BioSystems, 24 (1990) 61–81
Elsevier Scientific Publishers Ireland Ltd.

Note: conditions for evolutionary adaptability discussed at length in my monograph "Adaptability" (Plenum Press, 1983) especially ch. 6.

The geometry of evolution

Michael Conrad

Department of Computer Science, Wayne State University, Detroit MI 48202 (U.S.A.)

(Received August 7th, 1989)
(Revision received April 20th, 1990)

Some structures are more suitable for self-organization through the Darwin-Wallace mechanism of variation and selection than others. Such evolutionary adaptability (or evolvability) can itself evolve through variation and selection, either by virtue of being associated with reliability and stability or by hitchhiking along with the advantageous traits whose appearance it facilitates. In order for a structure to evolve there must be a reasonable probability that genetic variation carries it from one adaptive peak to another; at the same time the structure should not be overly unstable to phenotypic perturbations, as this is incompatible with occupying a peak. Organizations that are complex in terms of numbers of components and interactions are more likely to meet the peak-climbing condition, but less likely to meet the stability condition. Biological structures that are characterized by a high degree of component redundancy and multiple weak interactions satisfy these conflicting pressures.

Keywords: Evolution; Adaptive landscape; Stability; Complexity; Structuralism; Evolvability.

1. Structuralism and Darwinism

Recent years have seen spirited attacks and defenses of 'NeoDarwinism'. The whole complex of arguments is too difficult to characterize in a few short sentences. But it is probably adequate to say that what is being brought into question is the adequacy of the modern synthesis. Brian Goodwin (1985) has put the matter in a particularly succinct way. According to Goodwin the NeoDarwinian view treats the organism as a historical accident. Out of the materials provided by genetic variation the molding power of selection can create anything. Goodwin argues, instead, that structure is important; only certain forms are possible and selection chooses among these (cf. also Sibatani, 1985). The question, what are the possible structural forms, is equivalent to the question, what embryologies are possible?

In this author's opinion it is unfortunate that the term NeoDarwinism has become tagged with the anemia observed by Goodwin and with other anemias of some recent dominant trends in evolutionary thinking. There is no reason why the Darwinian framework should fail to accommodate and even precipitate new understandings of morphogenesis and hierarchical ecosystems biology (e.g. Conrad, 1983; Salthe, 1985). The tagging, I believe, is due to a misplaced proscription. Some self-styled NeoDarwinists have strongly rejected terms such as 'evolvability' . As a consequence the horse of variation and selection always pulls the cart of structure and can never be pushed by it.

Structure A, for example, could certainly have more evolutionary potentiality than Stucture B. Or structure A′ might perform the same function as A, yet have greater evolutionary adaptability. It would clearly be good for the evolutionary process if it could select A over B and A′ over A. But to some orthodox NeoDarwinists this smacks of group selection. And if one cannot talk about evolvability evolving, why should one be allowed to talk about the material structural basis of evolution at all? Not that our self-styled defender of Darwinism would deny a material

0303-2647/90/$03.50

basis of evolution. What the view seems to be is that this material basis is a combination of physics and history, but that no evolutionary considerations affect the history apart from those directly connected with the particular constraints necessary for the life of the individual organism. All the work of discovering these particular constraints can be attributed to the tremendous search power of variation and selection.

My purpose in this paper is to show that so far as evolvability is concerned the contrary structuralist view is just as Darwinian, if not twice as Darwinian. The picture is that a very special class of structures is particularly amenable to evolution, and that these are themselves selected through the Darwinian mechanism of variation and selection. The chief characteristic of this special class of structures is that it must satisfy at one and the same time two conflicting conditions. The first is that the organism be stable, that it sit in a developmental basin of attraction. This is more likely as the number of components in the organism and the number of interactions among them decreases, on the simple grounds that the chance of a valley occurring in the phase space of a system decreases with its dimensionality (May, 1973). The second condition is that the adaptive peaks corresponding to these basins of attraction should be close enough together to be connected by single genetic changes. But this is more likely as the number of components and interactions increases, since pathways between peaks in the adaptive peak space correspond to pathways between valleys in the basin space. The only way for a system to satisfy both conditions is to have many redundant components with multiple weak interactions. In this case the developmental system can have many genetically related homomorphic images. The extra components and weak interactions that allow for this special situation are costs to the individual organism; the structure that is most amenable to evolution will be functionally less effective from the thermodynamic point of view. Nevertheless the amenability-increasing structural features inevitably hitchhike along with the advantageous traits whose evolution they facilitate.

2. The classical adaptive landscape

Let us first recall the classical adaptive

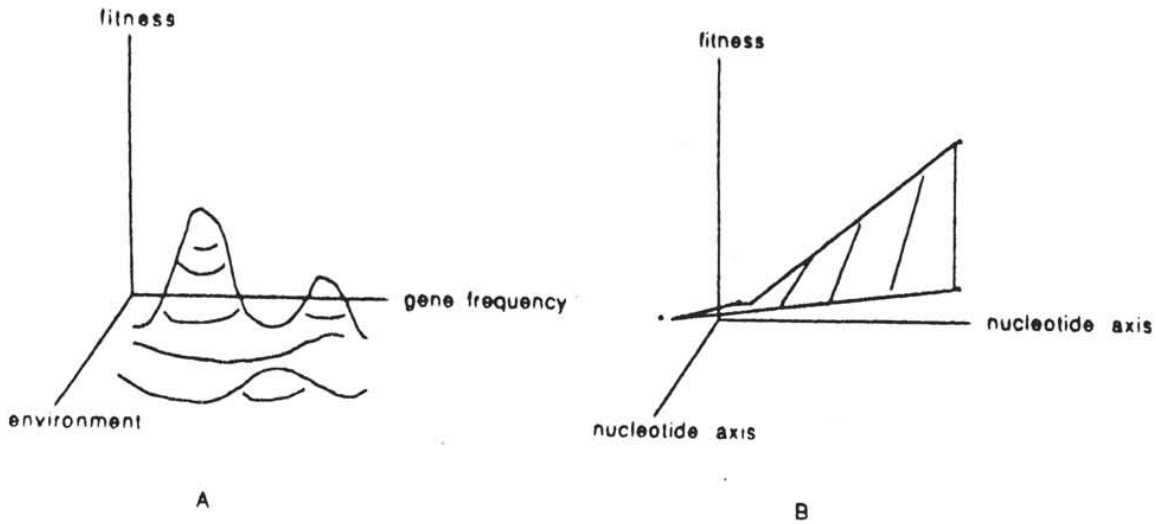

Fig. 1. Schematic picture of classical (A) and molecular (B) adaptive landscape. The classical landscape could have many gene and environment axes. The molecular landscape would have four nucleotide axes for each base position in DNA and also environment axes (omitted). The two axes in the illustration represent nucleotides at two different base positions and the black dots the four possible configurations. The lined plane schematically represents an imaginary sheet thrown over the points in the space.

landscape of Wright (1932). This is essentially an assignment of a performance (or fitness) measure to gene frequency (Fig. 1A). We can picture one axis for each gene locus and one axis for the fitness measure. To this we can add axes for each environmental variable, and we can add an axis for time. The fitness surface is a manifold with hyperdimensional peaks and valleys, and evolution may be pictured as a flow of particles (representing organisms) on this surface. Because of variation and selection the particles tend to flow uphill, towards the tops of the peaks. Or in some instances they may flow along shallow valleys between peaks, or along upward rising passes that connect peaks that are separated by valleys, possibly deep valleys, along most axes.

The particles, since they represent organisms, can disappear (die) or appear (be born). As a consequence only one or two particles need reach a new peak. Similarly if a peak moves, due to change in the environment, particles will be pushed off peaks, leading to particle death and in some cases extinction of whole populations. Evolutionary change in this case is necessary just to stay on the same moving peak.

As indicated above, we are here taking fitness as a measure of performance. Sometimes performance is defined in terms of relative contribution to the following generation (cf. Waddington, 1968). Unfortunately this definition has internal difficulties; for example, the relative contribution depends on the size of the population. Fitness in the original sense used by Darwin expressed a relationship between the organism and the environment. Those organisms whose traits actually fit the environment, somewhat like a key might fit to a lock, are more likely to reproduce. Such fitness relationships are in general too complex to capture with a scalar measure. For the present purposes we can simply regard our use of a scalar measure as a conceptual construct that is useful for analyzing evolutionary processes.

