



CHICAGO JOURNALS



Ecological Specialization and Adaptive Decay in Digital Organisms.

Author(s): Elizabeth A. Ostrowski, Charles Ofria, Richard E. Lenski

Source: *The American Naturalist*, Vol. 169, No. 1 (January 2007), pp. E1-E20

Published by: [The University of Chicago Press](#) for [The American Society of Naturalists](#)

Stable URL: <http://www.jstor.org/stable/10.1086/510211>

Accessed: 28/02/2011 15:09

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=ucpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and The American Society of Naturalists are collaborating with JSTOR to digitize, preserve and extend access to *The American Naturalist*.

<http://www.jstor.org>

Ecological Specialization and Adaptive Decay in Digital Organisms

Elizabeth A. Ostrowski,^{1,*} Charles Ofria,^{2,†} and Richard E. Lenski^{1,‡}

1. Ecology, Evolutionary Biology, and Behavior Program, Michigan State University, East Lansing, Michigan 48824;

2. Department of Computer Science and Engineering, Michigan State University, East Lansing, Michigan 48824

Submitted August 31, 2005; Accepted June 23, 2006;
Electronically published December 4, 2006

ABSTRACT: The transition from generalist to specialist may entail the loss of unused traits or abilities, resulting in narrow niche breadth. Here we examine the process of specialization in digital organisms—self-replicating computer programs that mutate, adapt, and evolve. Digital organisms obtain energy by performing computations with numbers they input from their environment. We examined the evolutionary trajectory of generalist organisms in an ecologically narrow environment, where only a single computation yielded energy. We determined the extent to which improvements in this one function were associated with losses of other functions, leading to organisms that were highly specialized to perform only one or a few functions. Our results show that as organisms evolved improved performance of the selected function, they often lost the ability to perform other computations, and these losses resulted most often from the accumulation of neutral and deleterious mutations. Beneficial mutations, although relatively rare, were disproportionately likely to cause losses of function, indicating that antagonistic pleiotropy contributed significantly to niche breadth reductions in this system. Occasionally, unused functions were not lost and even increased in performance. Here we find that understanding how the functions were integrated into the genome was crucial to predictions of their maintenance.

Keywords: adaptation, digital evolution, mutation accumulation, pleiotropy, regressive evolution, specialization.

* Corresponding author. Present address: Department of Ecology and Evolutionary Biology, Rice University, Houston, Texas 77005; e-mail: ostrowski@rice.edu.

† E-mail: ofria@msu.edu.

‡ E-mail: lenski@msu.edu.

Many theories about the origins and maintenance of biological diversity involve specialization and adaptive decay. Specialization describes the process by which organisms become highly adapted to a narrow range of environmental conditions and may be associated with adaptive decay—the loss of other traits, functions, or abilities that results in the evolution of narrow niche breadth. A tendency toward increased specialization is a defining feature of adaptive radiations, as it forms the underpinnings for niche partitioning and character displacement, which promote diversification (Simpson 1953; Schluter 2000).

The process of specialization can result in the loss of other functions in environments in which they are no longer useful, termed “adaptive decay.” For example, the transition from a free-living to a parasitic lifestyle is thought to involve not only adaptations that enable host exploitation but also extensive decay of other unused functions presumably necessary for survival outside the host, with parasites showing reduced or streamlined genomes relative to their free-living relatives (Andersson et al. 1998; Shigenobu et al. 2000; Cole et al. 2001; Ochman and Moran 2001).

Both specialization and adaptive decay have been documented in natural populations, but the underlying genetic mechanism remains unclear. Some have hypothesized that there are trade-offs, such that adaptation to one environment inevitably results in loss of adaptation to others (antagonistic pleiotropy hypothesis). Trade-offs may result from an energetic burden associated with maintaining or expressing unused functions or because improvements to a selected trait directly interfere with the functioning of an unselected trait. An alternative hypothesis is that the loss of specialized features results from relaxed selection, enabling mutations to accumulate in the portions of the genome that encode unused functions (mutation accumulation hypothesis). Which of these mechanisms predominates is important, insofar as they lead to different expectations as to the frequency of specialization and the types of circumstances that promote it. For instance, mutation accumulation requires that the genes that contribute to increased adaptation in alternative environ-

ments be distinct and that the environment be heterogeneous (in time or space), so as to give rise to the periods of relaxed selection that enable mutations to accumulate. Alternatively, antagonistic pleiotropy results in constraints that prevent organisms from being simultaneously adapted to many niches and does not require environmental heterogeneity, although it may be aided by it. To see why this is true, imagine a generalist organism that is adapted to use both resource A and resource B. If at some point all further improvements on resource A cause a fitness reduction on resource B and vice versa, then specialization can arise in the absence of any change in the environment. By contrast, mutation accumulation requires that selection be relaxed and thus implies that something previously under selection no longer is; for example, some resource may no longer be available. If fitness levels in two environments are considered to be two different traits (Falconer 1952), then antagonistic pleiotropy occurs when the same mutation improves fitness in one environment and reduces it in another. By contrast, in the case of mutation accumulation, the mutations that increase fitness in the novel environment and reduce it in the ancestral environment are distinct (Cooper et al. 2001). For this reason, if antagonistic pleiotropy is common, then the process of specialization will be closely tied to that of adaptive decay.

Despite long-standing interest in these two hypotheses, it has been difficult to distinguish between them. A large body of work on adaptive decay (also called regressive evolution) has focused on cave organisms (Jones et al. 1992; Jernigan et al. 1994; Jeffrey et al. 2003). These organisms often exhibit highly convergent and distinctive phenotypes, characterized by loss of pigmentation and reduced visual systems but with other sensory structures being highly developed, such as antennae. Mutation accumulation hypotheses suggest that the lack of light in caves resulted in relaxed selection and the accumulation of mutations that eventually led to the losses of pigmentation and visual sensory structures. Alternatively, antagonistic pleiotropy hypotheses posit that adaptation to low light levels led to highly specialized sensory structures and that the losses of other traits were a direct result of this adaptation, possibly because the maintenance of unused functions was costly. For example, Darwin (1859) hypothesized that eyes are costly to burrowing rodents because they are prone to infection and thus that their evolutionary loss may have been aided by natural selection. Although the extent to which regressive phenotypes reflect mutation accumulation or antagonistic pleiotropy has been a subject of great debate, a recent study of cave fish demonstrated linkage for quantitative trait loci (QTL) associated with both a regressive (eye size) and a constructive (body weight) trait, suggesting that either antagonistic pleiotropy or hitchhiking was responsible (Borowsky and

Wilkins 2002). In general, increased knowledge of the genetic basis of traits, as well as their evolutionary dynamics, is expected to shed light on the processes of mutation accumulation and antagonistic pleiotropy.

Experimental approaches provide an alternative to comparative approaches, allowing direct examination of the process of specialization and adaptive decay. A recent study of evolving populations of *Escherichia coli* examined the consequences of long-term adaptation to a simple environment for the evolution of catabolic niche breadth (Cooper and Lenski 2000). Replicate populations of *E. coli* were propagated for 20,000 generations in a medium containing only a single available carbon source, glucose. While the evolved populations were found to have increased ability to compete for and catabolize glucose relative to their ancestor, they also consistently evolved reduced ability to catabolize other resources. Moreover, the identities of these carbon sources were similar across independently evolved populations. This parallelism suggested that the losses resulted from antagonistic pleiotropy, that is, that the reduction in diet breadth had traded off with the improved ability to use glucose. Populations that evolved elevated mutation rates during the course of the experiment also did not show significantly greater losses, contrary to the expectations under mutation accumulation, further indicating that antagonistic pleiotropy was the primary mechanism.

Here we address the process of specialization in a very different medium—an evolving system composed of self-replicating computer programs that mutate, compete, and evolve in a computational environment. We examine the digital equivalent of diet breadth—the ability of these organisms to perform complex computations that enable them to garner energy from their environment. In this system, we can not only observe the pattern of evolution associated with specialization and adaptive decay, but we can also examine in detail the underlying genetic processes; that is, we can identify the specific mutations that result in losses of function and determine their fitness effects. We use this knowledge to distinguish between antagonistic pleiotropy and mutation accumulation by asking whether losses of function were the result of neutral or beneficial mutations. Whereas losses that result from mutations that are neutral in the selective environment constitute examples of mutation accumulation, those that result from mutations that are beneficial in the selective environment constitute examples of antagonistic pleiotropy.

Below, we give a brief introduction to the digital life system, Avida. Additional information is provided in appendixes A and B, including a schematic of a digital organism and a glossary of terms. A more detailed description of the system is available elsewhere (Wilke and Adami 2002;

Lenski et al. 2003; Ofria and Wilke 2004), and documentation is available online (<http://devolab.cse.msu.edu>).

The Avida System

The Avida system is a software platform wherein self-replicating computer programs (digital organisms) adapt and evolve in a computational environment. Each digital organism has a genome composed of a series of computer instructions, which, by default, are executed sequentially by a virtual central processing unit (CPU). However, some instructions permit jumps or loops; for example, replication generally involves the execution of a copy loop. Execution of a viable genome results in an organism copying itself, instruction by instruction, and on completion, dividing by binary fission to produce two organisms. If no empty space is available in the population, replication results in the replacement of another organism in the population. Thus, the faster a given organism produces offspring, the more likely its genotype is to persist and spread in the population over time.

Evolution occurs because the copy process is subject to random mutations, at a rate specified by the experimenter. Mutations can be point mutations, whereby one instruction in the genome is randomly replaced with another, or they can be insertions or deletions, which enable the genome to grow or shrink in length. Mutations are normally deleterious because they reduce the speed or efficiency with which an organism replicates; in the extreme, they are effectively lethal if they prevent an organism from being able to replicate altogether. Mutations that are beneficial increase the replication rate of the organism, either by improving the efficiency with which it produces copies of itself or else by enabling it to receive additional CPU cycles, which allow it to execute more instructions. CPU cycles can thus be thought of as energy in Avida: every instruction executed burns a CPU cycle, but organisms must execute instructions in order to replicate or perform other functions. Digital organisms will thus generally adapt in one of two ways. First, they may evolve to reduce the number of CPU cycles they require to produce an offspring, that is, to reduce their generation time. Alternatively, they may evolve to obtain more CPU cycles (more energy), which may allow them to produce more offspring.

