of π' between two adjacent vertices $\mathbf{x} \in \mathbb{A}^{L'}$ and $\mathbf{y} \in \mathbb{A}^{L'}$ in that list does not intersect π a third time in $\mathcal{A}^{L'}$. We call such a segment on $\mathcal{A}^{L'}$ an *arc* through \mathbf{x} and \mathbf{y} on π . An arc is said to be *trivial* if it is a segment of π itself. Immediately from the definition a trivial arc can only contain a single edge. We denote the number of non-trivial arcs which are accessible by $\tilde{Z}_{\mathbf{x}\mathbf{y}\pi}^L$. Each walk π' generates at most one arc through \mathbf{x} and \mathbf{y} on π . Also each walk π' is uniquely identified by the set of arcs it generates on π and π' is accessible on $\mathcal{A}^{L'}$ if and only if all of the arcs

it generates on π are accessible. Therefore we can bound for any valid walk π :

$$\tilde{Z}_{ab} \leq \prod_{\mathbf{x}, \mathbf{y} \in \mathcal{A}^{L'}} (1 + \tilde{Z}'_{\mathbf{x}\mathbf{y}\pi}). \quad (82)$$

With this we have

$$\hat{\mathfrak{R}} \leq \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{A}^{L'}} \sum_{\pi} I_{\mathbf{x}\mathbf{y}\pi} \mathbb{E} \left[\frac{X_{\pi}}{\mathbb{E}[\tilde{Z}_{ab}]} \ln(1 + \tilde{Z}'_{\mathbf{x}\mathbf{y}\pi}) \right] \quad (83)$$

where $I_{\mathbf{x}\mathbf{y}\pi}$ is an indicator variable which is 1 iff π contains \mathbf{x} and \mathbf{y} in $\mathcal{A}^{L'}$ and 0 otherwise.

Conditioned on the two fitness values $F_{\mathbf{x}}$ and $F_{\mathbf{y}}$, $\tilde{Z}'_{\mathbf{x}\mathbf{y}\pi}$ becomes independent of X_{π} since \mathbf{x} and \mathbf{y} are the only vertices whose fitness values influence both the quasi-accessibility of candidate arcs and π :

$$\hat{\mathfrak{R}} \leq \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{A}^{L'}} \sum_{\pi} I_{\mathbf{x}\mathbf{y}\pi} \mathbb{E} \left[\frac{X_{\pi}}{\mathbb{E}[\tilde{Z}_{ab}]} | F_{\mathbf{x}}, F_{\mathbf{y}} \right] \mathbb{E} \left[\ln(1 + \tilde{Z}'_{\mathbf{x}\mathbf{y}\pi} | F_{\mathbf{x}}, F_{\mathbf{y}}) \right] \quad (84)$$

For convenience we also assume that the conditioning of the fitness values $F_{(a,1)}$ and $F_{(b,1)}$ to a difference of β is contained in the outer expectation.

Currently $\tilde{Z}'_{\mathbf{x}\mathbf{y}\pi}$ is stochastically independent of X_{π} , but still explicitly dependent on π in the choice of candidate arcs that need to be counted. We can remove this dependence by loosening the restriction that included arcs must not be trivial and must not intersect π except at \mathbf{x} and \mathbf{y} . Doing so $\tilde{Z}'_{\mathbf{x}\mathbf{y}\pi}$ is upper bounded by $\tilde{Z}'_{\mathbf{x}\mathbf{y}}$, where the lack of third index indicates the loosened restriction. The resulting bound is not in general good enough for all choices of \mathbf{x} and \mathbf{y} in the sum. We will later revisit and adjust it for these special cases.

Because the logarithm is concave, the mean over it can be bounded by exchange of the two. Let $\downarrow \mathbf{x}$ be the projection of $\mathbf{x} \in \mathcal{A}^{L'}$ on the first component or equivalently the 1-section of $\mathcal{A}^{L'}$. Compared to all walks on $\mathcal{A}^{L'}$ generated from walks on \mathcal{A}^L from $\downarrow \mathbf{x}$ to $\downarrow \mathbf{y}$, arcs from \mathbf{x} to \mathbf{y} in $\mathcal{A}^{L'}$ are more restricted in the number of times vertices with projection $\downarrow \mathbf{x}$ or $\downarrow \mathbf{y}$ may or must be visited. Therefore the expectation over $\tilde{Z}'_{\mathbf{x}\mathbf{y}}$ may be bounded by the expectation over $\tilde{Z}_{\downarrow \mathbf{x} \downarrow \mathbf{y}}$.

$$\hat{\mathfrak{R}} \leq \mathbb{E} \left[\sum_{\mathbf{x}, \mathbf{y} \in \mathcal{A}^{L'}} \mathbb{E} \left[\frac{\sum_{\pi} I_{\mathbf{x}\mathbf{y}\pi} X_{\pi}}{\mathbb{E}[\tilde{Z}_{ab}]} | F_{\mathbf{x}}, F_{\mathbf{y}} \right] \ln \mathbb{E} [1 + \tilde{Z}_{\downarrow \mathbf{x} \downarrow \mathbf{y}} | F_{(\downarrow \mathbf{x}, 1)}, F_{(\downarrow \mathbf{y}, 1)}] \right] \quad (85)$$

Similarly all walks π through \mathbf{x} and \mathbf{y} can be separated into three segments from $(a, 1)$ to \mathbf{x} , from \mathbf{x} to \mathbf{y} and from \mathbf{y} to $(b, 1)$. Each walk is uniquely determined by these three segments and for any choice of these segments forming a valid walk, their accessibility is independent under the conditioning since valid

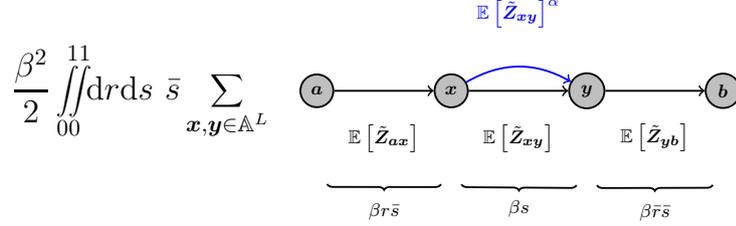


Fig. 8 Graphical representation of \mathfrak{K} 's bound given in eq. (87). The fitness range β is split into three segments corresponding to paths $\mathbf{a} \rightarrow \mathbf{x}$, $\mathbf{x} \rightarrow \mathbf{y}$ and $\mathbf{y} \rightarrow \mathbf{b}$, indicated by black arrows. Each segment contributes the given expectation value conditioned on the specified fitness difference between its endpoints into a product. The blue line represents the contribution of all arcs from \mathbf{x} to \mathbf{y} , which contribute the given α -dependent factor.

walks are selfavoiding in $\mathcal{A}^{L'}$. Taking all triples of walk segments from \mathbf{a} to $\downarrow \mathbf{x}$, from $\downarrow \mathbf{x}$ to $\downarrow \mathbf{y}$ and from $\downarrow \mathbf{y}$ to \mathbf{b} , all valid walks from \mathbf{a} to \mathbf{b} through any copy of the genotypes $\downarrow \mathbf{x}$ and $\downarrow \mathbf{y}$ in $\mathcal{A}^{L'}$ are generated. This allows together with the previous arguments for the bound

$$\mathfrak{K} \leq \mathbb{E} \left[\sum_{\mathbf{x}, \mathbf{y} \in \mathbb{A}^L} \frac{\mathbb{E}[\tilde{Z}_{\mathbf{a}\mathbf{x}} | F_{\mathbf{x}}, F_{\mathbf{y}}] \mathbb{E}[\tilde{Z}_{\mathbf{x}\mathbf{y}} | F_{\mathbf{x}}, F_{\mathbf{y}}] \mathbb{E}[\tilde{Z}_{\mathbf{y}\mathbf{b}} | F_{\mathbf{x}}, F_{\mathbf{y}}]}{\mathbb{E}[\tilde{Z}_{\mathbf{a}\mathbf{b}}]} \ln(1 + \mathbb{E}[\tilde{Z}_{\mathbf{x}\mathbf{y}} | F_{\mathbf{x}}, F_{\mathbf{y}}]) \right]. \quad (86)$$

