

The Hierarchical Fair Competition (HFC) Model for Parallel Evolutionary Algorithms

Jian Jun. Hu

hujianju@cse.msu.edu

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

Erik D. Goodman

goodman@egr.msu.edu

Genetic Algorithms Research and Applications Group
Michigan State University
2857 W. Jolly Rd., Okemos, MI 48864

Abstract -The HFC model for evolutionary computation is inspired by the stratified competition often seen in society and biology. Subpopulations are stratified by fitness. Individuals move from low-fitness subpopulations to higher-fitness subpopulations if and only if they exceed the fitness-based admission threshold of the receiving subpopulation, but not of a higher one. HFC's balanced exploration and exploitation, while avoiding premature convergence, is shown on a genetic programming example.

I. INTRODUCTION

One of the central problems in evolutionary computation is to combat premature convergence and to achieve balanced exploration and exploitation. In a traditional GA, selection pressure must not overwhelm the diversity-increasing operators (mutation and, to some extent, crossover) or premature convergence is likely to occur. As the evolutionary process goes on, the average fitness of the population gets higher and higher, and then only those new individuals with similarly high fitness tend to survive. New "explorer" individuals in fairly different regions of the search space usually have low fitness, until some local exploration and exploitation of their beneficial characteristics has occurred. So a standard EA tends to concentrate more and more of its search effort near several discovered peaks, and to get "stuck" in these local optima (we use here the language of continuous, real-valued function optimization, but more generally, the concept of "attractors" can instead be used). Many variations [1,2,3,4,5,6] on traditional GA's and especially many of the efforts on parallel EA's are aimed at addressing this problem, as described in the next section. In this paper, an observation about a strategy employed in some societal and biological systems to maintain different high fitness individuals in one population led to the discovery of the new approach presented in this paper, called the Hierarchical Fair Competition parallel model (HFC), which can efficiently combat the premature convergence of EA's.

A. Previous Work on the Premature Convergence Problem

Crowding [1] is proposed as a diversity-maintaining offspring replacement strategy. Replacing similar individuals tends to prevent the population from aggregating around one

peak, and thus extends the search area horizontally. However, genetic drift and sampling error often cause a GA to converge to one or two peaks of the search space, in spite of crowding [2].

Fitness sharing – a modification of the fitness-based selection operation [3] -- permits the formation of stable subpopulations centered on different peaks, thereby permitting parallel search at many peaks in the search space. The problem is that as the number of local optima in the search space grows (especially for GP), it becomes increasingly probable that the limited size of the population cannot accommodate all of them.

Typical multi-population models in EA explicitly maintain several subpopulations, each of which may hold individuals at or near one or more local optima. However, the feasible number of subpopulations may often be small compared to number of peaks that must be explored.

The injection island GA (iiGA) [4, 5] is a hierarchical, parallel model in which subpopulations are organized in a hierarchy with different representations at each level (often representing different resolutions of problem representation). The iiGA is similar in some respects (its hierarchical organization of the subpopulations) to our HFC model. The major difference is that the search space in iiGA is fundamentally divided into hierarchical levels according to the resolution (or type) of representation used at each level, in a predetermined way. For each subpopulation, especially for the highest-level subpopulation(s), there is still the risk of being trapped in a local optimum. The iiGA addresses the premature convergence problem primarily by using multiple subpopulations at each level, and by using mechanisms such as crowding within each subpopulation, rather than directly through the hierarchy.

The basic source of premature convergence of EA's comes from an explicit fact: selection pressure makes high fitness individuals reproduce quickly and thus supplant low-fitness individuals, some of which may, in fact, be more promising, but not yet fully exploited. This fact holds true even when we find search points near a global optimum, as long as they are not close enough to have high fitness relative to those near other, earlier-explored local optima. It is clear that in many cases, an EA doesn't "appreciate" good genetic material until it is placed in the right context. This difficulty

in appropriately rewarding good genes until they are assembled into a good genotype impedes search. The “unfair” competition contributes to the slow search progress of many EA’s when confronted with difficult, high-dimensionality, multi-modal problems.

