Does Evolutionary Plasticity Evolve?

Andreas Wagner

SFI WORKING PAPER: 1997-02-009

SANTA FE INSTITUTE

Does Evolutionary Plasticity Evolve?

Andreas Wagner

Yale University, Department of Biology
Center for Computational Ecology, OML327
P.O. Box 208104, New Haven, CT06520-8104
E-mail: waganda@doliolum.biology.yale.edu
Phone: (203)-432-9999

ABSTRACT

During the development of a multicellular organism from a zygote, a large number of epigenetic interactions take place on every level of suborganismal organization. This raises the possibility that the system of epigenetic interactions may compensate or "buffer" some of the changes that occur as mutations on its lowest levels, and thus stabilize the phenotype with respect to mutations. This hypothetical phenomenon will be called "epigenetic stability." Its potential importance stems from the fact that phenotypic variation with a genetic basis is an essential prerequisite for evolution. Thus, variation in epigenetic stability might profoundly affect attainable rates of evolution. While representing a systemic property of a developmental system, epigenetic stability might itself be genetically determined and thus be subject to evolutionary change. Whether or not this is the case should ideally be answered directly, i.e., by experimentation. The time scale involved and our insufficient quantitative understanding of developmental pathways will probably preclude such an approach in the foreseeable future. Preliminary answers are sought here by using a biochemically motivated model of a small but central part of a developmental pathway. Modeled are sets of transcriptional regulators that mutually regulate each others expression and thereby form stable gene expression patterns. Such gene expression patterns, crucially involved in determining developmental pattern formation events, are most likely subject to strong stabilizing natural selection. After long periods of stabilizing selection, the fraction of mutations causing changes in gene expression patterns is substantially reduced in the model. Epigenetic stability has increased. This phenomenon is found for widely varying regulatory scenarios among transcription factor genes. It is discussed that only epistatic (non-linear) gene interactions can cause such change in epigenetic stability. Evidence from paleontology, molecular evolution, development and genetics, consistent with the existence of

variation in epigenetic stability, is discussed. The relation of epigenetic stability to developmental canalization is outlined. Experimental scenarios are suggested that may provide further evidence.

Key words: developmental canalization, epigenetic system, epistatic gene interactions, gene networks, neutral mutations, transcriptional regulators

INTRODUCTION

The pathway from a zygote to an adult multicellular organism involves an immense number of components and interactions on each of several levels of sub-organismal organization. There are thousands of genes, each of which may express a spectrum of biologically active molecules, a multitude of regulatory interactions among those molecules, which jointly build the web of cellular organization, and finally an immense number of intercellular communication processes, which create the developmental coordination that produces a finely integrated adult organism. Taken together, they constitute the intricate fabric of the epigenetic system. This complexity invites the question of how changes on the lowest, submicroscopical levels of this production system are translated onto the macroscopic level of the phenotype. How do epigenetic interactions influence the effect of mutant genes and their gene products on the phenotype? Can they absorb or buffer some such effects? If yes, can such an ability to "protect" the phenotype from mutations be subject to evolutionary change? What direction would such change take? These and similar questions shall be addressed here. But why are they relevant in the first place? An essential prerequisite for evolution is (heritable) phenotypic variation, and the evolutionary potential of a group of organisms may often be tied to the availability of such variation. A large amount of phenotypic variation may increase the chances for the rare occurrence of superior phenotypes. The architecture of the epigenetic system will have to be recognized as an important determinant, facilitating or inhibiting, enabling or preventing organismal evolution, if it is found that different epigenetic architectures allow for different amounts of phenotypic variation. Related are many prominent issues, including those regarding the nature and frequency of neutral mutations, as well as the enigmatic maintenance of high amounts of heritable variation in natural populations. Some of them will be discussed in relation to

the findings reported below.

Why are these questions not ranked prominently in current evolutionary biology? Methodological reasons may be responsible. Essential for any buffering of mutation effects is a certain type of epigenetic interactions, namely non-linear (epistatic) interactions, as will be discussed in greater detail below. This type of interaction is most likely pervasive in nature, as Wright (1968) already notes. It represents the rule rather than the exception. However, its analysis and, in fact, the analysis of polygenic systems in general is notoriously difficult. Formal genetic studies of polygenic systems (for an example, see Milkman, 1960) provide invaluable information about patterns of gene interaction, but they may not have the resolving power required for the task at issue. Desirable would be a qualitative and quantitative understanding of all the steps involved in the developmental pathway leading to a phenotypic character of some model organism, an understanding that would require quantitative biochemical data of high accuracy. From such a system, a model could be formulated that allows to answer some of the above questions. Given available biochemical methods, such a model is unlikely to be available soon. An added complication of nonlinear models is that crucial properties of a model may change with subtle alterations in it. Robustness of results, for example to changes in model parameters, becomes a major issue. In spite of all these obstacles it seems worthwhile to seek a reference point for discussion, if one considers the importance of the above questions. In order to provide preliminary answers, I present a biochemically motivated model of an epigenetic process central to development.

The terms "stability of the epigenetic system to mutations" or simply "epigenetic stability" will be used to denote the kind of resilience considered here. Their intended meaning is that a system with high epigenetic stability buffers the phenotype effectively against mutations. The same amount of genetic variation leads to less phenotypic variation in a system with high stability than in one with low

stability. It is thus clearly related to "developmental canalization" in "epigenetic landscapes" (Waddington, 1957), representing the *genetic* aspect of canalization, as will be discussed in greater detail below. However, the term "epigenetic stability" will here be given preference over "genetic canalization," because it emphasizes the epigenetic nature of the phenomenon. It should be clearly distinguished from "phenotypic plasticity," the *environmental* aspect of canalization (e.g., Scheiner, 1993), as well as from "genomic plasticity," the spatiotemporal variability of genome structure (e.g., Gutman *et al.*, 1987, Leblond *et al.*, 1991).

Basic pattern formation phenomena in organismal development involve evolutionarily highly conserved proteins that regulate gene expression on the transcriptional level. Empirical evidence suggests that these transcriptional regulators frequently interact in a network-like fashion to establish patterns of gene expression that, in turn, determine basic *Bauplan* features of the organism (Ingham, 1988; McGinnis and Krumlauf, 1992). Their central role makes such networks and their constituent genes an important and popular subject of research in evolutionary and developmental biology (Ingham, 1988; Kappen *et al.*, 1989; Krumlauf, 1994; Lawrence and Morata, 1994; McGinnis *et al.*, 1990; McGinnis and Krumlauf, 1992; Olson, 1990; Rosenfeld, 1991). The establishment of a gene expression pattern by crossregulatory interactions within a network can be viewed as a small segment of a developmental pathway, and there may be many different ways by which regulatory interactions could establish one specific gene expression pattern (Carroll, 1990; Akam, 1989). The likelihood that a mutation affecting regulatory interactions changes the networks gene expression pattern can be viewed as a measure for the epigenetic stability of that segment of a developmental pathway. With this perspective in mind, a mathematical model for the regulatory interactions within gene networks is formulated and used below to address the following two questions. Are there differences in epigenetic stability among networks that produce the same

gene expression pattern. If so, can natural selection act on this variation? In other words, could epigenetic stability be an evolving rather than an inherent feature of a developmental pathway? The result is unambiguous and robust to changes in all model parameters: evolution towards high stability occurs if the gene expression pattern that a network produces is subject to stabilizing selection. Whether or not such a gain of epigenetic stability is a phenomenon of general biological importance is ultimately an empirical question. Potentially relevant paleontological, genetic and developmental evidence is discussed. General experimental scenarios are suggested that may provide further evidence.

THE MODEL

Only a fraction of the genes encoding transcriptional regulators are likely to be expressed in any given cell and during any given ontogenetic stage of an organism. Moreover, expression patterns of these genes may vary from cell to cell and from stage to stage. The model to be used here refers to the expression pattern of transcription factor genes only in one developmental stage and only in one cell or a body region that shares an expression pattern, such as, e.g., a set of nuclei in a part of a *Drosophila* blastoderm expressing a specific subset of segmentation genes (Ingham, 1988). A subset of N such genes, denoted as $(G_1, \ldots, G_N)$, whose products mutually regulate each other's expression on the transcriptional level will be referred to as a "network." Due to cross-coupling between regulatory pathways (Schüle and Evans, 1991), the number of regulatory proteins (transcriptional regulators and others) involved in most intracellular regulatory circuits is probably large. However, those circuits presumably relying to a large extent on transcriptional regulation may be quite small, involving $10 - 100$ or fewer genes, according to the available circumstantial evidence (Ingham, 1988; He et al., 1989; McGinnis and Krumlauf,

1992).

Regulation of transcription from RNA Polymerase II promoters in eucaryotes is a process in which DNA interacts with multi-protein complexes (Johnson and McKnight, 1989; Mermelstein, 1989). Many of the protein-protein and protein-DNA interactions involved are poorly understood and a considerable number of genes important for the process are probably not even cloned and characterized (Weinzierl *et al.*, 1993). Based on the available empirical data, it seems therefore unlikely that a good qualitative model – let alone a quantitative theory – of transcriptional regulation will be available in the near future. For these reasons and in order to arrive at an analytically and computationally tractable formalism, a number of simplifying assumptions will be used in the model presented here. It is assumed, (i) that expression of the genes in the network is regulated exclusively on the transcriptional level, (ii) that each gene of the network produces one and only one species of an active transcriptional regulator and, (iii) that enhancer elements mediating one regulator's effect on expression of the target gene act independently from enhancer elements for other regulators of the same gene.

