

Mass Extinction and the Simulated Annealing Algorithm

Douglas S. Robertson
15 Pavilion Drive
Gaithersburg, MD 20878

Received 23 December 1992, 22 April 1993

ABSTRACT: Classical Darwinian evolution processes involve a natural implementation of a mathematical optimization process, in this case the maximization of a quantity called "fitness." Like most optimization processes, evolution is subject to some well recognized mathematical difficulties, particularly the problem of a system becoming "stuck" in a local extremum of the fitness function instead of the global extremum. Recent developments in the mathematics of optimization processes provide some insight into both the nature of this difficulty and possible means of avoiding the problem.

* * * * *

Mass extinctions have been the subject of much recent debate, and most of the discussion has naturally centered on the causes of such extinctions. Here I wish to focus instead on the effects of the extinctions. It seems obvious that an evolutionary process that contains a mechanism that produces episodic mass extinctions will have different effects than one that lacks such a mechanism. Different ecosystems will be produced. But how will those ecosystems be different? What are the overall effects of mass extinctions on the evolutionary process? A consideration of the behavior of some simple but powerful mathematical algorithms may shed some light on these questions.

The broad class of mathematical algorithms that relate to this problem is the class that involves maximization or minimization of a function¹, particularly a function of many variables. Although there is a variety of methods available for solving these problems, virtually all of them share a common critical difficulty, the inability to distinguish between a local and a global extremum. If the function being explored has a large number of local extrema, then nearly all of these algorithms will converge on one particular extremum and ignore all of the others. In these circumstances the "global" extremum can be extremely difficult to locate. As Press *et al.*, put it, "Virtually nothing is known about finding global extrema in general." (1987, p 274). Even very large and nearby extrema (that are not necessarily global extrema) can be difficult to locate.

The problem of finding the extreme value of a function is of more than mathematical interest, because many of the maximization algorithms have analogues in the real world, and an understanding of the behavior of the mathematical algorithms can lend insight into the behavior of the analogous real systems, especially biological systems. The tendency of a biological system evolving under Darwinian natural selection to get "stuck" in a local maximum of a distribution of fitness values is a problem that has long been recognized. In a previous paper I addressed a variant of this problem, and showed how feedback loops could create a mechanism by which an individual species could avoid being "stuck" in a local maximum of the fitness curve (Robertson, 1991). This paper is concerned with a broader problem. Entire ecological systems that contain a large ensemble of interacting species may also exhibit a tendency to become stuck in a local maximum of an extremely complex objective function.

This objective function for an ecosystem should not be confused with the fitness function for individual species. It is a far more complex entity, one that depends in a highly complex manner on the fitness functions for all of the species present in the ecosystem. To clarify the distinction I will use the term "objective function" when referring to the quantity that is optimized for an entire ecosystem, and the term "fitness function" only when referring to the evolutionary behavior of an individual species.

A novel mathematical technique has recently been devised that deals with the problem of finding global extrema of complex objective functions in a manner that is simple, somewhat counter-intuitive, and startlingly effective, although it is not perfect. The algorithm goes by the name "simulated annealing" because it is a close analogue to solidifying and crystallizing processes that commonly occur in physical chemistry; as we will see it may have analogues in biological processes as well, analogues that are particularly helpful for understanding the effects of mass extinctions.

We can illustrate the working of the simulated annealing algorithm in two contexts, one mathematical, the other physical. The mathematical context involves the famous Traveling Salesman Problem (TSP). The physical problem is

¹There is no important difference between maximization and minimization problems--one can be converted into the other by changing a sign. They will be used interchangeably through this text. The function whose extremum is sought has a variety of names in different disciplines. In mathematics it is sometimes called the objective function; in economics the cost function; in evolutionary biology the fitness function. These terms may be used interchangeably.

as simple as a marble rolling on a complicated surface.