It appears that current EA’s have not addressed this unfair competition directly enough. What we need to do is to allow young but promising individuals (i.e., those in relatively new regions, which may ultimately give rise to high-fitness offspring, but which are currently not of high fitness) to “grow up” and, at an appropriate time, join in the cruel competition process and be kept for further exploitation or be killed (as appropriate) when they are demonstrated with some confidence to be bad. At the same time, we hope to maintain the already-discovered high fitness individuals and select from them even more promising individuals for exploitation without killing younger individuals. Holland has recently explored a mechanism that provided some protection for “new” individuals, in a different way, in his Cohort Genetic Algorithm [6].

The Hierarchical Fair Competition parallel model (HFC) provides a mechanism satisfying these seemingly conflicting requirements. It originated as a metaphor for the mechanism of hierarchical fair competition principles found in some societal and biological systems, which can maintain and foster potentially-high-fitness individuals (or, more accurately, progenitors of high fitness individuals) efficiently.

II. HIERARCHICAL FAIR COMPETITION IN SOCIETAL AND BIOLOGICAL SYSTEMS

Competition is widespread in societal and biological systems, but diversity remains large. After close examination, we find there is a fundamental principle underlying many types of competition in both societal and biological systems: the Fair Competition Principle.

A. The Fair Competition Principle in Societal Systems

In human society, competitions are often organized into a hierarchy of levels. None of them will allow unfair




FIGURE 1: In education system, low level students compete to get admission to higher level schools.

competition – for example, a young child will not normally compete with college students in a math competition. We use the educational system to illustrate this principle in more detail.

In the education system of China and many other developing countries, primary school students compete to get admission to middle schools and middle school students compete for spots in high schools. High school students compete to go to college and college students compete to go to graduate school [Fig. 1] (in most Western countries, this competition starts at a later level, but is eventually present, nonetheless). In this hierarchically structured competition, at each level, only individuals of roughly equivalent ability will participate in any competition; i.e., in such societal systems, only fair competition is allowed. This hierarchical competition system is an efficient mechanism to protect young, potentially promising individuals from unfair competition, by allowing them to survive, learn, and grow up before joining more intense levels of competition. If some individuals are “lost” in these fair competitions, they were selected against while competing fairly only against their peers. If we take the academic level as a fitness level, it means that only individuals with similar fitness can compete.




FIGURE 2: HFC model extends the search horizontally in search space and vertically in fitness dimension and kills bad individuals at appropriate times while allowing promising young individuals grow up continuously

An interesting phenomenon sometimes found in societal competitions is the “child prodigy.” A ten-year-old child may have some extraordinary academic ability. These prodigies may skip across several educational levels and begin to take college classes at a young age. An individual with sufficient ability (fitness) can join any level of competition. This also suggests that in subpopulation migration, we should migrate individuals according to their fitness levels, rather than according to “time in grade.”

With such a fair competition mechanism that exports high-fitness individuals to higher-level competitions, societal systems reduce the prevalence of unfair competition and the

unhealthy dominance or disruption that might otherwise be caused by “over-achieving” individuals.

B. The Fair Competition Principle in Biological Systems

It is somewhat surprising that in “cruel” biological/ecological systems, the fair competition principle also holds in many cases. For example, there are mechanisms that reduce unmatched or unfair competition between young individuals and mature ones. Among mammals, young individuals often compete with their siblings under the supervision of parents, but not directly with other mature individuals, since their parents protect them against them. When the young grow up enough, they leave their parents and join the competition with other mature individuals. Evolution has found the mechanism of parental care to be useful in protecting the young and allowing them to grow up and develop their full potentials. Fair competition seems to be beneficial to the evolution of many species.

III. THE HIERARCHICAL FAIR COMPETITION PARALLEL MODEL

Inspired by the fair competition principle and the hierarchical organization of competition within subpopulations in societal systems, we propose the Hierarchical Fair Competition parallel model (HFC), for genetic algorithms, genetic programming, and other forms of evolutionary computation.

In this model [Fig 3], multiple subpopulations are organized in a hierarchy, in which each subpopulation can only accommodate individuals within a specified range of fitnesses. The entire range of possible fitnesses is spanned by the union of the subpopulations’ ranges. Conceptually, each subpopulation has an admission buffer that has an **admission threshold** determined either initially (fixed) or adaptively. The admission buffer is used to collect qualified candidates, synchronously or asynchronously, from other subpopulations. Each subpopulation also has an **export threshold** (fitness level), defined by the admission threshold of the next higher-level subpopulation. Only individuals whose fitnesses are between the subpopulation’s admission threshold and export threshold are allowed to stay in that subpopulation. Otherwise, they are exported to the appropriate higher-level subpopulation. Exchange of individuals is allowed only in one direction, from lower-fitness subpopulations to higher-fitness subpopulations, but migration is not confined to only the immediately higher level.