In the model used here and motivated in a more formal way by Wagner (1994), a gene network is represented by a dynamical system whose state variables correspond to expression states of the network's genes. They are denoted as

$$\vec{S}(t) := (S_1(t), \ldots, S_N(t)), \tag{1}$$

where $S_i(t)$ is the expression state of the i-th gene at some time $t \geq 0$ during some developmental process in which the network acts. For reasons of computational simplicity, it is assumed that $S_i(t)$ can only assume two values, namely $(+1)$ and (-1), corresponding to a situation in which the gene G_i is expressed or not expressed, respectively, at time t. The gene expression state $\vec{S}(0)$ at time $t = 0$

is called the initial expression state. It can be conceptualized as being imposed onto the network by the products of one or more "upstream" genes that are not themselves part of the network. Such products might be extracellular signaling molecules, such as growth factors or differentiation signals, but also transcriptional regulators, e.g., a retinoic acid receptor acting on homeobox genes in a developing vertebrate limb. The boundary of a network is therefore somewhat arbitrary: a gene is defined as "upstream" by virtue of the fact that it regulates the expression of network genes, but is not regulated by these genes.

Starting from the initial gene expression pattern, $\vec{S}(0)$, cross- and auto-regulatory interactions among network genes cause the expression state to change. These changes are modeled by the set of difference equations

$$S_i(t+\tau) = \sigma[\sum_{j=1}^{N} w_{ij} S_j(t)] = \sigma[h_i(t)]. \tag{2}$$

Here, the expression state of gene G_i at time $t+\tau$, $S_i(t+\tau)$, is a function of a weighted sum, $h_i(t)$, of the expression state of all network genes at time t. $h_i(t)$ represents the sum of the regulatory effects of all network genes on gene G_i. $\sigma(x)$ is the sign function ($\sigma(x) = -1$ for $x < 0$, $\sigma(x) = +1$ for $x > 0$ and $\sigma(0) = 0$), and τ is a time constant characteristic for the process under consideration. Its value will depend on biochemical parameters such as the rate of transcription, or the time necessary to export mRNA into the cytoplasm for translation. The real constants w_{ij} represent the "strength" of regulatory interaction of the product of G_j with G_i, i.e., the degree of transcriptional activation ($w_{ij} > 0$) or repression ($w_{ij} < 0$) that the transcriptional regulator produced by gene G_j has on gene G_i. Such regulatory interactions are known to be mediated by regulatory (enhancer) DNA elements. In biological terms, individual w_{ij}'s can be thought of as some compound measure of the binding constant and the transcriptional activation (repression) ability of the

factor produced by G_j at the enhancer element that mediates its interaction with G_i. Alternatively, and in line with the structure of (2), w_{ij} can be thought of as a measure of the influence that the product of G_j has on G_i relative to other gene products. In this sense, it is the relative size of these constants that is relevant to the dynamics of (2). The i-th row of w, $w_i := \{w_{ij} | 1 \leq j \leq N\}$ corresponds to the entire enhancer of gene G_i with all regulatory DNA elements that affect the expression of G_i. The "connectivity matrix" $w = (w_{ij})$ that the constants w_{ij} define corresponds to all regulatory DNA elements relevant to regulatory interactions among network genes. Any non-zero diagonal element, $w_{ii} \neq 0$, corresponds to autoregulation of G_i by its own gene product (e.g., Regulski *et al.*, 1990; Sucov *et al.*, 1990). Some (or most) entries of w may be zero. The fewer non-zero entries w has, the fewer regulatory interactions exist among network genes. An important model parameter is therefore the fraction of entries different from zero, denoted by c ($c \in (0, 1)$), which will be called the "connectivity density" of the network. The discrete-time dynamical system (2) can also be viewed as the limiting case of a system of differential equations, in which concentrations of gene products, rather than binary (on-off) gene expression states change (Wagner, 1994), but computational limitations prohibited the use of such a system here. It should be noted that the structure of (2) is similar to "spin glass" (Binder and Young, 1986) or neural net (Amit 1989; Hopfield and Tank, 1986) type models of gene networks first introduced by Kauffman (1969, 1993). However, (2) is conceptually different from Kauffman's models in that a specific type of gene interaction, namely transcriptional regulation, is considered. Recently, models conceptually similar to (2) have been successfully used to describe and predict regulatory gene interactions in early *Drosophila* embryogenesis (Mjolsness *et al.*, 1991; Reinitz *et al.*, 1995; Reinitz and Sharp, 1995).

The dynamics of (2) will lead to the attainment of an equilibrium gene ex-

pression state, which may be a fixed point of (2) or a limit cycle. For reasons of tractability, only fixed point equilibria will be considered here. They are denoted by $\vec{S}(\infty)$. Below, the amount of time that a network takes to reach such an equilibrium when starting from $\vec{S}(0)$ will be of importance. This amount of time is also indicative of the length of the path to equilibrium in the space of all possible gene expression states. It will therefore be referred to as the pathlength to equilibrium. The genes expressed in the equilibrium state will affect the expression of genes outside ("downstream") of the network. Possible downstream genes may include structural genes or genes encoding proteins involved in signal transduction processes, but also transcriptional regulators that do not themselves regulate the expression of genes within the network (Budd and Jackson, 1991). Many experimental genetic studies (e.g., McGinnis and Krumlauf, 1992) suggest that deviations in the expression pattern of a gene network from the wild type pattern causes developmental perturbations that often lead to deleterious effects on the adult phenotype. In a biological population, such individuals would be eliminated by natural selection. These observations motivate the assumption that an optimal equilibrium gene expression state, denoted as $\vec{S}^{opt}(\infty)$, exists for networks acting in a developmental process. If a network attains an equilibrium state $\vec{S}(\infty)$ that is different from this optimal state, developmental perturbations will result and the fitness of the respective adult organism will be reduced. A networks equilibrium expression pattern may deviate from an optimal pattern for a variety of reasons, one of them being mutations. Mutations may affect a network in various ways. Mutations in downstream genes may alter the effect that a network's gene expression pattern has on these genes. Mutations in upstream genes may influence the initial state $\vec{S}(0)$, and mutations in transcription units of network genes will change the interaction pattern inside the network in a global and drastic way. It is, however, the matrix w defining the regulatory interactions within the network that represents its most interesting "organizational" properties, and

thus mutations in regulatory DNA regions, represented by changes in this matrix, will be the focus of this study.

The evolutionary scenario envisioned here involves a gene network acting in an ontogenetic process in each of the members of a population of organisms. It is assumed that both the initial gene expression pattern, $\vec{S}(0)$, and the optimal equilibrium gene expression pattern, $\vec{S}^{opt}(\infty)$, are the same for all organisms in that population. The organisms are subject to mutations of regulatory DNA regions, recombination among network genes, and genetic drift. They are also subject to stabilizing selection on the optimal gene expression state, $\vec{S}^{opt}(\infty)$. This requires a notion of an individual's fitness, which is modeled in the following way. A measure for the distance d between the equilibrium state $\vec{S}(\infty)$ attained by a network and the population's optimal equilibrium state is defined by $d[\vec{S}^{opt}(\infty), \vec{S}(\infty)] := \frac{1}{2} - \frac{1}{2N}\sum_{i=1}^{N} S_i(\infty)S_i^{opt}(\infty)$. This measure is known as the Hamming distance (Amit, 1989). It counts the number of expression states of individual genes that are different in the two states and normalizes it to the interval $(0, 1)$. Based on d, the fitness of an individual is then defined via a Gaussian fitness function as

$$\exp\left(-\frac{d[\vec{S}^{opt}(\infty), \vec{S}(\infty)]^2}{s}\right), \tag{3}$$

The parameter s ($s > 0$) represents the strength of selection, small values of s implying strong selection against deviations from the optimal state. Two principal questions will be asked in this scenario. First, consider only those networks within a population that attain an equilibrium expression state, $\vec{S}(\infty)$, that is identical to the optimal gene expression state $\vec{S}^{opt}(\infty)$. Is there variation in the stability of these networks to mutations, i.e., are there networks in which mutations are more likely to cause a change in $\vec{S}(\infty)$ than in others? Second, if such variation exists, is it heritable, and is it thus possible to select for high epigenetic stability?

To develop an analytical theory for this problem is a formidable task, given that dynamical systems like (2) in and by themselves pose difficult analytical problems (Amit, 1989, Binder and Young, 1986), and given the added dimension of a population level process acting on the networks. Therefore, a numerical approach is used here, in which a population of (initially) identical gene networks is generated from one "founder" network in a way detailed in the next section. The optimal gene expression state within the population is the equilibrium gene expression state of this "founder" network. Then, networks within the population are allowed to diverge by subjecting them to mutation and recombination, while exercising stabilizing selection on the optimal gene expression pattern, until the mean fitness of the population has reached a quasi-equilibrium. The evolution of epigenetic stability among members of this population is studied during this process.