Fitness, viewed as a relationship between organism and environment, is close to the concept of biological function. We can view any structure in biology in functional terms; but we must not forget that the concept of function is inherently ambiguous. A given structure could assume an indefinite number of potential functions. The dramatic transformations of function that have occurred in biological evolution are an expression of this fact. These structures are not obliged to adhere to the functions we initially ascribe to them. As a consequence, function is highly milieu dependent. It can change radically as the environment changes and as different peaks are occupied. When this occurs fitness and the peak structure of the landscape itself change. As unoccupied peaks become occupied the environment changes, and the peak structure of the space changes. Our adaptive space is really rather more of a trampoline than a stiff surface, and many of the transformations that occur in evolution are probably better understood in terms of an endogenously malleable fitness manifold than in terms of a fixed geometrical surface built up from a well founded fitness measure. For the purposes of the ensuing analysis, however, we need not emphasize this point.

3. The molecular landscape

Genes are sequences of nucleic acid bases, and the structure of our adaptive surface should be modified to accommodate this fact (Fig. 1B, cf. Conrad, 1979a, 1983). To do this, let each potential base position on the genome be represented by four axes, corresponding to the four possible bases that could occupy this position (A, T, G, or C). If A occupies the position, the A axis is assigned the value 1. Otherwise it is assigned the value 0, and similarly for the other three axes. One of the four axes will be assigned the value 1 if the potential position is occupied by a base, all positions will be assigned the value 0 if no base occupies the position, and under no circumstance will two or more of these four axes be assigned the value 1.

We will say that the dimensionality of a genotype is equal to the number of non-zero axes. Thus if a genome comprises 10^6 nucleotide bases it has dimensionality 10^6. We can also suppose that the number of axes is sufficiently large to accommodate any potential genome. Or alternatively, we can suppose that axes can always be added to accommodate any oversized genome. Strictly speaking some extra axes might be necessary to represent breaks between different parts of the genome (in particular between chromosomes). But the addition of such technical complications is not necessary here.

To the set of gene axes we add environment axes, a time axis, and a fitness axis. The dimensionality of the space associated with a single organism is equal to the sum total of axes to which values are assigned. As with the dimensionality of a single genotype we can admit extra gene axes whenever needed. Many different organisms can then be described in the space. Thus we can picture each of the existing organisms in an ecosystem as represented by a point in the space. Each organism will occupy a subspace of the whole space, and these subspaces will intersect.

Since our gene axes are discrete (value either 1 or 0) the fitness surface is not a continuous manifold. There will still be high points (on the fitness axis) corresponding to peaks, and low points, corresponding to valleys and gorges. If we imagine that a hyperdimensional sheet is tossed over our fitness axis we obtain a continuous landscape of peaks and valleys.

4. The phenotypic space

Associated with each possible genotype in the fitness space is a phenotype, or collection of organism traits. These traits are an expression of the organization of atoms and molecules comprising the organism at any given time. The collection of possible phenotypes constitute a space which should map to the genotype space. In general, any one genotype will map to an ensemble of phenotypes, since the genome is differently expressed in response to different environmental histories. Also, different genotypes may in some instances map into a single phenotype (because of degeneracy of the genetic code).

Let us consider more specifically how we can construct a phenotype space and how the structure of the phenotype space for the collection of possible phenotypes maps into the genotype for the collection of possible genotypes. The first step is to choose a set of entities in terms of which organisms are to be described. These could be atoms, such as hydrogen, carbon, oxygen, nitrogen, and various other elements. Or they can be atoms, molecules, and macromolecules. Or they might be types of biological cells. As the entities become more complex the number of types increases. If we work with cells we would have to classify them into general types; we would have virtually an indefinite number of entities if we chose each distinguishable cellular organization as an entity. If we choose to work with a large number of entity types we will see fewer interactions, since many interactions will be buried in the entities. If we have fewer entity types, we will see more interactions, since we will then expose the interactions within previously admitted entities. Interactions can either be strong or weak. In general it is useful to identify collections of elementary entities as a complex entity only if there are a substantial number of strong interactions among them. But how we choose the entities is to a considerable extent a matter of convenience, determined by how it is most useful to conceptualize the system for the purposes at hand. For now we can leave this choice open, since for the argument to be developed it will not in the first pass make any difference if we have more entity types and fewer interactions or fewer interactions and more entity types.

For definiteness, though, let us take the entities as the atoms and molecules of which the organism is composed. Consider first a

four dimensional space (x, y, z, and t), and assign space-time coordinates to each of the entities. We can also assign different colors to the different entities, in which case we will obtain a cloud of 'colored points'. But it is more convenient to convert this cloud to a single particle moving in a hyperdimensional space, with each atom or molecule assigned its own set of position and momentum coordinates. Strictly speaking this is not sufficient. Each of the atoms and molecules also has internal characteristics (e.g. its electronic structure and configuration of nuclei) which is pertinent to its dynamical development. But our space is merely to be viewed as a conceptual picture, not as an in principle complete physical description. We might also note that organisms are open systems, so in principle we should include particles that enter and leave the boundaries that define the organism at any given point in time. But for all practical purposes we can use the fact that most of these are identical particles, so we can think in terms of a set of axes smaller than would be necessary to describe a whole ecosystem. However, we can include in our space the same environmental axes we use in the genotype space.

Our hyperdimensional point will form a hyperdimensional trajectory in time. We can repicture this trajectory in terms of a phase (or state) space. The time axis is eliminated. The trajectory represents how the state (hyperdimensional point) at any instant of time is mapped into the state at the next instant of time. We can further picture the set of possible trajectories in state space as a flow of points, starting from all initial points allowable for the phenotype. This is the usual global picture of a dynamical system in phase space (Rosen, 1970). The flow will in general have some basins of attraction—equilibria or steady states, limit cycles, chaotic attractors (Fig. 2A). The organism, if it is to 'earn the right to persist' must occupy such a basin of attraction. If it is perturbed by an environmental event or by an internal fluctuation, it must either return to its stable point or trajectory, or jump to another acceptable basin of attraction. An organism could have multiple steady

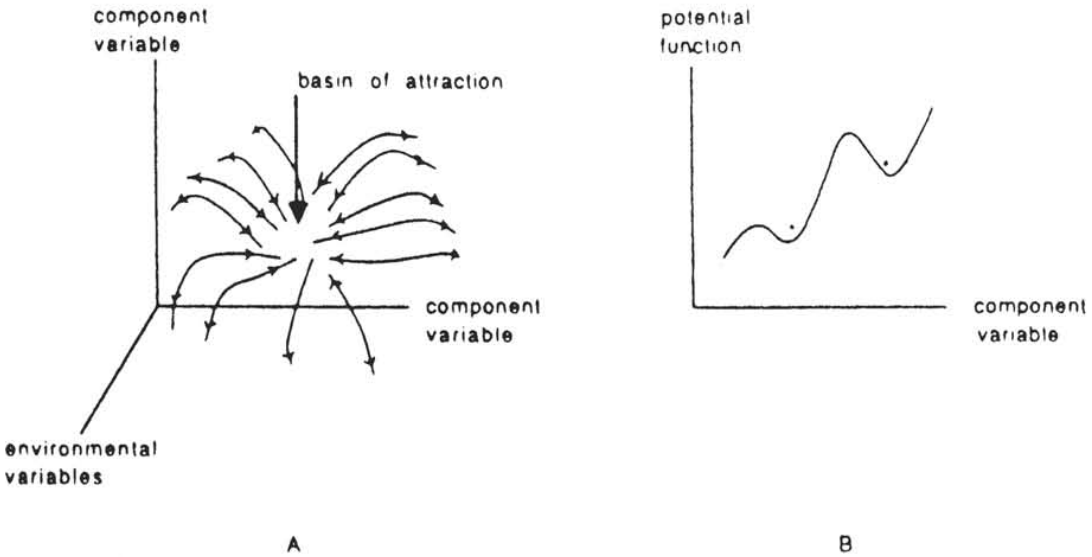

Fig. 2. Phenotypic spaces. (A) schematically represents a basin of attraction in a phase space picture. (B) illustrates a potential function picture. In principle the potential function picture can be associated with the phase space picture, though in most cases the association would be purely conceptual and notational. The two pictures do not correspond in this case, first because only one basin of attraction is illustrated in (A) and second because only one variable axis is represented in (B).

states, or basins of attraction, but under any circumstances it must have dynamics that are stable at some level of description if it is to survive. This does not preclude certain types of phenotypic instability from making an important contribution to fitness, in particular to adaptability (Conrad, 1983). Transitions between multiple steady states or chaotic (initial condition sensitive) dynamics are examples that mix aspects of instability and stability. The point is that the dynamics of the organism must be coherent and therefore stable so far as its overall function is concerned.