Digital organisms can obtain additional CPU cycles by performing bitwise logic functions on numbers they input from their environment. A correct computation provides an organism with CPU cycles above its initial allotment, which can then be put toward further execution of the genomic instructions, potentially resulting in an increased rate of replication. In our experiments, organisms receive an initial allotment of CPU cycles equal to their genome length. This makes evolution neutral with respect to ge-

nome size and prevents shorter organisms from having an inherent fitness advantage (for more details, see Lenski et al. 2003; Ofria et al. 2003). Thus, whether a given organism replicates faster depends only on whether the CPU cycles obtained offset the additional CPU cycles required to perform the computation. Organisms therefore evolve not only to perform computations but also to perform them as efficiently as possible.

The performance of computations represents a kind of a metabolism, in that the conversion of one or two numbers into another number provides the organism with energy. In Avida, the genomic instruction set includes one that is called *nand*; this instruction can allow digital organisms to perform the NAND (not and) logic function, provided that the instruction is properly coupled to input-output instructions to obtain the numbers and output the result. All other computations can be constructed by combining multiple *nand* instructions with various other instructions. In this way, digital organisms also resemble real computers, in that all computations performed by computers can be built out of combinations of NAND functions (sometimes referred to as *nand gates*).

The Equals Function

For the current study, we employed nine possible logic functions, eight of which require two inputs, that is, two binary numbers input from the environment. These nine bitwise logic functions are as follows: NOT, NAND, AND, OR-NOT, OR, AND-NOT, NOR, XOR, and EQU. Computation of these functions has been described elsewhere (Lenski et al. 2003), but for purposes of illustration, we describe in greater detail one of these functions, *equals* (EQU), which is the focus of the current study. EQU is a computation where, for any two inputs, the correct output contains a 1 (true) at every site where the bits are identical and a 0 (false) at every site where the bits are not identical. For example,

Input A:	0	1	0	1	1	0	1	1	1	0	0	1	...
Input B:	0	0	1	1	0	1	0	1	1	1	0	1	...
A EQU B:	1	0	0	1	0	0	0	1	1	0	1	1	...

Thus, in an environment where EQU is rewarded, an organism that inputs A and B and outputs the above number would receive additional CPU cycles. Because most computations require the coordination of multiple steps, digital organisms must store and manipulate intermediates or partial results. For example, the performance of EQU requires combining a minimum of five NAND functions and at least 19 instructions in total (Lenski et al. 2003). Finally, CPU cycle rewards are determined simply by comparing an organism's inputs with its output, such that selection is based on the phenotype, not the genotype.

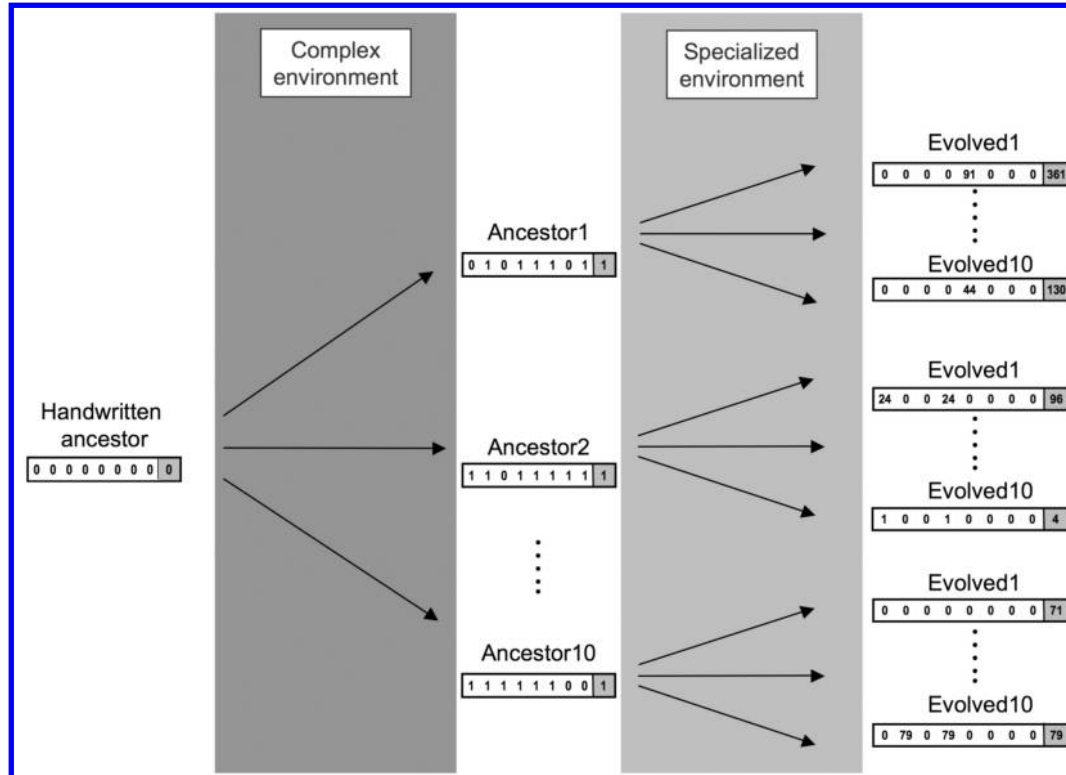


Figure 1: The evolution experiment had two phases: an initial period when replicate populations evolved in a complex environment with all functions rewarded, followed by a period of evolution in a specialized environment with only EQU rewarded. Evolved organisms represent the final dominant organism in each evolved population. Numbers shown beneath each organism specify its phenotype in terms of the number of times it performs each function during its life cycle, in the following order: NOT, NAND, AND, OR-NOT, OR, AND-NOT, NOR, XOR, and EQU. For clarity, the number that corresponds to an organism's performance of EQU has been shaded.

In a recent article, Lenski et al. (2003) examined the evolutionary origins of the EQU computation from an ancestral organism that could perform no functions. They found that the EQU function has the properties of a complex feature: its performance required the coordinated execution of numerous interacting parts. Moreover, its evolutionary emergence required that other, simpler functions also be rewarded; these simpler functions could then serve as building blocks for such a complex function. Here we expand on this work by examining specialization of the EQU function. Starting from generalist organisms (those that could perform a variety of computations, in addition to EQU), we examine their evolution in a narrow environment, where only EQU generates extra CPU cycles. We examine how these digital organisms adapt to their novel environment, the extent to which they evolve to be highly specialized, and the evolutionary processes that govern their transition from generalists to specialists. Finally, we describe an unexpected result: that some unrewarded functions were not lost and even increased in performance, despite evolution in an ecologically narrow environment.

Here, we take advantage of the remarkable transparency of digital systems, as well as the ability to probe the underlying genetic architecture of the traits in question, to develop a general understanding of the factors that drive the balance between niche breadth reduction and conservatism in this system.

Methods

Experimental Design

Two-Stage Evolution Experiment. In the first stage of the experiment, replicate populations evolved from a single handwritten ancestor that could self-replicate but that could not perform any logic functions (fig. 1). These populations evolved in an environment where the performance of any of the nine computations provided CPU cycles as rewards. These rewards were limited, however, to once per gestation cycle, such that organisms generally evolved to perform each function only once. (The gestation cycle is defined as the time from when an organism

executes the first instruction in its genome to when it produces an offspring.) Following 100,000 updates of evolution, an arbitrary unit of time in Avida corresponding to an average of 30 instructions executed per organism (see glossary, table B1), the dominant genotype was isolated from each population. These genotypes served as the generalist ancestors (subsequently denoted Ancestor1–Ancestor10) in the main experiment.

In the second stage of the experiment, replicate populations were founded from each generalist ancestor and evolved in an environment where only EQU yielded extra CPU cycles. The experiment consisted of 10 replicate populations for each of the 10 ancestors, for a total of 100 populations. The ancestors were generalists in that they could perform a wide variety of different logic functions, though they differed in the number and identity of the exact functions they performed (average = 7.3, range = 6–9 of nine possible logic functions). All ancestors, however, performed EQU exactly once per gestation cycle. The ancestors also varied in the number of instructions comprising their genomes, with the shortest having 59 instructions and the longest having 124 instructions (average = 99.7). All populations evolved for 100,000 updates, during which time they received a reward only for the EQU computation. In this new EQU-only environment, however, organisms received rewards every time they performed the EQU computation and output the appropriate result. Insertion, deletion, and point mutations occurred at rates of 0.01, 0.01, and 0.08 mutations per genome per replication, respectively. Population size was limited to a maximum of 3,600 organisms, and offspring were placed randomly in the population, such that the population was well mixed.

Examining the Line of Descent. To assess specialization and adaptive decay following evolution in the EQU-only environment, the most abundant genotype from each population was saved and assayed at the end of each experiment for its ability to perform each of the nine computations, including EQU. For each of these genotypes, we also determined its line of descent, which is the sequence of all genotypes leading back to the original ancestor. By looking along the line of descent, we identified pivotal genotypes where mutations arose that produced a loss of function. We then classified each mutation according to its fitness effect relative to the immediately preceding parent genotype: >1 was beneficial, 1 was neutral, and <1 was deleterious. For the purposes of distinguishing between antagonistic pleiotropy and mutation accumulation, we henceforth lump deleterious mutations with neutral ones and refer to them collectively as non-beneficial. The reason for doing so is that the antagonistic pleiotropy hypothesis specifically concerns beneficial mu-

tations, whereas mutation accumulation could encompass not only neutral mutations but also deleterious mutations that drift or hitchhike to fixation.

Generally, organisms along the line of descent differed from their immediate predecessors by a single mutation. Occasionally, however, they differed by two or more mutations. By default, we classified multiple mutations according to their fitness effect in combination. However, to ensure that these multiple mutational steps did not influence our results, we also analyzed our data without these multiple mutations. We also performed several additional sets of experiments to determine how our results were affected by our choice of particular parameter values. First, we repeated our experiments at higher and lower genomic mutation rates of 0.3 and 0.01, respectively, equal to threefold higher and 10-fold lower than those in our original experiments. To control for the effects of differential mutation supply, we performed additional experiments where we scaled the length of the experiments inversely to the mutation rate. Thus, experiments at a genomic mutation rate of 0.3 were run for 33,000 updates, and those at the 0.01 genomic mutation rate were run for 1,000,000 updates. Second, to determine whether the presence of energetic constraints on genome size could alter the balance between antagonistic pleiotropy and mutation accumulation, we repeated our experiments in an environment where there was strong selection to reduce genome length. To do so, we changed the baseline allotment of an organism's CPU cycles to be a constant (in this case, equal to 100) rather than proportional to its genome length. Under these conditions, organisms face strong selection to reduce their genome size to a minimum. Although there is little direct evidence that energetic constraints on genome size are common in biological organisms (Petrov and Hartl 2000; Sliwa and Korona 2005), the reduced genomes of some intracellular bacteria (Cole et al. 2001; Ochman and Moran 2001; Gómez-Valero et al. 2004) nevertheless suggest the potential importance of such constraints in the evolution of ecological specialization.