Little is known empirically about "typical" initial and equilibrium expression states, as well as about patterns of regulatory interactions within gene networks. Starting a population simulation with one type of network that has a prespecified initial, equilibrium state and connectivity matrix would therefore require many *ad hoc* assumptions about biologically significant network features. This is in part circumvented here by pursuing a statistical approach involving infinite sets E ("ensembles") of networks, each with its own initial state, equilibrium state and matrix w of regulatory interactions. Each network within an ensemble is the starting point of an evolution scenario as outlined above, during which evolution of stability in a population of organisms is monitored. This permits one to assess how sensitive results are to variations in network features. In this approach, the ensemble is a more suitable level to characterize network structure than the individual network, and network characterization is therefore statistical in nature. The following parameters are used to characterize a network ensemble. First, all networks within an ensemble have the same number of genes, N. Second, among networks within

an ensemble, the number of genes expressed in the initial state follows a binomial distribution $B(N, p)$. In other words, if one were to choose randomly (with uniform distribution) a network from the ensemble, the probability of finding a specific initial state $\vec{S}(0)$ with k expressed genes would be given by $\binom{N}{k} p^k (1-p)^{N-k}$. This also implies that the mean number of expressed genes in the initial state is given by Np. Expressed genes in the equilibrium state follow the same distribution, but are stochastically independent from those expressed in the initial state. Third, connectivity matrices w within the ensemble are defined by a probability distribution $\rho(w)$ of regulatory interaction strengths (specified below), with a mean fraction c of connectivities different from zero among networks within the ensemble. The number of these connectivities varies from network to network. Since a numerical analysis is carried out here, only finite samples of a network ensemble can be studied. Such samples are generated by a random search in the space of all possible gene networks with given N, p, c and ρ. The important question regarding robustness of results to variation in model features is addressed by numerically generating several ensemble samples with different values of N, p, c and ρ. As will be discussed below, results might also be sensitive to the specific way in which mutations are modeled. Therefore, two different models for introducing changes into connectivity matrices w are used to provide additional support for robustness of the results.

NUMERICAL METHODS

This section and the appendix describe how ensemble samples were generated, how evolution was simulated, and how epigenetic stability was assessed.

Generation of Individual Networks and of Ensemble Samples.

At the outset of this study are numerically generated finite samples of infinitely large network ensembles E. Each network in an ensemble (sample) has its own

initial state $\vec{S}(0)$, equilibrium state $\vec{S}(\infty)$, and connectivity matrix w. Sample sizes ranged from 150 to 300 networks, and are given along with results in the next section. Because sample members were generated independently from each other, it is sufficient to describe how one member was generated.

Figure 1 outlines the procedure used to generate individual ensemble members. First, two binary pseudorandom arrays in $\{-1, 1\}^N$, corresponding to $\vec{S}(0)$ and $\vec{S}^{opt}$, were generated. Individual entries of these arrays were stochastically independent both within each array and among arrays. Each entry was chosen according to the probabilistic rule $\mathbf{P}(S_i = 1) = p$, p being a real number in $(0, 1)$. It follows that the number of entries equal to $(+1)$ in each array is binomially $(B[N, p])$ distributed, with mean Np. Then, a $N \times N$ pseudorandom matrix $w = (w_{ij})$ with independently and identically distributed entries was generated. Individual entries of this matrix were different from zero with probability c, and entries different from zero followed a continuous probability distribution $\rho(w_{ij})$ (Gaussian and "reflected" gamma distribution, see appendix). Subsequently, network dynamics (2) was carried out using $\vec{S}(0)$ as initial state and w as connectivity matrix. If the network did not attain the equilibrium state $\vec{S}(\infty)$ as a fixed point, the matrix w was "discarded," a new matrix w was generated in the same way as the old one, and the network's dynamics was evaluated again. New matrices were generated in this way until a matrix had been found for which the network attained $\vec{S}(\infty)$ as a fixed point, or until 5×2^{2N} matrices had been generated, whichever came first. If no matrix had been found after 5×2^{2N} trials, the pair of states $\vec{S}(0)$ and $\vec{S}(\infty)$ was discarded, a new pair of state arrays $\vec{S}(0)$ and $\vec{S}(\infty)$ was generated in the same way as the old pair, and a stochastic search for a matrix w was carried out in the same way as for the old pair of states. Once a matrix was found for which (2) attained $\vec{S}(\infty)$ as a fixed point with $\vec{S}(0)$ as initial state, the triplet $(\vec{S}(0), \vec{S}(\infty), w)$ was considered a member of the ensemble sample. Further details of the sample generation procedure

are given in the appendix.

The following network features were kept constant while generating one ensemble sample: N, the number of network genes, p, the (mean) fraction of genes expressed in initial and equilibrium state, c the (mean) fraction of connectivities different from zero, and $\rho(w_{ij})$, the probability distribution of non-zero connectivities w_{ij} within the ensemble. However, for different ensemble samples, different values of these parameters were used ($4 < N < 10, 0.1 < p < 0.9, 0.4 < c < 1.0$, different distribution types for ρ, see appendix) to assess robustness of results to variations in these parameters.

Simulated Evolution. From each member of an ensemble thus generated, a population, i.e., an array of $P = 500$ identical copies of the ensemble member was generated. A simulated evolution process involving mutation, selection, genetic drift and, in some cases (see appendix), recombination was carried out for this population. Since an independent evolution simulation was carried out for each ensemble member, it is sufficient to describe this simulation for only one member. The following four processes were carried out in the given order. Each iteration of this sequence of processes was considered one generation of simulated evolution.

1. Recombination: In pairs of consecutive matrices in the population, starting with the first pair of matrices, rows were swapped with probability 0.5, corresponding to free recombination between genes and tight linkage among regulatory elements within a promoter. Note that randomness in the order of matrices is already implied by the selection algorithm used (step four).

2. Mutation: Each non-zero entry in each connectivity matrix of the population was replaced with probability $1/(cN^2)$ by a pseudorandom number distributed according to the same continuous probability density ρ as that used in the generation of the ensemble sample. This approach is essentially the house-of-cards assumption used in many models of population genetics (e.g., Turelli, 1985; Zeng

and Cockerham, 1993).

3. Fitness evaluation: Network dynamics (2) was evaluated numerically for each network in the population, using the same initial state $\vec{S}(0)$ for all networks. If a given network reached some fixed point $\vec{S}(\infty)$ after $3N$ or fewer time steps (see also appendix), its fitness was evaluated using (3) with $s = 0.1$. In this calculation, the reference (optimal) gene expression state $\vec{S}^{opt}(\infty)$ for all networks in the population was the equilibrium state of the network that had "founded" the population at the beginning of the simulation. If a network had not reached equilibrium after $3N$ time steps, its fitness was assigned the minimally possible value of $\exp(-1/s)$, thus making it very unlikely that the network survived the subsequent "selection" step.

4. Selection: Fitness of networks was normalized such that the maximum fitness in the population was equal to one. Then, a network was chosen at random and a pseudorandom number, r, with uniform distribution on $(0, 1)$ was generated. If r was smaller than the fitness of the individual, the individual "survived." This process was repeated until a new population of the same size as the old population had been generated, i.e., sampling of networks was carried out with replacement.

Assessment of Epigenetic Stability Before and After Simulated Evolution.

The stability of networks to changes in their connectivity matrices was assessed before evolution and after 400 generations of simulated evolution. Since epigenetic stability might depend on the kind of changes caused by a mutation, it is necessary to verify that results are not artifacts of the way mutations are modeled. For this reason, the sensitivity of networks to changes in their connectivity matrix was assayed by two different means, "mutation" and "orthogonal perturbation." A "mutation" consisted, as above, in the replacement of a non-zero connectivity by a random variate with the same probability distribution ρ as that used in the generation of the ensemble sample. The method of "orthogonal perturbation" is

described in detail in the appendix.

The same approach was taken to evaluate the stability of networks with respect to both types of perturbations, mutations and orthogonal perturbations. For each member of an ensemble sample, one perturbation was introduced and it was tested whether the network, starting at $\vec{S}(0)$, attained the same equilibrium state, $\vec{S}(\infty)$, as before perturbation. This process was repeated $(5000c)$ times with the original ensemble member, and the fraction of perturbations that led to the attainment of the stable equilibrium state was termed the stability of that network *before evolution* with respect to the kind of perturbation used. The corresponding assay after 400 generations of simulated evolution of the population generated from that ensemble member was carried out in a slightly different way. After step 4 ("Selection") in the last generation of the simulated evolution process, the number of networks in the population that attained the equilibrium state $\vec{S}^{opt}(\infty)$, N_{opt}, was determined. For each of these N_{opt} networks, sensitivity of the network to perturbation was assayed in the same way as that of the original ensemble member, except that $(c/N_{opt})5000$ perturbations were used *per network*, thus ensuring that the assay before and after evolution was based on approximately the same *total* number of perturbations. The fraction of perturbations that lead to a change in the trajectory of the networks, such that $\vec{S}^{opt}(\infty)$ was not attained, was taken as a measure for the (average) stability of the network to the respective perturbation *after evolution*. Strictly speaking, this measure is a feature of the population and not of an individual network (as opposed to the measure of stability before evolution), but since $\vec{S}(0)$ and $\vec{S}^{opt}(\infty)$ are identical for all members of a population, it will loosely be referred to as the stability of a network after evolution.

RESULTS

Epigenetic stability varies among gene networks and it can evolve. Figure 2a shows epigenetic stability before and after a simulated evolution process in networks of $N = 10$ genes with a density of regulatory interactions, c, equal to one, i.e., the expression of any given network gene is regulated by all other network genes. Values of stability shown on the abscissa of this figure demonstrate considerable variation in stability among networks within the ensemble sample, i.e., before evolution. Values on the ordinate demonstrate that stabilizing selection, mutation and recombination increase and homogenize epigenetic stability in the populations derived from individual ensemble members. In other words, individuals in the simulated populations have evolved the ability to "absorb" mutations. Their pattern of epigenetic interactions acts as a "buffer" that prevents the effects of most mutations from becoming visible as altered equilibrium gene expression states. Figure 2b shows that a qualitatively identical result holds if stability to "orthogonal perturbations" instead of stability to mutations is assayed, demonstrating that the observed phenomenon is not an artefact of a specific way of modeling mutations.