Further we use that the logarithm can be bounded from above by any power law $\ln(x) \leq \bar{\alpha}(x-1)^\alpha$ for $x \geq 1$, $0 < \alpha < 1$ and a constant $\bar{\alpha}$ depending on α . In particular $\bar{\alpha}$ as a function of α can be chosen so that it is bounded except around $\alpha = 0$, where it must diverge. Therefore, as long as we choose later any non-zero but constant α , the additional factor $\bar{\alpha}$ will not change the asymptotic order of \mathfrak{K} . All in all:

$$\mathfrak{K} \leq \mathbb{E} \left[\bar{\alpha} \sum_{\mathbf{x}, \mathbf{y} \in \mathbb{A}^L} \frac{\mathbb{E}[\tilde{Z}_{\mathbf{a}\mathbf{x}} | F_{\mathbf{x}}, F_{\mathbf{y}}] \mathbb{E}[\tilde{Z}_{\mathbf{x}\mathbf{y}} | F_{\mathbf{x}}, F_{\mathbf{y}}]^{1+\alpha} \mathbb{E}[\tilde{Z}_{\mathbf{y}\mathbf{b}} | F_{\mathbf{x}}, F_{\mathbf{y}}]}{\mathbb{E}[\tilde{Z}_{\mathbf{a}\mathbf{b}}]} \right] \quad (87)$$

The remaining expectation values depend only on the differences of the fitness values that they are conditioned on, not the actual placement of that difference. It is therefore convenient to use the variables s and r with $\bar{s} = 1 - s$ and $\bar{r} = 1 - r$ introduced previously, such that

$$F_{\mathbf{x}} - F_{\mathbf{a}} = \bar{s}r\beta \quad (88)$$

$$F_{\mathbf{y}} - F_{\mathbf{x}} = s\beta \quad (89)$$

$$F_{\mathbf{b}} - F_{\mathbf{y}} = \bar{s}\bar{r}\beta \quad (90)$$

with which the outer expectation value of \mathfrak{K} can be expressed as an integral over the unit square $(s, r) \in [0, 1]^2$ with a surface element $\frac{\beta^2 \bar{s}}{2} dr ds$, which through the factor $\frac{1}{2}$ already conditions on $F_{\mathbf{x}}$ and $F_{\mathbf{y}}$ being correctly ordered. We need to show that this integral is asymptotically bounded by a constant in order to show that \mathfrak{K} is asymptotically bounded by a constant from above as we intend.

The expectation value $\mathbb{E} [\tilde{\mathcal{Z}}_{ab}]$ may be bounded using eq. (61):

$$\mathbb{E} [\tilde{\mathcal{Z}}_{ab}] = \Theta(1) L e^{L \Gamma_{ab}(\beta)} \quad (91)$$

The same bound does not in general apply to the other expectation values uniformly over the integration domain due to the divergence of the constant term with vanishing fitness difference. For this reason, we split the integration region. For some sufficiently small constant $\epsilon > 0$ we will consider integration in the regions with $s \in [0, \epsilon]$ and $s \in [\epsilon, 1]$ separately and name the corresponding contributions to \mathfrak{K} accordingly with an index.

6.2 Case $s \in [\epsilon, 1]$

In the interval $[\epsilon, 1]$, s is bounded away from zero and therefore using eq. (61), the expectation values $\mathbb{E} [\tilde{\mathcal{Z}}_{\mathbf{x}\mathbf{y}} | \beta s]$ can be bounded uniformly by

$$\mathbb{E} [\tilde{\mathcal{Z}}_{\mathbf{x}\mathbf{y}} | \beta s] \leq \mathcal{O}(1) L e^{L \Gamma_{\mathbf{x}\mathbf{y}}(\beta s)} \quad (92)$$

For the remaining expectation values we follow the procedure used in the previous section and expand with eq. (63) to obtain a sum of locus-factorized terms. We name the special loci according to the walk segment's index. As we have already expanded the contributions for the second segment, only the first and third remain.

$$\mathfrak{K}_{[\epsilon, 1]} = \mathcal{O}(1) \bar{\alpha} L^\alpha \sum_{l_1, l_3=1}^L \mathbb{E} \left[\prod_{l=1}^L \mathfrak{F}_l \right] = \mathcal{O}(1) \bar{\alpha} L^\alpha \sum_{l_1, l_3=1}^L \mathbb{E} \left[\prod_{l=1}^L \mathfrak{G}_l \mathfrak{H}_l \right] \quad (93)$$

Again, the usual form for loci $l \notin \{l_1, l_3\}$ can be given through the exponential terms in the expectation values

$$\mathfrak{G}_l = \sum_{x_l, y_l \in A} e^{\Gamma_{a_l x_l}(\beta \bar{s} r) + (1+\alpha) \Gamma_{x_l y_l}(\beta s) + \Gamma_{y_l b_l}(\beta \bar{s} r) - \Gamma_{a_l b_l}(\beta)} \quad (94)$$

and again for loci $l \in \{l_1, l_3\}$ one or more of the exponential factors will be replaced by their derivatives. For the same reasons as used previously, in these cases \mathfrak{H}_l is uniformly bounded by a constant and therefore

$$\mathfrak{K}_{[\epsilon, 1]} = \mathcal{O}(1) \bar{\alpha} L^{2+\alpha} e^{L \bar{\alpha}} \quad (95)$$

where

$$\mathfrak{T} = \left\langle \ln \sum_{x_l, y_l \in \mathbb{A}} e^{\Gamma_{a_l x_l}(\beta \bar{s} r) + (1+\alpha) \Gamma_{x_l y_l}(\beta s) + \Gamma_{y_l b_l}(\beta \bar{s} r) - \Gamma_{a_l b_l}(\beta)} \right\rangle_l \quad (96)$$

\mathfrak{T} can be considered a function of β , s , r and α . At $\alpha = 0$, it is always 0 as can be verified by matrix multiplication. The first derivative towards α at $\alpha = 0$ is found to be exactly $\mathfrak{M}(s, r, \beta)$. It is therefore possible to bound

$$\mathfrak{T} = \alpha \mathfrak{M}(s, r, \beta) + \mathcal{O}(\alpha^2). \quad (97)$$

In the regular case, as β converges to $\hat{\beta}$, $\mathfrak{M}(s, r, \beta)$ is eventually bounded from above by a negative constant in the region $s \in [\epsilon, 1 - \epsilon]$, so that in this region the integrand falls exponentially quickly to zero for suitable choice of $\alpha > 0$, resulting in no asymptotic contribution to \mathfrak{K} . In the interval $s \in [1 - \epsilon, 1]$ we need to account for the boundary term at $s = 1$. At $s = 1$, \mathfrak{T} is exactly $\alpha \mathbf{\Gamma}_{ab}(\beta)$. At the candidate threshold function $\mathbf{\Gamma}_{ab}(\beta)$ simply evaluates to $-\frac{\ln L}{L}$ up to irrelevant higher orders in L . By assumptions for the regular case we also have that the derivative $\partial_\alpha \partial_s \mathfrak{T}$ is positive at $(s, \alpha, \beta) = (1, 0, \hat{\beta})$, so that

$$\mathfrak{T} \leq -\alpha \frac{\ln L}{L} + \mathcal{O}\left(\left(\frac{\ln L}{L}\right)^2\right) - c \bar{s} \alpha \quad (98)$$

for some $c > 0$. The term $-\alpha \frac{\ln L}{L}$ exactly compensates a factor L^α to the integrand of \mathfrak{K} and with a factor \bar{s} in the surface element of the integration, the contribution to \mathfrak{K} from $s \in [1 - \epsilon, 1]$ is then for suitably small constant $\alpha > 0$:

$$\mathfrak{K}_{[1-\epsilon, 1]} = \mathcal{O}(1) L^2 \int_{1-\epsilon}^1 ds \bar{s} e^{-cL\bar{s}\alpha} = \mathcal{O}(1). \quad (99)$$