Such stable dynamics can sometimes be described in terms of valleys of a potential surface (Fig. 2B). In general, it is not possible to construct a bona fide potential surface; nevertheless it remains useful to think in terms of a notational potential function. Basins of attraction correspond to valleys of the potential function, and we can think of peak climbing on the adaptive surface as equivalent to falling into some valley on the potential surface.

We can finally imagine a grand phenotype space with points (or flows) corresponding to all the phenotypes that correspond to possible genotypes in the fitness space. Each of these flows can then be coordinated to points on the fitness surface (or to more than one point if two genotypes code for equivalent phenotypes). Basins of attraction in the grand phenotype space correspond to peaks and other highlands in the fitness space. Deep valleys and gorges in the fitness space correspond to flows in the grand phenotype space that have no basin of attraction that corresponds to the living state. On some points along the environment axis the basins of attraction might become viable; but this would correspond to the valleys becoming highland at the corresponding point on the environment axes in the fitness space. We shall call non-viable flows unstable since their stable points, if any, fall into non-viable basins of attraction. The only points of the flow that could have acceptable fitness are unstable.

Finally we note that the dimensionality of the genotype space does not uniquely determine the dimensionality of the phenotype space. This depends on the 'rules' of development. A low dimensional genome could in principle specify a higher dimensional phenotype space than a higher dimensional genome. This depends in part on the redundancy of the genetic description and in part on the way the genetic description is used.

5. The evolvability criterion

Some fitness-scapes are better suited for evolution than others. If the terrain consists of isolated peaks that are separated by wide, deep valleys the chance of making the transition from one peak to another is very small; in fact, the chance of any peak becoming populated is small. If peaks are connected to each other by pathways which allow for continuous ascent, making the transitions will be easy. Strictly speaking these connected peaks are not peaks at all. They are peaks in most dimensions, but in one or a few dimensions there is an upward-running pathway connecting them. Evolution could also occur if bona fide peaks are connected by shallow valleys that are not extremely wide.

To make this condition somewhat more precise, let us suppose that a particular set of mutations must occur in order to make the transition. If these must occur simultaneously in order to avoid falling into an unacceptably deep valley, the evolution time scales as

$$T \sim 1/Ap^n$$

where A is the population size, p is the mutation probability, and n is the number of mutations. If the mutations can occur in a series of steps the evolution time becomes

$$T' < n/A'p$$

where A' is the smallest size reached by any population in the series. If $p = 10^{-10}$ (one mutation in 10^{10} base pairings) the pathway that requires a double mutation reduces the

rate of evolution by a factor of 10^{10} over the rate possible on the single step pathway. Clearly evolution is feasible in historical time only if peaks are connected by stepwise traversable pathways in at least one of the dimensions of fitness space.

The above criterion is not fully precise. We have ignored the fact that in a long genome, multiple mutations are likely to occur, apart from the desired ones. We have to add a factor of $(1 - p)^{m-n}$ to the denominator of both of the above expressions, where m is the number of bases in the genome. We should also add a factor to represent the time required to occupy the neighboring peak, which might become substantial if the peaks have similar fitness values. A precise calculation (Conrad, 1972b, 1978, 1983) shows that incorporation of both of these factors (including a wide variety of assumptions about the rate at which populations grow after mutation) hardly alters the picture. One step traversability of the landscape is still important. If the pathways between peaks are so particular that they must occur through only one sequence of mutations, then the genome cannot be too long, say much longer than 10^{10} bases. But in general if we are in a region of the landscape with few deep valleys and gorges many different mutations will be acceptable. As a consequence it is quite likely that the occurrence of two, three, four, or more simultaneous mutations will be acceptable. This does not mean that the landscape is so structured that peaks are separated by multiple mutational steps. It means that the peaks are so dense that multiple mutational events are acceptable, and in many cases are 'don't care' events.

The extreme case of such don't care events are scratch space genes, i.e. non-coding genes. These can drift in any direction without any effect on fitness. But the chance that a desired simultaneous mutation will be achieved in steps in this way is still small since it is likely that the remainder of the gene will drift in really undesirable directions.

The same considerations apply even if the mutations are not point mutations. Any type of genetic event might be involved – frameshift mutation, duplication, deletion, crossover, recombination. The hierarchical structure of the genome into genes and chromosomes allows for single step manipulations of blocks of nucleotides that have substantial functional independence. The necessary condition for evolutionary change to occur is that at least one such genetic event must be acceptable on its own. If evolutionary change requires two or more to occur simultaneously the evolution process will stagnate.

6. Genetic instability versus phenotypic stability

Now let us consider what is required for the genotype to sit on an adaptive peak, or at least to have an acceptable fitness level. The reasonable supposition is that the associated flow in phenotype space must be stable. Any condition necessary for stability in phenotype space should thus be a necessary condition for acceptable fitness in genotype space. The problem is that unqualified stability in phenotype space means isolated peaks in the fitness landscape, which conflicts with the evolvability criterion.

The situation is illustrated in Fig. 3. A and B represent basins of attraction in the phenotype space. These are stable in the sense that the organism will asymptotically return to either the A or B trajectory in response to a shift in any of the variables, provided the shift is not too large. Basins A and B would correspond to valleys of a potential surface. The corresponding peaks in the adaptive landscape (the fitness space) are also denoted by A and B. If the basins in the phenotype space are robustly stable to perturbation the corresponding peaks on the adaptive landscape are likely to be separated by wide deep valleys.

Now suppose that the dimensionality of the adaptive landscape is increased and that the dimensionality of the phenotypic space is also

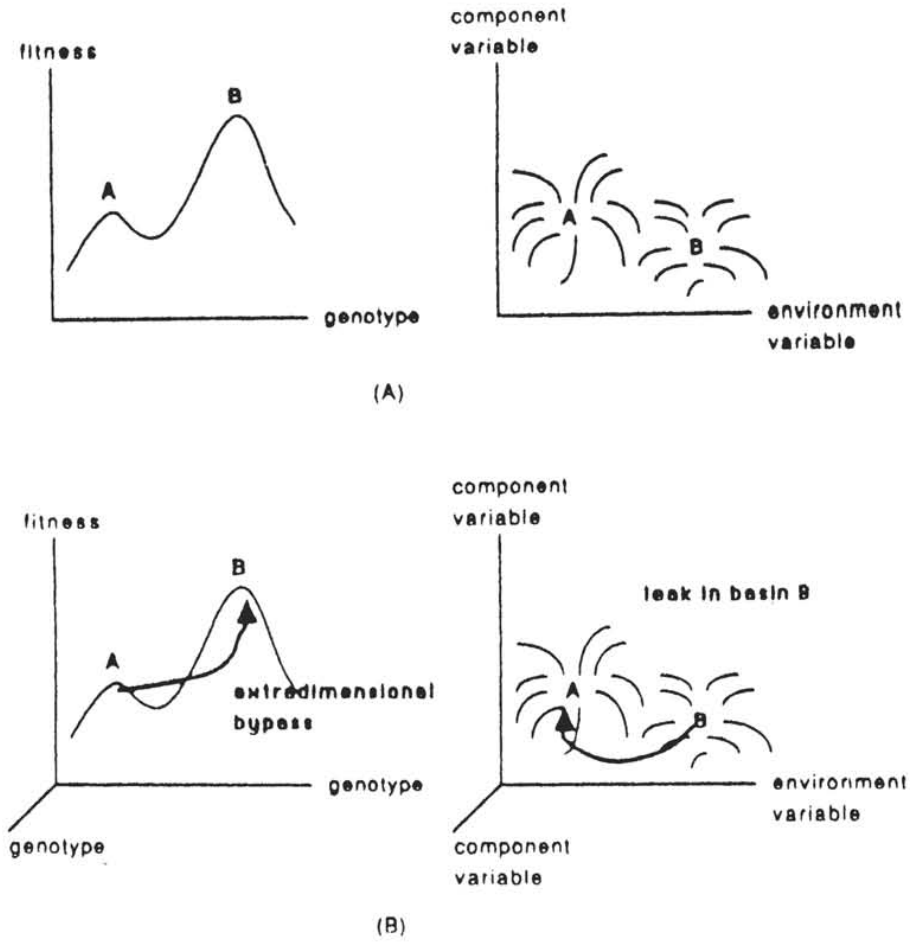

Fig. 3. Extradimensional bypass. The adaptive landscape fitness peaks A and B correspond to the basins of attraction A and B in phenotypic space. In the upper figures (A) the two peaks are separated, corresponding to the fact that the two basins are separated. In the lower figures (B) the addition of an extra component destabilizes basin B, allowing a transition to A. The appearance of the pathway from basin B to A corresponds to the appearance of an upward running extradimensional bypass from peak A to peak B in the adaptive landscape space. The two spaces do not precisely correspond since the environment axes are omitted from the adaptive landscapes, due to the obvious impossibility of representing more than three dimensions. The genotype axes would represent bundles of nucleotide axes and the addition of an extra genotype axis is used to indicate the increase in dimensionality of the space.

increased. This increase might be a direct consequence of an increase in the size of the genome, or it might be a consequence of the way in which the genes act. All else equal, the chance that A and B remain isolated basins decreases. The intuition is that the chance of a valley occurring in a high dimensional space (e.g. on a high dimensional potential surface) is less than in a low dimensional space. As a consequence a leak will develop in basin B and any organism located there will fall toward basin A. Actually, the two basins will merge and the organism will fall toward the point in A corresponding to the lowest potential. This merging of the basins will correspond to the appearance of an upward running pathway between peak A and peak B in the adaptive space. We will call this upward running pathway an extradimensional bypass, since the peaks remain isolated except in one of the newly added dimensions.