The TSP is easy to state. Given a set of points on a plane, find the shortest (closed) path that touches all of these points at least once. The name of the problem derives from its usefulness to a traveller who must visit a number of cities, who naturally wishes to expend the minimum possible time and effort.

The TSP has been called the most important unsolved simple problem in mathematics. Its solution would immediately allow the solution of an amazing variety of similar problems, all classified under the mathematical rubric of NP-complete problems (See Stewart, 1987, pp. 208-212 for a discussion of NP-complete problems). The TSP is incredibly easy to state and to understand, and, paradoxically, it is simultaneously both trivial and impossible to solve. It is trivial in that its solution is easily understood by a small child. You simply tabulate all of the possible paths that touch the given points, add up the length of each path, and select the shortest. But the problem is also impossible, because the trivial solution generally cannot be finished even by the fastest theoretical computers operating for geologic timescales. The difficulty is not solving the problem, but finding a solution in a reasonable amount of time.²

A simple algorithm for dealing with the TSP begins by taking any path that connects the points. The algorithm then proceeds by trying "small" changes to the path³. If the trial path is shorter than the original then the change is kept, and further changes are tested. This algorithm is simple and reasonably efficient, but it suffers from the classic problem of optimization algorithms, the tendency to get stuck in a local minimum. If there is a shorter path that requires a major change rather than a sequence of minor changes, this algorithm will never find it.

The simulated annealing algorithm provides a simple variation on this procedure that does a remarkably effective job in avoiding the problem. The trick is to accept all of the shorter paths, and *some* of the longer paths. The key lies in the decisions concerning which of the longer paths should be accepted. This is done randomly. The algorithm works by starting with a fairly high probability of accepting a longer path; it then proceeds to decrease this probability gradually. The details of the algorithm can be found in standard references (*e.g.*, Press *et al.*, 1987, pp. 326-334). The most surprising thing about the algorithm is not so much that steps are needed in the "wrong" direction (this is a mathematical consequence of the existence of local extrema) but that *random* steps in the "wrong" direction could be so effective in searching for the global extremum.

Press gives examples that illustrate the effectiveness of the algorithm, and says that: "For practical purposes, it [the algorithm] has effectively solved the famous traveling salesman problem . . ." (Press *et al.*, 1987, p 326). As noted above, the algorithm is called "simulated annealing" because the procedure of slowly decreasing the probability of accepting a longer path is strongly analogous to the process of growing a crystal by slowly cooling a melt. The level of probability of accepting a larger path is analogous to the temperature.

The second example that illustrates the functioning of an annealing algorithm involves a more physical situation. Consider a marble rolling on a surface. If the surface is smooth and shaped something like a bowl with only one minimum then the marble is a simple and efficient mechanism that will always find the global minimum. All you need is gravity, and a little friction. (Without friction, of course, the marble will never find the minimum.)

But suppose the surface is complicated. Particularly assume that the surface has a large number of dimples or holes (extrema), of widely varying depth. A marble placed on such a surface will generally roll to some nearby local minimum and stay there. There is no easy way to find a global minimum with a rolling marble. Again, as in the case of the TSP, progress can only be made if we allow a certain number of "uphill" steps.

To see how a simulated annealing approach would work in this case, we begin by classifying the depth of each hole according to the potential energy that would be released by the marble in dropping into the hole. The i^{th} hole would correspond to a potential energy P_i . Now suppose we are able to give random kicks to the marble, each kick giving the marble a constant kinetic energy ΔE . After a series of such kicks, the marble will generally not be trapped in any shallow hole, that is, any hole such that $P_i < \Delta E$. But it will be trapped permanently in the first hole it encounters with $P_i > \Delta E$. Now suppose the deepest hole has a potential P_{max} . If we start by giving the marble a sequence of kicks such that $\Delta E > P_{\text{max}}$, the marble generally will not settle down anywhere. We then gradually decrease ΔE , and when it becomes

²To illustrate the difficulty, suppose you have 25 "cities" to visit, and you have a computer that can calculate 1,000,000 possible paths in one second. Such a computer will exhaustively solve the problem in only 20 billion years. Next suppose that you have a really fast computer, one that could solve the 25-city problem in only one year. Such a computer could solve a 50-city problem in a mere 10^{40} years. It thus seems safe to assert that the most powerful computers, even in theory, cannot solve the general traveling salesman problem by exhaustive search.

