The Outlaw Method for Solving Multimodal Functions with Split Ring Parallel Genetic Algorithms

K. Burton Harvey and Chrisila C. Pettey
Department of Computer Science, Box 48 Middle Tennessee State University Murfreesboro, TN. 37132
kbharvey@phgtechnologies.com, cscbp@mtsu.edu
(615)898-2416

Abstract

The similarities between island model parallel genetic algorithms (PGAs) and natural islands invite the application of certain geobiological concepts to the field of evolutionary computation. Heretofore, island model PGAs have usually been used to find single solutions to unimodal functions. However, geobiology's split ring species phenomenon—which refers to the tendency of related environments to develop multiple, distinct species from the same genetic stock[6]—suggests that island model PGAs could be used to develop multiple, distinct solutions to multimodal functions. Though initial experiments demonstrate that the strong selection pressures exerted by island model PGAs overwhelm the split ring species phenomenon, the outlaw method is a technique for protecting the split ring species phenomenon so that it can develop multiple solutions despite these pressures.

1 Introduction

Genetic algorithms and parallel computation fit together neatly. In fact, John Holland postulated the marriage of genetic algorithms to parallel architectures as long ago as 1960[4]. Also, the biological process of natural selection which GA researchers strive to replicate is itself parallel and asynchronous, consisting of multiple isolated populations simultaneously exploring diverging paths of development[5]. The island model's similarity to these natural isolated population structures suggests a new tangent for GA exploration—the relationship between geography and evolutionary development[3]. In particular, could the split ring theory of biology be used to find multiple optima in a multimodal function? The remainder of this paper is a preliminary investigation of that possibility. In the next section, the split ring theory will be described. Subsequent sections will describe the split ring PGA and the outlaw method along with several experiments done to examine the proposed split ring solution to the problem of finding multiple optima.

2 The Split Ring Phenomenon

Island geobiologists use the term *split ring species phenomenon* to refer to the continuum of difference that may be observed among populations of organisms who share common ancestors but are separated across an island chain[6]. Populations on neighboring islands closely resemble each other genetically, while populations located on islands that are far apart are strikingly dissimilar. In fact, in most instances of the split ring phenomenon, the populations at either end of the chain are *reproductively isolated* from (unable to interbreed with) each other and different enough to be considered separate species. In the context of this phenomenon, "split" refers to the inability of populations on these end islands to interbreed and their resulting dissimilarity.

The continuum of difference develops gradually, and in response to the degrees of separation between the islands. Consider the five islands arranged in a split ring in Figure 1. Populations on neighboring islands A and B can regularly exchange members and interbreed, but populations on islands separated by greater distances—such as the populations on islands A and E—are unable to directly interbreed. Additionally, the differences between the islands themselves may cause differences between the populations; for instance, climactic dissimilarities between a northern island such as A and a southern island such as E could exert different selection pressures on their resident populations.

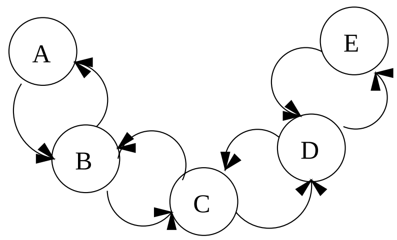

Figure 1: Split Ring of Islands

The split ring species phenomenon can be regarded as nature's attempt to develop multiple solutions (multiple species) in response to a single problem (the adaptation of a type of organism to an environment). Would it be possible to observe—or perhaps even catalyze—this type of multiple solution development in the context of island model PGAs? Specifically, could a split ring island model PGA isolate each of a bimodal function's two solutions on the islands at either end of the ring? In the past, Cavvicchio's "preselection," De Jong's "crowding," and Goldberg's "fitness sharing" have been used to tailor GAs to the solution of multimodal functions[2]. A split ring PGA might prove to be faster and easier to implement than these previous techniques, which necessitate the inclusion of GA-slowing evaluative procedures[2]. The following sections describe the initial experiments with a simple split ring PGA. Section 3 presents the implementation details. The remaining sections contain explanations of the tests done to determine the efficacy of using split ring PGAs and the enhancement to the split ring PGA known as the outlaw method.