The addition of the extra dimensional

bypass makes the organism sitting on peak A amenable to evolution — it can get to the better peak B. But unfortunately it reduces the fitness of B, since our assumption is that fitness requires stability in the phenotypic space.

Clearly we have arrived at a contradiction. A genetic-developmental organization must be slightly unstable to allow for evolution, but this is incompatible with the stability required for fitness. The resolution is not too difficult. It is only necessary to organize the phenotypic dynamics to be unstable to mutation and other genetic perturbation, but stable to the physiological class of perturbations.

7. Stability versus complexity

Let us look more carefully at how this conflict relates to the dimensionality of the space, or more generally, to the complexity of the system. For a system to be stable it must occupy a valley in all dimensions. As soon as this condition fails in one of the dimensions, the system will slide downward, either to another valley (usually in a lower dimensional space) or to extinction. Whenever an extra component is added to the system the dimensionality of a space increases. Thus systems with more components are less likely to be stable. Systems with fewer or weaker interactions among the components are also less likely to sit in a valley in all dimensions, hence less likely to be stable as well.

These relations between stability and complexity, first recognized by Gardner and Ashby (1970), have been analyzed in a particularly incisive way by May (1972, 1973). The systems considered by May are model ecosystems of the Lotka-Volterra type. Let C represent the probability that any pair of components in the system will interact, s the (common) average interaction strength, and m the number of components. For the present purposes May's main result is that in models in which the interaction structure is selected at random the probability of stability goes to zero as

$$s(mC)^{1/2} > 1$$

with the transition to instability with increase in s being extremely sharp for large m.

This result implies that the chance of a basin occurring in phenotype space decreases as the number of components in the organism increases, as the number of interactions among these components increases, and as the interactions become stronger. In grand phenotype space the lower dimensional basins are more likely to be the stable ones, hence more likely to meet a necessary condition for fitness. The addition of gratuitous complexity would appear quite likely to destabilize an organization and therefore to render it completely unfit. The species will either slide to extinction or self-simplify by sliding to a new basin with lower dimensionality and fewer interactions. Such self-simplification, if it occurs, would eliminate smooth pathways between the regions of the grand phenotype space corresponding to different phenotypes, where smooth means that one phenotype could slide into another in response to a small perturbation.

Lotka-Volterra systems are one instance of a broad class of dynamical models to which May's results apply (see also Hastings, 1982). We might also note that Wigner (1961) presented a somewhat similar argument for the unlikelihood of reproduction in quantum mechanical systems. Essentially as the number of components and interactions becomes large it requires very special constraints to maintain stability. May emphasizes that the result does not apply to all conceivable dynamical models or to all regions of parameter space. Certainly the range of possible phenotypic dynamics extends beyond the scope of the theorem. But in this case we can view the theorem as suggesting what sorts of special constraints might be at work if a system is in fact both complex and stable.

8. Stability versus evolvability

Let us briefly recapitulate the argument. In order to move from one fit position on the landscape to another there must be a 'smooth'

connecting path. This means a path that can be traversed in single steps, though in general the dimensionality of the path is such that each step may be accompanied by a number of side steps in different dimensions. The main point is that peaks have to be densely packed or connected in at least one dimension by smoothly climbable pathways. The reasonable intuition is that as the number of dimensions increases, the chance that there will be a connecting pathway also increases. To formalize this intuition we associate the genotype space with a phenotype space. This introduces dynamics, and we know from May's analysis that increase in dimensionality and interconnectedness of the phenotype does in fact increase the likelihood that it will be destabilized by perturbations of the genotype. Such instability is a permissive (necessary, not sufficient) condition for evolution, but a prohibitive condition for fitness. The likelihood that peaks will be connected in such a way that evolution can occur thus increases with the number of components and interactions in the phenotypic dynamics; but the likelihood that peaks will occur at all decreases.

It might be counterargued, as observed earlier, that the dynamics in phenotypic space need only be organized so that they are slightly unstable relative to genetic change (or perturbation) but stable relative to physiological perturbations. The problem is that one and the same phenotype must satisfy both pressures. The former drives it to complexity while the latter drives it to simplicity. But one and the same system cannot be both complex and simple.

9. Three principles of organization

Let us now consider what special dynamical features enable organisms to satisfy these conflicting drives. Three features are pertinent.

9.1. Compartmentalization

This in general increases the chance of stability (May, 1973). Here compartments are defined as blocks of components that interact mostly among themselves in terms of number or strength of interactions. The effect is to reduce the ramification of perturbation (Conrad, 1983). This allows for some channeling of the effect of mutation and other genetic change on specific aspects of phenotypic dynamics. It also serves to block off the effects of physiological disturbance. As a consequence it increases physiological stability and at the same time keeps the effect of genetic instability within bounds.

9.2. Component redundancy

This means the presence of components that are essentially functionally equivalent. If some of the components are removed the system as a whole will nevertheless not be noticeably altered. Redundancy is a form of compartmentalization since similar components are grouped by virtue of their interactions with other groups of components, possibly similar within the groups as well. The difference is that in this case it is not the interactions among the components in the group (or block) which are the main point; and in fact such interactions may be less important than with components in a second group. As with compartmentalization proper, however, the effect is to prevent the ramification of perturbation from one block to another. As a consequence redundancy increases stability in a way which is absolutely contrary to the Gardner-Ashby-May theorem. The probability of stability increases as the number of components and interconnections increases. The reason is that the interconnections are not random (May, 1973; see also Conrad, 1972a, 1983).

Redundancy serves to buffer the effect of mutation and other genetic perturbation on phenotypic dynamics. The effect of a genetic perturbation can be distributed over a larger number of elements, hence expressed as a more gradual and graceful alteration of phenotypic features critical for function. Compartmentalization in the strict sense also

serves this function. Such gradualism means that adaptive peaks that are metrically close in the fitness space correspond to basins of attraction that are metrically close in the phenotypic space, and that a high density of peaks corresponds to a high density of basins. The reason is that similar structures and dynamical behaviors are likely to have similar fitness, therefore to afford more traversible pathways in the fitness space. Radical changes in the phenotypic dynamics in response to genetic events with a significant likelihood of occurrence could conceivably lead to occupation of new peaks or major increases in fitness; but in general major changes in a complex system are completely non-viable.

Redundancy and the associated multiple steady states in phenotypic space afford stability to the physiological class of perturbations as well as buffering the genetic class. So it is precisely the type of organization needed to reconcile evolvabilty with fitness.

9.3. Multiple weak interactions

Multiple weak interactions are a form of redundancy. If one connection is broken the operation of the system is still supported by other connections. The advantage of weak interactions is that they allow for gradual transformation of function. If one or a few of the weak interactions is broken off, the relationship between two parts can be altered gradually. If the interactions are all strong the breaking either leads to a major change in the relationship or, if the interactions are redundant in the strictest sense, leave it exactly the same. Redundancy thus plays a key role in mutation buffering.

From the standpoint of May's theorem weak interactions provide the best way of compromising genetic instability and phenotypic stability. When m and C become large in the expression $s(mC)^{1/2}$ the only way to maintain a good chance of stability in a randomly constructed organization is to make s small. If the component redundancy is organized in such a way that stability increases with the number of components, this is not necessarily true. But it is still true that small s is necessary for gradual transformation of function. From the standpoint of buffering the effect of mutation on the phenotypes it is best for C and m to be large and for s to be small.

Many other factors bear on biological organization. We mention two that are closely connected to the principles mentioned above.