Statistical Analyses

Performance of the EQU Function. To determine how fitness and the performance of EQU varied depending on the ancestor, we performed two one-way ANOVAs. These analyses were performed using PROC GLM in SAS, with ancestor designated as a random effect. In the first ANOVA, we used the log relative fitness of evolved populations as the response variable, where each evolved population's fitness is relative to that of its own ancestor. In the second ANOVA, the response variable was the number of times EQU was performed in the numerically dominant

genotype isolated from each evolved population. Because variances were heterogeneous across ancestors, we performed the ANOVAs as nonparametric Kruskal-Wallis tests. These tests were performed in SAS using PROC NPAR1WAY.

Antagonistic Pleiotropy versus Mutation Accumulation. To examine the relative contributions of antagonistic pleiotropy and mutation accumulation, we totaled the number of beneficial and nonbeneficial mutations per ancestor across the 10 replicate experiments at those steps where some unused function was lost. Because nonbeneficial mutations are typically more common than beneficial mutations, even in the line of descent (Lenski et al. 2003), we also assembled a baseline calculation of the relative proportion of these two mutation types over the course of evolution by totaling their numbers over the line of descent as a whole, irrespective of whether they were associated with a loss of function. To determine whether the ratio of beneficial to nonbeneficial mutations was significantly higher among those mutations that caused a loss of function (which would provide support for the antagonistic pleiotropy hypothesis), we performed a Fisher's exact test that compared the number of beneficial versus nonbeneficial mutations causing a loss of function with the number that did not. To assess the statistical significance of the contingency tables, we used the right-tailed P value of a Fisher's exact test, where a low P value would indicate that beneficial mutations were significantly over-represented among mutations causing losses of function. The analyses were performed in SAS, using PROC FREQ and the Fisher Exact option.

Results

Specialization and Adaptive Decay in the EQU-Only Environment

We consider three components of specialization. First, we examine the extent to which populations evolve increased performance of EQU, where the performance is determined as the total number of times an organism outputs the result of the EQU function per reproductive cycle. Second, because organisms can make improvements in the efficiency of their EQU performance without increasing the number of times it is performed, we also consider the degree to which fitness increased in the EQU-only environment. Third, we examine the extent of adaptive decay, that is, the extent to which unrewarded functions were lost during evolution in the EQU-only environment, leading to the evolution of narrow niche breadth.

With regard to the first of these criteria, we find that evolved populations had greatly improved performances

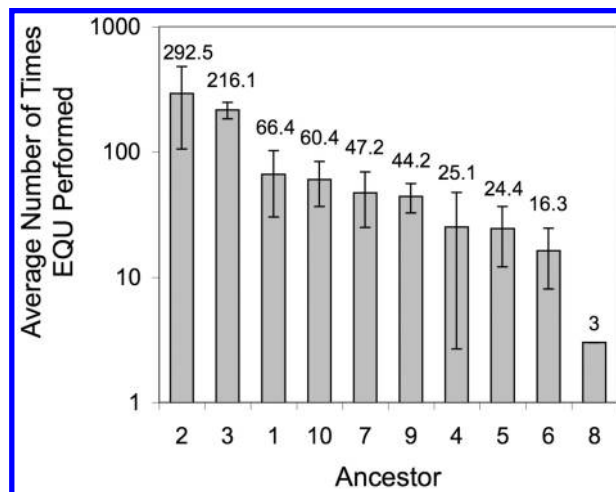


Figure 2: Average number of times EQU is performed per life cycle by evolved organisms, arranged from highest to lowest for each of the 10 ancestors. The ancestors all performed EQU only once, and each bar represents the mean across 10 replicate populations evolved from that ancestor. Values are plotted on a logarithmic scale, and error bars represent 1 SE. For clarity, the average value is also written above each bar.

of EQU. Whereas all ancestors performed EQU only once per reproductive cycle, most evolved organisms performed it tens or even hundreds of times (fig. 2). Interestingly, the magnitude of this improvement depended strongly on the ancestor (Kruskal-Wallis $\chi^2 = 42.41$, $df = 9$, $P < .0001$). There was also variation in the performance of EQU among replicate populations evolved from the same ancestor. For example, in five of 100 populations (three derived from Ancestor1 and one each from Ancestor9 and Ancestor10), the performance of EQU did not increase above the ancestral level. However, when averaged over the 10 replicate populations, organisms evolved from Ancestor1 had the third highest performance of EQU overall (fig. 2). Although EQU performance did not increase in these five populations, at least one beneficial mutation fixed in every population, indicating that some adaptation to the EQU-only environment occurred in all populations. In these five populations, fitness improved by reductions to generation time, although the fitness increases in these populations were small compared with populations that evolved increased performance of the EQU function. Such variation in the extent of fitness improvement among replicate populations evolved from the same ancestor indicates that the chance occurrence of different mutations in replicate populations was an important component of specialization in this system. Similarly, because the generalist ancestors themselves evolved from the same handwritten ancestor (fig. 1), differences in outcome that were contin-

gent on the generalist ancestor also demonstrate the importance of chance events at an earlier stage of evolution.

While EQU performance did not increase in every population, fitness universally improved in the EQU-only environment (fig. 3A). Once again, the magnitude of the increase depended greatly on the ancestor. Figure 3A shows the fitness trajectory of the populations over time, averaged over the 10 replicates for each ancestor. While all populations increased in fitness, there was substantial heterogeneity in the magnitude of this improvement, and again the ancestor had a highly significant effect (Kruskal-Wallis $\chi^2 = 32.45$, $df = 9$, $P = .0002$).

Evolution of Niche Breadth Reductions

The loss of unrewarded functions was not universal and also varied across ancestors. Retention of a particular function is indicated by a black cell in figure 4. Only seven of 100 populations retained only EQU and lost all unused functions; these seven populations were distributed across six different ancestors (fig. 4). As expected, there was a significant correlation between EQU performance and fitness in the EQU-only environment (Spearman's $\rho = 0.81$, $P = .005$). However, there was no overall relation-

ship between increased EQU performance and the magnitude of niche breadth reductions (Spearman's $\rho = 0.33$, $P = .35$). For example, populations evolved from Ancestor3 tended to maintain a relatively broad niche (fig. 4) and yet were the second-highest performers of EQU (fig. 2). The decline in niche breadth over time is plotted in figure 3B. Each line represents the average for a different ancestor, and the colors for each population are the same as those used for the fitness trajectories in figure 3A.

Population Genetic Processes Underlying the Evolution of Reduced Niche Breadth

Where functions were lost, we were interested in determining whether losses were caused by mutation accumulation or antagonistic pleiotropy. To address this question, we determined the fitness effect of every mutation that resulted in a loss of function along the line of descent. If mutations causing functions to be lost were neutral or deleterious in the EQU-only environment, it would indicate that mutation accumulation was responsible for losses of function. Similarly, if the mutations leading to losses of function were beneficial, it would indicate that antagonistic pleiotropy was responsible. Note that there




Figure 3: A, Fitness trajectory of populations in the EQU-only environment. Each line represents the average of 10 replicate evolution experiments for each of 10 different generalist ancestors. Fitness is expressed as the log ratio of values for evolved organisms relative to their own ancestor, such that all populations start at 0. B, Reduction in niche breadth during evolution in the EQU-only environment. Niche breadth was calculated as the proportion of organisms performing each function at a given time and summed over all functions. Each line represents the average of 10 replicates for each of the 10 ancestors.