3 Implementation Details

Unlike diffusion model PGAs, which require hard-to-find massively parallel processing computers, island model PGAs can be implemented on common single processor machines running readily available software. The experiments described in this paper were performed using a cluster if IBM-compatible PCs equipped with Intel Pentium 100 MHz processors and the Linux Red Hat 5.0 operating system. The PCs were linked together with a freeware clustering package called PVM (Parallel Virtual Machine). PVM creates an invisible middle layer between a PCs' operating system and its application programs, masking the participating computers' individual architectures and effectively combining them into one large multiprocessing virtual machine.

For all experiments in this paper roulette wheel selection and elitism was used on each island. Each island had a mutation rate of .001 and a crossover rate of 0.6. In the split ring PGA, each island determines its left and right neighbors in the ring. After the island GAs initialize their populations and enter their main loops, they communicate chromosomes to their left and right neighbors. To introduce a "split" into the island model PGA described above, the last process in the ring should not communicate with its right neighbor, and the first process in the ring should not communicate with its left neighbor. Thus the split is between the first and last islands.

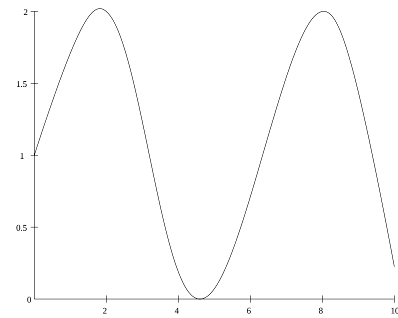

Figure 2: Function 1

4 The Split Ring PGA

In order to determine a viable communication rule and whether or not the split ring PGA has the potential for exhibiting the split ring phenomenon, two experiments were performed on the bimodal function $f(x) = sin(x) + 1$ for $0 \leq x \leq 10$ (see figure 2). The optima for this function are located at $x = \pi/2$ and $x = 5\pi/2$. For this function, individuals were 10 bit chromosomes. The two experiments contrasted each island communicating its best individual to each of its neighbors vs. communicating a random individual to each of its neighbors. Both experiments used 10 islands, each of which communicated with their neighbors each generation. Since these initial experiments were simply to determine if the Split Ring PGA was feasible, the islands ran for only 10 generations. The experiments were run 10 times each and averaged. In this and all subsequent experiments roulette wheel selection was used, the mutation rate was .001, the crossover rate was .6, the population size per island was 250 individuals, each individual was ten bits in length and was scaled to represent the numbers from 0 to 10.23, and elitism was used.

In these experiments, the ring using best individual communication succeeded in isolating different optima at oposite ends in only three of the ten runs. In the other seven runs, the populations of those "end is-

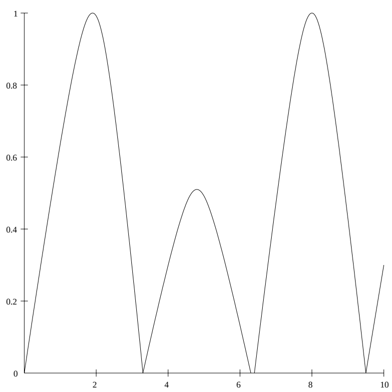

Figure 3: Function 2

lands" either developed to the same optimum, had a mixture of both of the optima, or found neither of the optima. On the other hand, the random individual communication ring was better. In seven of the ten runs, the islands at opposite ends of the split ring developed different optima. Inspection of the three unsuccessful runs revealed that while the end islands' populations had not yet converged on their respective optima, in a few more generations they possibly could have.

From these experiments it appeared that a split ring PGA using random individual communication was indeed capable of finding multiple optima for a multimodal solution. So, in order to more thoroughly test the split ring PGA, a new function was devised (see figure 3). For Function 2, $f(x) = sin(x)$ for $0 \leq x \leq \pi$, $f(x) = [sin(x - \pi)/2]$ for $\pi \leq x < 2\pi$, and $f(x) = sin(x)$ for $2\pi \leq x \leq 10$. Function 2 thus has two optima at $\pi/2$ and $5\pi/2$ and a sub- optimal solution located exactly between the two true optima at $3\pi/2$. Of interest was how would intermediary islands react to the presence of this sub-optimum? The best case scenario was a chain of three islands in which the first island found the left optimum, the middle found the sub-optimum, and the right island found the last optimum.