The first is that evolutionary flexibility should be a maximum when the number of possible variations on the organism is greatest (Conrad and Hastings, 1985). The variations must either be due to alterations in the initial state of the system or alterations in its interaction structure (the connections of its components). According to the binomial theorem the number of possible variations on the interaction structure of m components, N, is greatest when one half of the connections are turned on ($N = m^2/2$ and on the average $C = 1/2$). Substituting into $s(mC)^{1/2}$ yields $s(m/2)^{1/2}$. Since m should be large if N is to be large, this again suggests the great importance of small s. Biological organizations that are well suited for evolution should have multiple weak interactions, tending to fifty percent interactivity. If the chance of stability is to be high the strength of the interactions should be small. If many of the components and interactions contribute to a stabilizing redundancy structure, the interaction strength should still be small since this is most effective from the standpoint of buffering.

The second factor has to do with functional efficiency. If we simulate a physical system we have to consider every interaction that significantly contributes to its behavior. According to current force laws this could mean up to m^2 interactions in a system comprising m particles. The behavior may or may not have significance from the point of view of performing a function. In the case of an organism, this means it may or may not have significance from the point of view of contributing to its survival and reproduction. Suppose that the number of interactions

that actually contribute to the performance of the function is h. We can then define efficiency as h/m^2. As $h \to m^2$ the computational difficulty of simulating the function with a conventional computer increases.

Efficiency is equal to 1 when $h = m^2$. This is the best efficiency possible. If we want high evolutionary flexibility h should equal $m^2/2$, giving an efficiency of 1/2. From the standpoint of achieving high efficiency and high evolutionary flexibility it is best for an organism to operate on the principle of multiple weak interactions.

10. Non-programmability and homomorphism

Multiple weak interactions have a fundamental implication for developmental biology. We can picture the processes of development as described by some sort of mathematical map. The initial state of the organism and its environment is mapped into the final state of the organism and environment. Naturally we focus our attention on the organism, and treat the environment as a set of externally driven parameters, and as an energy source and heat bath. The term 'developmental program' is sometimes used to refer to this map. A computer program is a map. But for development, the usage is somewhat metaphorical, since programs in the strict sense are a highly restricted type of map. They can be thought of as rules (or tables) that specify the next state of a system given its present state and input. The symbols are distinct and finite in number, and the transformations from state to state occur on a discrete time scale. The sequence of transformations may be called the execution sequence, and so far as is known any physically realizable map can be simulated (in the sense of computing the input/output behavior) by writing the appropriate tables. This commonly held tenet of today's computer science is sometimes called the strong form of the Turing-Church thesis. For the present purposes, however, the truth or falsity of this thesis makes no difference. We have raised the point only to underline the metaphorical status of the 'developmental program' concept.

The salient feature of machines that run computer programs is that they are programmable. This is possible because each component in a digital computer is built up from components whose function is definitively specified by a manageably small user's manual. As a consequence we can always map a computer program, expressed as a table, into the physical structure of a machine, that is, into the states of its components and their connectivity. We can of course always extend the table, by adding a new state, without in any way altering the previously specified transformation. In the case of the actual machine this means that we can always add components and connections without in any way altering the user's manual specification of the components already in place. For this to be possible the number of interactions that affect the behavior of a component cannot increase as the size of the system increases. The engineer must organize matters so that the maximum number of interactions in which a component can participate scales as a constant, independent of the size of the system. Systems that have this property will be called *structurally programmable*. All digital computers are programmable in the structural sense; they may also be programmable at the interpretive level if they are initially wired up to realize a program that can read and follow any other program.

Now recall that for best efficiency and evolvability about half the interactions should be turned on in an organism ($N = m^2/2$). Thus the optimum number of interactions per component is $N/m = m/2$. The number of interactions that define the behavior of a component increases as the number of components in the organism increases. Thus, an organism cannot be structurally programmable.

What this means, first of all, is that even if we use the metaphorical idea of a developmental program, we must be very careful not

to carry along with this the computer science concept of programmability. Digital computers execute programs, and they are programmable. We can think of an organism's development being generated by a metaphorical program; but the organism is not programmable, and its program of development is not written into the states of its components and their pattern of connection according to a manageably small user's manual. If it could be so written, each component would have a functional specification that is independent of the number of components in the organism. But in the regime of multiple weak interactions this is not the case. Ultimately the only user's manual would be the equations of physics, and these in effect have an infinite number of entries (since they involve continuous dynamics).

What about the DNA base sequence? This also cannot be a computer program in any way reminiscent of a digital computer program. Even if we adhere to the metaphor of a developmental program we cannot identify this with the genetic description of the organism. The genes are better viewed as knobs which are used to alter the developmental program. This involves continuous dynamical processes embedded in a sequence of bifurcations from one developmental stage to another. Some of the genotypic knobs control dynamic parameters in each stage (e.g. rate constants of various biochemical and cellular processes), and others control the space-time order of the bifurcations (the bifurcation parameters). Environmental factors also exert an influence, either on the expression of the genes or on their action.

Why is this good for evolution? Computer programs are notoriously fragile. A single change in the code is rarely acceptable. Changes in parameters of the program might be acceptable. But blind changes in the syntax affect the structure of the execution sequence and are hardly ever acceptable. As a consequence the adaptive surface of a computer program is extremely rugged, with peaks separated by deep, wide gorges. Evolution of computer programs by variation and selection is unworkable. The same would be true of biological evolution if DNA were like a digital computer program. Most changes would alter the structure of the execution sequence and would lead to teratologic behavior. But in fact changes in DNA are much more like changes in the parameters of a dynamic process. Some of the changes modify the structure of the developmental execution sequence, that is, of the space-time pattern of turning genes on and off, and of turning various biochemical and biomechanical processes on and off. But the vast majority of the changes leave the overall order of development essentially the same and modify the emerging form by topologically distorting the dynamics. As a consequence mutation buffering is possible. If the distortability of the dynamics is not sufficiently gradual it is always possible to add redundancies that increase the gradualism. The adaptive landscape becomes increasingly well suited for evolution.

The whole process can be viewed in terms of homomorphic realizations of the developmental map (Conrad, 1970). The developmental program of a mouse and an elephant may be essentially similar; but by modifying some structural and regulatory genes we may obtain a different realization of this program. The two realizations are homomorphic in the same sense that two houses built with different sized bricks from the same blueprint are homomorphic. The two systems may be 'coded' into one another by 'relabeling' them with different components.

The process can also be viewed in terms of structural stability. A structurally stable system is one whose qualitative features are invariant to perturbation (Thom, 1970). This is precisely what genetic buffering achieves. Different regimes that are equivalent up to structural stability are in fact topologically homomorphic (strictly speaking diffeomorphic).

This is not to say that all evolution can be reduced to gradual change in given forms, or structures. We can presume that a number of basic forms are possible, corresponding to the basic morphological plans in living systems. The evolution of these basic forms may have involved syntactic changes — basic changes in the space-time order of the developmental program. But once these basic structures appeared, the radiative epoch of evolution began. The main process here is the generation of homomorphic images of the basic forms. These homomorphic realizations are suited to different environmental conditions, and hence occupy different high points on the adaptive landscape. For the complex life plans, such as the vertebrates, the transformation of realizations probably proceeds in small increments; in simpler forms, such as the plants, larger increments may occur in a single step and at the same time have a good chance of being viable.

We can finally note that this distinction between significant and gradual change may be related to Gel'fand and Tsetlin's (1962) idea of a well-organized function. Structures described by such a function admit manipulation through two groups of parameters, those that led to a significant change in performance evaluation and those that led to relatively small changes. The idea was developed in the context of control problems, in particular animal locomotion, but it is evident that if 'evaluation' is replaced by 'fitness' and 'small change' is associated with gradualism, the type of structures that would have good controllability would also have good evolvability. The origin of basic forms, of phyla, would involve changes in the gross control parameters, while the adaptive radiation of these forms would involve the generation of alternative realizations (or homomorphic images) via the fine control parameters.

11. Genetic buffering mechanisms

Elsewhere the author has described how organizational structures suitable for effective evolution manifest themselves at different levels of biological organization (Conrad, 1983). Here we shall briefly review examples of genetic buffering mechanisms at the level of structural genes, at the level of genetic regulatory mechanisms, and at the level of polygenic organization.