Results are robust to changes in model parameters.

Important parameters of the model include the number of genes in a network, N, and the density of regulatory interactions within the network, c. Figures 3a-c show the effect of varying the connectivity density c on the evolution of stability. It can be seen that, as c decreases from one, both mean stability after evolution and variability of stability before evolution decrease. Figure 4a summarizes the results of Figures 2a and 3a-c, showing mean and standard deviation of stability before and after evolution. "Error" bars in this figure indicate differences in variation of stability before and after evolution, and not only "significance" of differences in mean stability. Figure 4b shows results corresponding to those in Figure 4a, but for "orthogonal perturbation" instead of mutation, demonstrating again that the phenomenon observed here is not an artefact of a specific way of modeling mutations.

Figure 5 shows results analogous to Figure 4, but here the number of genes, N, is varied, while the connectivity density c is kept constant at $c = 1$. In sum, it can be said that both a decrease in the density of regulatory connections and in the number of network genes lead to a reduced increase in epigenetic stability during evolution, while not changing the qualitative finding that such an increase takes place. Varying two further important characteristics of network structure, the mean number of expressed genes, Np, and the distribution type, ρ, of individual connectivities yields increases in stability comparable in magnitude to those reported above (results not shown).

A decrease in (i) the number of regulatory interactions per gene or (ii) the number of network genes causes a decrease in stability to mutations after evolution (Figures 4 and 5). This behavior can be qualitatively understood from the structure of the network model (2). The sum of all regulatory influences on gene G_i at time t, $h_i(t) = \sum_{j=1}^{N} w_{ij} S_j(t)$, determines G_i's expression state at time $t + \tau$. Changing the value of individual regulatory connections w_{ij} by mutation will change the value and, sometimes, the sign of $h_i(t)$. It can be shown (A. Wagner, 1994), that the probability that such an event changes the sign of $h_i(t)$ is proportional to the probability that the equilibrium gene expression pattern of a network is changed. The probability of a sign change in $h_i(t)$, in turn, is clearly determined by the number of non-zero products $w_{ij} S_j$ that contribute to this sum. The larger the number of these contributions, the smaller the effect that an individual mutation, i.e., a change in one connectivity w_{ij}, has on the whole sum, $h_i(t)$. The number of non-zero products $w_{ij} S_j$ in $h_i(t)$ depends on (i) the number of genes in the network and (ii) the density of regulatory connections, i.e., the number of non-zero w_{ij}'s in the network. In other words, as network size or connectivity density is decreased, the influence that individual mutations have on $h_i(t)$ increases, and, thus, the likelihood that a mutation changes the equilibrium gene expression pattern increases as well.

The number of regulatory connections and the number of network genes set an upper bound for epigenetic stability that can not be exceeded. A network with many connections or many genes can absorb higher mutational pressure than a network with few connections or few genes.

The absence of a rigorous analytical understanding of this system does not preclude the possibility to characterize networks of high and low stability phenomenologically. What happens in a population of gene networks during the evolution of high stability? The following numerical analysis provides a partial answer. It is based on networks with $N = 10$ genes that are densely connected, i.e., $c = 1$. The phenomena observed for theseparameter values are representative of and exemplify those observed for different values, which are therefore not shown.

Some pathlengths occur more frequently than others. The pathlength of a member network of the ensemble E is defined as the number of time steps of (2) that the network takes to get from $\vec{S}(0)$ to $\vec{S}(\infty)$. Figure 6 shows the distribution of pathlengths in an ensemble sample of 1000 networks, i.e., before evolution, depicting the relative frequency of each pathlength. It is clear that not all pathlengths occur with equal frequency. Networks with pathlengths 2 to 5, for example, occur more frequently than networks with pathlength 10 to 12. Since the network search algorithm used to generate the sample is unbiased with respect to pathlength, one can conclude that networks with certain (short) pathlengths occupy a much larger volume in the space of all possible genotypes, than networks with other (longer) pathlengths. This observation invites the question whether epigenetic stability and pathlength are correlated.

Epigenetic stability is coupled to pathlength. Figure 7a shows a diagram of epigenetic stability vs. pathlength in a sample of E, i.e., before evolution. It is clear that there is a correlation between pathlength and epigenetic stability. The

estimated Pearson correlation regression coefficient is $r \approx -0.79$. Figure 7b is included to demonstrate that this feature is not likely to be an artefact of the way mutations are modeled ($r \approx -0.49$). These results imply that mutations are more likely to have deleterious effects in a network with a long path to equilibrium than in a network with a short path. This raises the possibility that networks with high pathlength are replaced with networks of lower pathlength during evolution, since the latter are less likely to mutate to suboptimal variants.

Evolutionary forces change pathlengths in a population. Figure 8a shows mean and standard deviation of pathlength during 400 generations of simulated evolution in a population of networks with $N = 10$ and $c = 1$. In generation zero, this population consisted of 500 identical networks, all of which were copies of one network with pathlength 10. It is obvious that a significant reduction in pathlength occurs in less than 200 generations of network evolution. Figure 8b shows the distribution of pathlengths before and after evolution in a sample of E. It demonstrates that a process similar to that shown in Figure 8a occurs also on the ensemble level.

From the results presented so far, one can also infer that the evolution of pathlengths can not fully account for the evolution of high stability. To see that, consider a set of networks with a given pathlength before evolution, e.g., networks from Figure 8a ($N = 10$, $c = 1$) with pathlength 3. It is evident from the figure that these networks show variation in epigenetic stability (the coefficient of variation is $\sigma_x/\bar{x} \approx 0.129$). This demonstrates that there are differences in stability among ensemble members with a given pathlength. Further, the mean pathlength of networks after evolution in Figure 8b ($N = 10$, $c = 1$) is 2.46, a value close to 3. Variation in epigenetic stability for networks after evolution is low, as can be seen from Figure 2a. The coefficient of variation in stability after evolution for networks

in this figure is $\sigma_x/\bar{x} \approx 0,026$, much lower than that for networks with pathlength 3 before evolution. Thus, stabilizing selection eliminates not only pathlengths with low stability, but also networks with a given pathlength and low stability.

Evolution of stability and recombination. A question that has not yet been addressed concerns the influence of recombination on the evolution of stability. All results shown so far involved free recombination among network genes. Is recombination required for the observed phenomena? The answer is no. For example, Figure 9a shows the evolution of mean and standard deviation of pathlength during 400 generations of simulated evolution in a population of networks ($N = 10$ and $c = 1$). It is completely analogous to Figure 8a, except that no recombination occurred in the population of Figure 9a. Similarly, Figure 9b shows results on the ensemble level, analogously to Figure 9b but without recombination. It is clear from the figures that the presence or absence of recombination makes little difference for the evolution of pathlengths. Epigenetic stability before and after evolution is also virtually identical for the two scenarios, and to avoid further redundancy no separate graphs are shown for the asexual case.

The mere observation that selection is effective in the presence of recombination suggests that recombination leaves important network properties unchanged. This may seem surprising, given that recombination causes profound "systemic" changes in a network. One example for such a property is pathlength to equilibrium. One might suspect that pathlength, being an "emergent" network feature depending on the interaction of all network genes, would be randomized by the random shuffling of the genes involved in recombination. However, the observation that mean pathlength in a population changes under the influence of selection suggests that pathlength is a "heritable" feature of such networks. That this is indeed true is demonstrated by the numerical example in Figure 10. Shown is a plot of

mean parental pathlength vs. mean offspring pathlength in generation 200 of a simulated evolution process initialized by an ensemble member with pathlength 10 (heritability $h^2 \approx 0.66$, $P < 0.01$).

In sum, it can be said that stabilizing selection increases epigenetic stability in the model of gene networks used here and it does so regardless ofnetwork architecture and the mode (sexual/asexual) of reproduction. Mutations in networks with low stability and high fitness produce many genotypes with low fitness. However, they also produce some genotypes with high fitness and higher stability than the original network. Networks with low fitness are eliminated, but networks with high fitness and high stability accumulate, since they produce fewer suboptimal variants. Thus, networks with high stability and high fitness replace networks with low stability and high fitness in a population. In this model, stability is associated with short paths to equilbrium gene expression states of a gene network. The selection process at work acts in part on these pathlengths.

DISCUSSION

Networks of transcriptional regulators guide important developmental pattern formation processes in various organisms, and the ubiquity of transcriptional regulation suggests that important developmental functions of such gene networks still await discovery in many other organisms. Most significantly, gene networks are involved in the patterning of characters of the organismal body plan. These characters are highly conserved and show little naturally occurring variation. In this paper, it is investigated whether such conservation may in some cases reflect an evolved resilience, termed "epigenetic stability," of a developmental pathway to mutation. Since this question is difficult to address experimentally, a mathematical model for gene networks was used here to provide preliminary answers. Using this model, it

was found that epigenetic stability varies among networks of transcriptional regulators. In some gene networks, almost all mutations have a phenotypic effect, whereas in others almost none do. The following phenomenon is observed for all except those very few networks who have exceptionally high stability to begin with. In a simulated population of organisms in which a network acts and in which selection favors a specific, "optimal" gene expression pattern, epigenetic stability is substantially increased in the course of an evolution process involving many generations of mutation and selection. What happens during evolution? Epigenetic stability is a property of the epigenetic system, i.e., the pattern of epigenetic interactions represented here by regulatory interactions between transcriptional regulators and the genes that encode for them. Subjected to a certain amount of mutational pressure, a network with low stability will produce a larger number of deleterious mutants (i.e., gene networks with expression patterns other than the optimal pattern) than one with high stability. Occasionally, some mutant genotypes are produced that are phenotypically neutral (i.e., they leave the gene expression pattern unchanged), but have higher epigenetic stability than the original genotype. Thus, a quite indirect process is at work, in which genotypes with low epigenetic stability are eliminated from a population due to the large fraction of maladapted offspring they produce. A reorganization of the epigenetic system towards high stability to mutations takes place.