6.3 Case $s \in [0, \epsilon]$

For the integration interval $[0, \epsilon]$ we will fix $\alpha = 1$ and since we cannot apply the simple bound to the expectation $\mathbb{E}[\tilde{\mathcal{Z}}_{xy} | \beta s]^{1+\alpha}$ used before uniformly in this region, we will expand it using the sum form of the expectation value. Since $1 + \alpha = 2$ now, there will effectively be two additional sums resulting from this, for which we label the corresponding locus variables l_{21} and l_{22} . Again, we bring the contribution into the form

$$\mathfrak{K}_{[0, \epsilon]} = \mathcal{O}(1) L^{-1} \sum_{l_1, l_{21}, l_{22}, l_3} \mathbb{E}[\mathfrak{F}_l] = \mathcal{O}(1) L^{-1} \sum_{l_1, l_{21}, l_{22}, l_3} \mathbb{E}[\mathfrak{G}_l \mathfrak{H}_l] \quad (100)$$

Here, since all expectation values in the numerator of \mathfrak{R} were expanded into sums, only a single factor L^{-1} remains from the expectation value in its denominator. The usual form \mathfrak{G}_l is unchanged from the region $[\epsilon, 1]$ except for the choice $\alpha = 1$. As before \mathfrak{H}_l can be bounded by a constant for all special l , but this will turn out not to be sufficient here. Suppose we used such a bound, then we would obtain

$$\mathfrak{R}_{[0,\epsilon]} \leq \mathcal{O}(1)L^3 \mathbb{E} [e^{L\mathfrak{T}}] \quad (101)$$

where \mathfrak{T} is unchanged from the previous integration region except for the choice $\alpha = 1$. At $s = 0$ only terms with $x_l = y_l$ can contribute to \mathfrak{T} and so it becomes 0 by matrix multiplication. The first derivative towards s can be formed directly, using that derivatives of matrix exponentials correspond to multiplication with the matrix exponent. Using that $(A)_{x_l y_l} \mathbb{I}_{x_l y_l} = 0$ since the allele graph is simple, the derivative evaluates exactly to $-\beta \mathbf{\Gamma}'_{ab}(\beta)$, so that:

$$\mathfrak{T} = -\beta \mathbf{\Gamma}'_{ab}(\beta)s + \mathcal{O}(s^2) \quad (102)$$

As $\beta \mathbf{\Gamma}'_{ab}(\beta)$ is strictly positive and bounded away from zero asymptotically, this shows that ϵ can always be chosen such that \mathfrak{T} is negative for $s \in (0, \epsilon]$ with negative first derivative at $s = 0$. Consequently the integration at the boundary is of the form

$$\mathfrak{R}_{[0,\epsilon]} \leq \mathcal{O}(1)L^3 \int_0^\epsilon ds e^{-cLs} = \mathcal{O}(L^2) \quad (103)$$

for some constant $c > 0$. In contrast to the boundary at $s = 1$ the surface element does not contribute here and does not yield an additional factor L^{-1} . The naive bound shown above is not sufficient and two powers of L remain that we have to suppress.

To cancel these factors, we need to bound the terms \mathfrak{H}_l more carefully around small s instead of applying uniform constant bounds. In particular it would be sufficient to show that these terms introduce at least two factors s into the integrand, since the integration over $s^n e^{-cLs}$ would result in a value of order L^{-1-n} instead of just L^{-1} . Depending on the choices of distances between \mathbf{x} and \mathbf{y} and the choices of the special loci it is possible to provide these two factors. However not all combinations of these choices yield such a factor. The problematic cases will however turn out to be marginal in the sense that they only apply to a fraction $\frac{1}{L}$ or $\frac{1}{L^2}$ of the summands in the sums over special loci. Each such factor L^{-1} offsets the need for one s factor in the integrand, allowing the total contribution to \mathfrak{R} to still be constant. In the following we need to list all of the relevant combinations and show their contributions of s orders.

The method of bounding \mathfrak{H}_l is to consider the small- s behavior of the factors $(e^{s\beta A})_{x_l y_l}$ and $(Ae^{s\beta A})_{x_l y_l}$ which appear in it. In particular, depending

on the distance we have

$$(e^{s\beta A})_{x_l y_l} = \begin{cases} 1 + \mathcal{O}(s^2) & x_l = y_l \\ \mathcal{O}(s^{d_{x_l y_l}}) & x_l \neq y_l \end{cases} \quad (104)$$

$$(Ae^{s\beta A})_{x_l y_l} = \begin{cases} \mathcal{O}(s) & x_l = y_l \\ \mathcal{O}(s^{d_{x_l y_l} - 1}) & x_l \neq y_l \end{cases} \quad (105)$$

and due to the bounded degree of the graph, all of these bounds are uniform.

Using the bounds above, we can obtain the necessary factors of s . First consider the case of all four special loci distinct. We then have for the two loci l_{21} and l_{22} :

$$\mathfrak{H}_{l_{2i}} = \frac{\sum_{x_{l_{2i}}, y_{l_{2i}} \in \mathbb{A}} (e^{\beta \bar{s} r A})_{a_{l_{2i}} x_{l_{2i}}} (e^{\beta s A})_{x_{l_{2i}} y_{l_{2i}}} (Ae^{\beta s A})_{x_{l_{2i}} y_{l_{2i}}} (e^{\beta \bar{s} r A})_{y_{l_{2i}} b_{l_{2i}}}}{\sum_{x_{l_{2i}}, y_{l_{2i}} \in \mathbb{A}} (e^{\beta \bar{s} r A})_{a_{l_{2i}} x_{l_{2i}}} [(e^{\beta s A})_{x_{l_{2i}} y_{l_{2i}}}]^2 (e^{\beta \bar{s} r A})_{y_{l_{2i}} b_{l_{2i}}}} \quad (106)$$

From eq. (104) we can see that for all distances

$$(e^{\beta s A})_{x_l y_l} (Ae^{\beta s A})_{x_l y_l} = \mathcal{O}(s) \quad (107)$$

and therefore each of $\mathfrak{H}_{l_{21}}$ and $\mathfrak{H}_{l_{22}}$ contribute at least one factor s , resulting in a sufficient contribution of s^2 as explained above. If not all of the four loci are distinct the form of \mathfrak{H}_{2i} will be different. However, the only modifications are in the placement of derivatives of matrix exponentials. As long as still $l_{21} \neq l_{22}$, the relevant terms which are small around $s = 0$, namely the exponentials for the second walk segment, remain unchanged.

Therefore the remaining cases are for $l_{21} = l_{22}$, for which we will write l_2 . This equality reduces the number of summands to consider by a factor L^{-1} as discussed before and consequently we need to find only one factor s . In particular if either l_1 or l_3 are equal to l_2 as well, then the weight of these cases is reduced by another factor L^{-1} , so that no s is required anymore. Therefore, we can focus only on the case where l_1, l_2 and l_3 are all distinct. For this case the contribution of locus l_2 is

$$\mathfrak{H}_{l_2} = \frac{\sum_{x_{l_2}, y_{l_2} \in \mathbb{A}} (e^{\beta \bar{s} r A})_{a_{l_2} x_{l_2}} [(Ae^{\beta s A})_{x_{l_2} y_{l_2}}]^2 (e^{\beta \bar{s} r A})_{y_{l_2} b_{l_2}}}{\sum_{x_{l_2}, y_{l_2} \in \mathbb{A}} (e^{\beta \bar{s} r A})_{a_{l_2} x_{l_2}} [(e^{\beta s A})_{x_{l_2} y_{l_2}}]^2 (e^{\beta \bar{s} r A})_{y_{l_2} b_{l_2}}} \quad (108)$$

Following again eq. (104), the numerator is of order $\mathcal{O}(s^2)$ except if $d_{x_{l_2} y_{l_2}} = 1$, in which case there is a zeroth order contribution. The latter case requires additional considerations to resolve.

First, we consider the subcase with $d_{\mathbf{x}\mathbf{y}} \geq 2$. In this case it is possible that $d_{x_{l_2} y_{l_2}} = 1$, but if this is the case we always have another locus l' with $d_{x_{l'} y_{l'}} \geq 1$. Following the separation of edges in the previous section, we can handle one such locus as a special locus in exchange for another sum over of order L . However, the $x_{l'} \rightarrow y_{l'}$ factor contributions in $\mathfrak{H}_{l'}$'s numerator will then always be $[(e^{\beta s A})_{x_{l'} y_{l'}}]^2$ without any derivatives since $l' \neq l_2$. From eq.