11.1. Structural genes

These are the genes that code for structural and enzymatic proteins. The situation is schematically illustrated in Fig. 4. The sequence of bases in DNA is translated to a sequence of amino acids. This in turn folds on the basis of weak interactions among many of the amino acids. Prominent examples include van der Waal's interactions, hydrogen bonds, coordination bonds, hydrophobic interactions, and disulphide bonds. Some features of the three-dimensional shape are particularly important for function — recognition sites, binding sites, control sites. Other features of the shape are quite unimportant from the functional point of view. The essence of buffering is to absorb part of the effect of mutation and other genetic variation in these noncritical regions and thereby to express them in terms of more graded changes in the critical regions. Graded here means that the amount of change in any step is in many cases small. This does not mean that big changes do not occur and that such changes may not in some instances be significant. But if a system is to be tuned to perform a function it is advantageous if the tuning can proceed in small steps. Furthermore, since shape features are not scalar properties, gradualism allows for exploring a large variety of neighboring possibilities.

The following mechanisms buffer the effect of gene changes on the three-dimensional shape and function of proteins (Conrad, 1979b, 1983).

(1) *Redundancy of weak bonding.* More weak bonds are like more springs in a mat-

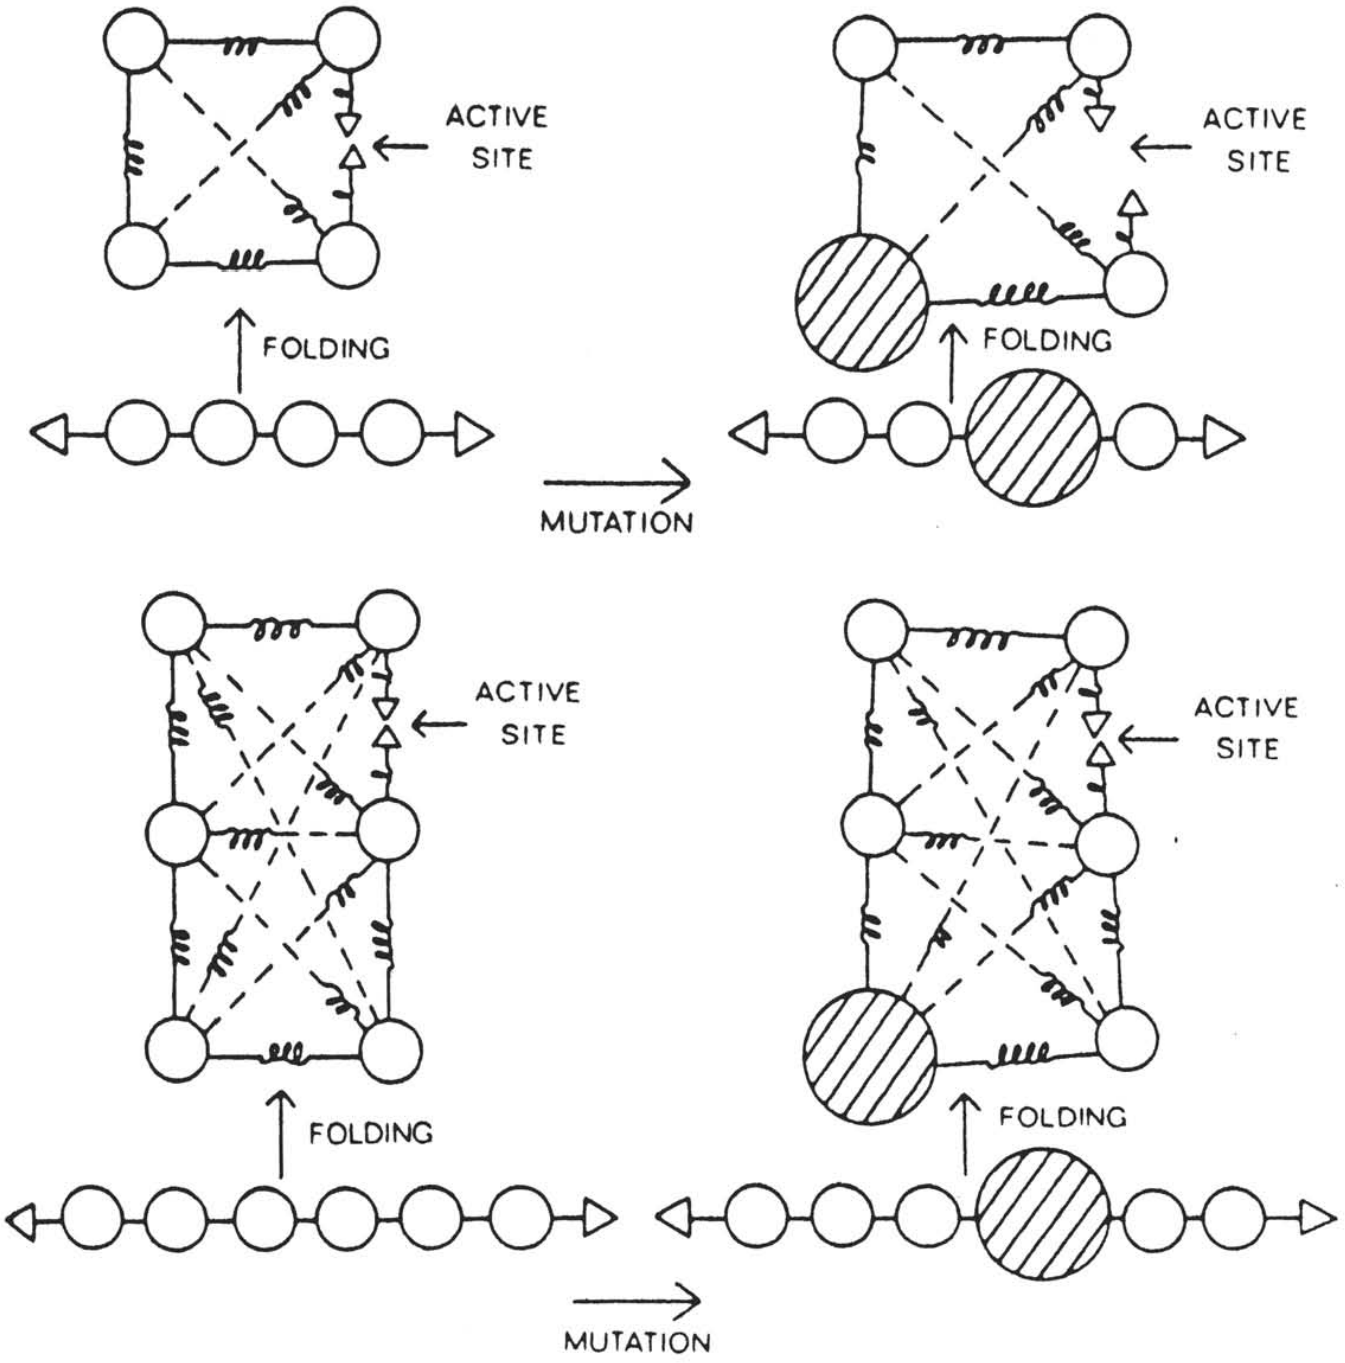

Fig. 4. Mutation buffering model. Balls and triangles represent amino acids. The solid springs represent strong (covalent) bonds, while the dashed springs represent weak bonds responsible for protein folding. Increasing the number of amino acids and the number of weak interactions is somewhat like making a spring mattress with more springs: it serves to absorb the effects of mutation (represented by a switch from a small to a large ball) on features of the shape critical for function (represented by the distance between triangles at the active site). The increase in the number of amino acids corresponds to an increase in the dimensionality of the adaptive landscape (see Fig. 3), and the buffering effect of this increase corresponds to a traversable extra dimensional bypass from one fit form to a form of higher fitness (adapted from Conrad, 1979b, 1983).

tress. The features of the mattress crucial for comfort are altered less in response to one spring breaking as the number of springs increases.

(2) *Redundant amino acids.* The incorporation of extra amino acids allows for greater redundancy of weak bonding and is in general equivalent to the addition of more capacity for absorbing the effects of mutation. Redundancies that preserve the sequence of steps in folding are particularly important.

(3) *Redundancy of amino acid types.* This means the use of amino acids with close structural analogs. If mutation to a close structural analog is possible gradual variation of the protein shape is certainly possible as well. Hydrophobicity is one important factor in determining the replaceability of amino acids (Volkenstein, 1979).

(4) *Specific organizational formats.* Some proteins have special organizational features that amplify the number of possible variations on structure. The structure of the immunoglobin molecule (with its highly modifiable claws) is an example.

The protein could also have organizational features that amplify the response to mutation. Buffering mechanisms can be thought of as analogous to a fine control on a microscope. A coarse control makes it possible to find a general region of operation quickly, but a fine control is necessary for thoroughly exploring the region. It is possible to construct regions of an amino acid sequence (using proline for example) that provide a coarse control capability, and other regions that use redundancy to provide a fine control capability. A combination of coarse and fine control is optimal from the standpoint of evolvability.