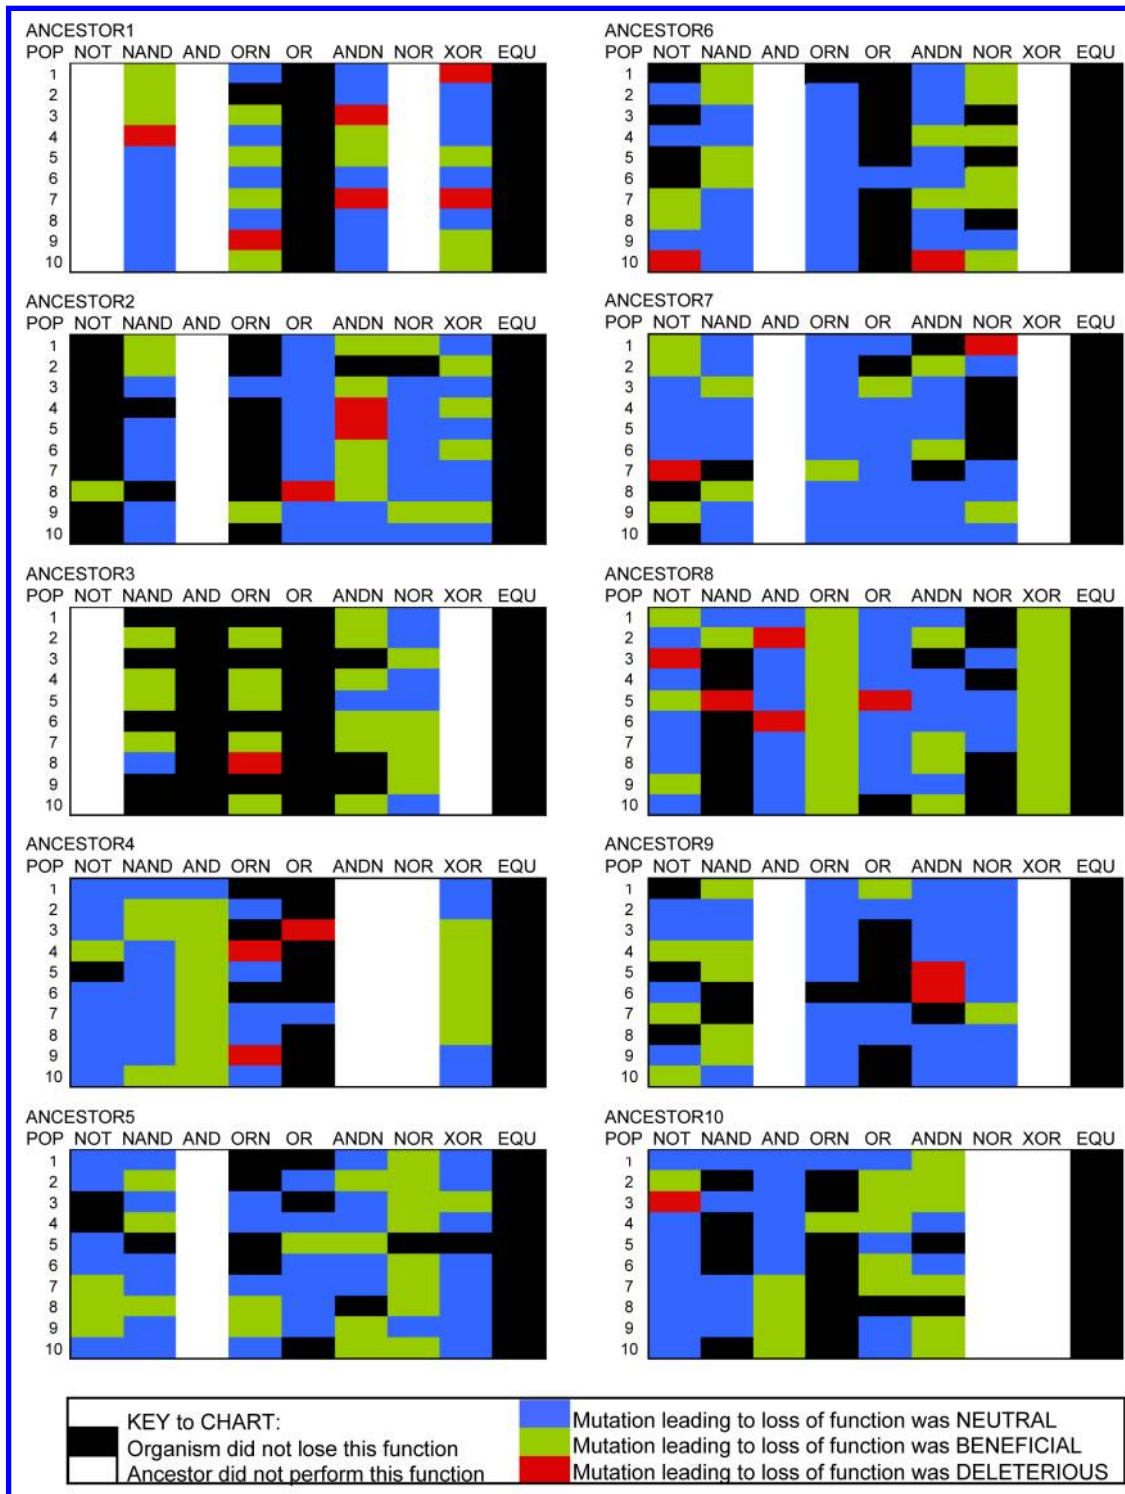


Figure 4: Outcome of evolution in the specialized EQU-only environment for all 100 populations, arranged by ancestor. Each row specifies the final dominant genotype from one evolved population, and the colors indicate which functions were lost and the type of mutation (beneficial, neutral, or deleterious) that caused the loss of function. In some cases, a single mutation led to the loss of multiple functions at once, such that the colors of the blocks are not necessarily independent in every row.

were at least two ways in which a mutation causing a loss of function could be beneficial. First, it may be beneficial because the instructions that encode EQU also encode other functions, such that mutations that enhance EQU performance interfere with these other functions; this would constitute a classic example of pleiotropy. Second, mutations causing losses of function could also be beneficial because they reduce the energy spent performing useless tasks, thereby increasing fitness in the EQU-only environment. This also constitutes antagonistic pleiotropy, in the sense that a single mutation improves fitness in one environment but reduces it in another. For our purposes, we did not distinguish between these two explanations: all mutations that were simultaneously beneficial in the EQU-only environment and resulted in a loss of function were interpreted as support for the pleiotropy hypothesis. Our classification scheme thus captures two categories of explanation: those mutations that fix via selection for improved performance in an EQU-only environment and those mutations that fix by genetic drift or by hitchhiking alongside beneficial mutations.

The mutations leading to losses of function are shown for all populations in figure 4, arranged by ancestor. Because a single mutation occasionally led to the simultaneous loss of multiple functions, cells are not necessarily independent of one another. In addition, because we are interested in understanding the niche breadth of the final derived organisms and the mutations that led to that state, we do not consider cases where a function was lost but subsequently regained. Thus, we examine only the mutations causing losses of function if the function was absent at the end of the experiment. In cases where a function was lost, regained, and subsequently lost again, we consider only the final loss of function. This methodology is most likely conservative with respect to detecting antagonistic pleiotropy, as earlier mutations (when adaptation is most rapid) are more likely to be beneficial than later mutations.

For eight of 10 ancestors, the beneficial to nonbeneficial ratio was higher among mutations causing losses of function than among those that did not. In seven of these eight cases, the Fisher's exact test was highly significant (table 1). In the two cases where the mutations causing losses of function were predominantly neutral or deleterious (Ancestor2 and Ancestor7), the differences were quite small. In these cases, the left-hand P values of the Fisher's exact tests, which would test for overrepresentation of neutral or deleterious mutations among mutations causing losses of function, were nonsignificant ($P = .07$ and $.53$ for Ancestor2 and Ancestor7, respectively).

To determine overall support for the antagonistic pleiotropy hypothesis, we combined the P values from the 10 separate tests into a single P value, using a method developed by Fisher (Sokal and Rohlf 1995, p. 194). This

Table 1: Comparison of mutations associated with losses of function relative to their overall proportion along the line of descent

Ancestor	Beneficial or not	Number of mutations		P_B		P
		Loss	Not	Loss	Not	
Anc1	Beneficial	12	301	.308	.176	.034
	Not	27	1,407			
Anc2	Beneficial	15	592	.306	.424	.966
	Not	34	803			
Anc3	Beneficial	20	496	.714	.370	<.001
	Not	8	844			
Anc4	Beneficial	19	233	.396	.261	.032
	Not	29	661			
Anc5	Beneficial	22	303	.386	.205	.002
	Not	35	1,178			
Anc6	Beneficial	14	268	.326	.183	.020
	Not	29	1,195			
Anc7	Beneficial	10	282	.200	.208	.610
	Not	40	1,076			
Anc8	Beneficial	28	201	.431	.172	<.001
	Not	37	968			
Anc9	Beneficial	10	204	.213	.168	.265
	Not	37	1,010			
Anc10	Beneficial	17	394	.386	.232	.017
	Not	27	1,302			

Note: Results of Fisher's exact tests comparing the loss of functions due to beneficial versus nonbeneficial (not; neutral or deleterious) mutations. The left-hand side shows the contingency table for each of the 10 ancestors. In each case, the number of mutations was summed over 10 replicate populations. Loss and not categories refer to the number of mutations that were associated with a loss of function or not, respectively. P_B indicates the proportion of mutations that were beneficial. P is the probability associated with the right tail of a Fisher's exact test—in other words, the probability of seeing, by chance alone, as much or more overrepresentation of beneficial mutations among loss-of-function mutations.

overall test is also highly significant ($P < .001$). For all of our tests, we used a right-tailed P value, where a significant result indicates that beneficial mutations were disproportionately represented among mutations causing losses of function. Significant overrepresentation of neutral or deleterious mutations is not expected because, unlike beneficial mutations, which can become overrepresented as a result of selection, there is no analogous mechanism to cause the overrepresentation of neutral or deleterious mutations. Consistent with this expectation, none of the P values based on the left tail was significant (data not shown). Also, using a two-tailed P value (and noting that the two tails are not symmetric in a Fisher's exact test) produced the same statistically significant results. The only exception was the test for Ancestor1; in this case, the P value increased from .034 (one tailed) to .053 (two tailed).

To determine how energetic constraints on genome size might affect the balance between antagonistic pleiotropy and mutation accumulation, we repeated our experiments

in an environment where organisms received a baseline CPU cycle allotment that was constant rather than proportional to genome length. In such an environment, all else being equal, decreases in genome size will confer a fitness benefit; however, organisms also receive additional CPU cycles for performing rewarded computational functions. These experiments show that, as expected, genomes are significantly shorter in organisms evolved in this environment (Kruskal-Wallis $\chi^2 = 46.6$, $n = 200$, $P < .0001$). In addition, the presence of energetic constraints on genome size shifted the balance further in favor of antagonistic pleiotropy; the right tail of the Fisher's exact test was significant for nine out of 10 ancestors, indicating that beneficial mutations were disproportionately represented among mutations causing losses of function (data not shown). Although selection to reduce genome size in this system is probably stronger than that operating in natural systems (since a single instruction deletion constitutes a proportionately larger decrease in genome size than does a single base pair deletion), these findings nevertheless confirm that energetic constraints on genome size can increase the importance of ecological specialization arising from antagonistic pleiotropy.

Finally, although we expect that most mutations that cause losses of function will be deleterious in the ancestral environment, we verified this expectation by measuring the fitness effect of each loss-of-function mutation in the ancestral environment. Of the 470 mutations that caused a loss of function in our experiments (i.e., fig. 4, all blue, red, or green cells), 448 of them were indeed deleterious in the ancestral environment. Moreover, excluding those few mutations that did not reduce fitness in the ancestral environment had no effect on the statistical significance of our results. Overall, our findings indicate that, where losses of function occurred, they were disproportionately likely to be caused by a beneficial mutation, and therefore, antagonistic pleiotropy was an important factor in driving the decay of unrewarded functions.

Steps with Multiple Mutations

Because some steps along the line of descent occasionally included multiple mutations (i.e., a derived genotype differed from its immediate parent by more than a single mutation), we sought to determine whether these multiple mutations made a significant contribution to losses of function. We found that multiple mutations accounted for approximately 6.2% of all genotypic steps along the line of descent and for approximately 4.1% of mutations causing losses of function. Thus, multiple mutations were, if anything, underrepresented among the mutations causing losses of function. The Fisher's exact tests comparing losses of function due to beneficial versus nonbeneficial muta-

tions (table 1) were largely unaffected by the exclusion of multiple mutations (data not shown). The statistical significance of the results differed only for Ancestor1, which became nonsignificant once these mutational steps were excluded.

Niche Breadth Reductions at Higher and Lower Mutation Rates

Our initial experiments were performed at a genomic mutation rate of 0.1 for 100,000 updates. To assess the generality of these results, we repeated our experiments at significantly higher and lower mutation rates of 0.3 and 0.01, respectively. As expected, niche breadth usually declined more rapidly with increasing mutation rate (fig. 5). However, it was not obvious whether the faster decay of niche breadth was a result of the greater overall mutation supply or whether mutation rate disproportionately affected losses of function by altering the relative importance of beneficial and nonbeneficial mutations. For example, in asexual organisms, increasing the mutation rate is expected to increase the fixation of nonbeneficial mutations to a greater extent than beneficial mutations because, at high mutation rates, beneficial mutations will more often arise in different lineages that interfere with each other's fixation, a phenomenon termed "clonal interference" (Muller 1932; Gerrish and Lenski 1998; de Visser et al. 1999; Orr 2000). The relative roles of antagonistic pleiotropy and mutation accumulation might thus be altered by changes to the mutation rate.