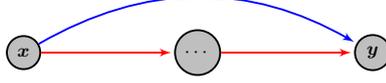


Fig. 9 Distance-1 case for $\tilde{Z}'_{xy\pi}$. The dots represent an arbitrary acyclic subgraph. The blue path is always accessible, assuming that the initial and final fitness values for x and y are correctly ordered, which we enforce through integral bounds. The focal path is also conditioned on being accessible. Assuming that π is the focal red path and π' the blue path, then in $\tilde{Z}'_{ab\pi'}$ all paths except π' are counted, while in $\tilde{Z}'_{ab\pi}$ at least one accessible path (π) is excluded, but more paths that are not arcs may also be excluded. Therefore $\tilde{Z}'_{ab\pi} \leq \tilde{Z}'_{ab\pi'}$ under the stated conditioning.

(104), such a factor results in a factor s^2 compensating the additional L sum as well as the required s factor to the integrand.

The only remaining case is then $d_{xy} = 1$. For this case the contribution to \mathfrak{K} is indeed not bounded as we require. However, this contribution turns out to be an overcounting issue introduced by our loosening of the arc restrictions on $\tilde{Z}'_{xy\pi}$. Specifically, if $d_{xy} = 1$, we will enforce the restriction that $\tilde{Z}'_{xy\pi}$ should not count the direct walk segment $x \rightarrow y$ if π is taking this direct step. Since the direct step is always accessible given that F_x and F_y are ordered correctly, this segment contributes exactly 1 to the expectation value $\mathbb{E}[\tilde{Z}_{xy}|\beta s]$, which we can therefore subtract from it. This is possible even if π does not use this direct step since $\tilde{Z}'_{xy\pi} \leq \tilde{Z}'_{xy\pi'}$ if π does not use the trivial arc, but π' does (Figure 9). With this modification the value of \mathfrak{H}_{l_2} becomes

$$\mathfrak{H}_{l_2} = \frac{\sum_{x_{l_2}, y_{l_2} \in \mathbb{A}} (e^{\beta \bar{s} r A})_{a_{l_2} x_{l_2}} (A e^{\beta s A})_{x_{l_2} y_{l_2}} \left((A e^{\beta s A})_{x_{l_2} y_{l_2}} - 1 \right) (e^{\beta \bar{s} r A})_{y_{l_2} b_{l_2}}}{\sum_{x_{l_2}, y_{l_2} \in \mathbb{A}} (e^{\beta \bar{s} r A})_{a_{l_2} x_{l_2}} [(e^{\beta s A})_{x_{l_2} y_{l_2}}]^2 (e^{\beta \bar{s} r A})_{y_{l_2} b_{l_2}}} \quad (109)$$

Since $(A)_{x_{l_2} y_{l_2}} = 1$, the leading order in the numerator is now $\mathcal{O}(s)$, which is sufficient to obtain a bounded contribution to \mathfrak{K} .

All in all, the total contributions to \mathfrak{K} are bounded in L at the candidate threshold function for the regular type, implying that there is a constant $C > 0$, such that

$$\liminf \mathbb{P}[\tilde{Z}_{ab} \geq 1] \geq C \quad (110)$$

implying the (weak) critical point given in our main statement. It remains to improve this bound from non-zero C to 1, which we expect can be done as mentioned in the Introduction.

Acknowledgements

This work was supported by DFG within SPP 1590 *Probabilistic structures in evolution*. We thank David Augustin for his contributions to the early stages of this project, and Alexander Drewitz for discussions.

References

1. Alon, N., Spencer, J.: *The Probabilistic Method*. 2nd ed., Wiley, New York (2000).
2. Altenberg, L.: Fundamental properties of the evolution of mutational robustness. Preprint arXiv:1508.07866 (2015)
3. Berestycki, J., Brunet, É., Shi, Z.: The number of accessible paths in the hypercube. *Bernoulli* **22**, 653–680 (2016)
4. Berestycki, J., Brunet, É., Shi, Z.: Accessibility percolation with backsteps. *ALEA, Lat. Am. J. Probab. Math. Stat.* **14**, 45–62 (2017)
5. Carneiro, M., Hartl, D.L.: Adaptive landscapes and protein evolution. *Proc. Nat. Acad. Sci. USA* **107**, 1747–1751 (2010)
6. Crona, K., Greene, D., Barlow, M.: The peaks and geometry of fitness landscapes. *J. Theor. Biol.* **318**, 1–10 (2013)
7. Franke, J., Klözer, A., de Visser, J.A.G.M., Krug, J.: Evolutionary accessibility of mutational pathways. *PLoS Comp. Biol.* **7**(8), e1002134 (2011)
8. Gillespie, J.H.: Molecular evolution over the mutational landscape. *Evolution* **38**, 1116–1129 (1984)
9. Hegarty, P., Martinsson, A.: On the existence of accessible paths in various models of fitness landscapes. *Ann. Appl. Probab.* **24**, 1375–1395 (2014)
10. Kauffman, S., Levin, S.: Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology* **128**(1), 11–45 (1987)
11. Kingman, J.F.C.: A simple model for the balance between selection and mutation. *Journal of Applied Probability* **15**(1), 1–12 (1978)
12. Kistler, N., Schertzer, A.: Undirected polymers in random environment: Path properties in the mean field limit. Preprint arXiv:2012.04076 (2020)
13. Krug, J.: Accessibility percolation in random fitness landscapes. In: *Probabilistic Structures in Evolution*, ed. by E. Baake and A. Wakolbinger (EMS Press, 2021)
14. Li, L.: Phase transition for accessibility percolation on hypercubes. *J. Theor. Prob.* **31**, 2072–2111 (2018)
15. Martinsson, A.: Accessibility percolation and first-passage site percolation on the un-oriented binary hypercube. Preprint arXiv:1501.02206 (2015)
16. Martinsson, A.: Unoriented first-passage percolation on the n-cube. *Ann. Prob.* **26**, 2597–2625 (2016)
17. Martinsson, A.: First-passage percolation on Cartesian power graphs. *Ann. Prob.* **46**, 1004–1041 (2018)
18. Nowak, S., Krug, J.: Accessibility percolation on n -trees. *Europhys. Lett.* **101**, 66,004 (2013)
19. Orr, H.A.: The population genetics of adaptation: the adaptation of DNA sequences. *Evolution* **56**, 1317–1330 (2002)
20. de Visser, J.A.G.M., Krug, J.: Empirical fitness landscapes and the predictability of evolution. *Nature Reviews Genetics* **15**, 480–490 (2014)
21. Weinreich, D.M., Watson, R.A., Chao, L.: Sign epistasis and genetic constraint on evolutionary trajectories. *Evolution* **59**, 1165–1174 (2005)
22. Wu, N.C., Dai, L., Olson, C.A., Lloyd-Smith, J.O., Sun, R.: Adaptation in protein fitness landscapes is facilitated by indirect paths. *eLife* **5**, 16,965 (2016)
23. Zagorski, M., Burda, Z., Waclaw, B.: Beyond the hypercube: evolutionary accessibility of fitness landscapes with realistic mutational networks. *PLoS Comp. Biol.* **12**(12), e1005218 (2016)

Chapter 7

Discussion

NK fitness landscapes are a simple approach to model fitness landscapes of varying ruggedness, fitness distribution and interaction structure between loci. Nonetheless many of their structural properties can be characterized by few parameters of the NK model, the mean size of neighborhoods k and the mean degree r in the neighborhood structure.

In particular for uniform- and regular NK structures these two parameters largely determine the qualitative behavior of the landscape.

At constant values of these parameters accessibility of distant genotypes drops quickly to zero, while the expected number of such accessible paths given accessibility increases factorially in L . This can be contrasted with the HoC model in which the expected number of accessible paths never increases faster than exponentially and, dependent on the particular arrangement of the allele graph, accessibility between distant genotypes is often asymptotically non-zero, although one might intuitively expect HoC landscapes to be more rugged than fixed- k NK landscapes.