Each subpopulation can have the same or different size and running parameters. However, considering that there are often more low-fitness peaks than high-fitness peaks, we tend to allocate larger population sizes or more subpopulations to lower fitness levels, to provide extensive exploration; and we tend to use higher selection pressure in higher-fitness-level subpopulations to ensure efficient exploitation. As it is often easier to make a big fitness jump in a lower level subpopulation, we often end up using larger fitness ranges for

lower level subpopulations, and smaller ranges for high-level subpopulations [Fig. 2], but, of course, that depends on the properties of the fitness landscape being explored. The critical point is that the whole range of possible fitness must be spanned by the union of the ranges of all levels of subpopulations. Of course, the highest-level subpopulation(s) need no export threshold (unbounded above) and the lowest-level subpopulation(s) need no admission threshold (unbounded below).

Exchange of individuals can be conducted synchronously after a certain interval or asynchronously as in

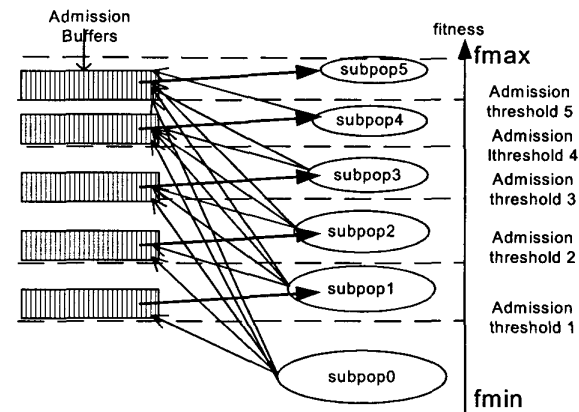


FIGURE 3: In HFC model, subpopulations are organized in a hierarchy with ascending fitness level. Each subpopulation accommodates individuals within a certain fitness range determined by the admission thresholds

many parallel models. At each moment of exchange, each individual in each subpopulation is examined, and if it is outside the fitness range for its subpopulation, it is exported to the admission buffer of a subpopulation with an appropriate fitness range. When a new candidate is inserted into an admission buffer, it can be inserted into a random position or inserted by sorting. After export, each subpopulation imports the appropriate number of qualified candidates from its admission buffer into its pool. Subpopulations (especially at the base level) fill any spaces still open after emptying their admission buffers by generating new individuals at random to fill the spaces left by the exported individuals.

The number of levels in the hierarchy or number of subpopulations (if each level has only one subpopulation) can be determined initially or adaptively. In the static HFC model, we can manually decide into how many levels the fitness range will be divided, the fitness thresholds, and all other GA parameters. In the dynamic HFC model, we can dynamically change the number of levels, number of subpopulations, size of each subpopulation, and admission and export fitness thresholds. The benefit of the dynamic HFC model is that it can adaptively allocate search effort according to the characteristics of the search space of the problem to be solved, thereby searching more efficiently (initial research on

adaptation of admission thresholds is in preparation for reporting elsewhere). However, even “coarse” setting of the parameters in a static HFC model has yielded major improvement in search efficiency over current EA’s on example problems.

Another useful extension to HFC used here is to introduce one or more *sliding* subpopulations, with dynamic admission thresholds that are continually reset to the admission threshold of the level in which the current best individual has been found. Thus, these subpopulations provide additional search in the vicinity of the advancing frontier in the hierarchy.