To address this issue, we repeated our experiments but this time scaled their duration inversely to the mutation rate. Because our initial experiments were run at a 0.1 genomic mutation rate for 100,000 updates, we reran the high mutation rate (0.3) experiments for 33,000 updates and the low mutation rate (0.01) experiments for 1,000,000 updates. Scaling the runs in this way should control for differential mutation supply; this prediction was verified by examining the number of genotypes along the line of the descent, which was found to be similar across treatments (mean: low = 147.1, medium = 141.9, high = 139.4; all pairwise comparisons not statistically significant).

Our results show that increasing the mutation rate while holding mutation supply constant tends to decrease the number of beneficial mutations that cause losses of function per experiment (least squares means: low = 1.92, medium = 1.67, high = 1.38). A two-way parametric ANOVA based on the number of beneficial mutations causing losses of function found a significant effect of ancestor ($F = 12.41$, $df = 9, 270$, $P < .0001$), mutation rate ($F = 5.60$, $df = 2, 18$, $P = .013$), and their interaction ($F = 1.66$, $df = 18, 270$, $P = .046$). As expected, losses of function resulting from nonbeneficial mutations

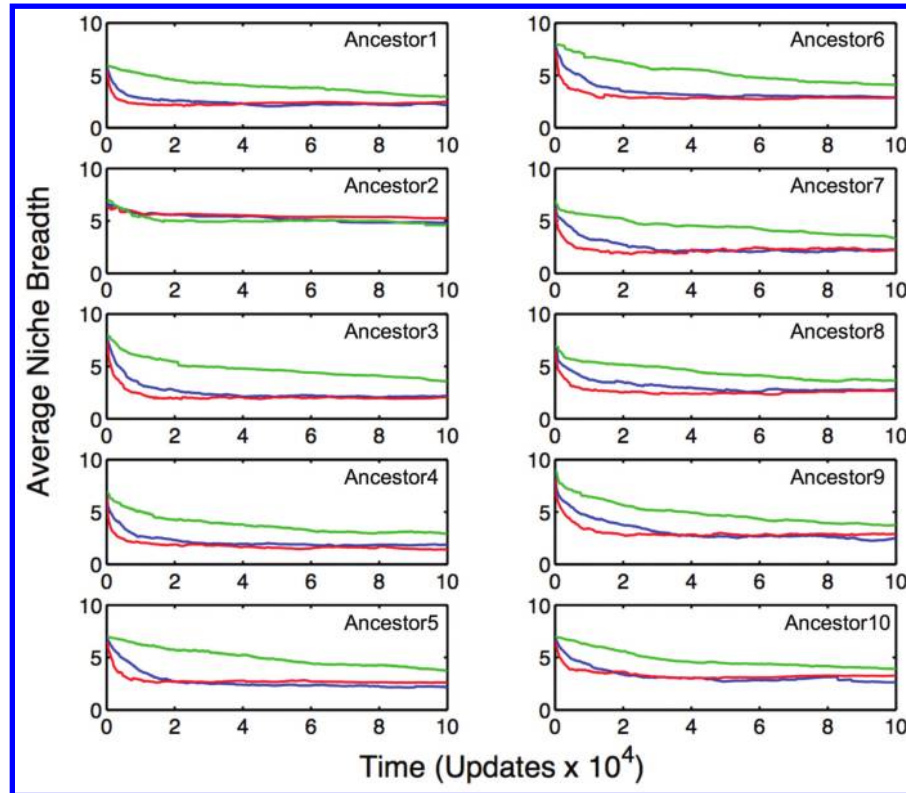


Figure 5: Decline in average niche breadth over time as a function of mutation rate. Average niche breadth was calculated as the mean niche breadth of the 10 replicate populations derived from each ancestor. Niche breadth was calculated as the proportion of organisms performing each function at a given time point and summed over all functions. *Green* = low, 0.01 genomic mutation rate; *blue* = medium, 0.1 genomic mutation rate; *red* = high, 0.3 genomic mutation rate.

showed the opposite pattern (least squares means: low = 2.76, medium = 3.03, high = 3.20). Ancestor and mutation rate were again both statistically significant ($F = 20.09$, $df = 9, 18$, $P < .0001$ and $F = 3.70$, $df = 2, 18$, $P = .045$, respectively). The interaction between mutation rate and ancestor, however, was nonsignificant ($F = 1.08$, $df = 18, 270$, $P = .376$) for these mutations. Finally, a paired t -test comparing the proportion of loss-of-function mutations that were beneficial at low and high mutation rates showed the proportion to be significantly higher at lower mutation rates, as hypothesized (two-tailed $P = .002$, based on table 2).

We can also ask whether beneficial mutations remain overrepresented among mutations causing losses of function at higher and lower mutation rates. These data are presented in table 2, which shows the proportion of beneficial mutations, relative to the total, that were associated or not associated with a loss of function. At all three mutation rates, beneficial mutations were usually present in greater proportions among mutations causing losses of function than among those that did not. For all mutation

rates, the proportion of beneficial mutations causing losses of function was higher in descendants of eight out of 10 ancestors, although the identities of these eight ancestors varied across the treatments. Many of these differences were significant when we performed the Fisher's exact tests to examine the number of beneficial versus nonbeneficial mutations causing losses of function or not, as we also saw for our earlier analysis at the genomic mutation rate of 0.1 (table 1). While the pattern at higher and lower mutation rates is qualitatively similar to that for the medium mutation rate, somewhat fewer tests were significant at both extremes. In general, however, the pattern was similar, despite large changes to the mutation rate, indicating that antagonistic pleiotropy was an important contributor to niche specialization at all three mutation rates.

As before, we also analyzed our results to determine whether they were affected by the presence of multiple mutations in some steps. At the low mutation rate, multiple mutations comprised 2.3% and 0% of steps along the line of descent and losses of function, respectively. The conclusions of the Fisher's exact tests that compared the

Table 2: Proportion of mutations along the line of descent that were beneficial as a function of mutation rate

Ancestor	Genomic mutation rate								
	Low (.01)			Medium (.1)			High (.3)		
	Loss	Not	<i>P</i>	Loss	Not	<i>P</i>	Loss	Not	<i>P</i>
1	.162	.195	NS	.308	.176	*	.175	.190	NS
2	.321	.333	NS	.306	.424	NS	.288	.252	NS
3	.862	.321	***	.714	.370	***	.500	.256	*
4	.596	.484	NS	.396	.261	*	.426	.199	***
5	.545	.267	***	.386	.205	**	.354	.363	NS
6	.362	.251	NS	.326	.183	*	.341	.200	*
7	.192	.186	NS	.200	.208	NS	.192	.175	NS
8	.431	.259	**	.431	.172	***	.441	.148	***
9	.396	.210	**	.213	.168	NS	.227	.226	NS
10	.381	.249	*	.386	.232	*	.250	.242	NS

Note: Comparison of the proportion of mutations substituted on the line of descent that were beneficial among those causing losses of function versus those that did not, at three different mutation rates. Notice that beneficial mutations are generally present in higher proportions among mutations causing losses of function. Asterisks indicate the significance of the associated Fisher's exact test, which compared the number of beneficial versus nonbeneficial (neutral or deleterious) mutations that caused losses of function (loss) versus those that did not (not). NS = not significant.

* $P < .05$.

** $P < .01$.

*** $P < .001$.

fitness effects of mutations causing losses of function to those that did not were generally unaffected by these mutations, with the exception of Ancestor4, for which the test became significant once steps with multiple mutations were excluded. At the high mutation rate, multiple mutations comprised 7.1% of all mutations along the line of descent and 10.9% of mutations causing losses of function. The statistical significance of all of the Fisher's exact tests at the high mutation rate were unchanged by the exclusion of these mutations.

Functional Genetic Explanations for Niche Conservatism

A striking feature of these experiments is the extent to which some functions were repeatedly retained across replicate populations started from the same ancestor (fig. 4, columns of black cells). For example, all 10 populations evolved from Ancestor1 invariably retained OR, while those evolved from Ancestor3 always kept both AND and OR. There are at least two explanations for the maintenance of unrewarded functions. One possibility is that there may have been insufficient mutational pressure to cause losses of function. While this effect may be expected to be random with respect to functions, some functions present larger targets for mutations because they require more instructions to encode, and thus they may be lost more consistently. To test whether mutational pressure was

strong enough to lead to decay of functions, we ran additional experiments with one ancestor, Ancestor3, for which derived populations had decayed the least on average over the course of their evolution. These experiments were identical to the original experiment, except that no functions—including even EQU—were rewarded. In 10 replicate experiments starting from Ancestor3, every function was lost, showing that insufficient mutational pressure could not explain the failure for losses of function to occur.

A second possibility is that these functions were maintained because their performance was coupled to that of EQU—in other words, due to pleiotropy. One line of evidence that pleiotropy was often responsible for the maintenance of some functions is that their performance, despite not being rewarded, often increased during evolution in the EQU environment and, in many cases, in proportion to that of EQU. Figure 6 shows the phenotype of evolved organisms from three different ancestors. Correlations comparing the performance of retained functions to that of EQU are consistently significant (Ancestor10, OR-NOT: $r = 0.72$, $df = 6$, $P = .044$; Ancestor1, OR: $r = 0.81$, $df = 8$, $P = .005$; Ancestor2, NOT: $r = 0.98$, $df = 7$, $P < .0001$; Ancestor2, OR-NOT: $r = 0.90$, $df = 6$, $P = .002$).

Given that different functions appeared to be coupled in their performance, we wanted to see whether we could understand the mapping between genotype and phenotype that gave rise to these correlations. In other words, rather than merely observing that some genotypes retained more unused functions than others, we wanted to understand the origins of this evolutionary pattern by investigating the relationship among different functions in the ancestral genome. To do so, we systematically assessed all 10 ancestral genomes for the extent to which knocking out a given instruction affected the performance of each function. The resulting “genotype-phenotype map” allowed us to infer the regions of the genome that encoded each function an organism performs, as well as the overlap in these regions.