For the expected number of local maxima the particular choices of fitness distribution and NK structure do have some influence on the asymptotic leading order behavior, even if restricted to uniform- and regular NK structures, although the variation seems to be limited to a finite power in the per-genotype probability for well-behaved fitness distributions. Only for fitness distributions with extremely long tails my considerations suggest a possibly stronger deviation towards fewer local maxima. Although not presented in fully here, I believe that this qualitatively different behavior can be proven.

Interestingly, it seems that among all uniform- and regular NK structures, the block structures is optimal for the expected number of local maxima in the sense that no other such structures seem to result in a asymptotically significantly higher expected number of local maxima. I am not aware of any proof of this.

The results for the NK model are qualitatively independent of the particular allele graph chosen, as long as it is considered fixed, however in the HoC model the

choice of allele graph may have more significant ramifications. In particular allele graphs with distant alleles result in pairs of genotypes which are asymptotically never accessible, while especially for high-degree allele graphs of small diameter even distant genotypes often are asymptotically accessible with finite probability. The presented results allow to make this determination explicitly for many choices of allele graphs, however the remaining cases in which the conditions on Martinsson's function are not satisfied and the expected number of accessible paths doesn't "tell the truth", remains an open problem.

Appendix A

Proofs and other calculations

A.1 Expected number of local optima for normal-distributed uniform and regular NK landscapes

This is a more precise implementation of the strategy in,^[23] attached as chapter 5. Here I assume a standard normal distribution for individual fitness values.

The probability that a genotype is a local fitness maximum is the probability that the fitness difference to all one-mutant steps is positive. Given that all fitness contributions are independently standard-normal distributed, these differences will be jointly normal distributed with zero mean. The covariance matrix A will have diagonal values of $2r$ for an r -uniform NK landscape, accounting for differences of two fitness values on all r partial landscapes contributing to the associated mutation. There is a contribution of 1 to each off-diagonal term A_{ij} for each partial landscape containing loci i and j . These are the only partial landscapes contributing to the covariance between the two mutations associated with i and j , since the only possible shared fitness contribution in the fitness difference is the fitness of the partial landscape at the focal genotype.

Given k -uniformity in addition to regularity, the column and row sums are therefore all the same with value $r(k+1)$. The smallest eigenvalue of this matrix satisfies $\lambda_- \geq r$ and by the Perron-Frobenius theorem its largest eigenvalue is $r(k+1)$ with corresponding eigenvector the constant vector.

The probability for the genotype to be a local optimum is then

$$P_{\text{opt}} = \left\langle \prod_{i=1}^L \Theta(x_i) \right\rangle \quad (\text{A.1})$$

where x is distributed as explained above. We can rescale x to $x' = \frac{x}{\lambda_-}$. Fur-

thermore separating x' into a sum of two jointly normal random vectors, the first independent with variances 1, we have

$$P_{\text{opt}} = \left\langle \prod_{i=1}^L \Phi(y_i) \right\rangle \quad (\text{A.2})$$

where Φ is the standard normal cdf and y is distributed jointly normal with covariance matrix B with diagonal values $\frac{2r}{\lambda_-} - 1$ and off diagonal values $B_{ij} = \frac{A_{ij}}{\lambda_-}$. It's smallest eigenvalue is zero and its largest eigenvalue is $\lambda'_+ = \frac{r(k+1)}{\lambda_-} - 1$. The integration variable may be shifted by a vector c , which given the PDF of the joint normal distribution results in

$$P_{\text{opt}} = e^{-\frac{1}{2}\langle c|B^{-1}|c \rangle} \left\langle e^{-\langle y|B^{-1}|c \rangle} \prod_{i=1}^L \Phi(y_i + c_i) \right\rangle \quad (\text{A.3})$$

where y is still distributed as before. Technically the matrix inverse does not exist, this is however not relevant for the considerations here, since we will consider only shifts without contribution of the zero eigenvector. Alternatively one may consider λ_- to be offset by some small ϵ from the real smallest eigenvalue, so that B is positive-definite.

Choosing c as a constant vector and using q for its entries, we have then

$$P_{\text{opt}} = e^{-\frac{q^2}{2\lambda'_+}L} \left\langle \prod_{i=1}^L e^{-\frac{q}{\lambda'_+}y_i} \Phi(y_i + q) \right\rangle \quad (\text{A.4})$$

because the constant vector is the Frobenius-Perron eigenvector for B .

Bounds on this expression can now be obtained by application of Jensen's inequality and by considering the maximum argument of the expectation.

For the maximum argument we can consider the factors for each entry of y independently. Each entry has the form

$$\exp\left(-\frac{q}{\lambda'_+}y_i + \ln \Phi(y_i + q)\right) \quad (\text{A.5})$$

A shift of y_i back by q modifies this to

$$\exp\left(\frac{q^2}{\lambda'_+} - \frac{q}{\lambda'_+}y_i + \ln \Phi(y_i)\right) \quad (\text{A.6})$$

Choosing roughly $q = \sqrt{2 \ln \lambda'_+}$, as $\lambda'_+ \rightarrow \infty$, the term $\frac{q}{\lambda'_+}$ tends towards 0. The maximal value of y_i therefore will diverge to infinity in the same limit, maximizing

$\ln \Phi(y_i)$. Expanding $\ln \Phi(y_i)$ yields

$$\ln \Phi(y_i) = - \left(\frac{1}{y_i \sqrt{2\pi}} + \mathcal{O}\left(\frac{1}{y_i^2}\right) \right) e^{-\frac{1}{2}y_i^2} \quad (\text{A.7})$$

$$\partial_{y_i} \ln \Phi(y_i) = \left(\frac{1}{\sqrt{2\pi}} + \mathcal{O}\left(\frac{1}{y_i}\right) \right) e^{-\frac{1}{2}y_i^2} \quad (\text{A.8})$$

and so by comparison with the linear term in the exponent, the maximum is reached at

$$y_i = \sqrt{2 \ln \left(\frac{\lambda'_+}{q \sqrt{2\pi}} + \mathcal{O}(1) \right)} \quad (\text{A.9})$$

where the order of the remaining terms can be estimated from the leading correction of the expansion of $\partial_{y_i} \ln \Phi(y_i)$. This yields a maximum for the expectation's argument and the overall upper bound on the probability

$$P_{\text{opt}} \leq \exp \left(L \left(\frac{q^2}{2\lambda'_+} - \frac{q \sqrt{2 \ln \left(\frac{\lambda'_+}{q \sqrt{2\pi}} + \mathcal{O}(1) \right)}}{\lambda'_+} \right) \right) \quad (\text{A.10})$$

Here I dropped the $\ln \Phi(y_i)$ term which is always negative for simplicity. Choosing $q = \sqrt{2 \ln \lambda'_+}$, this bound yields

$$P_{\text{opt}} \leq \exp \left(-L \frac{\ln \lambda'_+ - \frac{1}{2} \ln \ln \lambda'_+ - \ln(2\sqrt{\pi}) + o(1)}{\lambda'_+} \right) \quad (\text{A.11})$$

For the lower bound, instead of searching for the maximum of the expectation's argument, we apply Jensen's inequality to the exponential:

$$P_{\text{opt}} \geq \exp \left(L \left(-\frac{q^2}{2\lambda'_+} + \langle \ln \Phi(y_i + q) \rangle \right) \right) \quad (\text{A.12})$$

The linear term in y_i does not contribute since y_i 's expectation is zero. The term $\langle \ln \Phi(y_i + q) \rangle$ is the same for all i due to the regularity assumption. Only the marginal distribution of y_i is relevant for this expectation, which is a normal distribution with variance $B_{ii} = \frac{2r}{\lambda_-} - 1$. λ_- is at least r and therefore this variance is asymptotically bounded. At $y_i \rightarrow -\infty$, the term $\ln \Phi(y_i + q)$ falls to negative infinity only quadratically, so that the expectation always exists. With this, assuming q is diverging, the contributions to the expectation may be pointwise asymptotically expanded to

$$\langle \Phi(y_i + q) - 1 \rangle \quad (\text{A.13})$$

which evaluates to

$$\Phi\left(\frac{q}{\sqrt{1+B_{ii}}}\right) - 1 \quad (\text{A.14})$$

Choosing $q = \sqrt{2(1+B_{ii}) \ln \lambda'_+}$ then yields the overall bound

$$P_{\text{opt}} \geq \exp\left(-L(1+B_{ii})\frac{\ln \lambda'_+ + \mathcal{O}(1)}{\lambda'_+}\right) \quad (\text{A.15})$$

The maximal value for λ_- is $2r$, in which case B_{ii} becomes 0 and λ'_+ becomes $\frac{k-1}{2}$. For this constellation the resulting bounds become

$$-\frac{2 \ln k + \mathcal{O}(1)}{k} \leq \frac{\ln P_{\text{opt}}}{L} \leq -\frac{2 \ln k - \ln \ln k + \mathcal{O}(1)}{k} \quad (\text{A.16})$$

In particular the leading order term is tight.

