Evolution of a 60 Decibel Op Amp Using Genetic Programming

Forrest H Bennett III[1]   John R.Koza[2]

David Andre[3]   Martin A. Keane[4]

1) Visiting Scholar, Computer Science Department,
Stanford University, Stanford, California 94305 USA
fhb3@slip.net

2) Computer Science Department, Stanford University

3) Computer Science Department, University of California, Berkeley, California

4) Econometrics Inc., 5733 West Grover
Chicago, IL 60630 USA

Abstract: Genetic programming was used to evolve both the topology and sizing (numerical values) for each component of a low-distortion, low-bias 60 decibel (1000-to-1) amplifier with good frequency generalization.

1. Introduction

In nature, complex structures are designed by means of evolution and natural selection. This suggests the possibility of applying the techniques of evolutionary computation in order to automate the design of complex structures.

The problem of circuit synthesis involves designing an electrical circuit that satisfies user-specified design goals. A complete design of an electrical circuit includes both its topology and the sizing of all its components. The *topology* of a circuit consists of the number of components in the circuit, the type of each component, and a list of all the connections between the components. The *sizing* of a circuit consists of the component value(s) of each component.

Evolvable hardware is one approach to automated circuit synthesis. Early pioneering work in this field includes that of Higuchi, Niwa, Tanaka, Iba, de Garis, and Furuya (1993a, 1993b); Hemmi, Mizoguchi, and Shimohara (1994); Mizoguchi, Hemmi, and Shimohara (1994); and the work presented at the 1995 workshop on evolvable hardware in Lausanne (Sanchez and Tomassini 1996).

The design of analog circuits and mixed analog-digital circuits has not proved to be amenable to automation (Rutenbar 1993). CMOS operational amplifier (op amp) circuits have been designed using a modified version of the genetic algorithm (Kruiskamp 1996; Kruiskamp and Leenaerts 1995); however, the topology of each op amp was one of 24 topologies based on the conventional human-designed stages of an op amp. Thompson (1996) used a genetic algorithm to evolve a frequency discriminator on a Xilinx 6216 reconfigurable processor in analog mode.

2. Genetic Programming

Genetic programming is an extension of the genetic algorithm described in John Holland's pioneering *Adaptation in Natural and Artificial Systems* (1975).

The book *Genetic Programming: On the Programming of Computers by Means of Natural Selection* (Koza 1992) provides evidence that genetic programming can solve, or approximately solve, a variety of problems. See also Koza and Rice 1992. Genetic programming (GP) starts with a primordial ooze of randomly generated computer programs composed of the available programmatic ingredients and then applies the principles of animal husbandry to breed a new (and often improved) population of programs. The breeding is done in a domain-independent way using the Darwinian principle of survival of the fittest, an analog of the naturally-occurring genetic operation of crossover (sexual recombination), and occasional mutation. The crossover operation is designed to create syntactically valid offspring programs (given closure amongst the set of ingredients). Genetic programming combines the expressive high-level symbolic representations of computer programs with the near-optimal efficiency of learning of Holland's genetic algorithm. A computer program that solves (or approximately solves) a given problem often emerges from this process.

The book *Genetic Programming II: Automatic Discovery of Reusable Programs* (Koza 1994a, 1994b) describes a way to evolve multi-part programs consisting of a main program and one or more reusable, parameterized, hierarchically-called subprograms.

Recent work in the field of genetic programming is described in Kinnear 1994, Angeline and Kinnear 1996, Koza, Goldberg, Fogel, and Riolo 1996.

Gruau's *cellular encoding* (1996) is an innovative technique in which genetic programming is used to concurrently evolve the architecture, weights, thresholds, and biases of neurons in a neural network.

3. Evolution of Circuits

Genetic programming can be applied to circuits if a mapping is established between the kind of rooted, point-labeled trees with ordered branches used in genetic programming and the line-labeled cyclic graphs encountered in the world of circuits.