The first series of experiments with Function 2 was an attempt to conclusively determine whether the "send the best chromosome" or the "send a random chromosome" communication strategy is best for uncovering both optima of a bimodal function. To make the results more conclusive it was decided to increase the number of subpopulations and the number of generations. First a control experiment was investigated. In the control experiment an unsplit ring of 15 islands each communicating their best individual bidirectionally was monitored over 100 generations. The phenomenon which was observed could best be described as "peer pressure." In the first generations of the run, optima were scattered evenly across the islands; for example, although the majority of island 0's chromosomes might bear closest resemblance to the first optimum, its neighbor, island 1, might have a larger percentage of chromosomes closest to the sub-optimum, while the next island might have a percentage closest to optimum 2. As time goes by, however, this variation vanishes. In fact, one can observe the islands' influence on each other. If an island favoring the second optima is located between two islands favoring the first, the optimum chromosomes which the left and right islands submitted to their mutual neighbor gradually brings his population composition to resemble theirs. This pattern of influence happened simultaneously over the entire ring of islands, bringing it gradually to homogeneity and losing one of the two optima. There were hold-out groups of islands who refused to give up their optima for a long time, but by the 100th generation, even those islands took on the semblance of the other islands in the ring. As one might expect, the speed with which the islands reached homogeneity was directly proportional to their migration rate; the more often islands sent each other their best chromosomes, the more quickly the nonconforming islands gave up their unpopular optima in favor of those of the majority.

These experiments were repeated with a ring of the same size that communicated unidirectionally. Interestingly, peer pressure could be seen here too, but only in one direction. In other words, an island could cause its right neighbor to change optima. The last islands in the ring would be the first to achieve homogeneity. Then the other islands in the population would achieve homogeneity. Again, tinkering with the migration intervals showed that more frequent communication caused the ring to converge more quickly.

Next, another ring of 15 islands—these sending a random chromosome—were tested. On the basis of the experiments with Function 1, it was expected that this ring would exhibit the split ring species phenomenon. It did—but only for the first 15 generations. After that, one of the end islands forsook its optimum and began to develop the same one as was being developed by the island on the opposite end of the ring. Gradually, each of the intermediary islands also converged to this optimum. By the fiftieth generation, one of the optimum chromosomes had disappeared completely from the ring, never to be seen for the rest of the program's execution.

These experiments tended to show that a "send a random" chromosome strategy is better than a "send the best" chromosome strategy for uncovering multiple optima. However, a split ring PGA employing either strategy apparently, in a very short time, will lose one of the optima and converge completely towards the other optima. Since the split ring phenomenon does tend to show up in the "random communication split ring PGA" in the earlier generations, perhaps the number of islands in the ring has a bearing on the ability of the split ring PGA to find the solutions.

In order to determine what effect, if any, the number of intermediary islands had on the isolation of different solutions at different ends of the split ring, experiments were run with split ring PGAs with 1-15 islands for one hundred generations each. The results were disappointing. Although the split ring species phenomenon was sometimes evinced in the first few generations of the PGA, after 20 or so generations, it was gone. Furthermore, if the ring of islands did avoid homogeneity and found both optima, those optima were rarely located on the islands at the end of the chain. Also, rings with fewer islands and rings with more islands tended to converge to homogeneity at approximately the same rate.

Disappointed by the split ring PGA's apparent inability to find Function 2's optima and maintain them for a reasonable length of time, the experiments on Function 1 were performed again allowing the islands to evolve for 100 generations rather than just for 10. Once again, although the beginnings of the split ring phenomenon were evident in the first few generations of a split island model PGA, those beginnings were quickly covered over as the PGA progressed.

One can only hypothesize about why the split ring PGA fails to develop both optima. High selection pressure in general, though, is probably responsible. In nature, there are an infinite number of unique variables which diffuse strong selection pressure. Perhaps what is needed is an addendum to the PGA designed to encourage divergent evolution and protect nonconformity as soon as it occurs. Such an addendum would latch on to the split ring species phenomenon evident in the early stages of a split ring PGA, and preserve it through the rest of the generations. The outlaw method which is described in the next section is one possible addendum to a split ring PGA.