On the opposite side of the spectrum, the smallest eigenvalue is only $\lambda_- = r$. In this case $B_{ii} = 1$ and $\lambda'_+ = k$. The bounds are not tight in this case yielding

$$-\frac{2 \ln k + \mathcal{O}(1)}{k} \leq \frac{\ln P_{\text{opt}}}{L} \leq -\frac{\ln k - \frac{1}{2} \ln \ln k + \mathcal{O}(1)}{k} \quad (\text{A.17})$$

Since we did not actually use the fact that λ_- is the smallest eigenvalue, rather than just a positive value not larger than the smallest eigenvalue, this bound applies in general to all cases.

The bound also applies in the case of k growing with L , as all L dependencies are explicit. Therefore with $\alpha = \frac{k}{L}$ constant:

$$C_1(\alpha L)^{-\frac{2}{\alpha}}(1 + o(1)) \leq P_{\text{opt}} \leq C_2(\alpha L \sqrt{\ln L})^{-\frac{1}{\alpha}} \quad (\text{A.18})$$

where C_1 and C_2 are positive constants independent of α and L . I expect the additional logarithmic correction on the upper bound to be an artifact of a suboptimal bound and not to be actually achievable.

A.2 Mean number of local optima in regular and uniform NK models with sum-dominating distribution

In this section I obtain bounds on the expected number of local optima for NK structures under an extremely heavy-tailed distribution. The derivation will assume the complete graph on two vertices as allele graph.

Specifically I consider a specific sequence ϕ_N of sum-dominating distributions, as follows. Let θ_n be the standard uniform distributions centered at c_n with $c_n = 1 + \sum_{i < n} (c_i + \frac{1}{2})$ and $c_1 = \frac{1}{2}$. Now let the distribution ϕ_N be the distribution obtained by drawing a value from θ_n with n an integer chosen uniformly between 1 and N .

Drawing elements from ϕ_N and taking N to infinity, the probability of the same n being drawn twice becomes zero and so asymptotically almost surely for all subsets of the chosen elements a sum of the elements of the subset multiplied with ± 1 is dominated by the largest element. This means that the signed subset sums will be positive if the sign of the largest element is positive and negative otherwise.

Whether a genotype is a local maximum is only dependent on the signs of sums of fitness differences. As a consequence this distribution allows, in the limit $N \rightarrow \infty$, to reduce the problem to a property of the ranking of the individual fitness values.

With the sum-dominating distribution with sufficiently large N , we can now consider the possible orderings of the individual fitness contributions and how they affect the sign of individual mutation effects on the landscape.

The relevant fitness contributions can be labeled by a pair (i, l) . Here i is the partial landscape to which they belong. If the fitness contribution is contributing to the mutant obtained by a mutation on a given locus, then l is the index of that locus. If the fitness contribution is contributing to the focal genotype, then $l = -1$.

In each mutation, only the sign of the largest fitness contribution matters. That means, for each locus l , the mutation on l decreases fitness if and only if all (i, l) for all i are smaller than $(i, -1)$.

Given an ordering of the individual fitness value contributions from highest to lowest, we can use the following algorithm to determine whether all mutations decrease fitness and therefore result in a local maximum:

- Initially, for all loci l , s_l is unset.
- Iterate in the given order through fitness contributions (i, l) .
 - If $l = -1$:
 - * For all l in the neighborhood i , set s_l .
 - If $l \neq -1$:
 - * For all l in the neighborhood i :
 - If s_l is not yet set, exit the algorithm, the genotype is not a local maximum.

Successively this algorithm sets the values s_l indicating the mutations which are known to be decreasing fitness based on the highest fitness contributions considered so far. The algorithm exits if one of the signs is found to be positive, indicating absence of a local maximum. If the algorithm terminates without explicit exit, all s_l will be set and the genotype is a local maximum. Once a sign has been found other fitness values cannot influence it anymore, since the earlier fitness contribution determining its sign dominates all possible sums of remaining fitness contributions.

The algorithm can be terminated early, once all s_l are found. An alternative point-of-view for the algorithm is to consider only the $(i, -1)$ contributions and their ordering first. A $(i, -1)$ contribution is reached in the algorithm, only if none of the (i, l) contributions with $l \neq i$ and s_l unset has been reached before and if it reached, it sets all s_l with l in the NK neighborhood i . The probability that $(i, -1)$ is the first among the (i, l) is simply $\frac{1}{k'_i+1}$ where k'_i is the number of unset s_l with l in i .

$$P_{\text{opt}} = \left\langle \exp \left(- \sum_i \ln(k'_i + 1) \right) \right\rangle \quad (\text{A.19})$$

where the average is over all orderings of NK neighborhoods and the sum is over all indices corresponding to NK neighborhoods.

Trivial bounds on P can be obtained by the fact that NK blocks chosen with $k'_i = 0$ do not contribute to the exponent and that $\sum_i k'_i = 0$. Worst-case all non-zero k'_i are 1, in which case there must be L of them. Best case all k'_i are k , in which case there are $\frac{L}{k}$ of them. This implies the bounds

$$\exp(-L \ln 2) \leq \lim P_{\text{opt}} \leq \exp \left(-L \frac{\ln(k+1)}{k} \right) \quad (\text{A.20})$$

The lower bound is the trivial lower bound for any fitness landscape and the upper bound corresponds to the value of the BN structure for any distribution, indicating again that the BN structure seems to be optimal in terms of the number of local maxima.

Given the bounds above, in the limit of L growing sufficiently faster than k , it is expected that $\lim P_{\text{opt}}$ decreases exponentially. As a consequence it is possible that only a small probability mass contributes to the expectation in eq. (A.19). This may be used to optimize the bounds obtained so far.

The upper trivial bound requires that, except for already fully covered NK blocks, each chosen NK block covers exactly $\frac{L}{k}$ loci without overlap, meaning that up to duplicates, the chosen edges form a *perfect cover* of the NK structure. An (*edge*) *cover* of a hypergraph is a subset of its edges such that each vertex is contained in at least one of the edges in the subset. A *perfect (edge) cover* of a

k -uniform hypergraph with L vertices is a cover which contains exactly $\frac{L}{k}$ edges, the minimal possible value. The (*edge*) *covering number* of a hypergraph is the size of its smallest cover. We write ν for the covering number divided by the number of vertices.

The BN structure is the only one in which all possible choices yield perfect covers. However, it is often possible to find covers which are close to perfect in k -uniform r -regular NK structures. In any case, by selection of a random subset of edges it can be seen that^[1]

$$\nu \leq \frac{\ln k + 1}{k} \quad (\text{A.21})$$

Although it is generally unlikely that a random ordering of the edges starts with a minimal cover, the probability that it happens is at least $\binom{L}{L\nu}^{-1} \geq e^{-L\nu \ln \nu^{-1}}$ and furthermore, the above bound on ν can be obtained by a random uniform choice of edges with finite probability, making it unnecessary to account for the likelihood of one particular ordering. Depending on the order in which the edges of the cover are chosen, the distribution of k'_i differs, however by Jensen's inequality

$$\langle \ln(k'_i + 1) \rangle_i \leq \ln(\langle k'_i \rangle + 1) = \ln(\nu^{-1} + 1) \quad (\text{A.22})$$

and so

$$\lim P_{\text{opt}} \geq \exp(-2L\nu \ln(\nu^{-1} + 1) - 1) \quad (\text{A.23})$$

$$\lim P_{\text{opt}} \geq \exp\left(-L \frac{\ln k + 1}{k} \ln\left(\frac{k}{\ln k + 1} + 1\right) + \mathcal{O}(1)\right) \quad (\text{A.24})$$

$$= \exp\left(-L \left(\frac{(\ln k)^2}{k} (1 + o(1))\right) + \mathcal{O}(1)\right) \quad (\text{A.25})$$

The second inequality yields exponential growth factors which are smaller than $\ln 2$ for all $k \geq 2$, implying that with the sum-dominating distribution all non-trivial uniform and regular NK structures have an exponentially growing expected number of local optima.