Developmental biology provides the motivation for this mapping. The starting point of the growth process used herein is a very simple embryonic electrical circuit. The embryonic circuit contains certain fixed parts appropriate to the problem at hand and certain wires that are capable of subsequent modification. An electrical circuit is progressively developed by applying the functions in a circuit-constructing program tree to the modifiable wires of the embryonic circuit (and, later, to both the modifiable wires and other components of the successor circuits).

The functions are divided into four categories:

(1) connection-modifying functions that modify the topology of circuit (starting with the embryonic circuit), and

(2) component-creating functions that insert components into the topology of the circuit,

(3) arithmetic-performing functions that appear in arithmetic-performing subtrees as argument(s) to the component-creating functions and specify the numerical value of the component, and

(4) automatically defined functions in function-defining branches.

Each branch of the program tree is created in accordance with a constrained syntactic structure. Branches are composed from construction-continuing subtree(s) that continue the developmental process and arithmetic-performing subtree(s) that determine the numerical value of the component. Connection-modifying functions have one or more construction-continuing subtrees, but no arithmetic-performing subtrees. Component-creating functions have one construction-continuing subtree and typically have one arithmetic-performing subtree. This constrained syntactic structure is preserved by using structure-preserving crossover with point typing (Koza 1994a).

3.1. The Embryonic Circuit

The developmental process for converting a program tree into an electrical circuit begins with an embryonic circuit.

Figure 1 shows a one-input, one-output embryonic circuit that serves as a test harness for the evolving circuits. VSOURCE is the incoming signal. VOUT is the output signal. There is a fixed 1,000 Ohm load resistor RLOAD and a fixed 1,000 Ohm source resistor RSOURCE. Because we are evolving an amplifier, there is also a fixed 1,000,000 Ohm feedback resistor RFEEDBACK, a fixed 1,000 Ohm balancing source resistor RBALANCE_SOURCE, and a fixed 1,000,000 Ohm balancing feedback resistor RBALANCE_FEEDBACK. This arrangement limits the possible amplification of the evolving circuit to the 1000-to-1 ratio (which corresponds to 60 dB) of the feedback resistor to the source resistor.
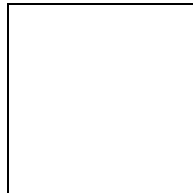

Figure 1 Embryonic circuit.

There are three modifiable wires Z0, Z1, and Z2 arranged in a triangle so as to provide connectivity between the input, the output, and the balancing resistors. All of the above elements (except Z0, Z1, and Z2) are fixed and are not modified during the developmental process. At the beginning of the developmental process, there is a writing head pointing to (highlighting) each of the three modifiable wires. All development occurs at wires or components to which a writing head points.

The top part of this figure shows a portion of an illustrative circuit-constructing program tree. It contains a resistor-creating R function (labeled 2 and described later), a capacitor-creating C function (labeled 3), and a polarity-reversing FLIP function (labeled 4) and a connective LIST function (labeled 1). The R and C functions cause modifiable wires Z0 and Z1 to become a resistor and capacitor, respectively; the FLIP function reverses the polarity of modifiable wire Z2.

## 3.2.	Component-Creating Functions

Each circuit-constructing program tree in the population contains component-creating functions and connection-modifying functions.

Each component-creating function inserts a component into the developing circuit and assigns component value(s) to the component. Each component-creating function has a writing head that points to an associated highlighted component in the developing circuit and modifies the highlighted component in some way. The construction-continuing subtree of each component-creating function points to a successor function or terminal in the circuit-constructing program tree.

Space does not permit a detailed description of each function that we use herein. See Koza, Andre, Bennett, and Keane (1996), and Koza, Bennett, Andre, and Keane (1996a, 1996b, 1996c, 1996d, 1996e) for details.

The two-argument resistor-creating R function causes the highlighted component to be changed into a resistor. The value of the resistor in kilo-Ohms is specified by its arithmetic-performing subtree.

The arithmetic-performing subtree of a component-creating function consists of a composition of arithmetic functions (addition and subtraction) and random constants (in the range –1.000 to +1.000). The arithmetic-performing subtree specifies the numerical value of a component by returning a floating-point value that is, in turn, interpreted, in a logarithmic way, as the value for the component in a range of 10 orders of magnitude (using a unit of measure that is appropriate for the particular type of component involved). The floating-point value is interpreted as the value of the component as described in Koza, Andre, Bennett, and Keane (1996, 1997).