A. Some Characteristics of the HFC Model

- 1) Exchange of individuals is one-directional, in a hierarchical exchange structure permitting migration from lower-fitness subpopulations to higher-fitness ones. While at each fitness level, the low-fitness offspring of high-fitness parents (created by mutation or recombination) will be replaced by immigrants or eliminated by selection.
- 2) The balance between exploration and exploitation is maintained by the use of subpopulations at different fitness levels. Low-fitness subpopulations tend to explore new search areas and send their promising offspring to higher-fitness subpopulations for exploitation.
- 3) The HFC model maintains a large number of high-fitness individuals without threatening to eliminate lower-fitness (but perhaps promising) individuals. Thus possibly promising new search locales can persist long enough to be appropriately exploited.
- 4) The HFC model quickly captures superior offspring and moves them to a place where they are free to compete with, and be recombined with, each other. This produces an effect similar to the elitism often used in multi-objective evolutionary computation, such as NSGAII or SPEAII (Zitzler *et al.*, 2000), in which superior individuals are also kept separately. At that level, we can control the intensity of selection to determine the tradeoff between exploitation of those high-fitness individuals and exploration in their neighborhoods.
- 5) With multiple subpopulations organized hierarchically according to fitness, the HFC model allows one to look across a multimodal fitness landscape and to see relatively stable regions at each of several fitness levels. Each level may include bands of individuals with similar fitness from many different peaks.
- 6) HFC can be regarded as maintaining several niches in the vertical fitness dimension rather than in a horizontal dimension as in traditional niching techniques. It does not use genotype or phenotype distance to form niches; however, if desired, multiple niches at each level can

also be maintained, through multiple subpopulations at each level or by using crowding, fitness sharing, or other such techniques. The intensity of search at each fitness level can be independently controlled.

IV. EXAMPLE PROBLEMS

The HFC model with Genetic Programming (HFC-GP) has been applied to a real-world analog circuit synthesis problem that was first pursued using GP without HFC [8]. In this problem, an analog circuit is represented by a bond graph model [9] and is composed of inductors (I), resistors (R), capacitors (C), transformers (TF), gyrators (GY), and Sources of Effort (SE). Our task is to synthesize a circuit, including its topology and sizing of components, to achieve specified performance. The objective is to evolve an analog circuit with response properties characterized by a pre-specified set of eigenvalues. By increasing the number of eigenvalues specified, we can define a series of synthesis problems of increasing difficulty, in which premature convergence problems become more and more significant when traditional GP methods are used.

Circuit synthesis by GP is a well-studied problem that generally demands large computational power to achieve good results. Since both topology and the parameters of a circuit affect its performance, it is easy to get stuck in the evolution process. Koza usually uses a population size from 30,000 to 640,000 in his experiments [10].

A. Experiments on an Analog Circuit Synthesis Problem

Four circuits with increasing difficulty are to be synthesized, with eigenvalue sets as specified in Table 1.

Circuits were evolved with single-population GP, multiple-population GP and HFC-GP. The GP parameter for the single population GP is shown in cell (1, 2) of Table 2.

TABLE 1 TARGET EIGENVALUES

Problem 1: 6-eigenvalue problem
$-2 \pm 3.3i, -7.5 \pm 4.5i, -3.5 \pm 12.0i$
Problem 2: 8-eigenvalue problem
$-2 \pm 3.3i, -7.5 \pm 4.5i, -3.5 \pm 12.0i, -3.4 \pm 12.0i$
Problem 3: 10-eigenvalue problem
$-2 \pm 3.3i, -7.5 \pm 4.5i, -3.5 \pm 12.0i, -3.4 \pm 12.0i, -10.0 \pm 8.0i$
Problem 4: 12-eigenvalue problem
$-2 \pm 3.3i, -7.5 \pm 4.5i, -3.5 \pm 12.0i, -3.4 \pm 12.0i, -10.0 \pm 8.0i, -1.5 \pm 3.0i$

Additional parameters for the multi-population GP were shown in cell (2, 2) of Table 2. A one-way ring migration topology was used.

The parameters for HFC-GP were the same with multi-population GP except that the ring migration is replaced with the HFC scheme. For this problem, the range of raw fitness values was [0.5, 1.0], so we defined a fitness admission threshold for each subpopulation (one subpopulation per level, in this case) as shown in cell (3, 2) of

Table 2. Subpopulation 15 was used as a “sliding” subpopulation to aggressively explore the fitness frontier.

First, it is important to notice that these problems exhibit a very high degree of epistasis, as a change in the placement of any pair of eigenvalues has a strong effect on the location of the remaining eigenvalues. Eigenvalue placement is very different from “one-max” or additively decomposable optimization problems, and constitutes an increasingly difficult sequence of problems with the problem order. The performance of each of the three GA approaches was assessed on four problems of increasing difficulty. Each experiment was run ten times, and the average of the results is reported in Fig.4, where three GP methods are indicated by

HFC-GP: Hierarchical Fair Competition Model for GP

OnePop: Single population GP

MulPop: multi-population GP (ring topology)

From Figure 4, it is impressive to see that in all four problems, HFC-GP achieves dramatically better performance vis-à-vis best of run, and that the improvement is more dramatic on the more difficult problems. The superior performance at the initial generations may result from the rapid combination of superior individuals, in a single subpopulation, relative to the ring parallel GA. Yet convergence in the HFC is much slower than in the single- and multi-population GP runs. In fact, we observe relatively steady improvement during the runs for this set of problems.