Further, if there exists a *almost-perfect cover*, meaning a cover of size $\frac{L}{k} + o(1)$, then

$$\lim P_{\text{opt}} \geq \exp\left(-2L \frac{\ln k}{k} (1 + o(1))\right) \quad (\text{A.26})$$

and if there exists a *decent cover*, meaning a cover of size $\mathcal{O}\left(\frac{L}{k}\right)$, then

$$\lim P_{\text{opt}} \geq \exp\left(-L\Theta\left(\frac{\ln k}{k}\right)\right) \quad (\text{A.27})$$

So far, I am not aware of any k -uniform regular NK structure and fitness distribution for which this growth rate is not $\mathcal{O}\left(\frac{\ln k}{k}\right)$. This poses the question whether there is such a structure with the sum-dominating fitness distribution. The possibility remains since the covering number may be of order $L\frac{\ln k}{k}$. In fact Alon et al.^[1] show that for all $k \leq r \leq e^{4k}$ there exist k -uniform r -regular simple hypergraphs with $\nu \geq C_\nu \frac{\ln k}{k}$ for some universal constant $C_\nu > 0$. *Simple* here means that the co-degree of all pairs of edges is at most 1, meaning that edges never intersect in more than one locus.

This alone is not sufficient to establish a non- $\mathcal{O}\left(\frac{\ln k}{k}\right)$ behavior. It still needs to be shown that the expectation of $\exp(-\sum_i \ln(k'_i + 1))$ must be small if ν as above. In particular there are different orderings in which the cover can be chosen.

However, the considerations above seem to be an interesting indicator that it is indeed possible to achieve a uniform- and regular NK landscape with an asymptotically lower expected number of local maxima than

$$\mathbb{E}[N_{\max}] = A^L \exp\left(-L\Theta\left(\frac{\ln k}{k}\right)\right) \quad (\text{A.28})$$

Appendix B

Bibliography

- [1] N. Alon, B. Bollobás, J. H. Kim, and V. H. Vu, “Economical covers with geometric applications,” *Proceedings of the London Mathematical Society* **86**, 273–301 (2003).
- [2] P. Baldi and Y. Rinott, “On normal approximations of distributions in terms of dependency graphs,” *The Annals of Probability* **17**, 1646–1650 (1989).
- [3] N. Beerenwinkel, L. Pachter, and B. Sturmfels, “Epistasis and shapes of fitness landscapes,” *Statistica Sinica*, 1317–1342 (2007).
- [4] J. Berestycki, É. Brunet, and Z. Shi, “Accessibility percolation with back-steps,” *Latin American Journal of Probability and Mathematical Statistics* **14**, 45–62 (2017).
- [5] J. Berestycki, É. Brunet, and Z. Shi, “The number of accessible paths in the hypercube,” *Bernoulli* **22**, 653–680 (2016).
- [6] A. Bird, “Perceptions of epigenetics,” *Nature* **447**, 396 (2007).
- [7] A. Bray and M. Moore, “Metastable states in spin glasses with short-ranged interactions,” *Journal of Physics C: Solid State Physics* **14**, 1313 (1981).
- [8] K. Crona, D. Greene, and M. Barlow, “The peaks and geometry of fitness landscapes,” *Journal of theoretical biology* **317**, 1–10 (2013).
- [9] É. Danchin, A. Charmantier, F. A. Champagne, A. Mesoudi, B. Pujol, and S. Blanchet, “Beyond dna: integrating inclusive inheritance into an extended theory of evolution,” *Nature reviews genetics* **12**, 475–486 (2011).
- [10] J. A. G. De Visser, T. F. Cooper, and S. F. Elena, “The causes of epistasis,” *Proceedings of the Royal Society B: Biological Sciences* **278**, 3617–3624 (2011).
- [11] J. A. G. De Visser and J. Krug, “Empirical fitness landscapes and the predictability of evolution,” *Nature Reviews Genetics* **15**, 480–490 (2014).

- [12] B. Derrida, “Random-energy model: an exactly solvable model of disordered systems,” *Physical Review B* **24**, 2613 (1981).
- [13] R. Durrett and V. Limic, “Rigorous results for the nk model,” *Annals of probability*, 1713–1753 (2003).
- [14] S. F. Edwards and P. W. Anderson, “Theory of spin glasses,” *Journal of Physics F: Metal Physics* **5**, 965 (1975).
- [15] S. N. Evans and D. Steinsaltz, “Estimating some features of NK fitness landscapes,” *The Annals of Applied Probability* **12**, 1299–1321 (2002).
- [16] L. Ferretti, B. Schmiegelt, D. Weinreich, A. Yamauchi, K. Yutaka, F. Tajima, and A. Guillaume, “Measuring epistasis in fitness landscapes: the correlation of fitness effects of mutations,” *Journal of theoretical biology* **396**, 132–143 (2016).
- [17] L. Ferretti, D. Weinreich, F. Tajima, and G. Achaz, “Evolutionary constraints in fitness landscapes,” *Heredity* **121**, 466–481 (2018).
- [18] R. A. Fisher, “Xxi.—on the dominance ratio,” *Proceedings of the royal society of Edinburgh* **42**, 321–341 (1923).
- [19] H. Flyvbjerg and B. Lautrup, “Evolution in a rugged fitness landscape,” *Physical Review A* **46**, 6714 (1992).
- [20] J. H. Gillespie, “Molecular evolution over the mutational landscape,” *Evolution*, 1116–1129 (1984).
- [21] P. Hegarty and A. Martinsson, “On the existence of accessible paths in various models of fitness landscapes,” *The Annals of Applied Probability* **24**, 1375–1395 (2014).
- [22] J. Huxley, “*Evolution. the modern synthesis.*,” (1942).
- [23] S. Hwang, B. Schmiegelt, L. Ferretti, and J. Krug, “Universality classes of interaction structures for nk fitness landscapes,” *Journal of Statistical Physics* **172**, 226–278 (2018).
- [24] S. Kauffman and S. Levin, “Towards a general theory of adaptive walks on rugged landscapes,” *Journal of theoretical Biology* **128**, 11–45 (1987).
- [25] S. A. Kauffman and E. D. Weinberger, “The nk model of rugged fitness landscapes and its application to maturation of the immune response,” *Journal of theoretical biology* **141**, 211–245 (1989).
- [26] J. F. Kingman, “A simple model for the balance between selection and mutation,” *Journal of Applied Probability* **15**, 1–12 (1978).

- [27] N. Kistler and A. Schertzer, “Undirected polymers in random environment: path properties in the mean field limit,” arXiv preprint arXiv:2012.04076 (2020).
- [28] N. Kistler, A. Schertzer, and M. A. Schmidt, “Oriented first passage percolation in the mean field limit, 2. the extremal process,” *The Annals of Applied Probability* **30**, 788–811 (2020).
- [29] S. Klavzar, “Some new bounds and exact results on the independence number of cartesian product graphs,” *Ars Combinatoria* **74**, 173–186 (2005).
- [30] F. D. Klironomos, J. Berg, and S. Collins, “How epigenetic mutations can affect genetic evolution: model and mechanism,” *BioEssays* **35**, 571–578 (2013).
- [31] R. Knippers, *Molekulare genetik* (Georg Thieme Verlag, 2006).
- [32] L. Li, “Phase transition for accessibility percolation on hypercubes,” *Journal of Theoretical Probability* **31**, 2072–2111 (2018).
- [33] V. Limic and R. Pemantle, “More rigorous results on the kauffman–levin model of evolution,” *The Annals of Probability* **32**, 2149–2178 (2004).
- [34] C. A. Macken and A. S. Perelson, “Protein evolution on rugged landscapes,” *Proceedings of the National Academy of Sciences* **86**, 6191–6195 (1989).
- [35] A. Martinsson, “Accessibility percolation and first-passage site percolation on the unoriented binary hypercube,” arXiv preprint arXiv:1501.02206 (2015).
- [36] A. Martinsson, “First-passage percolation on cartesian power graphs,” *The Annals of Probability* **46**, 1004–1041 (2018).
- [37] A. Martinsson, “Unoriented first-passage percolation on the n-cube,” *The Annals of Applied Probability* **26**, 2597–2625 (2016).
- [38] P. A. P. Moran, *The statistical process of evolutionary theory* (Clarendon Press, 1962).
- [39] V. Mustonen and M. Lässig, “From fitness landscapes to seascape: non-equilibrium dynamics of selection and adaptation,” *Trends in genetics* **25**, 111–119 (2009).
- [40] S. Nowak and J. Krug, “Analysis of adaptive walks on nk fitness landscapes with different interaction schemes,” *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P06014 (2015).
- [41] H. A. Orr, “The population genetics of adaptation on correlated fitness landscapes: the block model,” *Evolution* **60**, 1113–1124 (2006).
- [42] H. A. Orr, “The population genetics of adaptation: the adaptation of dna sequences,” *Evolution* **56**, 1317–1330 (2002).