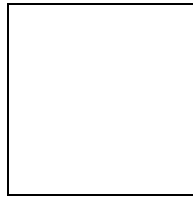Figure 2 shows a modifiable wire Z0 connecting nodes 1 and 2 of a partial circuit containing four capacitors.

Figure 2 Modifiable wire Z0.

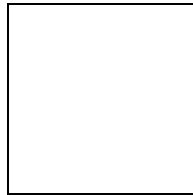Figure 3 shows the result of applying the R function to the modifiable wire Z0 of figure 2.

Figure 3 Result of applying the R function.

Similarly, the two-argument capacitor-creating C function causes the highlighted component to be changed into a capacitor. The value of the capacitor in nano-Farads is specified by its arithmetic-performing subtree.

The one-argument `Q_D_PNP` diode-creating function causes a diode to be inserted in lieu of the highlighted component, where the diode is implemented using a PNP transistor whose collector and base are connected to each other. The `Q_D_NPN` function inserts a diode using an NPN transistor in a similar manner.

There are also six one-argument transistor-creating functions (called `Q_POS_COLL_NPN`, `Q_GND_EMIT_NPN`, `Q_NEG_EMIT_NPN`, `Q_GND_EMIT_PNP`, `Q_POS_EMIT_PNP`, `Q_NEG_COLL_PNP`) that insert a transistor in lieu of the highlighted component. For example, the `Q_POS_COLL_NPN` function inserts a NPN transistor whose collector is connected to the positive voltage source.

The three-argument transistor-creating `Q_3_NPN` function causes an NPN transistor (model `Q2N3904`) to be inserted in place of the highlighted component and one of the nodes to which the highlighted component is connected. The `Q_3_NPN` function creates five new nodes and three new modifiable wires. There is no writing head on the new transistor. Similarly, the three-argument transistor-creating `Q_3_PNP` function causes a PNP transistor (model `Q2N3906`) to be inserted.

Figure 4 shows the result of applying the `Q_3_NPN0` function, thereby creating transistor Q6 in lieu of modifiable wire Z0 of figure 2.
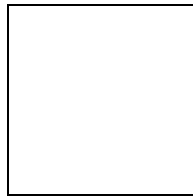

Figure 4 Result of applying the `Q_3_NPN0` function.

3.3. Connection-Modifying Functions

Each connection-modifying function in a program tree points to an associated highlighted component and modifies the topology of the developing circuit.

The one-argument polarity-reversing `FLIP` function attaches the positive end of the highlighted component to the node to which its negative end is currently attached and vice versa. After execution of the `FLIP` function, there is one writing head pointing to the component.

The three-argument `SERIES` division function creates a series composition consisting of the highlighted component, a copy of it, one new modifiable wire, and two new nodes (each with a writing head).

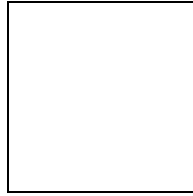Figure 5 illustrates the result of applying the `SERIES` division function to resistor R1 from figure 3.

Figure 5 Result after applying the `SERIES` function.

The four-argument `PSS` and `PSL` parallel division functions create a parallel composition consisting of the original highlighted component, a copy of it, two new wires, and two new nodes. Figure 6 shows the result of applying `PSS` to the resistor `R1` from figure 3.

There are six three-argument functions (called `T_GND_0`, `T_GND_1`, `T_POS_0`, `T_POS_1`, `T_NEG_0`, `T_NEG_1`) that insert two new nodes and two new modifiable wires and make a connection to ground, positive voltage source, or negative voltage source, respectively. Figure 7 shows the results of applying the `T_GND_0` function to the circuit of figure 3.
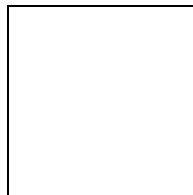

Figure 6 Result of the `PSS` parallel division function.