An example of a genotype-phenotype map, constructed for Ancestor1, is shown in figure 7. Each row of the map represents one of the instructions in the genome, starting from the first (top row) to the last instruction (bottom row). Taking each site in the genome in turn, we replaced the instruction present at that site with a null instruction, called nop-X, and then tested the ability of the resulting knockout mutant to perform logic functions. Organisms were tested only for functions that the unmutated “wild-type” organism had itself performed. Each column of the map denotes a different logic function that could be performed by the unmutated organism, and the cells are colored as follows. White means that when the instruction in the corresponding row is replaced with a null instruc-

Ancestor 10								
NOT	NAND	AND	OR-N	OR	AND-N	NOR	XOR	EQU
0	0	0	0	0	0	0	0	71
0	32	0	32	0	0	0	0	62
0	0	0	66	0	0	0	0	261
0	9	0	0	0	0	0	0	31
0	17	0	17	0	16	0	0	17
0	16	0	16	0	0	0	0	46
0	0	0	1	0	0	0	0	1
0	0	0	24	24	19	0	0	24
0	0	0	12	0	0	0	0	12
0	79	0	79	0	0	0	0	79
correlation with EQU: 0.72								
Ancestor 1								
NOT	NAND	AND	OR-N	OR	AND-N	NOR	XOR	EQU
0	0	0	0	91	0	0	0	361
0	0	0	1	1	0	0	0	1
1	0	0	0	1	0	0	0	3
0	0	0	0	101	0	0	0	101
0	0	0	0	32	0	2	0	62
0	0	0	0	1	0	0	0	2
0	0	0	0	1	0	0	0	2
0	0	0	0	1	0	0	0	1
0	0	0	0	1	0	0	0	1
0	0	0	0	44	0	0	0	130
correlation with EQU: 0.81								
Ancestor 2								
NOT	NAND	AND	OR-N	OR	AND-N	NOR	XOR	EQU
24	0	0	24	0	0	0	0	96
158	0	0	239	0	76	158	0	316
488	0	0	0	0	0	0	0	1952
1	48	0	49	0	0	0	0	196
1	0	0	1	0	0	0	0	4
11	0	0	11	0	0	0	0	64
58	0	0	58	0	0	0	0	173
0	1	0	1	0	0	0	0	4
29	0	0	0	0	0	0	0	116
1	0	0	1	0	0	0	0	4
correlation with EQU: 0.98 (NOT) 0.90 (OR~)								

Figure 6: Functions that were not lost were correlated with the performance of EQU. Each row represents the phenotype of a different evolved organism (one from each of the 10 replicate populations) for three illustrative ancestors. Numbers show the number of times the organism performs each logic function per life cycle. While many functions were lost (indicated by 0), those that were not lost show a correlated increase in their performance with EQU. Correlation coefficients for the performance of these functions with EQU are indicated below each table. Correlations were calculated only between pairs of data where the function in question had been maintained (i.e., if the value in the table was >0) in over half the replicate populations, as indicated by shading.

tion, there is no effect on the function in the corresponding column. Colored cells indicate that replacing the corresponding instruction with a null instruction resulted in a loss of that function, and thus these cells correspond to the areas of the genome that encode the different functions. Among the colored cells, red cells represent the subset of the instructions required for any given function that are also required for the EQU function. For any other

function, a mixture of red and blue therefore indicates only partial overlap with the instructions that encode EQU. Logic functions that lack any blue coloring, such as the function OR, indicate that there are no sites in the genome that can be mutated to cause the loss of that function and still maintain the EQU function. We therefore expect that functions such as these might rarely be lost in an EQU-only environment, owing to the absence of any mutational target that does not also affect EQU. Consistent with this expectation, populations evolved from this ancestor in an EQU-only environment never lost the ability to perform OR but usually lost all other unnecessary functions (fig. 4, upper left). Analysis of the genotype-phenotype map for this ancestor thus implies that the differential overlap in the encoding of the various functions with that of the EQU function is at least partly responsible for their differential maintenance during evolution in an EQU-only environment.

To assess this relationship more generally, we used these genotype-phenotype maps for all the ancestors to identify the genomic regions that corresponded to each logic function. We then assessed the extent to which each of these functions overlapped with the EQU function and calculated the number of nonoverlapping instructions. We then determined how many times (out of a possible 10) each function was actually lost during evolution in an EQU-only environment. The relationship between these two measures is presented in figure 8 and shows that functions that overlap completely with EQU (i.e., those with 0 nonoverlapping instructions) were most likely to be maintained, but the probability of maintenance drops rapidly as the number of nonoverlapping instructions increases. This result provides compelling support for our hypothesis that the integration of these functions in the genome played an important role in maintaining certain unused functions during evolution in the EQU-only environment. Specifically, it demonstrates that niche breadth evolution in this system was driven not only by the selective environment in which these organisms evolved but also by the way in which their genotypes mapped onto their phenotypes, that is, by their genetic architecture.

Discussion

Two distinct population genetic mechanisms are thought to promote the evolution of ecological specialization, as reflected in a narrow niche breadth. One entails the accumulation of mutations that are neutral or deleterious in a novel environment owing to relaxed selection on unused functions. In the other, fitness improvements in a novel environment may come at the expense of other traits, leading to trade-offs and losses of function. Trade-offs can occur if the same genes contribute to two or more




Figure 7: Genotype-phenotype map showing the instructions that encode each function in Ancestor1. Each row in the map represents a single instruction, starting from the first instruction in the organism’s genome (*top row*) to the final instruction (*bottom row*). Each column represents a different function performed by the ancestor, and the coloring indicates what happens to the performance of that function when a given instruction is knocked out (replaced by a null instruction). *White* = knocking out the instruction does not affect performance of the function; *blue* or *red* = knocking out the instruction causes the function to be lost. Red blocks indicate the subset of instructions required for a given function that, when knocked out, also cause the loss of EQU. Note that every instruction in this organism that knocks out OR also knocks out EQU, whereas this pattern does not hold for NAND, OR-N, AND-N, or XOR.

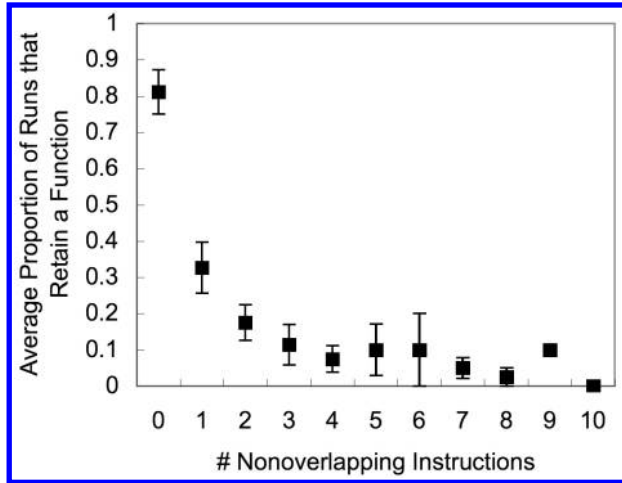


Figure 8: Association between the number of nonoverlapping instructions (required to perform some function but not required for the EQU function) and the average proportion of times (out of a possible 10) that the function was maintained during evolution in the EQU-only environment. Error bars represent 1 SE.

traits, such that mutations that improve one may worsen others. Trade-offs can also arise, even when traits do not share a genetic basis, if the maintenance of unselected or weakly selected traits entails an energetic burden.

Here, we describe the evolution of ecological specialization in digital organisms. Starting from a set of generalist ancestors, each of which could perform a wide variety of logic computations, we examined their adaptation to a novel environment where only a single computation was directly selected. A benefit of examining the process of specialization in digital organisms is that we can precisely trace the mutational steps leading from the generalist ancestor to the evolved specialist, which allows us to examine in detail the mutations that lead to losses of function along the way.

Our results revealed significant heterogeneity in the magnitudes of fitness improvements in the novel EQU-only environment, with different populations evolving to perform EQU to different extents depending on the ancestor used to initiate the experiment. All of the ancestors performed EQU once per gestation cycle at the start of the experiment, and, in a few cases (five out of 100), evolved organisms did not increase their performance of EQU above the ancestral level. In most cases, however, organisms evolved to perform the function tens or even hundreds of times per gestation cycle. The evolved organisms also varied in the extent to which their niche breadth became narrower, with very few populations (only seven out of 100) evolving to become pure EQU specialists that could not perform any other functions.

Examination of the mutations that led to losses of function allowed us to quantify the relative importance of mutation accumulation and antagonistic pleiotropy. These data showed that, in absolute terms, more losses of function were caused by neutral or deleterious mutations than by beneficial mutations. Yet, when we standardized for the greater numbers of nonbeneficial substitutions along the lines of descent, beneficial mutations were disproportionately associated with losses of function. Although the proportion of losses of function that could be attributed to beneficial mutations was generally higher at lower mutation rates, the proportion of beneficial substitutions overall was also higher. For this reason, changes to the mutation rate had little effect on the results of the Fisher's exact tests. Nevertheless, functions were more often lost due to the accumulation of neutral and deleterious mutations, and this effect increased with increasing mutation rate.

The finding that lower mutation rates permit the fixation of proportionally more beneficial mutations suggests that some kind of interference is occurring at the higher mutation rates, although it is not clear whether the interference arises from deleterious or beneficial mutations. At high mutation rates, beneficial and deleterious mutations may often arise on the same background, limiting fixation to those beneficial mutations of large effect (Peck 1994; Orr 2000; Johnson and Barton 2002). High mutation rates can also lead to interference among beneficial mutations that arise in different clonal lineages (Gerrish and Lenski 1998; de Visser et al. 1999). Distinguishing between these two alternatives in evolving digital populations is a subject for further study.

One surprising result of these experiments was how, in particular ancestors, certain functions were often maintained in the absence of direct selection for their performance. Examination of the genetic architecture of the ancestors revealed that overlap in the genetic instructions that encode the different functions was a good predictor of their maintenance during evolution in the EQU-only environment. Not only were these functions maintained but also their performance often increased in parallel with that of EQU, resulting in unexpected positive correlations between certain traits across populations evolved from the same ancestor. Because we know that there was no direct selection on these functions, their maintenance is more analogous to that of a vestigial trait rather than the outcome of selection operating on two traits simultaneously. Wright (1977, p. 428) and Lande (1978) both suggested that useless or even slightly detrimental functions might be retained over long periods of time, owing to their pleiotropic relationships to characters under direct selection. Nevertheless, it would be interesting to examine the consequences of these genetically integrated traits in the event that selection were to operate on them in opposing di-

rections (Beldade et al. 2002), and we will examine this possibility in future work.