5 The Outlaw Method

Traditionally, processes in island model PGAs communicate good chromosomes to each other in the hope that the single best answer will be found by one of

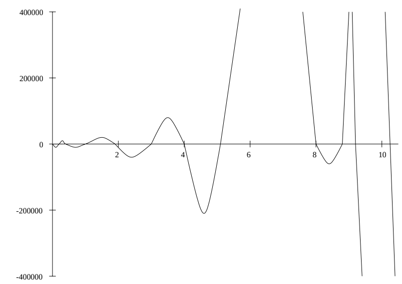

Figure 4: Function 3

them in a short amount of time. In other words, the goal of their combined efforts is computational speedup and their communication is cooperative in purpose. But as shown by the previous experiments, the traditional approaches to communication taken by island model PGAs eventually eradicate all variation from the ring of processes—even if that approach is "send a random chromosome." If PGAs are to find multiple solutions to multimodal functions in an efficient and reliable fashion, perhaps the purpose of their communication methods needs to be re-thought.

For instance, what would be the effect of making communication competitive in nature? Specifically, what if an island communicated chromosomes to its neighbor not in the hope that the neighbor would find the communicated chromosomes useful, but to prohibit the neighbor from developing those chromosomes? This philosophy would be something akin to the toddler's creed or the top dog's rule, "It's mine, and don't you dare take it!" Such a restriction might force each island in the chain to come up with its own unique solutions to the problem at hand. One such restrictive method is the *Outlaw Method*. Four types of outlaw methods are considered in this paper. For all outlaw methods, communication was unidirectional. Unidirectional communication establishes a "pecking order" among the islands; the first island in the ring gets to develop whichever solution it stumbles across, while the last island must avoid all of the answers which its predecessors in the ring have discovered. With bidirectional communication, chaos would ensue as adjacent islands claimed the same outlaw chromosome and sent it to each other. Also, unidirectional communication reduces run time because it generates only half as many messages[1]. In order to test these outlaw methods, a third function was created (see figure 4). Function 3 was a polynomial with 12 roots at 0, .2, .3, 1.1, 1.9, 2.8, 4.0, 5.2, 8.0, 8.6, 9.1, and 10.2. It was decided that the PGA would search for roots rather than for maxima. To establish control, a split ring PGA was applied to this function. Twelve islands were used in the ring. The PGA was run for 100 generations five times, each time with a different random number seed. In four out of the five runs, the split ring PGA found seven of the roots. In the other run only six roots were found. In all five runs each island found a root, but the solutions were not unique. Also, adjacent islands often developed the same solutions.

5.1 Outlaw Method 1

In the first version of the outlaw method, each island sends its best chromosome to its right neighbor. An island considers the chromosome which it receives from its left neighbor during the communication phase to be an "outlawed chromosome," and sets about replacing all copies of the outlaw chromosome found in its population with random chromosomes. This removal process occurs once per generation, after selection, crossover, mutation and communication have already taken place. The removal process is intended to thwart the redundant development of answers already found by other islands further left in the chain and to encourage the development of new chromosomes. In all five runs Outlaw Method 1 uncovered eight roots. Pleasingly, adjacent islands did not develop the same solutions. However, four islands did develop the same solutions as more distant islands. Outlaw Method 1 appears to be an improvement over the original split ring method because of the slight improvement in results and because of the definite improvement in communication time.

5.2 Outlaw Method 2

In Outlaw Method 1 it bacame apparent that each island needed to know not only what solution its immediate left neighbor was developing, but what solutions its more distant neighbors were developing as well. However, it was necessary to develop a method that would instill this knowledge in the islands without opening up direct communication paths between them. It was noticed that solution development in the Outlaw Method 1 environment was a very dynamic process. Often, an island which had just developed a particular solution had to forego that solution when alerted in the communication phase of the development of that solution by another island. This news would force the island to search for a new solution, and this of course would cause the island to send a new "outlaw chromosome" to its right neighbor, perhaps robbing that island of its claimed solution, too. This tendency of outlawed chromosomes to "trickle down" from one island to the next over the course of several generations could be used to endow each island with a memory of chromosomes against which it had been warned. Given this memory, islands might not attempt to repeatedly re-develop redundant solutions.

In the second outlaw method, each island maintains a log of outlawed chromosomes which it has received from its left neighbor. The size of this log is related to the position of the island in the ring; islands further down the ring have larger logs to accommodate the greater number of chromosomes which they have to avoid developing. The log is filled gradually, with one chromosome per communication phase. It is hoped that the log will fill with chromosomes representing the best efforts of islands to the left in the chain, or at least with strong variations on them. This log is implemented by means of a circular queue. During the removal process, each chromosome in the population is compared with each chromosome in the outlaw queue; if a chromosome in the population matches one or more of the chromosomes in the outlaw queue, that chromosome is replaced by a random chromosome.