³It is not necessary to specify here what is meant by a "small" change. Different algorithms use different choices. For a discussion, see Press *et al.*, p 328.

less than P_{\max} there will be a good probability that the marble will eventually become trapped in the deepest hole. Thus the importance of decreasing ΔE *slowly*: the marble must have sufficient time to find the right hole at each energy level. Further decreases in ΔE will allow the marble to "explore" lesser holes near the bottom of a deep hole, and thereby have a good chance of finding the overall global minimum.

Notice that this procedure is not guaranteed to find the global extremum--only to have a certain probability of finding it. But absolute mathematical guarantees are hard to come by in the real world. In practice, the annealing method is found to work well on real problems, well enough that it is of great practical use in dealing with practical problems similar to the TSP, such as the design of the interconnections among the circuit elements on integrated circuit chips⁴.

The relevance of this discussion to the evolution of a complex ecosystem lies in the fact that the ecosystem may also have a tendency to stick in a local extremum of a complex objective function. Variation and natural selection by themselves will tend to find such local extrema. Darwinian evolution may therefore be adequate to find the extremum of a very simple and smooth objective function, but it may well have difficulty in finding the global extremum (or even a large and nearby local extremum) of a very complex objective function. In pure Darwinian evolution there is no obvious way around this difficulty, no ready analogue of the process of occasionally accepting a poorer solution in hope of finding a different, (higher) maximum in the objective curve. In other words, individuals with low fitness values in a given ecosystem will not be selected simply because there is a chance that in some other ecosystem they might have a very high fitness value. Thus in order to explore a complicated objective function thoroughly and find global extrema (if they exist) or even to locate the larger local extrema, we will need to have an analogue to the procedure of "kicking" the marble.

It may be possible to find such an analogue in processes that are essentially external to the biological system. An ecosystem may be disturbed by mass extinctions caused by such an external process, such as an asteroid impact or a major climate or sea-level change, for example. The extinction process can provide the needed analogue to a random "kick" (to an entire ecosystem) from outside the system. The organisms that survive following the extinction will generally evolve toward some new and probably different maximum in the objective function. As with the simulated annealing algorithm, there is no guarantee that the system will in fact find a higher maximum, or even a different one, only a certain probability that it will. There is, however, a good chance that if the new maximum is higher, that it will be more resistant to processes that cause mass extinctions, just as the marble that happened to land in a deep extremum will resist small "kicks." It may thus endure for a longer time, and we may have a higher probability of observing it.

How do we know that the extinction will take the system lower on the objective function, away from the local maximum? If the system is truly at such a local extremum, then the answer to this question is almost a mathematical tautology: any change at all relative to a local maximum will take the system lower on the function, by definition.

This procedure will only be important if the objective function is complicated enough to have a broad spectrum of local extrema. Whether it is or not is a question that needs to be investigated in more detail than can be done in this brief essay. There are, however, several reasons to think that it must possess such complexity. In the first place, extremely complex objective functions are commonly found in very simple problems such as the TSP. Although it is not impossible, there is no general reason to expect that systems that are far more complicated (such as biological ecosystems) will have simpler objective functions. More tangibly, if the objective function of an ecosystem were very simple, then we would expect that the system would tend to return to the same extremum, *i.e.*, the same ecosystem would tend to re-evolve following a disturbance such as a mass extinction, just as a marble in a simple, bowl-shaped surface will return to the same point after being disturbed. But, plainly, ecosystems do not tend to repeat themselves after such events. Kauffman and Levin (1987) give a detailed development of some of the implications of the objective functions having this level of complexity. For this discussion I am assuming that the function is sufficiently complex that the problem of a system becoming trapped in a local extremum is serious.