The three-argument `PAIR_CONNECT_0` and `PAIR_CONNECT_1` functions enable distant parts of a circuit to be connected together. The first `PAIR_CONNECT` to occur in the development of a circuit creates two new wires, two new nodes, and one temporary port. The next `PAIR_CONNECT` to occur (whether `PAIR_CONNECT_0` or `PAIR_CONNECT_1`) creates two new wires and one new node, connects the temporary port to the end of one of these new wires, and then removes the temporary port.

The one-argument `NOP` function has no effect on the highlighted component; however, it delays activity on the developmental path on which it appears in relation to other developmental paths in the overall program tree.

The zero-argument `END` function causes the highlighted component to lose its writing head. The `END` function causes its writing head to be lost – thereby ending that particular developmental path.

The zero-argument `SAFE_CUT` function causes the highlighted component to be removed from the circuit provided that the degree of the nodes at both ends of the highlighted component is three (i.e., no dangling components or wires are created).
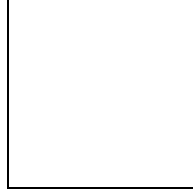
Figure 7 Result of applying the `T_GND_0` function.

4.   Preparatory Steps

Our goal in this paper is to evolve a design for a 60 decibel amplifier with low distortion and low bias. Before applying genetic programming to circuit synthesis, the user must perform seven major preparatory steps, namely

(1) identifying the embryonic circuit that is suitable for the problem,

(2) determining the architecture of the overall circuit-constructing program trees,

(3) identifying the terminals of the to-be-evolved programs,

(4) identifying the primitive functions contained in the to-be-evolved programs,

(5) creating the fitness measure,

(6) choosing certain control parameters (notably population size and the maximum number of generations to be run), and

(7) determining the termination criterion and method of result designation.

The feedback embryo of figure 1 is suitable for this problem.

Since the embryonic circuit has three writing heads – one associated with each of the result-producing branches – there are three result-producing branches (called `RPB0`, `RPB1`, and `RPB2`) in each program tree. We decided to include two one-argument automatically defined functions (called `ADF0` and `ADF1`) in each program tree and that there would be no hierarchical references among the automatically defined functions. Thus, there are two function-defining branches in each program tree. Consequently, the architecture of each overall program tree in the population consists of a total of five branches (two function-defining branches and three result-producing branches) joined by a `LIST` function.

The function sets are identical for all three result-producing branches of the program trees. The terminal sets are identical for all three result-producing branches.

For the three result-producing branches, the function set, $\mathcal{F}_{\text{ccs-rpb}}$, for each construction-continuing subtree is

$\mathcal{F}_{\text{ccs-rpb}}$ = {`ADF0`, `ADF1`, `R`, `C`, `SERIES`, `PSS`, `PSL`, `FLIP`, `NOP`, `T_GND_0`,
  `T_GND_1`, `T_POS_0`, `T_POS_1`, `T_NEG_0`, `T_NEG_1`,
  `PAIR_CONNECT_0`, `PAIR_CONNECT_1`, `Q_D_NPN`, `Q_D_PNP`,
  `Q_3_NPN0`, ..., `Q_3_NPN11`, `Q_3_PNP0`, ..., `Q_3_PNP11`,
  `Q_POS_COLL_NPN`, `Q_GND_EMIT_NPN`, `Q_NEG_EMIT_NPN`,
  `Q_GND_EMIT_PNP`, `Q_POS_EMIT_PNP`, `Q_NEG_COLL_PNP`}.

For the three result-producing branches, the terminal set, $\mathcal{T}_{\text{ccs-rpb}}$, for each construction-continuing subtree consists of

$\mathcal{T}_{\text{ccs-rpb}}$ = {`END`, `SAFE_CUT`}.

For the three result-producing branches, the function set, $\mathcal{F}_{\text{aps-rpb}}$, for each arithmetic-performing subtree is,

$\mathcal{F}_{\text{aps-rpb}} = \{+, -\}$.

For the three result-producing branches, the terminal set, $\mathcal{T}_{\text{aps-rpb}}$, for each arithmetic-performing subtree consists of

$\mathcal{T}_{\text{aps-rpb}} = \{\leftarrow\}$,

where $\leftarrow$ represents floating-point random constants from -1.0 to $+1.0$.