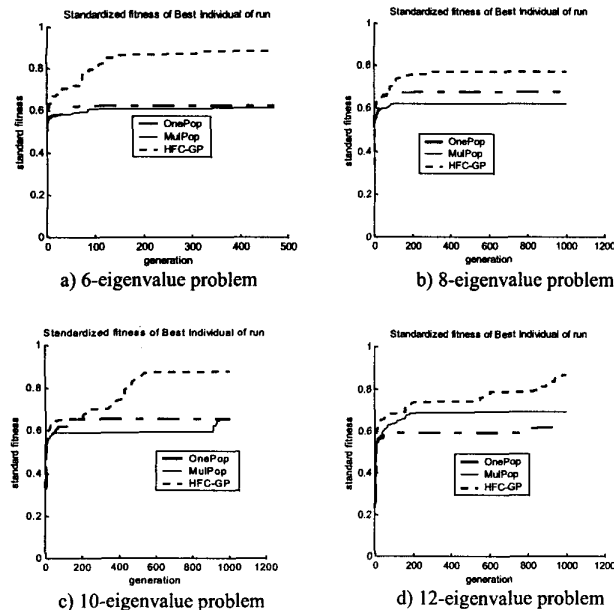


FIGURE 4. Fitness of Best Individual To Date vs. Generation: dashdot (OnePop), solid (MulPop), dashed (HFC)

V. HFC BEHAVIOR AND POTENTIAL EXTENSIONS

A. Observation of HFC's Behavior

- 1) Takeover of subpopulations by high-fitness immigrants

is largely avoided, since only immigrants of the appropriate fitness range are allowed to enter a subpopulation and continuously there are new immigrants coming from lower levels with competitive fitness. If a new immigrant's offspring turn out to be extraordinarily fit, they immediately emigrate to a yet-higher level subpopulation, where they again compete fairly.

- 2) HFC provides another mechanism for maintaining diversity. In a classical GA/GP, a balance must be maintained between the diversity-decreasing effect of selection and the diversity-increasing effects of mutation and some forms of crossover. Addition of HFC can shift this balance to allow maintaining greater diversity with lower rates of mutation and/or crossover. Alternatively, greater selection pressure can be applied without leading to premature convergence, while speeding exploitation of high-fitness individuals.
- 3) In HFC, the tendency of “standard” EA's to narrow their search fairly rapidly to the earliest-discovered regions of relatively high fitness is countered. This allows more thorough exploration around new individuals (usually with low-fitness) that may contain ultimately valuable genetic material that might be discarded by standard EA's. Because low-fitness individuals are not forced out of the lower levels by competition from higher-fitness individuals, they continue to explore the space widely, feeding promising new search regions to higher-fitness subpopulations as they are found.
- 4) HFC often discovers new peaks by building a smooth “path” for potential progenitors of a potential global optimum solution to go up through low-fitness subpopulations to one or more subpopulations of the highest-fitness level. The continual insertion of new random individuals into the lower-level subpopulations reduces the chance that HFC search will get “stuck” at a local optimum and helps it explore new search areas. HFC thus employs a type of multi-start or reinitialization on a continual basis.
- 5) While individuals that are only locally optimal eventually tend to disappear from any given subpopulation, they tend to be explored (and their better offspring to migrate upward) before being eliminated, since their competitors are of similar fitness.

B Potential Extensions of HFC

- 1) Asynchronous export and import of individuals in HFC allow for easy parallelization compared to other diversity maintaining techniques that require calculation of population-level similarity measures, such as fitness-sharing methods.

- 2) The fitness range of each fitness level in the hierarchy can be evolved during the search process. Then we can utilize the problem characteristics to adaptively distribute the search effort (yielding an adaptive HFC model).
- 3) This model is compatible with other existing techniques to improve the search ability of EA's. Existing multi-population and parallel EA models can easily be adapted to the HFC model – only the migration rules and subpopulation topology need to be changed.