The results of the experiments with Outlaw Method 2 indicated that Method 2 is somewhat effective at prohibiting islands from developing the same solutions, but utterly fails to encourage the islands to develop new solutions. The split ring with Outlaw Method 2 found 8 roots in four of the five runs, but in all five runs, the two islands at the end of the ring failed to develop any solution at all. Outlaw Method 2 exerts increasing amounts of "negative" selection pressure on islands farther down the ring. Subsequent islands have increasingly less freedom; not only must they develop solutions, but they must develop solutions not developed by any island before them in the ring. Although they have the smae amount of time as the first islands in which to develop answers, those last islands must frequently throw out their work and start over again at the beginning because an island before them in the ring gave up its root and chose their root.

5.3 Outlaw Method 2 With Spin

In an effort to help the last islands in the PGA come up with a solution, a supplemental procedure called *spin* was devised. When spin is implemented, any island which has to remove a large number of chromosomes from its population during the remove outlaw phase will perform the select, crossover, mutate, and remove outlaw phase several times in rapid succession without communicating with its neighboring island. Spin is intended to give islands which have been sent "back to the drawing board" a chance to catch up with their more privileged neighbors. In all five runs of Outlaw Method 2 with spin, nine of the islands found unique roots. The other three islands found roots uncovered by other islands. Outlaw Method 2 with spin shows potential for being the better of the four split ring PGAs discussed so far. It seems to more consistently find more solutions than the split ring PGA or Outlaw Method 1 or Outlaw Method 2, and it has half the communication of the original split ring PGA. It does, however, require more storage than the other split ring PGAs.

5.4 Outlaw Method 3

Though the preceding versions of the outlaw method had proven useful in improving a split ring PGA's chances of uncovering all the optima in a multimodal function, they had not achieved the striking results of the non-redundant discovery of all twelve solutions in Function 3. Given the possibility that finding all twelve solutions is more important than actual run time of the PGA, a third outlaw method was devised. As in Outlaw Method 2, each island has a log with a size dependent on the island's position in the ring. During the communication phase, however, an island broadcasts its best chromosome not only to its immediate right neighbor, but to every island to its right to the end of the ring. To keep this increase in communication in perspective, consider the number of "receives" done by the nodes in each of the methods during each generation. Since sends are typically non-blocking, only receives can slow down a process. The maximum number of receives done by an island each generation in the original split ring PGA is 2. The maximum number of receives done by an island each generation for Outlaw Methods 1 and 2 is 1. For Outlaw Method 3 the maximum number of receives would be one less than the number of islands. As was mentioned previously, this increase in communication time may be relatively unimportant if all solutions can be found—particularly if the number of islands is fairly small (say 5-15).

The results of applying Outlaw Method 3 to function 3 were impressive. In four of five runs, all twelve of the roots were discovered. In the other run, eleven roots were discovered, while one island found no solution.

6 Future Research and Conclusions

While the experiments presented here suggest that split ring island model PGAs without the outlaw method may not be particularly suited to finding multiple optima, the experiments are admittedly preliminary. As with all GAs, there are multiple parameters to be set. Several modifications to these experiments would be worthy of further investigation. For instance, since geographically separate islands have many different environmental factors, the islands in the split ring PGA could have different selection techniques and/or different mutation rates and/or different crossover rates.

Differing mutation rates could also be effective in the outlaw methods. Because islands farther down the ring may find it more difficult to find a solution, increasing the mutation rate for each successive island might give them more freedom with which to explore the search space.

As for other improvements to the outlaw methods, it makes sense to remove redundant answers from an island's population. But should those removed individuals be replaced with completely random chromosomes? After all, such replacement chromosomes are "random" only in that they have been selected thoughtlessly; they may or may not contain schemata possessed by the very chromosomes which they replace, and in time might develop into the same outlawed individuals. In fact, one can imagine that an "evolutionary loop" outlawing certain individuals causes the same individuals to be deleted from an island's population only to re-evolve and be deleted again and again. Given random replacement, it is even conceivable that a deleted chromosome could be replaced by an exact replica of itself. Perhaps a more thoughtful consideration of how replacement chromosomes should be chosen might prove productive.