The terminal sets are identical for both function-defining branches (automatically defined functions) of the program trees. The function sets are identical for both function-defining branches.

For the two function-defining branches, the function set, $\mathcal{F}_{\text{ccs-adf}}$, for each construction-continuing subtree is

$\mathcal{F}_{\text{ccs-adf}} = \mathcal{F}_{\text{ccs-rpb}} - \{\texttt{ADF0}, \texttt{ADF1}\}$.

For the two function-defining branches, the terminal set, $\mathcal{T}_{\text{ccs-adf}}$, for each construction-continuing subtree is

$\mathcal{T}_{\text{ccs-adf}} = \mathcal{T}_{\text{ccs-rpb}}$.

For the two function-defining branches, the function set, $\mathcal{F}_{\text{aps-adf}}$, for each arithmetic-performing subtree is,

$\mathcal{F}_{\text{aps-adf}} = \mathcal{F}_{\text{aps-rpb}} = \{+, -\}$.

For the two function-defining branches, the terminal set, $\mathcal{T}_{\text{aps-adf}}$, for each arithmetic-performing subtree consists of

$\mathcal{T}_{\text{aps-adf}} = \{\leftarrow\} \approx \{\texttt{ARG0}\}$,

where $\texttt{ARG0}$ is the dummy variable (formal parameter) of the automatically defined function.

The evaluation of fitness for each individual circuit-constructing program tree in the population begins with its execution. This execution applies the functions in the program tree to the embryonic circuit, thereby developing the embryonic circuit into a fully developed circuit. A netlist describing the fully developed circuit is then created. The netlist identifies each component of the circuit, the nodes to which that component is connected, and the value of that component. Each circuit is then simulated to determine its behavior. The 217,000-line SPICE (Simulation Program with Integrated Circuit Emphasis) simulation program (Quarles et al. 1994) was modified to run as a submodule within the genetic programming system.

An amplifier can be viewed in terms of its response to a DC input. An ideal inverting amplifier circuit would receive a DC input, invert it, and multiply it by the amplification factor. A circuit is flawed to the extent that it does not achieve the desired amplification; to the extent that the output signal is not centered on 0 volts (i. e., it has a bias); and to the extent that the DC response of the circuit is not linear.

Thus, for this problem, we used a fitness measure based on SPICE's DC sweep. The DC sweep analysis measures the DC response of the circuit at several different DC input voltages. The circuits were analyzed with a 5 point DC sweep ranging from -10 millvolts to $+10$ MV, with input points at -10 MV, -5 MV, 0 MV, $+5$ MV,

and +10 MV. SPICE then simulated the circuit's behavior for each of these five DC voltages.

Fitness is then calculated from four penalties derived from these five DC output values. Fitness is the sum of the amplification penalty, the bias penalty, and the two non-linearity penalties.

First, the amplification factor of the circuit is measured by the slope of the straight line between the output for –10 MV and the output for +10 MV (i.e., between the outputs for the endpoints of the DC sweep). If the amplification factor is less than the target (which is 60 dB for this problem), there is a penalty equal to the shortfall in amplification.

Second, the bias is computed using the DC output associated with a DC input of 0 volts. There is a penalty equal to the bias times a weight. For this problem, a weight of 0.1 is used.

Finally, the linearity is measured by the deviation between the slope of each of two shorter lines and the overall amplification factor of the circuit. The first shorter line segment connects the output value associated with an input of –10 MV and the output value for –5 MV. The second shorter line segment connects the output value for +5 MV and the output for +10 MV. There is a penalty for each of these shorter line segments equal to the absolute value of the difference in slope between the respective shorter line segment and the overall amplification factor of the circuit.

Many of the circuits that are randomly created in the initial random population and many that are created by the crossover and mutation operations are so bizarre that they cannot be simulated by SPICE. Such circuits are assigned a high penalty value of fitness (10^8).