11.2. DNA structure and genetic regulation

The pertinent fact here is that DNA is not an ideal double helix (Sasisekharan and Pattabiraman, 1978). Handedness and conformation within a given handedness is influenced by base sequence and is milieu dependent. A plausible assumption is that the variability of DNA conformation helps facilitate the action of transcription and translation enzymes, and that it plays an important role in the regulation of gene expression.

Two types of buffering are possible (Conrad, 1985; Conrad et al., 1986). The first is readout buffering. In this case redundant DNA isolates the conformation of coding regions from conformational strains in regulatory regions. Introns and some other redundant DNA may serve as readout buffers. The second is evolutionary buffering. Here redundant DNA can enhance the 'tunability' of conformation-dependent regulation of gene expression. By absorbing conformational strain the number of conformational gradations that can be achieved increases. Repetitive DNA may serve as an evolutionary buffer. Readout buffers would also serve as evolutionary buffers.

DNA buffering, since it is based on redundancies in the base sequence, entails an increase in the dimensionality of the adaptive landscape. With such buffering, sequence-dependent changes in DNA conformation are capable of playing as important a role in the diversification of metazoan organizations as amino acid sequence changes played in the diversification of proteins and microorganisms.

11.3. Multigenic organization

Many mechanisms contribute here (see Conrad, 1983). One prominent example is polygenic inheritance. Genes control the rates of reactions by determining the properties of enzymes, and one mode of evolution involves transformation of enzyme properties. A second mode involves changes in the concentration of enzymes. This could be achieved through regulatory mechanisms or through altering the number of genes that code for that enzyme.

The hierarchical organization of the genome into genes, blocks of co-acting genes, and chromosomes allows for higher level genetic operations, such as cross-over and recombination. The enormous context dependence in computer programs makes cross-over infeasible as a mechanism of evolution. In general it is not possible to evolve two segments of code independently and then combine them. The exception is in production systems, or rule-based programming. The rules (but not those of the inference engine) can be added and deleted separately. To some extent genetic organization has this independence property. Many genes may affect one trait, and one trait may affect many genes. But to the extent that genes act in parallel rather than independently they can often evolve independently and then be combined to yield specific useful characteristics. This type of organization, involving components with a high degree of mutual independence, may not be as efficient from the energetic point of view as a highly integrated system, but it is more suitable to self-organization through the Darwinian mechanism.

Other evolution facilitating redundancies that could be cited are in the hormonal system and in the immune system. The occurrence of first and second messengers is a redundancy mechanism that allows for relatively independent evolution of multicellular controls and intracellular responses. Special compartmental organizations also allow for a high degree of evolutionary and developmental flexibility. Add on growth mechanisms are particularly important here, such as the open growth system of plants and segmented organization in animals. In complex organisms, such as vertebrates, segmented organization is combined with allometric control. The multigenic organization is such that a variety of traits are influenced in a coherent way, probably through the response to hormonal controls. The organization appears to involve a combination whose action has enough independence to allow for effective evolutionary search, but whose integrated effects are coherent enough to produce a wide variety of homomorphic images. The essence of the organization is redundancy-buffered instability to genetic perturbations and redundancy-enhanced stability to phenotypic perturbations (in particular, structural stability of development).

12. Evolution of evolvability

Let us now recall why the concept of evolvability is controversial. Some evolutionists argue that natural selection can act only on properties that are advantageous to the individual (e.g. Williams, 1986). Evolvability is advantageous to the species. Do not, therefore, let the concept of evolvability mix into biological thinking.

This dictum is wrong on two counts. The first is that some mutation buffering redundancies are in fact advantageous to the individual organism. Some of the redundancies that confer stability on the phenotypic dynamics also serve to buffer the effect of genetic change. Readout buffering redundancies in DNA are a direct advantage to the organism and provide evolutionary buffering at the same time. The immune system provides an example of a different kind. This is to some extent an internalized evolutionary system (Jerne, 1955) and as a consequence buffering mechanisms that facilitate the ontogenetic evolution of immunoglobin molecules are an advantage to the individual organism.

The second count is that mutation buffering and other evolution facilitating mechanisms can accumulate even if they are a tax from the standpoint of the individual organism. Mutation buffering redundancies can always be added or deleted in a gradual way, and as a consequence every population will exhibit some variation in this respect. Evolutionary advances are more likely to emanate from this portion of the population. When they occur, the evolution facilitating redundancies will hitchhike along with the advantageous traits whose appearance they facilitate.

Hitchhiking has been proposed as a process

influencing the evolution of mutation rates (Strobeck et al., 1976). The notion of parasitic DNA also subsumes a hitchhiking effect, except that there is no causative relation between the parasitic and advantageous component of the DNA. According to the mutation buffering concept, however, so-called parasitic DNA may in many cases be facilitating in this respect. If it could evolve in the absence of such facilitation it could certainly evolve, and in an accelerated way, in its presence.

The point may be illustrated with a simple model. Recall our formula for the evolution time, $T \sim 1/Ap^n$. Suppose that a population is isolated atop an adaptive peak and that to move to some other adaptive peak two simultaneous mutations are necessary. It might be that it is isolated from all other peaks by a two mutation gap; or that it is isolated from some peaks in this way but not from others. There are two possibilities for jumping the gap. The first is to wait for a double mutation. The waiting time will be

$$T \sim 1/Ap^2$$

The second is to proceed through m evolution enhancing mutations and then to make v fitness enhancing mutations to the new peak (or rather to the corresponding peak in the higher dimensional space). The waiting time is now

$$T' < (v + m)/A'p$$

where we recall that A' is the smallest size reached by any population in the series. The relative slowdown due to the requirement for double mutation is at best

$$\frac{T}{T'} \approx \frac{A'}{(v + m)Ap}$$

Even if v and m are large and A' small as compared to A this ratio will be dominated by p. The hitchhiking mode of evolution, which involves forging an extradimensional pathway between the two peaks, is approximately 10^{10} faster than is the double mutation mode (assuming a mutation probability of $p = 10^{-10}$).

Our concept of hitchhiking and buffering can also be extended to the ecosystem level of organization (Conrad, 1988). The term perturbation buffering is more appropriate here. Ecological systems pass through a series of developmental (or successional) stages. Each stage corresponds to a basin of attraction, but in general with some leak in it that causes the system to move to the next stage. The final, or climax stage, is the most stable. Suppose that the community reaches basin K. Basin K' is deeper. But there is no easily traversable pathway leading from K to K'. Many changes in the structure of the community would have to take place simultaneously. The system will either stay in K or run off to a third basin that is no more stable than K. In general the community will be broken up into many subcommunities, and some of these will have more redundancy than others in terms of numbers of species and in terms of energy and communication channels among these species. These redundancies increase the chance that there will be an extra dimensional pathway from K to K'. The redundancy rich subcommunity will then flow into K' and will eventually displace the subcommunities remaining in K. In this way the succession facilitating redundancies will hitchhike along with the movement to the next stage of succession that they facilitate. As a consequence ecological communities will in many cases develop in the direction of a large number of somewhat functionally redundant species, with many weak energetic and informational interactions among these species (Conrad, 1983).

13. The principle of self-complication

Our arguments about the evolution of evolvability, at genetic, organismic and ecological levels of organization, can be summed up in a *principle of self-complication*. At the

level of the gene and organism the principle may be stated thus: the complexity of biological organization increases because (buffered) dynamic instability in response to genetic variations is advantageous to evolutionary self-organization. At the ecological level, buffered instability to perturbations emanating either from the environment or from other organisms in the community are advantages from the standpoint of successional stabilization.

The principle of self-complication contrasts with what has been termed the principle of *self-simplification*. Some authors (e.g. Levins, 1970; May, 1973) have argued that complex systems, because they are unstable, will self-simplify. This is a reasonable assumption, except for those special cases in which the structure of complexity confers extra stability. Our analysis suggests that complication in terms of redundant components and weak interactions will in general facilitate the achievement of stability and that biological organization is a consequence of self-complicating as well as self-simplifying processes.