Our results show that there was no single function that was always retained with EQU; rather, the identity of the retained functions varied depending on the particular ancestor. For example, organisms evolved from Ancestor3 failed to lose AND and OR, whereas organisms evolved from Ancestor2 often failed to lose NOT and OR-NOT. Moreover, because the ancestors all shared the same historical environment, these differences in outcome reflect stochasticity in the origins of each ancestor's unique genetic architecture—a genetic architecture that influenced the subsequent trajectory of evolution in the EQU-only environment. Where multiple functions were maintained (e.g., Ancestor3), it would be interesting to explore whether they had been built on each other sequentially. One could imagine, for instance, that EQU evolved from AND and that AND evolved from OR, and so on. Of course, the construction of complex functions out of simpler ones—a process that contributes to the emergence of pleiotropy in this system—also occurs in natural systems (Jacob 1977; Nilsson and Pelger 1994; Meléndez-Hevia et al. 1996; Chen et al. 1997; Dean and Golding 1997). Thus, investigations into the form and direction of pleiotropy in nature might be informed through a consideration of the evolutionary history of the traits in question.

The importance of genetic integration for the maintenance of unrewarded functions in highly specialized environments led to substantial variation in the niche breadth of evolved organisms, with some organisms evolving very narrow specialization and others maintaining their niche breadth at about half their ancestral levels. This result raises interesting questions about the relative long-term success of these organisms in a fluctuating environment, where antagonistic pleiotropy and mutation accumulation may continually degrade functions that might become necessary again at some later time (Kawecki 2000). Organisms with highly integrated genetic architectures would potentially prosper in such environments, whereas those with greater modularity might do better in a more stable environment, particularly if genetic correlations among traits were found to constrain the optimization of each trait individually. These predictions do not differ from existing theories about the kinds of environments that select for generalist versus specialist species, with the former predicted to emerge in a temporally heterogeneous environment and the latter when there is environmental constancy (Levins 1968; Futuyma and Moreno 1988; Kassen 2002). However, this perspective emphasizes the role of genetic architecture in mediating these transitions rather than selection as the sole determinant of niche breadth.

Although our results show that most losses of function resulted from neutral or deleterious mutations, we note

that several factors are expected to alter the relative contributions of antagonistic pleiotropy and mutation accumulation to ecological specialization. Our results indicate that an important consideration is the relative likelihood that a given mutation will be beneficial or not: increases to the mutation rate, in combination with asexual reproduction, permitted the fixation of more neutral and deleterious mutations, increasing losses of function that result from these mutation types. Several authors have recently considered the role of deleterious mutations in adaptation, and this work has led to a reevaluation of how mutation rate alters the rate of adaptation in asexual organisms (Orr 2000; Johnson and Barton 2002; Wilke 2004). One of the difficulties encountered by this work is the complexity of the process, which requires modeling many competing lineages and evaluating nonequilibrium conditions, making it difficult to derive exact solutions. Digital systems may prove to be a suitable testing grounds for some of the hypotheses generated by this work, particularly because of the ease of estimating parameters that are difficult to measure in biological systems, such as the rate of occurrence and fixation of beneficial and deleterious mutations (see also Rozen et al. 2002; Sanjuan et al. 2004). Moreover, mutation-accumulation explanations for specialization often assume that the relevant mutations are either conditionally or weakly deleterious because unconditionally deleterious mutations have difficulty attaining fixation except in small populations (Kawecki 1994; MacLean et al. 2004). In asexual organisms, however, deleterious mutations can hitchhike to fixation alongside beneficial mutations. Given that many extreme examples of adaptive decay involve bacteria (Cole et al. 2001; Ochman and Moran 2001; Wernegreen et al. 2002), the potential role of deleterious mutations needs to be considered more carefully. In asexual organisms, niche breadth reductions could be occurring by both antagonistic pleiotropy (fixation of beneficial mutations) and mutation accumulation (via increased fixation of deleterious mutations), with the fixation of the former predisposing that of the latter through hitchhiking. This result suggests that ecological specialization, by this coupling of mutation accumulation and antagonistic pleiotropy, may occur more readily in asexual organisms.

Ecological theories of niche specialization predict that organisms will evolve to match the heterogeneity of their environment (Levins 1968; Via and Lande 1985; Scheiner 1993; Kassen 2002). Our results show that environmental constancy can, in fact, drive the evolution of niche breadth reductions, with all organisms evolving niche breadths that were narrower than that of their ancestors. However, substantial diversity in niche breadth was observed among independently evolved organisms despite their evolution in identical environments. While the extent to which traits are encoded by the same or different genes has sometimes

been taken into consideration when predicting the relative importance of antagonistic pleiotropy and mutation accumulation in driving specialization (Futuyma and Moreno 1988; Fry 1993; Kawecki 1994, 1998), the degree to which suites of traits may be genetically integrated has not often been considered with regard to the maintenance of functions across environments (but see Rausher 1988).

Although trade-offs are widely thought to promote the evolution of ecological specialization, the requisite negative correlations have often not been forthcoming (Jaenike 1990; Fry 1996; Agrawal 2000). The failure to detect negative correlations has led to a developing body of work that focuses on alternative explanations for the evolution of ecological specialization (Kawecki 1994, 1998; Whitlock 1996). In addition, Rausher (1988) suggested that trade-offs may not always be detected. For example, studies of diet breadth in phytophagous insects often employ host species that are already part of the natural diet. If severe trade-offs exist, then those host species are more likely to have been excluded from the diet previously, and the observed niche breadth will consist only of hosts for which there was either little or no conflict. Our results are consistent with this hypothesis, with trade-offs often quickly leading to losses of function, leaving mostly positive correlations among the remaining functions.

Ultimately, of course, experiments with digital organisms cannot tell us what processes are actually at work in

any given natural system—that is an empirical question that cannot be addressed by any model system, digital or otherwise. However, digital systems provide a novel way of assessing the logic that underlies many evolutionary theories, especially where complex interactions limit the opportunity for purely theoretical analysis. Our results show that ecological specialization occurs in digital organisms and, moreover, that some of the same patterns that have complicated simple theories of niche breadth in natural systems, such as the apparent paucity of trade-offs and an excess of positive correlations, also emerge here. Finally, digital systems offer the ability to connect patterns to processes and thus allow investigations of causal mechanisms more directly than is possible in any other system, enabling tests of existing hypotheses as well as the development of new ones that can in turn be tested in other systems.

Acknowledgments

We thank J. Conner, T. Cooper, K. Gross, A. Jarosz, and two anonymous reviewers for helpful comments on an earlier version of this manuscript. This work was supported by grants from the National Science Foundation to R.E.L. and C.O. and by support from the Quantitative Biology and Modeling Initiative at Michigan State University to E.A.O.

APPENDIX A
Schematic of a Digital Organism in Avida

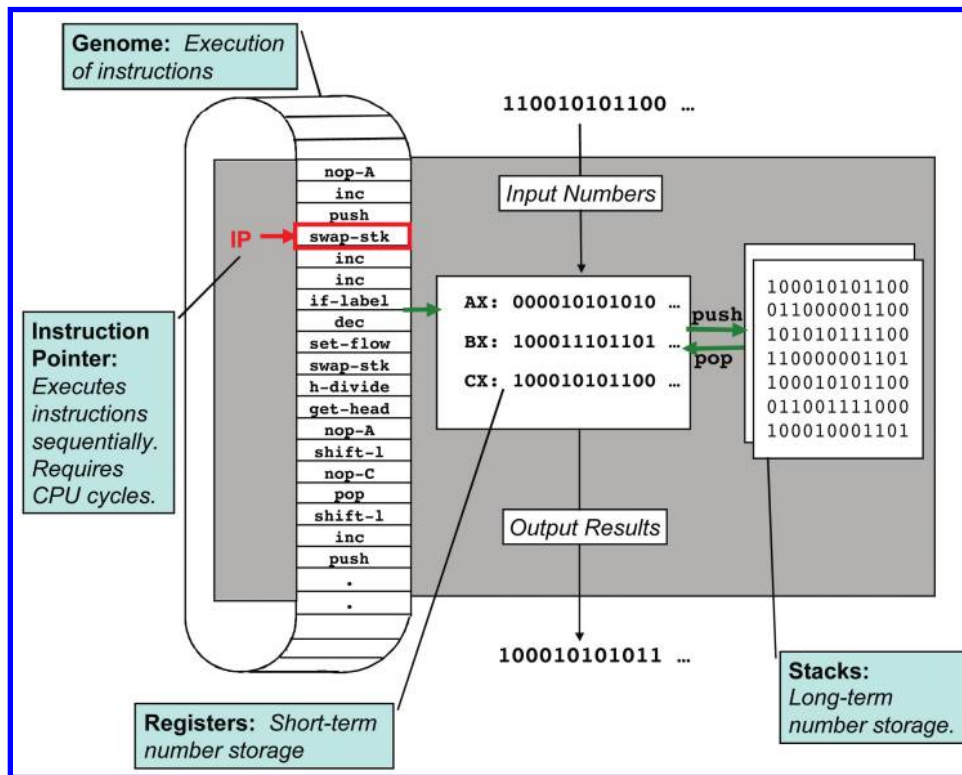


Figure A1: A digital organism consists of a genome (computer program), three registers, two stacks, and four heads (one of which is the instruction pointer [IP]). Execution of the program requires central processing unit cycles, and the point of execution is indicated by the location of the IP. An input-output instruction enables an organism to input binary numbers into the registers and output the results of computations. Most genomic instructions operate directly on the numbers in the registers, although the push and pop instructions cause numbers in the registers to be pushed onto the stack or popped off of the stack, respectively. The stacks are thus primarily used for storing numbers, whereas the registers are used to manipulate them.

APPENDIX B
Glossary of Terms

Table B1: Glossary of terms

Term	Definition
CPU	Central processing unit. All organisms have the instructions in their genomes executed by a virtual CPU. A mutation that causes an organism to have more CPU cycles (i.e., to have its genome executed faster than others) is generally beneficial.
Digital organism	A virtual computer, consisting of a genome (a computer program) and its associated hardware. The hardware consists of the CPU, which processes the instructions in the genome, two stacks, and three registers, which are used for storing, retrieving, and manipulating numbers. Each organism also has an instruction pointer (IP), which points to the next instruction to be executed in the genome, and read, write, and flow heads, which are used to specify positions in memory, such as in the copy process or for jumping and looping.