- [43] A. S. Perelson and C. A. Macken, “Protein evolution on partially correlated landscapes,” *Proceedings of the National Academy of Sciences* **92**, 9657–9661 (1995).
- [44] P. C. Phillips, “Epistasis—the essential role of gene interactions in the structure and evolution of genetic systems,” *Nature Reviews Genetics* **9**, 855–867 (2008).
- [45] F. J. Poelwijk, D. J. Kiviet, D. M. Weinreich, and S. J. Tans, “Empirical fitness landscapes reveal accessible evolutionary paths,” *Nature* **445**, 383–386 (2007).
- [46] F. J. Poelwijk, S. Tănase-Nicola, D. J. Kiviet, and S. J. Tans, “Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes,” *Journal of theoretical biology* **272**, 141–144 (2011).
- [47] Z. R. Sailer and M. J. Harms, “Detecting high-order epistasis in nonlinear genotype-phenotype maps,” *Genetics* **205**, 1079–1088 (2017).
- [48] B. Schmiegelt, “Sign epistasis networks,” MA thesis (Universität zu Köln, 2016).
- [49] B. Schmiegelt and J. Krug, “Accessibility percolation on cartesian power graphs,” 2021.
- [50] B. Schmiegelt and J. Krug, “Evolutionary accessibility of modular fitness landscapes,” *Journal of Statistical Physics* **154**, 334–355 (2014).
- [51] D. Sherrington and S. Kirkpatrick, “Solvable model of a spin-glass,” *Physical review letters* **35**, 1792 (1975).
- [52] R. J. Sommer, “Phenotypic plasticity: from theory and genetics to current and future challenges,” *Genetics* **215**, 1–13 (2020).
- [53] P. F. Stadler, “Fitness landscapes,” in *Biological evolution and statistical physics* (Springer, 2002), pp. 183–204.
- [54] D. L. Stein, *Spin glasses and biology*, Vol. 6 (World Scientific, 1992).
- [55] I. G. Szendro, M. F. Schenk, J. Franke, J. Krug, and J. A. G. De Visser, “Quantitative analyses of empirical fitness landscapes,” *Journal of Statistical Mechanics: Theory and Experiment* **2013**, P01005 (2013).
- [56] F. Tanaka and S. Edwards, “Analytic theory of the ground state properties of a spin glass. i. ising spin glass,” *Journal of Physics F: Metal Physics* **10**, 2769 (1980).
- [57] M. B. Taylor and I. M. Ehrenreich, “Higher-order genetic interactions and their contribution to complex traits,” *Trends in genetics* **31**, 34–40 (2015).

- [58] T. Uller, A. P. Moczek, R. A. Watson, P. M. Brakefield, and K. N. Laland, “Developmental bias and evolution: a regulatory network perspective,” *Genetics* **209**, 949–966 (2018).
- [59] V. G. Vizing, “The cartesian product of graphs,” *Vycisl. Systemy* **9**, 33 (1963).
- [60] J. W. Weibull, *Evolutionary game theory* (MIT press, 1997).
- [61] E. D. Weinberger, “Local properties of kauffman’s n-k model: a tunably rugged energy landscape,” *Physical review A* **44**, 6399 (1991).
- [62] D. M. Weinreich, Y. Lan, C. S. Wylie, and R. B. Heckendorn, “Should evolutionary geneticists worry about higher-order epistasis?” *Current opinion in genetics & development* **23**, 700–707 (2013).
- [63] D. M. Weinreich, R. A. Watson, and L. Chao, “Perspective: sign epistasis and genetic constraint on evolutionary trajectories,” *Evolution* **59**, 1165–1174 (2005).
- [64] C. O. Wilke, C. Ronnewinkel, and T. Martinetz, “Dynamic fitness landscapes in molecular evolution,” *Physics Reports* **349**, 395–446 (2001).
- [65] S. Wright, “The roles of mutation, inbreeding, crossbreeding and selection in evolution,” in *Proceedings of the sixth international congress of genetics*, Vol. 1 (1932), pp. 356–366.
- [66] S. Wright, “Evolution in mendelian populations,” *Genetics* **16**, 97 (1931).
- [67] S. Wright, “The distribution of gene frequencies in populations,” *Proceedings of the National Academy of Sciences of the United States of America* **23**, 307 (1937).
- [68] M. Zagorski, Z. Burda, and B. Waclaw, “Beyond the hypercube: evolutionary accessibility of fitness landscapes with realistic mutational networks,” *PLoS computational biology* **12**, e1005218 (2016).

Appendix C

Teilpublikationen

I included the following articles published in peer-reviewed journals or as preprints to this thesis and detail my contributions to these articles as follows:

- S. Hwang et al., “Universality classes of interaction structures for nk fitness landscapes,” *Journal of Statistical Physics* **172**, 226–278 (2018)

The results and wordings of sections “Mathematical Background and Definitions”, “Accessible Pathways” and appendix C are primarily contributions by myself with supportive and editorial contributions by the co-authors. My contributions to other sections of the article were to varying degrees editorial and supportive on obtaining of results.

- B. Schmiegelt and J. Krug, “Accessibility percolation on cartesian power graphs,” 2021

Results were obtained primarily by myself under guidance by the co-author. With the exception of parts of the introductory section and editorial influence, the wording of the article is my own.

Appendix D

Erklärung nach Promotionsordnung

Nach § 7 Absatz 8 Satz 1 der Promotionsordnung der Mathematisch-Naturwissenschaftlichen Fakultät der Universität zu Köln vom 12. März 2020 gebe ich die folgende Erklärung im Wortlaut der Ordnung ab:

Hiermit versichere ich an Eides statt, dass ich die vorliegende Dissertation selbstständig und ohne die Benutzung anderer als der angegebenen Hilfsmittel und Literatur angefertigt habe. Alle Stellen, die wörtlich oder sinngemäß aus veröffentlichten und nicht veröffentlichten Werken dem Wortlaut oder dem Sinn nach entnommen wurden, sind als solche kenntlich gemacht. Ich versichere an Eides statt, dass diese Dissertation noch keiner anderen Fakultät oder Universität zur Prüfung vorgelegen hat; dass sie - abgesehen von unten angegebenen Teilpublikationen und eingebundenen Artikeln und Manuskripten - noch nicht veröffentlicht worden ist sowie, dass ich eine Veröffentlichung der Dissertation vor Abschluss der Promotion nicht ohne Genehmigung des Promotionsausschusses vornehmen werde. Die Bestimmungen dieser Ordnung sind mir bekannt. Darüber hinaus erkläre ich hiermit, dass ich die Ordnung zur Sicherung guter wissenschaftlicher Praxis und zum Umgang mit wissenschaftlichem Fehlverhalten der Universität zu Köln gelesen und sie bei der Durchführung der Dissertation zugrundeliegenden Arbeiten und der schriftlich verfassten Dissertation beachtet habe und verpflichte mich hiermit, die dort genannten Vorgaben bei allen wissenschaftlichen Tätigkeiten zu beachten und umzusetzen. Ich versichere, dass die eingereichte elektronische Fassung der eingereichten Druckfassung vollständig entspricht.