Actually our whole discussion has used a rather naive definition of complexity. Many other definitions exist. According to the algorithmic definition of Chaiten and Kolmogoroff, the complexity of a pattern can be represented by the length of the shortest computer program that can generate the pattern (see Chaiten, 1977). A truly random (not pseudorandom) pattern is thus the most complex. Redundancy means that some of the features of the pattern are related to each other by a rule. Thus our principle of self-complication has a self-simplifying aspect when looked at from the point of view of the Chaiten-Kolmogoroff definition. The Chaiten-Kolmogoroff complexity of an evolutionary system would increase less in the course of evolution than would the complexity as measured by the number of components and interactions. Evolutionary systems would move toward some situation intermediate between order and randomness. From the point of view of constructing scientific theories this is of course the most complex (difficult) region. Pure randomness, no matter how complex from the standpoint of Chaiten-Kolmogoroff, lends itself to probabilistic models; while highly ordered situations lend themselves to group theory. The organizations that are best suited to evolution are precisely those that are most ill suited to the classical standards of scientific description.

14. The adaptive landscape reconsidered

Why does evolution work? The reason is *not* to be found solely in the magic optimizing power of variation and selection. It is as much due to the organizational structure that undergoes the variation. Evolution works because this organization is amenable to evolution, and because this amenability itself increases in the course of evolution.

In terms of the adaptive landscape the picture is that Darwinian systems move to regions of the fitness surface that are increasingly well suited to hill climbing through the mechanisms of variation and selection. This in general means moving to high dimensional regions. The chance of extradimensional pathways increases as the redundancy of components and weak interactions increases.

Why should there be a peak structure at all? We can imagine some basic structures — the ontogenetic plans corresponding to the different phyla, for example. All the different forms that appear in the course of evolution could well be homomorphic realizations of these basic forms. From the purely developmental point of view the adaptive landscape is essentially a small collection of easily climbable peaks, with populations occupying different locations on these peaks. The problem with this picture is that it ignores ecology, and the selective forces immanent in the interactions among different species. As populations climb the peaks they modify its structure, turning some locations into valleys and others into peaks. From the genetic and developmental point of view these valleys correspond to well formed organisms; from

the ecological point of view they are excluded. In some cases a continuum of forms may be allowed, yet distinct forms may become separated by deep wide valleys. In order to understand the evolution it is necessary to picture the development of the peak structure. At each stage of the evolution, populations are either tracking moving peaks or passing from peak to peak along traversable pathways. In the mature community the selective forces may be such that the intermediate pathways are transformed to deep, wide valleys. But this does not mean that the system at any point in the past had to solve the problem of traversing these valleys. What it means is that the high suitability of biological organization to evolution allows populations to climb upward running pathways rapidly when they appear, and in general to keep up with the changing landscape.

There is some connection here with the red queen hypothesis of van Valen (1973). Populations must keep evolving to stay in place. Our version is that the structure of organisms must be evolvable enough to evolve fast enough to follow the unfolding adaptive landscape. There is also a connection to the punctuated equilibrium concept of Gould and Eldredge (1977). The equilibrium situation corresponds to a stable landscape. The punctuated situation corresponds to fast tracking of an unfolding landscape. This fast tracking is not, however, due to bizarre nonlinearities in the developmental process. This would not provide a robust way of keeping up with a fast changing peak structure. Rather it is due primarily to mutation buffering, and is thus dependent on gradual transformation of structure and function in response to genetic change. We can recall our image of gross and fine controls on a microscope. Some features of the structure may allow for a wide sweep of variations; but a fine tuning capability is essential and likely to be associated with extradimensional pathways. Finally we should mention the connection to neutralism. If the organizational structure allows for buffering of mutation and other genetic change there is bound to be a great deal of neutral or quasi-neutral variation. Thus if mutation buffering is essential for the effective operation of variation and selection, neutral phenomena will be an inevitable concomitant. Neutralism is better interpreted as a condition for the effectiveness of selection rather than as a phenomenon which implies its irrelevance.

The picture is thus one in which search and structure are inextricably tied together. Search works because the structures are suitable. The structures are suitable because search leads to them. This is why a structuralist view can be just as Darwinian as orthodox NeoDarwinism, and actually twice as Darwinian.

Acknowledgement

This work was supported in part by National Science Foundation Grant IRI 87-02600.

References

Chaiten, G.J., 1977, Algorithmic information theory. IBM J. Res. Dev. 21, 350–359.

Conrad, M., 1972a, Stability of foodwebs and its relation to species diversity. J. Theor. Biol. 34, 325–335.

Conrad, M., 1972b, Information processing in molecular systems. Curr. Modern Biol. (now BioSystems) 5, 1–4.

Conrad, M., 1974, Molecular automata, in: Physics and Mathematics of the Nervous System, M. Conrad, W. Güttinger and M. Dal Cin (eds.) (Springer-Verlag, Heidelberg) pp. 419–430.

Conrad, M., 1978, Evolution of the adaptive landscape, in: Theoretical Approaches to Complex Systems, R. Heim and G. Palm (eds.) (Springer-Verlag, Heidelberg) pp. 147–169.

Conrad, M., 1979a, Bootstrapping on the adaptive landscape. BioSystems 11, 167–182.

Conrad, M., 1979b, Mutation-absorption model of the enzyme. Bull. Math. Biol. 41, 387–405.

Conrad, M., 1983, Adaptability (Plenum Press, New York).

Conrad, M., 1985, The mutation-buffering concept of biomolecular structure. J. Biosci. 8, 669–679.

Conrad, M., 1988, The ecosystem as an existential computer, in: Progress in Systems Ecology, B.C. Patten and S. Jorgensen (eds.), (Prentice Hall, Englewood Cliffs, NJ) in press.

Conrad, M. and Hastings, H., 1985, Scale change and the emergence of information processing primitives. J. Theoret. Biol. 112, 741–755.

Conrad, M., Brahmachari, S. and Sasisekharan, V., 1986, DNA structural variability as a factor in gene expression and evolution. BioSystems 19, 123–126.

Gardner, M.R. and Ashby, W.R., 1970, Connectance of large dynamical (cybernetic) systems: critical values for stability. Nature 228, 784.

Gel'fand, I.M. and Tsetlin, M.L., 1962, Some methods of control for complex systems. Russian Math. Surv. 17, 95–116.

Goodwin, B.C., 1965, Developing organisms as self-organizing fields, in: Mathematical Essays on Growth and the Emergence of Form, P.L. Antonelli (ed.) (University of Alberta Press, Edmonton, Canada) pp. 185–200.

Gould, S.J. and Eldredge, N., 1977, Punctuated equilibrium: the tempo and mode of evolution reconsidered. Paleobiology 3, 115–151.

Hastings, H.M., 1982, The May-Wigner stability theorem. J. Theor. Biol. 97, 155–166.

Jerne, N.K., 1955, The natural selection theory of anti body formation. Proc. Natl. Acad. Sci. USA 41, 849–857.

Levins, R., 1970, Complex systems, in: Towards a Theoretical Biology: Drafts, C.H. Waddington (ed.) (Edinburgh University Press, Edinburgh) pp. 73–88.

May, R.M., 1972, Will a large complex system be stable? Nature 238, 413–414.

May, R.M., 1973, Stability and Complexity in Model Ecosystems (Princeton University Press, Princeton, NJ).

Rosen, R., 1970, Dynamical System Theory in Biology (John Wiley, New York).

Sasisekharan, V. and Pattabiraman, N., 1978, Structure of DNA predicted from stereochemistry of nucleoside derivatives. Nature 275, 159–162.

Salthe, S.N., 1985, Evolving Hierarchical Systems (Columbia University Press, New York).

Sibatani, A., 1985, Molecular biology: a structuralist revolution. Rev. Biol. 78, no. 3, 373–397.

Strobeck, C., Maynard-Smith, J. and Charlesworth, B., 1976, The effects of hitchhiking on a gene for recombination. Genetics 82, 547–558.

Thom, R., 1970, Topological models in biology, in: Towards a Theoretical Biology, Vol. 3, C.H. Waddington (ed.) (Edinburgh University Press, Edinburgh) pp. 89–116.

Van Valen, L., 1973, A new evolutionary law. Evol. Theor. 1, 1–30.

Volkenstein, M., 1979, Mutations and the value of information. J. Theor. Biol. 80, 155–169.

Waddington, C., 1968, The basic ideas of biology, in: Towards a Theoretical Biology, 1. Prolegomena, C.H. Waddington (ed.) (Aldine, Chicago) pp. 1–41.

Wigner, E.P., 1961, The Probability of the Existence of Self-Reproducing Units, in: The Logic of Personal Knowledge (Routledge and Kegan Paul, London).

Williams, G.C. (ed.), 1966, Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought (Princeton University Press, Princeton, NJ).

Wright, S., 1932, The roles of mutation, inbreeding, crossbreeding, and selection in evolution. Proc. Sixth Int. Congr. Genet. 1, 356–366.