Results are robust to changes in model features. Are the findings just described likely to apply to real gene networks? The two issues relevant to this question are robustness of results to changes in network architecture and to relaxation of assumptions made in the model. Network architecture is specified by model parameters, such as the number of genes in a network, the number of transcriptional regulators influencing the expression of any given gene, the pattern of regulatory interactions and the number of genes that are transcriptionally active. It is reassur-

ing that the results summarized above are qualitatively independent of variations in these aspects of network architecture. Interestingly, the largest increase in stability is seen in network types with large numbers of regulatory interactions. The observed phenomenon does not seem to be an artefact of the way in which mutations are modeled, since two quite different ways of introducing changes in networks yield qualitatively identical results. Further, the mode of reproduction (sexual reproduction with recombination or asexual reproduction) seems to be of little significance. The evidence presented implies that epigenetic stability can be heritable in sexually reproducing populations. Also, note that a conscious decision to use a haploid model was made here, since in a haploid model no relevant genetic variation can be hidden in relations of dominance and recessiveness among alleles. It is not clear *a priori* whether diploidy would further increase the magnitude of the effects observed here, since evolution of specific dominance relations among alleles seems possible in a model like this. However, it is unlikely that diploidy would diminish these effects. A further important simplifying assumption was that the regulatory influences of different transcription factors act additively in determining the activation state of the gene whose expression they regulate, an assumption reflected by the additive superimposition of regulatory contributions in (2). Because non-additivity is responsible for the phenomena reported here, a change in this aspect of the model would most likely not eliminate the observed effects. They should rather become more pronounced.

Epistasis and the evolution of epigenetic stability. Figure 11a illustrates a metaphor that may be useful to visualize the results discussed above. In this figure, shaded regions correspond to optimal genotypes within the rectangular box representing all possible genotypes, i.e., "genotype space." In terms of the model, the matrix w of regulatory interactions within a gene network is represented by a point in the space of all possible matrices (genotypes) w. The matrix w of a

network that attains its optimal gene expression pattern is represented by a point on one of the shaded "islands." Mutations in regulatory DNA sequences change connectivity matrices, such that the matrix w' generated from a matrix w by mutation occupies a different location in genotype space. The mean displacement in genotype space caused by a mutation is a measure for the extent of this change. An evolving population of organisms can be viewed as a moving cloud of points in genotype space, and one factor determining the extension of this cloud is the mean displacement caused by a mutation event. Consider, for example, population A in Figure 11a, which consists of some optimal and some suboptimal genotypes (the mean displacement caused by a mutation is indicated by the circle drawn around one of the individuals in the population). Evolution of high epigenetic stability takes place when such a cloud moves away from the edge and towards the center of an island. After evolution of high stability (e.g., population B in Figure 11a) mutations generate fewer suboptimal genotypes. Population C illustrates that the topography of the regions of optimal genotypes and the extent of change caused by mutations jointly determine whether evolution of stability is possible. In relation to the extension of population C, the island it occupies is too small to allow evolution of epigenetic stability. The image presented here is clearly only a caricature of a real, high-dimensional genotype space, since high-dimensional spaces may have complicated and counterintuitive topological features. It does, however, allow one to illustrate that epistatic (non-linear) gene interactions are essential for the evolution of epigenetic stability. Consider a system analogous to the one considered here, but in which a phenotype is determined by additive contributions from some genotype w. In this case, the islands of Figure 11a would be replaced by a hyperplane in genotype space (Figure 11b), i.e., an infinitely thin set of zero volume. From the point of view of an evolving population, this set shows no local differences in epigenetic stability that are "visible" to the population. Evolution of epigenetic

stability becomes impossible here.

A time scale for the evolution of epigenetic stability? Among the factors that may influence the evolution of epigenetic stability are (i) population sizes, (ii) the strength of stabilizing selection, and (iii) the number of genes and the gene interactions involved (here, only a minute segment of a developmental pathway was modeled). Given that these factors may vary widely and that they represent in general unknown variables, it is difficult to estimate a characteristic time scale for the evolution of stability. However, a crude calculation for the network model used here serves to show that any increase in stability for a comparable real gene network is likely to take place on a macroevolutionary time scale. Consider a network with $N = 10$ genes, a density of regulatory interactions of $c = 0.2 - 1.0$, and assume that one regulatory enhancer element corresponds to 20-100 base pairs of DNA. If mutations affecting these elements occur at a rate of approximately 10^{-9} per base pair per generation, and if one takes into account the mutation rates used for simulated evolution here, as well as the time course of evolutionary change in the model (Figure 8a), one arrives at a range of $10^7 - 2.5 \times 10^8$ generations during which most of the change in stability would take place.

Paleontological data and the evolution of stability. Indirect evidence may be required to establish the biological significance of the results obtained here, because the time scales involved are potentially large. At first sight, paleontological data might seem to provide such evidence. The rapid radiation of morphs associated with the origin of many taxa is often followed by evolutionary stasis. This might hint towards a process analogous to the one described here. However, many lines of evidence suggest that major radiations are associated with new ecological opportunities (e.g., Erwin *et al.*, 1987, Futuyma, 1986, Jablonski and Bottjer, 1990, Simpson, 1944, Valentine and Erwin, 1987), whereas only circumstantial and anecdotal evidence for the importance of genetic factors is available (for an exam-

ple involving the biogeography of sea snakes and vampire bats see Futuyma [1986, pp256]). Thus, although the potential importance of genetic factors is discussed and acknowledged in the literature (Erwin *et al.*, 1987, Jablonski and Bottjer, 1990, Valentine and Erwin, 1987), conclusive evidence is lacking, and paleontological data may be ill-suited to discriminate between ecological and genetic factors.

Why are neutral mutations neutral? Selection against genotypes with low epigenetic stability is responsible for the evolution of stability. This finding, if generalizable, adds another facet to the important role that natural selection plays in organismal and molecular evolution (e.g., Gillespie, 1991). However, the way in which selection acts here is unusual, in that selection itself is responsible for the evolution of genotypes in which mutations are likely to be neutral. This "selected" or "evolved neutrality" also raises a fundamental question about the nature of neutral mutations. Are mutations phenotypically neutral because they have no effect on the activity of a gene product (such as an enzyme activity), or are they neutral because the effect they have on a gene product is buffered by epigenetic interactions? In the following, these hypothetical types of neutrality will be referred to as molecular neutrality and epigenetic neutrality, respectively. The difference between them is well illustrated by the network model.Here, mutation events do not correspond to mutations in protein coding sequences, but to mutations in the binding sites of proteins at an enhancer. In general, these mutations are far from molecularly neutral. In fact, in the mutation model used here, the strength of binding of a protein at an enhancer site before mutation is completely uncorrelated to that after mutation. However, in a gene network with high epigenetic stability, most such molecularly non-neutral mutations are phenotypically neutral. Thus, here epigenetically neutral mutations are mainly responsible for epigenetic stability.

How important could epigenetic neutrality be in comparison to molecular neutrality? Insight into the structure-function relationships of proteins suggests that

protein function may be very tolerant to amino acid substitution (e.g., Bowie *et al.*, 1990 and references therein). Thus, molecularly neutral mutations will represent an important, and maybe the predominant class of phenotypically neutral mutations. Much less is known in this regard about mutations in regulatory DNA sequences, a type of mutation that has probably been important in organismal evolution. Also, molecularly non-neutral mutations with biochemical effects strong enough to overwhelm the buffering capacity of the epigenetic system will most likely be possible in any developmental pathway. Moreover, the relative importance of molecular and epigenetic neutrality may crucially depend on the specific developmental pathway and/or biochemical pathway considered. Especially in comparatively simple genetic systems, such as viruses, there may be little leeway for epigenetic neutrality.

What kind of evidence is there and what kind of evidence would demonstrate the importance of epigenetic neutrality? First, note that carefully designed studies that test for the tolerance of proteins to amino acid substitutions have been carried out (e.g., Bowie *et al.*, 1990 and references therein). However, these studies test primarily for phenotypical neutrality, since they identify non-neutral mutations by the suboptimal phenotypes they produce or by evolutionary invariable amino acid residues in proteins. They do not distinguish between epigenetic and molecular neutrality. Such a distinction could be made by experimentsshowing that there are phenotypically neutral mutations that are not molecularly neutral, quite tedious experiments, to be sure, since biochemical activities of mutant gene products would have to be assayed mutant by mutant. The fact that some gene products, especially regulatory proteins, have different functions in different developmental contexts is not likely to facilitate this type of analysis. Fortunately, other lines of evidence provide relevant information as well, although much more is needed before definite conclusions will be possible. Graur and Li (1988) report that some serine protease inhibitors in mammals evolved faster at functionally important sites than at func-

tionally neutral sites. While an unusual amino acid composition at the functional sites is invoked in the author's explanation for this apparent anomaly, a different reading might be possible. If epigenetic stability for mutations in these proteins is high, some of the substitutions may have represented epigenetically neutral but molecularly non-neutral mutations. The rate variations in the "molecular clock" induced by phenotypically neutral nucleotide substitutions provides further hints. Variation in generation time among taxa has been invoked in an explanation of this phenomenon, but it cannot account for all variation in substitution rates (Kimura, 1987 and references therein). Kimura (1987) touches upon the phenomenon at issue here by proposing the "alteration of the selective constraint of each molecule (due to change of internal molecular environment)" (pp24) as one possible explanation. He analyzes two candidate examples, guinea pig insulin and opossum hemoglobin α, whose patterns of variation in substitution rates can not be accounted for by variations in generation time.