Mass extinctions can be thought of as a form of super-variation, different from the normal variation of Darwinian evolution principally in that super-variation acts to change whole systems of species at once, rather than to change individual species. We should not be too surprised to find that *random* super-variations are effective (given that ordinary random variation is also effective) in providing the raw material on which natural selection can act. Natural selection is the important non-random factor here. The two important facts about super-variation are, first, that it can act in directions other than toward the local maximum of the objective function (as can ordinary variation), and second, that it can sometimes move the system far enough along the objective function to escape some particular local extremum.

In his recent book describing the difficulties posed by the organisms fossilized in the Burgess shale, Gould spends some time debating whether or not there is any pattern to the species that survive mass extinctions (Gould, 1989, pp. 305-

⁴Recall the difficulty of solving the TSP for numbers in the range of 20-50 cities. In contrast, the simulated annealing algorithm routinely solves similar problems involving the interconnections of hundreds of thousands of components on integrated circuit chips.

308). If in fact mass extinction is a form of super-variation, then the problem disappears. We do not expect ordinary variation of species to be other than a random effect, and no more should be expected of super-variation.

We have described here a system with two levels of variation: the usual one caused by random mutations and sexual shuffling of genetic material, and the second caused by mass extinction processes. There is no particular reason that the situation could not be more complicated, with several different levels or types of super-variation. For example there could be a variety of different mechanisms (of varying efficiency) for producing mass extinctions. Such a multi-level system could well prove extremely efficient for exploring very complex objective functions.

Mass extinctions are probably not the only mechanism capable of kicking an ecosystem out of a local extremum in the objective function. If the extremum is sufficiently shallow, ordinary species variation might be enough to spill the system over a low "watershed." Kauffman and Levin describe the possibility of an evolutionary "long jump," *i.e.*, a single mutation that has large effects on an organism (Kauffman and Levin, 1987, p 33), and Gould describes a similar mechanism he refers under the heading of "hopeful monsters" [Gould, 1982, pp. 186-193]. These mechanisms may provide yet another type of "super-variation," another mechanism by which ecosystems can avoid being trapped in local extrema of the objective function.

Another aspect of the simulated annealing algorithm may be relevant to biological problems. In simulated annealing processes the broad outlines of the final solution generally appear early in the process--the ball is trapped in the vicinity of the deepest hole early on. Large variations away from the final solution are seen only very early in the process. Press describes an example of this behavior in a variant on the TSP that involves a penalty for the connecting path when it crosses certain boundaries that are somewhat whimsically termed "rivers." Solutions that involve more than the minimum number (two) of river crossings appear only at the very beginning of the process. The final number of crossings appears fairly early in the solution process, and tends not to change through successive iterations.

This behavior is perhaps clearest in the case of the rolling marble, wherein many portions of the complex surface are explored only in the early part of the solution before the marble is trapped in the very deep holes. Similarly, in the evolution of biological systems, the variety of possibilities that is "tested" is larger in earlier times, and quickly decreases with time. Gould refers to this as a pattern of "maximal early breadth." He says "Paleontologists have long recognized the Burgess pattern of maximal early disparity in conventional groups of fossils with hard parts. The echinoderms provide our premier example." (Gould, 1989, p 310-311). This pattern should be recognized as the natural consequence of a system that is evolving under something like an annealing algorithm.

Mass extinctions in the fossil record may lack one critical feature of the simulated annealing algorithm--there is no clear "cooling" process. That is, there is no obvious mechanism by which the mass extinctions mechanisms would be made to decrease gradually with time--they may in fact not decrease at all. Thus the process may resemble periodic re-melting rather than annealing. Nevertheless, by analogy with the rolling marble problem, a system that is periodically "kicked" may eventually get to a deeper extremum, and get there faster, than a system that is not kicked at all. There has been little mathematical analysis of such an "annealing" process without cooling. It seems intuitively, however, that such a system, even without cooling, will still spend more time in the vicinity of global extrema of the objective function than will a system that lacks the extinction processes. This assertion may need more detailed, quantitative analysis.