VI. CONCLUSIONS AND FUTURE WORK

Based on our analysis of the premature convergence problem in EAs, we believe that the premature convergence problem can to some extent be attributed to unfair (or unbalanced) competition in the selection process of evolution. While competition is essential to the evolutionary process, we observed that in many societal and biological systems, the negative effects of that competition are often modulated by structuring or stratifying it. We therefore developed the Hierarchical Fair Competition parallel model (HFC) for evolutionary algorithms, based on the belief that fair or stratified competition is often beneficial in evolution.

While allowing the sorts of “horizontal” extension of search typically provided by multiple subpopulations and niching methods within subpopulations, the HFC mandates a vertical stratification of search according to fitness, forcing high-fitness individuals generated in any subpopulation to migrate immediately (although asynchronously) to subpopulations containing only individuals of similar fitness levels. This has several beneficial effects – quickly exploiting promising individuals via crossover and mutation while protecting low-fitness individuals from being eliminated before their traits are explored, and dramatically reducing takeover and convergence at all levels.

The HFC model can thus successfully balance exploration and exploitation, while avoiding premature convergence. Experiments on a series of difficult, highly epistatic real-world problems demonstrate the effectiveness of the HFC model in improving significantly both the search speed and the quality of the best solution found.

The authors have already demonstrated (to be reported elsewhere) that the HFC model, with its fair competition principle, is also effective with genetic algorithms, and the extension to any other type of evolutionary algorithm that maintains a population of solutions at any time is clear. It should be useful in any situations where premature convergence is an issue.

Acknowledgment

The first author would like to thank his friends who gave him much inspiration in this work, and whose names, coincidentally, are included in the name of the model (HFC). His previous advisor, Professor Shuchun Wang of Beijing University of Aeronautics & Astronautics of China,

encouraged him to use societal and biological models in AI research. This work was supported by the National Science Foundation under contract DMI 0084934. The authors also acknowledge the assistance of Dr. Kisung Seo, Prof. Ronald C. Rosenberg and Zhun, Fan in defining and exploring the example problem presented here.

References

- [1] K. A. De Jong. An Analysis of the Behavior of a Class of Genetic Adaptive Systems. PhD thesis. University of Michigan. Dissertation Abstracts International 36(10), 5410B. 1975.
- [2] D. E. Goldberg and P. Segrest. “Finite Markov Chain Analysis of Genetic Algorithms,” *Proc. Second International Conf. on Genetic Algorithms*, pp. 1-8, 1987.
- [3] D. E. Goldberg and J. Richardson, “Genetic Algorithms with Sharing for Multimodal Function Optimization,” *Proc. Second International Conf. on Genetic Algorithms*, pp. 41-49, 1987.
- [4] S.-C. Lin, E. Goodman, and W. Punch, “Coarse-Grain Parallel Genetic Algorithms: Categorization and New Approach,” *IEEE Conf. on Parallel and Distrib. Processing*, Nov., 1994.
- [5] D. Eby, R. C. Averill, E. Goodman, and W. Punch, “Optimal Design of Flywheels Using an Injection Island Genetic Algorithm,” *Artificial Intelligence in Engineering Design, Analysis and Manufacturing*, 13, p. 389-402, 1999.
- [6] Holland, J. H., “Cohort Gas and Hyperplane-Defined Functions,” *Evolutionary Computation*, 8(4), pp.372-391, 2000.
- [7] B. Manderick, P. Spiessens, “Fine-Grained Parallel Genetic Algorithms,” in J. D. Schaffer, ed., *Proc. Third Internat. Conf. on Genetic Algorithms*, Morgan Kaufmann, San Mateo, CA, p. 428-433, 1989.
- [8] K. Seo, E. Goodman, and R. Rosenberg, “First Steps toward Automated Design of Mechatronic Systems Using Bond Graphs and Genetic Programming,” *Proc. Genetic and Evolutionary Computation Conf. - 2001*, July 7-11, Morgan Kaufmann Publishers, San Francisco, p. 189, 2001.
- [9] Z. Fan, J. Hu, K. Seo, E. Goodman, R. Rosenberg, and B. Zhang, “Bond Graph Representation and GP for Automated Analog Filter Design,” *Genetic and Evolutionary Computation Conference Late Breaking Papers*, pp. 81-86, 2001.
- [10] J. R. Koza, F. H. Bennett III, D. Andre, and M. A. Keane, *Genetic Programming III: Darwinian Invention and Problem Solving*, Morgan Kaufmann Publishers, San Francisco, 1999.