Epigenetic stability and the evolution of developmental pathways. How would epigenetic stability be affected if a change in a developmental pathway leading to some (quantitative) phenotypic character occurs? Note that there is ample evidence for such variation. Homologous characters in different taxa can be the product of different developmental pathways (e.g., G.P. Wagner, 1994). To illustrate the effect of such variation, it is useful to consider results from the gene network model in conjunction with the "island" metaphor of optimal genotypes (Figure 11a). In terms of the model, a modification in a developmental pathway consists in a change in the developmentally optimal gene expression pattern of a network, a change that is compensated by subsequent developmental events. This implies that the set of islands defining the optimal genotypes in genotype space (Figure 11a) has changed as well. As a result, epigenetic stability may be reduced, in which case the process of reorganizing the pathway's system of epigenetic interac-

tions towards high epigenetic stability has to begin anew. This scenario generates a testable hypothesis. Consider a (quantitative) phenotypic character subject to stabilizing selection, and the developmental pathway leading to this character. If a phylogenetically recent change in this pathway occurred in a taxon, the new pathway should have lower epigenetic stability than the ancestral pathway, and one should be able to observe more genetic character variation in the new taxon than in an ancestral taxon. A similar line of reasoning can be applied, if mutants affecting a developmental pathway produce variant phenotypes that are outside the normal range of phenotypic variation in a population, but still viable and fertile. If populations are established from such deviant genotypes, such as in an artificial selection experiment, more phenotypic variation with a genetic component should be observed in the respective character(s). Rendel (1979) states that this is a general feature of mutants, exposing "hidden genetic differences between members of a uniform population" (pp 140) and he gives concrete examples, one of which concerns the number of secondary mustacial whiskers in the house mouse. Their number is usually 19, with little variation. In mutants of the sex-linked gene *Tabby*, the number of whiskers is substantially reduced and varies considerably. That this variation has a genetic component is shown by the effectiveness of directional selection on the number of whiskers (Rendel, 1979, Dun and Fraser, 1958). A second, analogous example concerns changes in the number of scutellar bristles caused by the mutant *scute* of *Drosophila melanogaster* (Rendel, 1979).

The phenomena observed here are clearly related to "developmental canalization" in "epigenetic landscapes" (Waddington, 1957). Canalization can be conceptually partitioned into two components, an environmental component that is responsible for the buffering of a developing organism against influences from the external and the internal ("developmental noise") environment, and a genetic component that determines to what extent genetic variation is transformed into phe-

notypic variation. While the environmental component may well be genetically determined, only the genetic component, often termed "genetic canalization", relates to epigenetic stability. In fact, "genetic canalization" and "epigenetic stability" are synonymous, but the latter term was used here, because it emphasizes the epigenetic nature of the phenomenon. Unfortunately, most available experimental data pertains to the environmental aspect of canalization (Rendel, 1979; Waddington, 1957). However, Stearns and Kawecki (1994) were able to demonstrate recently that differences in epigenetic stability exist among several life history traits in *Drosophila melanogaster*. As one would predict, they find a high degree of canalization in traits that are closely correlated with fitness.

Some experimental evidence discussed here argues for variation in epigenetic stability as well as for its evolution. This evidence is far from conclusive, however, and much empirical work would be needed to provide firm empirical footing for the proposed mechanism. Fortunately, this mechanism generates a series of testable predictions that will be useful in assessing the biological significance of epigenetic stability and its evolution.

ACKNOWLEDGMENTS

The author would like to thank L.W. Buss, R. Vaisnys and G.P. Wagner for constructive criticism and invaluable suggestions on earlier drafts of the manuscript. This paper is communication number 18 from the Center for Computational Ecology at the Yale Institute for Biospheric Studies.

LITERATURE CITED

Akam, M. 1989. Making stripes inelegantly. Nature 341: 282-283.

Amit, D.J. 1989. Modeling brain function. Cambridge University Press, New York.

Binder, K. and A.P. Young. 1986. Spin glasses: experimental facts, theoretical concepts and open questions. Reviews of Modern Physics 58: 801-976.

Bowie, J.U., J.F. Reidhaar-Olson, W.A. Lim and R.T. Sauer. 1990. Deciphering the message in protein sequences: tolerance to amino acid substitution. Nature 247: 1306-1310.

Budd, P.S. and I.J. Jackson. 1991. What do the regulators regulate? First glimpses downstream. Trends in Genetics 7:74-76.

Carey, M., Y.-S. Lin, M.R. Green and M. Ptashne. 1990. A mechanism for synergistic activation of a mammalian gene by GAL4 derivatives. Nature 345: 361-364.

Carroll, S.B. 1990. Zebra pattern in fly embryos: activation of stripes or repression of interstripes? Cell 60: 9-16.

Chisaka, O. and M.R. Capecchi. 1991. Regionally restricted developmental defects result from targeted disruption of the mouse homeobox gene hox-1.5. Nature 350: 473-479.

Dun, R.B. and A.S. Fraser. 1958. Selection for an invariant character - "vibrissa number" - in the house mouse. Nature 181: 1018-1019.

Elredge, N. and J. Cracraft. 1980. Phylogenetic patterns and the evolutionary process. Columbia University Press, New York.

Erwin, D.H., J.W. Valentine and J.J. Sepkoski Jr. 1987. A comparative study of diversification events: the early paleozoic versus the mesozoic. Evolution 41: 1177-1186.

Falconer, D. 1981. Introduction to quantitative genetics. Longman, London.

Futuyma, D.J. 1986. Evolutionary biology. Sinauer Associates, Sunderland, MA.

Gillespie, J.H. 1991. The causes of molecular evolution. Oxford University Press, New York.

Graur, D. and W.-H. Li. 1988. Evolution of protein inhibitors of serine proteinases: positive darwinian selection or compositional effects? Journal of Molecular Evolution 28: 131-135.

Gutman, G.A., R.M. Besta, M.B. Frank and P.R. Baverstock. 1987. Duplication of J-K genes within the genus Rattus. Immunogenetics 26: 14-20.

He, X., M.N. Treacy, D.M. Simmons, H.A. Ingraham, L.W. Swanson and M.G. Rosenfeld. 1989. Expression of a large family of POU-domain regulatory genes in mammalian brain development. Nature 340: 35-42.

Hopfield, J.J. and D.W. Tank. 1986. Computing with neural circuits: a model. Science 233: 625-633.

Ingham, P.W. 1988. The molecular genetics of embryonic pattern formation in *Drosophila*. Nature 335: 25-32.

Jablonski, D. and D.J. Bottjer. 1990. The ecology of evolutionary innovations: the fossil record. Pp. 253-288. *in* M.H. Nitecki, ed. Evolutionary Innovations. University of Chicago Press, Chicago.

Johnson, P.F. and S.L. McKnight. 1989. Eukaryotic transcriptional regulatory proteins. Annual Review of Biochemistry 58: 799-839.

Kappen, C., K. Schughart and F.H. Ruddle. 1989. Two steps in the evolution of Antennapedia-class vertebrate homeobox genes. Proceedings of the National Academy of Sciences, USA 86: 5459-5463.

Kauffman, S.A. 1969. Metabolic stability and epigenesis in randomly connected

nets. Journal of Theoretical Biology 22: 437-467.

Kauffman, S.A. 1993. The origins of order. Oxford University Press, New York.

Kimura, M. 1987. Molecular evolutionary clock and the neutral theory. Journal of Molecular Evolution 26: 24-33.

Krumlauf, R. 1994. Hox genes in vertebrate development. Cell 78: 191-201.

Lawrence, P.A. and G. Morata. 1994. Homeobox genes: their function in *Drosophila* segmentation and pattern formation. Cell 78:181-189.

Leblond, P., P. Demuyter, J.-M. Simonet and B. Decaris. 1991. Genetic instability and associated genome plasticity in *Streptomyces ambofaciens* pulsed field gel electrophoresis. Evidence for large DNA alterations in a limited genomic region. Journal of Bacteriology 173: 4229-4233.

LeBowitz, J.H., R.G. Clerc, M. Brenowitz and P.A. Sharp. 1989. The Oct-2 protein binds cooperatively to adjacent octamer sites. Genes and Development 3: 1625-1638.

Lufkin, T., A. Dierich, M. LeMeur, M. Mark, and P. Chambon. 1991. Disruption of the Hox-1.6 homeobox gene results in defects in a region corresponding to its rostral domain of expression. Cell 66: 1105-1119.

McGinnis, W. and R. Krumlauf. 1992. Homeobox genes and axial patterning. Cell 68: 283-302.

McGinnis, N., M.A. Kuziora and W. McGinnis. 1990. Human hox-4.2 and *Drosophila* deformed encode similar regulatory specificities in *Drosophila* embryos and larvae. Cell 63: 969-976.

Mermelstein, F.H., O. Flores and D. Reinberg. 1989. Initiation of transcription by RNA Polymerase II. Biochimica et Biophysics Acta 1009: 1-10.