There is one possible mechanism for "cooling" the mass extinction process, although it relies on a number of controversial or untested assumptions. If asteroid impacts are a major (not necessarily exclusive) cause of such extinctions, the frequency of these impacts should decrease with time as these celestial bodies are swept up by the major planets. It is clear from studies of craters on the Moon that there was a significant gradient in the frequency of impacts after the first billion years or so of the Earth's history (see, *e.g.*, Short, 1978, pp. 94-96). Whether or not the rate of change is large enough to cause significant "cooling" is something that will require further study. This process may be taking the word "slow" to an extreme.

We should not lose sight of the fact that the biological systems that are evolving under Darwinian natural selection are almost infinitely more complex than the simple mathematical models (*e.g.*, the TSP) we have explored here. In particular, it has been suggested that the fitness function for individual organisms is controlled by the existence of other organisms (see, *e.g.*, Van Valen, 1980; Robertson, 1991). Thus a mass extinction event would not only knock a system out of a local extremum in the objective function, it would vastly alter the objective function itself, thereby producing a wildly non-linear situation that could enormously amplify the effects of the extinction itself. In other words, the objective function (or the complex surface that the marble is rolling on) is not only extremely complex but is time-variable as well, and the time variations can be both rapid and extreme where mass extinctions are involved. Thus the general problem of the effects of mass extinctions may prove to be far more complicated than would be implied by the simple example of a marble rolling on a complex surface that is not time-variable. The non-linearities introduced by time variability lead directly to feedback loops and chaotic behavior, as outlined in Robertson, (1991). The mathematical complexities inherent in such chaotic behavior suggest that the patterns seen in biological evolution may in the long run be no more predictable than the weather.

Darwinian evolution upset the earlier notions of special creation of species. It showed how species could be produced by the mechanistic process of natural selection. But some Darwinists created a comfortable myth of their own, the myth that present-day species, including of course the human race, were the natural, perhaps inevitable, product of a process of slow, steady, inexorable improvement of organism design by variation and selection, a process in which organisms slowly found their way along a smooth, well-behaved fitness curve. This concept underlies the idea of evolutionary "gradualism" contrasted to more recent ideas such as Gould and Eldridge's "punctuated equilibrium." But "gradualism" almost certainly requires smooth and well behaved fitness functions. If we allow the possibility that the function is not smooth or well-behaved in the mathematical sense, but instead has fractal-like characteristics and is strongly time-variable, then species we observe today may represent only the present, fascinating pattern created by a chaotic process. Present ecosystems might be but one beautiful loop on pattern of Mandelbrot-like complexity, or more precisely, on a pattern so complex that it makes the Mandelbrot set look smooth and simple by comparison.

References:

- Gould, S.J., 1982, *The Panda's Thumb*, New York: W.W. Norton.
Gould, S.J., 1989, *Wonderful Life*, New York: W.W. Norton.
Kauffman, S., and S. Levin, 1987, Toward a General Theory of Adaptive Walks on Rugged Landscapes, *J. Theor. Biol.*, **128**, 11-45.
Press, W.H., B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, 1986, *Numerical Recipes*, Cambridge University Press.
Robertson, D.S., Feedback Theory and Darwinian Evolution, *Journal of Theoretical Biology*, **152/4**, 469-484, 1991.
Short, N.M., 1975, *Planetary Geology*, Englewood Cliffs, NJ: Prentice-Hall.
Stewart, I., 1987, *The Problems of Mathematics*, New York, Oxford University Press.
Van Valen, L.M., July, 1980, Evolution as a Zero-Sum Game for Energy, *Evolutionary Theory*.