Another drawback of the current system of outlaw removal is its deterministic nature. While it is good to remove most of the outlawed chromosomes from an island's population, it might be better to leave in a few to serve as intermediate steps in the development of different solutions.

Because it has no restrictions imposed upon its solution by its neighbor's outlaw method, the first island in a ring often succeeds in discovering and refining a solution long before the later islands find and refine their solutions. Sometime after this solution is found, the first island may find another. This fickleness can have disastrous effects on the later islands in the ring, especially if the new solution which the first island has chosen was already claimed by another island in the ring; that island is robbed of ists solution and must find a completely new solution. As the robbed island searches for a new solution, it will pass a new outlawed chromosome to its right neighbors, possibly disturbing them as well. This chain reaction unsettles a ring of islands and can set it back many generations. To prevent this occurrence, islands could be equipped with locking mechanisms. After an island's population converges to a solution of predetermined quality, an internal trigger in the island could cause it to shut down its selection, crossover, and mutation phases, thereby preventing the island from abandoning the solution which it has found in favor of another one. For the rest of the ring's life, that locked island would only relay outlawed chromosomes.

Split ring PGAs employing the outlaw method show potential for being applicable in many areas. For example, in the past, PGAs have been applied to the solution of the Traveling Salesman Problem. A split ring PGA with the outlaw method could simultaneously develop many routes from one node to another. In the event of an unforeseen contingency, such as the failure of a critical router or the blocking of a runway by snow, users could choose a solution which does not involve the failed node. This course of action would effectively allow the problem to be solved by computer without reprogramming the problem to incorporate the contingency.

The area of artificial intelligence is a second example of the potential use of the outlaw method. Sentient beings incorporate multimodality in their planning processes. When planning, they think, "If I can't do A because of B, then I will do C," etc. When turned to artificial intelligence applications, a split ring PGA with the outlaw method could mimic that type of planning for multiple unforeseen contingencies.

One interesting occurrence in Outlaw Method 2 was that adjacent islands exhibited a tendency to develop solutions next to each other in the search space. It was as if an island which was discouraged by its left neighbor from developing a certain solution chose to develop the next best thing. Perhaps the outlaw method might also serve to automatically sort solutions from best to worst.

The work presented here suggests that the split ring phenomenon is evident at least in the earliest generations of a split ring PGA. What is needed is an enhancement that will encourage the split ring PGA to maintain this phenomenon. As an example, when split ring PGAs are enhanced with the outlaw method, they show potential for simultaneously developing multiple—if not all—solutions to multimodal functions. The success of the outlaw method warrants investigation of other enhancements which could be as successful as the outlaw method with perhaps less algorithm complexity and fewer communications.

7 Acknowledgements

The split ring species phenomenon was first brought to the attention of a computer science graduate student named John Clark by Dr. Deborah Clark, his wife and a professor in the biology department at MTSU, in the Fall of 1997. The authors want to thank John Clark for telling us about the split ring species phenomenon, and for his contribution to the experiments with Function 1.

References

[1] Cohoon, J.P., S.U. Hegde, W.N. Martin, and D. Richards. "Puctuated Equilibria: a Parallel Genetic Algorithm." In *Genetic Algorithms and their Applications: Procedings of the Second International Conference on Genetic Algorithms*, John Grefenstette, ed. Lawrence Erlbaum and Associates. Hillsdale, New Jersey: 1987.

[2] Goldberg, David E. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley: 1989.

[3] Goldberg, David E. "Zen and the Art of Genetic Algorithms." In *Proceedings of the Third International Conference on Genetic Algorithms*, John Grefenstette, ed. Lawrence Erlbaum and Associates. Hillsdale, New Jersey: 1987.

[4] Holland, J.H. "Iterative Circuit Computers." In *Proceedings of the 1960 Western Joint Computer Conference*, pp. 259-265.

[5] Manderick, Bernard and Piet Spiessens. "Fine-Grained Parallel Genetic Algorithms." In *Proceedings of the Third International Conference on Genetic Algorithms*, John Grefenstette, ed. Lawrence Erlbaum and Associates. Hillsdale, New Jersey: 1987.

[6] Moritz, Craig, C.J. Schneider and D.B. Wake. "Evolutionary Relationships within the Ensatina Eschscholtzii Complex Confirm the Ring Species Interpretation." *Systematic Biology*, March 1992, 41:273-291.