Milkman, R.D. 1960. The genetic basis of natural variation. I. Crossveins in

Drosophila melanogaster. Genetics 45: 35-48.

Mjolsness, E., D.H. Sharp and J. Reinitz. 1991. A connectionist model of development. Journal of Theoretical Biology 152: 429-453.

Olson, E. 1990. MyoD family: A paradigm for development. Genes and Development 4: 1454-1461.

Regulski, M., S. Dessain, N. McGinnis and W. McGinnis. 1991. High-affinity binding sites for the Deformed protein are required for the function of an autoregulatory enhancer of the Deformed gene. Genes and Development 5: 278-286.

Reinitz, J. and D. Sharp. 1995. Mechanism of formation of eve stripes. Mechanisms of Development 49:133-158.

Reinitz, J., E. Mjolsness and D.H. Sharp. 1995. Cooperative control of positional information in *Drosophila* by *bicoid* and maternal *hunchback*. Journal of Experimental Zoology 271:47-56.

Rendel, J.M. 1979. Canalization and selection. Pp 139-156 *in* J.N. Thompson Jr. and J.M. Thoday, eds. Quantitative genetic variation. Academic Press, New York.

Rosenfeld, M.G. 1991. POU-domain transcription factors: pou-er-ful developmental regulators. Genes and Development 5: 897-907.

Scheiner, S.M. 1993. Genetics and evolution of phenotypic plasticity. Annual Review of Ecology and Systematics 24: 35-68.

Simpson, G.G. 1944. Tempo and mode in evolution. Columbia University Press, New York.

Sokal, R.R. and F.J. Rohlf. 1981. Biometry. Freeman, New York.

Stearns, S.C. and T.J. Kawecki. 1994. Fitness sensitivity and the canalization of life history traits. Evolution 48: 1438-1450.

Sucov, H.M., K.K. Murakami, and R.M. Evans. 1990. Characterization of an

autoregulated response element in the mouse retinoic receptor type β gene. Proceedings of the National Academy of Science, USA 87: 5392-5396.

Struhl, K. and S. Oliviero. 1991. Synergistic transcriptional enhancement does not depend on the number of acidic activation domains bound to the promoter. Proceedings of the National Academy of Science, USA 88: 224-228.

Turelli, M. 1985. Effects of pleiotropy on predictions concerning mutation-selection balance for polygenic traits. Genetics 111: 165-195.

Valentine, J.W. and D.H. Erwin. 1987. Interpreting great developmental experiments: the fossil record. Pp 71-107 *in* C.A. Raff and E.C. Raff, eds. Development as an evolutionary process. Alan R. Liss, New York.

Waddington, C.H. 1957. The strategy of the genes. Macmillan, New York.

Wagner, A. 1994. Evolution of gene networks by gene duplications: a mathematical model and its implications on genome organization. Proceedings of the National Academy of Science, USA 91: 4387-4391.

Wagner, G.P. 1994.
Homology and the mechanisms of development. Pp 273-299 *in* B.K. Hall, ed. Homology: the hierarchical basis of comparative biology. Academic Press, San Diego.

Weinzierl, R.O.J., B.D. Dynlacht and R. Tjian. 1993. Largest subunit of *Drosophila* transcription factor IID directs assembly of a complex containing TBP and a coactivator. Nature 362: 511-517.

Wright, A.P.H. and J.-A. Gustafsson. 1991. Mechanism of synergistic transcriptional transactivation by the human glucocorticoid receptor. Proceedings of the National Academy of Science, USA 88: 8283-8287.

Wright, S. 1968. Evolution and the genetics of populations. Volume 1. University of Chicago Press, Chicago.

Zeng, Z.-B. and C.C. Cockerham. 1993. Mutation models and quantitative genetic variation. Genetics 133: 729-736.

APPENDIX

Generation of Individual Networks and Ensemble Samples. Connectivity matrices $w = (w_{ij})$ were generated as $N \times N$ pseudorandom matrices of independently and identically distributed entries w_{ij}, where the probability distribution $(1 - c)\delta(w_{ij}) + c\rho(w_{ij})$ was used for each entry.

Here, ρ denotes a continuous probability distribution, $\delta(x)$

denotes the Dirac delta function and

c is a real constant $c \in (0, 1)$ which determines the mean number, cN^2, of

entries of w that are different from zero. Two types of

probability distributions, ρ, were used: Gaussian

($\rho(x) = 1/(\sqrt{2\pi}\sigma) \exp[-x^2/(2\sigma^2)]$,

$x \in \Re$) or "reflected" Gamma ($\rho(x) = [2\Gamma(a)]^{-1} e^{-|x|}|x|^{a-1}$, $x \in \Re \setminus \{0\}$, $a > 0$),

because they cover a wide range of qualitatively

different shapes of symmetric distributions. Note that because the dynamics (2) is invariant to

multiplication

of w with an arbitrary scalar $x \neq 0$, the search for ensemble members can be restricted to a

bounded

area, e.g., a N^2-dimensional ball defined by $\sum_{i,j=1}^{N} w_{ij}^2 < \omega \in \Re^+$, and

still

sample all relevant areas of matrix space in a uniform way. This is what the algorithm of matrix

generation

used here achieves, since the probability distribution used in generating individual connectivities is

symmetric around the origin.

Throughout the search procedure used to find individual ensemble members, the network dynamics (2)

was

carried out for $3N$ time steps or until a fixed point had been reached, whichever came first. This

particular

value was chosen because test simulations had demonstrated that networks with a larger number of

time steps to equilibrium are rare for $N \leq 10$ (see, e.g., Fig. 6). The maximum number of matrix

trials (5×2^{2N}) used for a particular pair of states $\vec{S}(0)$, $\vec{S}(\infty)$ was chosen

because

exploratory simulations (not shown) had demonstrated that, 5×2^{2N} is at least two orders of

magnitude greater than the mean number of trials required to find a matrix for any given pair of states.

Simulated Evolution. To study the influence of recombination on the phenomena under

consideration,

two different kinds of simulated evolution scenarios were explored, a

scenario of sexual reproduction with recombination in which steps one

through four from numerical methods ("simulated evolution") were carried out, and an

asexual scenario in which step one was omitted. Any difference between these scenarios will be most

obvious in a situation where individual genes are unlinked. This is the reason

why free recombination between genes was used in step one.

All simulations were carried out for 400 generations, because test simulations (not shown) had demonstrated that a quasi-equilibrium in mean fitness is attained in less than 400 generations. This quasi-equilibrium is characterized by the disappearance of any trends in the evolution of mean fitness. Less heuristic criteria for the attainment of quasi-equilibrium, such as a decrease of temporal variation in mean fitness below a certain threshold turned out to be impractical, since the value of such a threshold would have depended on the individual simulation. Mean fitness during the simulations was characterized by a slight drop from its initial value of 1 until quasi-equilibrium was attained. Mean fitness values in quasi-equilibrium depended on the individual simulation, but were always greater than 0.93 (s=0.1).

Assessment of Epigenetic Stability. Here, the process of "orthogonal perturbation" is described. While having little biological appeal, it differs from mutation in that it is non-local, causing simultaneous changes in all entries of w. For what follows, it is useful to consider the connectivity matrix $w = (w_{ij})$ as a vector $v = (w_{11}, \ldots, w_{1N}, w_{21}, \ldots, w_{2N}, w_{N1}, \ldots, w_{NN}) := (v_1, \ldots, v_{N^2})$ in the vector space $\Re^{N^2}$. From this vector v, a vector v' is generated by

the perturbation.

Note that the dynamics of (2) is invariant to multiplication

of v with any real number $x \neq 0$. This implies that only a vector

v' with a

component orthogonal to v may cause a change in trajectory of (2), and thus only such a vector has the potential to affect the equilibrium gene expression pattern attained. In other words, one can restrict the analysis of the space of connectivity matrices to a $N^2 - 1$ dimensional manifold, the surface of a ball defined by

$v_1^2 + \ldots + v_{N^2}^2 = r > 0$. A natural measure of the

magnitude of a perturbation in v is therefore given by the angle α

between v and

v', defined as

$$\alpha = \arccos \frac{\langle v, v' \rangle}{\|v\| \|v'\|},$$

where $\langle v, v' \rangle := \sum_{i=1}^{N^2} v_i v_i'$

denotes the Euclidian inner product and $\|v\| := \sqrt{\langle v, v \rangle}$ is

the associated norm. This angle is used as an order

parameter for the magnitude of the perturbation in v. A vector v'

with a predetermined angle α to v was generated in the following

way. First, a vector orthogonal to v, $v^{\perp}$, was generated. Its

entries v_1 through v_{N^2-1} were chosen as identically and independently

distributed Gaussian [N(0,1)] pseudorandom variates. Then, $v_{N^2}^{\perp}$

was evaluated as

$$v_{N^2}^{\perp} = -\frac{\sum_{i=1}^{(N^2-1)} v_i v_i^{\perp}}{v_{N^2}},$$

and from that

$$v' = v + \tan\alpha \|v\| \left(\frac{v^{\perp}}{\|v^{\perp}\|} \right).$$

All directions in $\Re^{N^2}$ will be sampled by v' with equal probability, subject to the constraint that v' has a fixed angle α to v.