Table B1 (Continued)

Term	Definition
EQU	A logic function, where two binary inputs are compared and the correct output is a 1 if the input bits are the same and a 0 if the bits are different. In this system, EQU is actually a bitwise EQU, in that the correct output is the computation of EQU across all 32 bits for the two inputs. Performance of EQU requires, at a minimum, combining the outputs of five different NAND statements, in coordination with various other instructions.
Genome	Sequence of instructions that may contain information for making duplicate copies of the genome and for interactions with the environment. Execution of the instructions in a properly functioning genome leads to the production of an offspring.
Gestation time	Number of instructions executed, and hence CPU cycles required, to produce an offspring. Gestation time is generally a multiple of genome length but varies as a function of the efficiency of the copy process and the number of loops in execution.
Instruction	Units that comprise the genome. Each site in the genome is one of 26 possible instructions. Instructions not present in an ancestral genome may be introduced into the genomes of descendants via mutation.
Logic function	Computations based on binary inputs. Organisms may evolve to perform bitwise logic functions based on numbers they input from the environment.
Mutations	Mutations can be point mutations, where one instruction is randomly replaced with another during the copy process. Mutations can also be insertion or deletion mutations, causing genomes to grow or shrink in length. The rates of point, insertion, and deletion mutations are specified by the experimenter.
NAND	One of the 26 possible instructions in the genome. Also a core logic function; all other logic functions can be built from combinations of NANDs.

Literature Cited

- Agrawal, A. A. 2000. Host range evolution: adaptation and trade-offs in fitness of mites on alternative hosts. *Ecology* 81:500–508.
- Andersson, S. G. E., A. Zomorodipour, J. O. Andersson, T. Sicheritz-Ponten, U. C. M. Alsmark, R. M. Podowski, A. K. Naslund, A. S. Eriksson, H. H. Winkler, and C. G. Kurland. 1998. The genome sequence of *Rickettsia prowazekii*. *Nature* 396:133–143.
- Beldade, P., K. Koops, and P. M. Brakefield. 2002. Developmental constraints versus flexibility in morphological evolution. *Nature* 416:844–847.
- Borowsky, R., and H. Wilkens. 2002. Mapping a cave fish genome: polygenic systems and regressive evolution. *Journal of Heredity* 93:19–21.
- Chen, L., A. L. DeVries, and C.-H. C. Cheng. 1997. Evolution of antifreeze glycoprotein gene a trypsinogen gene in Antarctic notothenioid fish. *Proceedings of the National Academy of Sciences of the USA* 94:3811–3816.
- Cole, S. T., K. Eiglmeier, J. Parkhill, K. D. James, N. R. Thomson, P. R. Wheeler, N. Honore, et al. 2001. Massive gene decay in the leprosy bacillus. *Nature* 409:1007–1011.
- Cooper, V. S., and R. E. Lenski. 2000. The population genetics of ecological specialization in evolving *Escherichia coli* populations. *Nature* 407:736–739.
- Cooper, V. S., A. F. Bennett, and R. E. Lenski. 2001. Evolution of thermal dependence of growth rate of *Escherichia coli* populations during 20,000 generations in a constant environment. *Evolution* 55:889–896.
- Darwin, C. 1859. *On the origin of species by means of natural selection*. J. Murray, London.
- Dean, A. M., and G. B. Golding. 1997. Protein engineering reveals ancient adaptive replacements in isocitrate dehydrogenase. *Proceedings of the National Academy of Sciences of the USA* 94:3104–3109.
- de Visser, J. A. G. M., C. W. Zeyl, P. J. Gerrish, J. L. Blanchard, and R. E. Lenski. 1999. Diminishing returns from mutation supply rate in asexual populations. *Science* 283:404–406.
- Falconer, D. S. 1952. The problem of environment and selection. *American Naturalist* 86:293–298.
- Fry, J. D. 1993. The general vigor problem: can antagonistic pleiotropy be detected when genetic covariances are positive? *Evolution* 47:327–333.
- . 1996. The evolution of host specialization: are trade-offs overrated? *American Naturalist* 148(suppl.):S84–S107.
- Futuyma, D. J., and G. Moreno. 1988. The evolution of ecological specialization. *Annual Review of Ecology and Systematics* 19:207–233.
- Gerrish, P. J., and R. E. Lenski. 1998. The fate of competing beneficial mutations in an asexual population. *Genetica* 102/103:127–144.
- Gómez-Valero, L., A. Latorre, and F. J. Silva. 2004. The evolutionary fate of nonfunctional DNA in the bacterial endosymbiont *Buchnera aphidicola*. *Molecular Biology and Evolution* 21:2172–2181.
- Jacob, F. 1977. Evolution and tinkering. *Science* 196:1161–1166.
- Jaenike, J. 1990. Host specialization in phytophagous insects. *Annual Review of Ecology and Systematics* 21:243–273.
- Jeffrey, W. R., A. G. Strickler, and Y. Yamamoto. 2003. To see or not to see: evolution of eye degeneration in Mexican blind cavefish. *Integrative and Comparative Biology* 43:531–541.
- Jernigan, R. W., D. C. Culver, and D. W. Fong. 1994. The dual role

- of selection and evolutionary history as reflected in genetic correlations. *Evolution* 48:587–596.
- Johnson, T., and N. H. Barton. 2002. The effect of deleterious alleles on adaptation in asexual populations. *Genetics* 162:395–411.
- Jones, R., D. C. Culver, and T. C. Kane. 1992. Are parallel morphologies of cave organisms the result of similar selection pressures? *Evolution* 46:353–365.
- Kassen, R. 2002. The experimental evolution of specialists, generalists, and the maintenance of diversity. *Journal of Evolutionary Biology* 15:173–190.
- Kawecki, T. J. 1994. Accumulation of deleterious mutations and the evolutionary cost of being a generalist. *American Naturalist* 144: 833–838.
- . 1998. Red Queen meets Santa Rosalia: arms races and the evolution of host specialization in organisms with parasitic lifestyles. *American Naturalist* 152:635–651.
- . 2000. The evolution of genetic canalization under fluctuating selection. *Evolution* 54:1–12.
- Lande, R. 1978. Evolutionary mechanisms of limb loss in tetrapods. *Evolution* 32:73–92.
- Lenski, R. E., C. Ofria, R. T. Pennock, and C. Adami. 2003. The evolutionary origin of complex features. *Nature* 423:139–144.
- Levins, R. 1968. *Evolution in changing environments*. Princeton University Press, Princeton, NJ.
- MacLean, R. C., G. Bell, and P. B. Rainey. 2004. The evolution of a pleiotropic fitness trade-off in *Pseudomonas fluorescens*. *Proceedings of the National Academy of Sciences of the USA* 101:8072–8077.
- Meléndez-Hevia, E., T. G. Waddell, and M. Cascante. 1996. The puzzle of the Krebs citric acid cycle: assembling the pieces of chemically feasible reactions, and opportunism in the design of metabolic pathways during evolution. *Journal of Molecular Evolution* 43:293–303.
- Muller, H. J. 1932. Some genetic aspects of sex. *American Naturalist* 66:118–138.
- Nilsson, D.-E., and S. A. Pelger. 1994. A pessimistic estimate of the time required for an eye to evolve. *Proceedings of the Royal Society B: Biological Sciences* 256:53–58.
- Ochman, H., and N. A. Moran. 2001. Genes lost and genes found: evolution of bacterial pathogenesis and symbiosis. *Science* 292: 1096–1098.
- Ofria, C., and C. O. Wilke. 2004. Avida: a software platform for research in computational evolutionary biology. *Journal of Artificial Life* 10:191–229.
- Ofria, C., C. Adami, and T. C. Collier. 2003. Selective pressures on genomes in molecular evolution. *Journal of Theoretical Biology* 4:477–483.
- Orr, H. A. 2000. The rate of adaptation in asexuals. *Genetics* 155: 961–968.
- Peck, J. R. 1994. A ruby in the rubbish: beneficial mutations, deleterious mutations and the evolution of sex. *Genetics* 137:597–606.
- Petrov, D. A., and D. L. Hartl. 2000. Pseudogene evolution and natural selection for a compact genome. *Journal of Heredity* 91: 221–227.
- Rausher, M. D. 1988. Is coevolution dead? *Ecology* 69:898–901.
- Rozen, D. E., J. A. G. M. de Visser, and P. J. Gerrish. 2002. Fitness effects of fixed beneficial mutations in microbial populations. *Current Biology* 12:1040–1045.
- Sanjuan, R., A. Moya, and S. F. Elena. 2004. The distribution of fitness effects caused by single-nucleotide substitutions in an RNA virus. *Proceedings of the National Academy of Sciences of the USA* 101:8396–8401.
- Scheiner, S. M. 1993. Genetics and evolution of phenotypic plasticity. *Annual Review of Ecology and Systematics* 24:35–68.
- Schluter, D. 2000. *The ecology of adaptive radiation*. Oxford University Press, Oxford.
- Shigenobu, S., H. Watanabe, M. Hattori, Y. Sakaki, and H. Ishikawa. 2000. Genome sequence of the endocellular bacterial symbiont of aphids *Buchnera* sp. APS. *Nature* 407:81–86.
- Simpson, G. G. 1953. *The major features of evolution*. Columbia University Press, New York.
- Sliwa, P., and R. Korona. 2005. Loss of dispensable genes is not adaptive in yeast. *Proceedings of the National Academy of Sciences of the USA* 102:17670–17674.
- Sokal, R. R., and F. J. Rohlf. 1995. *Biometry*. Freeman, New York.
- Via, S., and R. Lande. 1985. Genotype-environment interaction and the evolution of phenotypic plasticity. *Evolution* 39:505–522.
- Wernegreen, J. J., A. B. Lazarus, and P. H. Degnan. 2002. Small genome of *Candidatus Blochmannia*, the bacterial endosymbiont of *Camponotus*, implies irreversible specialization to an intracellular lifestyle. *Microbiology* 148:2551–2556.
- Whitlock, M. C. 1996. The Red Queen beats the jack-of-all-trades: the limitations on the evolution of phenotypic plasticity and niche breadth. *American Naturalist* 148(suppl.):S65–S77.
- Wilke, C. O. 2004. The speed of adaptation in large asexual populations. *Genetics* 167:2045–2053.
- Wilke, C. O., and C. Adami. 2002. The biology of digital organisms. *Trends in Ecology & Evolution* 17:528–532.
- Wright, S. 1977. *Evolution and the genetics of populations*. Vol. 3. *Experimental results and evolutionary deductions*. University of Chicago Press, Chicago.

Associate Editor: Richard Gomulkiewicz
Editor: Michael C. Whitlock