The mutation algorithm and this algorithm of orthogonal perturbation are related, since replacement of individual connectivities within networks by independent random variates will result in a mean angular displacement from the original array of connectivities. An analytical estimate of this mean displacement can be obtained in the following way. Consider a member of E, represented by a vector (connectivity matrix) v and assume that the correlations among the entries of v are weak and that the individual entries follow a normal distribution, $N(0, \sigma^2)$, with mean zero and variance σ^2. Then, $(1/\sigma^2)\|v\|^2$ is distributed as $\chi^2(N^2)$. Its mean is N^2. Thus, $(1/\sigma)\|v\|$ will be of order N. Assume now that v' is the vector generated from v by replacing one of its entries with an $N(0, \sigma^2)$ distributed random variable. The quantity $(1/2\sigma^2)\|v - v'\|^2$ is distributed as $\chi^2(1)$. Thus, $(1/\sigma)\|v - v'\|$ is of order $\sqrt{2}$. The maximum angle between v and v' is achieved if $v - v'$ is orthogonal to v, in which

case

$$\tan\alpha = \frac{\|v - v'\|}{\|v\|}$$

If N is sufficiently large, the angular displacement by mutation will be small, such that $\tan\alpha \approx \alpha$. Therefore, the mean angular displacement by mutation should be of order $\sqrt{2}/N$ or smaller. This approximation will become more accurate as N increases, because of the increase in possible orthogonal directions in high dimensional spaces. In fact, numerical estimates obtained by Monte Carlo simulations (Fig. A1) demonstrate that even for $N < 10$ the actual mean angular displacement induced by mutation remains within a factor two of this analytical estimate. Whenever an orthogonal perturbation was carried out in the simulations, its angle was chosen as $\sqrt{2}/N$.

Heritability of Pathlengths.

The pathlength of a network that reaches some fixed point is defined as the number of time steps that (2) takes to reach this fixed point after starting from $\vec{S}(0)$. During the process of simulated evolution including recombination, the heritability of pathlengths in the population was estimated by parent-offspring regression (Falconer, 1981) in the following way. In a population of networks ("individuals") the average pathlength (mid-parent value) of each pair of prospective "parents", i.e., pair of networks that are destined to undergo recombination as described above, was evaluated. The networks were then subjected to recombination and mutation and the mean pathlength of the pair of "offspring" networks produced from every parental pair was evaluated. If at

least one of the networks in a pair of parents and their corresponding pair of offspring had not attained an equilibrium state, the whole "mating" was excluded from the calculation. A linear regression of mean parental pathlength vs. mean
offspring pathlength was carried out for the remaining matings. A
significance test (Sokal and Rohlf, 1981) of the resulting regression
coefficient was carried out as well. Note that the procedure also allows for networks that attain an equilibrium state other than $\vec{S}^{opt}(\infty)$.

FIGURE CAPTIONS

Fig. 1. Numerical generation of network ensemble members. See "numerical methods" for details.

Fig. 2. Evolution of epigenetic stability. $N = 10$, $c = 1$.

a) Epigenetic stability as measured by sensitivity of networks to mutation.

b) Epigenetic stability as measured by sensitivity of networks to orthogonal perturbation. Numerical values represent the fraction of mutation/orthogonal perturbation events that leave the simulated stable gene expression pattern unchanged.

Results are based on an ensemble sample of 300 networks. "Before evolution" indicates the stability of

each member of the sample, whereas "after evolution" corresponds to the

mean stability in the population derived from it after 400 generations of

simulated evolution, as described in "numerical methods". Further parameters: $\rho(w_{ij})$ is N(0,1)

distributed; $p = 0.5$; population size: 500 networks.

Fig. 3. Evolution of stability for networks with varying densities c of regulatory interactions.

a) $c = 0.8$; b) $c = 0.6$; c) $c = 0.4$.

Numerical values represent the fraction of mutation events that leave the simulated stable gene

expression pattern unchanged. "Before evolution" indicates the stability

of each member of the sample, whereas "after evolution" corresponds to the

mean stability in the population derived from it after 400 generations of

simulated evolution, as described in "numerical methods". Results are based on a sample of 200, 200 and

198 networks for a), b) and c), respectively. Further parameters: $N = 10$, $\rho(w_{ij})$ is N(0,1)

distributed; $p = 0.5$; population size: 500 networks.

Fig. 4. Evolution of epigenetic stability for networks with varying densities c of regulatory

interactions.

a) Epigenetic stability as measured by sensitivity of networks to mutation. Values are

calculated from the simulation results shown in Figures 2a, 3a, 3b, and 3c.

b) Epigenetic stability as measured by sensitivity of networks to orthogonal

perturbation. Left and right column in each pair of columns shows mean stability of a network ensemble

sample "before evolution" and mean stability of all populations derived from it "after evolution",

respectively.

The length of error bars is equal to one standard deviation. They indicate differences in variation of stability before and after evolution, and not only

"significance" of differences in mean stability. Results shown are based on numerically generated

ensemble samples of at least 175 networks.

Further parameters: $N = 10$; $\rho(w_{ij})$ is N(0,1); $p = 0.5$; population size: 500 networks.

Fig. 5. Evolution of epigenetic stability for networks with varying numbers of genes N.

a) Epigenetic stability as measured by sensitivity of networks to mutation.

b) Epigenetic stability as measured by sensitivity of networks to orthogonal

perturbation. Left and right column in each pair of columns shows mean stability of a network ensemble

sample "before evolution" and the mean stability of all populations derived from it "after evolution",

respectively.

The length of error bars is equal to one standard deviation. They indicate differences in variation of stability before and after evolution, and not only

"significance" of differences in mean stability.

Results shown are based on numerically generated ensemble samples of at least 175 networks.

Further parameters: $c = 1$; $\rho(w_{ij})$ is N(0,1); $p = 0.5$; population size: 500 networks.

Fig. 6. Pathlengths are unequally distributed. Relative frequencies q of different pathlengths in a numerically generated sample of 1000 networks from an ensemble with

$N = 10$ and $c = 1$. Dots correspond to relative frequencies, length of bars to one standard deviation (calculated as $\sqrt{q(1-q)/1000}$). Pathlength 0 occurs if and only if $\vec{S}(0) = \vec{S}(\infty)$. Note the significant differences in relative frequencies

for different pathlengths. Further parameters: $\rho(w_{ij})$ is N(0,1);

$p = 0.5$.

Fig. 7. Correlation of pathlength to epigenetic stability. Results are based on an ensemble sample of size 300. "Pathlength" indicates the pathlength of each member of the ensemble, "Stability" indicates its epigenetic stability before evolution based on a) mutation (estimated Pearson correlation coefficient $r \approx -0.79$, significantly different from zero at $\mathbf{P} \ll .01$) and b) orthogonal perturbation ($r \approx -0.49$ at $\mathbf{P} \ll .01$).

Network parameters: $N = 10$; $c = 1$; $\rho(w_{ij})$ is N(0,1); $p = 0.5$.

Fig. 8. Evolution of pathlengths. a) Change of mean (dots) and standard deviation (length of bars) of pathlength in a population of 500 networks derived from one ensemble member with pathlength 10. The population was subjected to 400 generations of mutation, recombination and selection. Note the rapid decrease in mean pathlength. b) Pathlength of each member of an ensemble sample (225 networks) before evolution vs. mean pathlength in a population (derived from the respective ensemble member) after 400 generations of evolution as in a). Network parameters: $N = 10$, $c = 1$, $\rho(w_{ij})$ is N(0,1); $p = 0.5$.

Fig. 9. Evolution of pathlengths in an "asexual" population.
a) Change of mean (dots) and
standard deviation (length of bars) of pathlength in a population of 500 networks derived from one ensemble member with pathlength 10. The population was subjected to 400 generations of mutation and selection, but *no* recombination (cf. Figure 8).
Note the rapid decrease in mean pathlength. b) Pathlength of each member of an ensemble sample (113 networks) before evolution vs. mean pathlength in a population (derived from the respective ensemble member) after 400 generations of evolution as in a). Network parameters: $N = 10$, $c = 1$, $\rho(w_{ij})$ is N(0,1); $p = 0.5$.

Fig. 10. Pathlengths are heritable. Results shown were obtained in generation 200 of a simulated evolution process initialized with one member of a network ensemble ($N = 10$, $c = 1$) with pathlength 10. The abscissa shows the mean pathlength of two "parent" networks that were subsequently subjected to recombination, the ordinate axis represents the mean pathlength

of the two "offspring" networks obtained from them. Numbers to the right side of each dot indicate the absolute frequency with which a pair of values occurred. The estimated regression coefficient (heritability) is 0.66 at $P < 0.01$. Further parameters: $\rho(w_{ij})$ is N(0,1); $p = 0.5$.

Fig. 11. Epistatic gene interactions are necessary for the evolution of epigenetic stability. Shaded regions represent sets of optimal genotypes in "genotype space". Clouds of points represent populations of organisms (genotypes). The circle around one individual (genotype) in population A in a) indicates the mean distance by which a genotype is displaced in genotype space after a mutation event. See text for details.

Fig. A1. Mean angular displacement caused by mutation. The solid line represents the analytically obtained upper bound for the mean angular displacement $\sqrt{2}/n$, as a function of dimension n. Note that n in this figure corresponds to the squared number of genes, N^2. For each n, Monte Carlo estimates of the mean angular displacement were obtained as follows. A vector $(x_1, x_2, \ldots, x_n)$ of independently and identically distributed (N(0,1)) pseudorandom variates was generated, and its angle to the vector $(x'_1, x_2, \ldots, x_n)$, where x'_1 is a stochastically independent N(0,1) pseudorandom variate, was evaluated. Dots show the mean angular displacement obtained from 10^5 independently generated pairs of vectors. The lengths of the corresponding bars are equal to one standard deviation of displacement. Note that the mean displacement stays close to the theoretical upper bound and approaches it as n is increased.