The population size, M, was 640,000. The crossover percentage was 89%; the reproduction percentage was 10%; and the mutation percentage was 1%. A maximum size of 300 points was established for each of the branches in each overall program. The other minor parameters were the default values in Koza 1994a (appendix D).

This problem was run on a medium-grained parallel Parsytec computer system consisting of 64 80 MHz Power PC 601 processors arranged in a toroidal mesh with a host PC Pentium type computer. The distributed genetic algorithm was used with a population size of $Q = 10,000$ at each of the $D = 64$ demes. On each generation, four boatloads of emigrants, each consisting of $B = 2\%$ (the migration rate) of the node's subpopulation (selected on the basis of fitness) were dispatched to each of the four toroidally adjacent processing nodes. See Andre and Koza 1996 for details.

5.   Results

The best circuit (figure 11) from generation 0 achieves a fitness of 986.1 and has nine transistors, three capacitors, and two resistors (in addition to the five resistors of the feedback embryo).

About 45% of the circuits of generation 0 cannot be simulated by SPICE. However, the percentage of unsimulatable programs drops to 8% by generation 1. Moreover, this percentage does not exceed 2% after generation 16 and does not exceed 1% after generation 58.

The best circuit (figure 12) from generation 49 achieves a fitness of 404.0 and has 16 transistors, no capacitors, and four resistors (in addition to the five resistors of the feedback embryo).

The best circuit (figure 13) from generation 109 has 22 transistors, no capacitors, and 11 resistors (in addition to the five resistors of the feedback embryo). It achieves a fitness of 0.178. Its circuit-constructing program tree has 40, 98, and 27 points, respectively, in its first, second, and third result-producing branches and 33 and 144 points, respectively, in its first and second automatically defined functions.

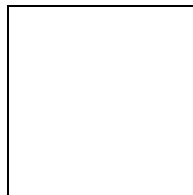**Figure 8 DC sweep of best circuit from generation 109.**

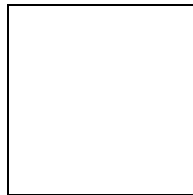**Figure 9 Time domain behavior of best of generation 109.**
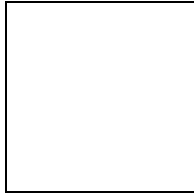
Figure 10 AC sweep for best circuit from generation109.

Figure 11 Best circuit from generation 0.
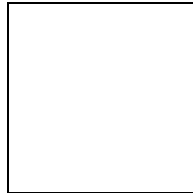

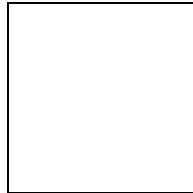**Figure 12 Best circuit from generation 49.**

Figure 13 Best circuit from generation 109.

The amplification of an op amp can be measured from the DC sweep (figure 8). The amplification in decibels is 20 times the common logarithm of the ratio of the change in the output divided by the change of the input (i.e., 20 millivolts here) . The amplification is 60 dB here (i.e., 1,000-to-1 ratio).   There is a bias of 0.2 volts. Notice the linearity of the DC sweep in this figure.

Figure 9 shows the time domain behavior of the best circuit from generation 109. The vertical axis shows voltage and ranges from –10 volts to +10 volts. The input is the 10 millivolt sinusoidal signal; however, this sinusoidal input signal appears as a nearly straight line because of the scale. At 1,000 Hz, the amplification is 59.7 dB; the bias is 0.18 volts; and the distortion is very low (0.17%).

The amplification of an op amp can also be measured from an AC sweep. Figure 10 shows that the amplification for the best circuit of generation 109. The amplification

at 1,000 Hz is 59.7 dB. The flatband gain is 60 dB and the 3 dB bandwidth is 79, 333 Hz.

We then tested whether the genetically evolved 22-transistor best circuit from generation 109 provided more than 60 dB amplification when embedded in a test harness that is appropriate for testing amplification of up to 80 dB. The amplification is 80.15 dB when SPICE's DC sweep is applied. The amplification is 77.79 dB at 1,000 Hz in the time domain. When the AC sweep is applied, the circuit delivers over 80 dB of amplification from 1 Hz to 36 Hz; the circuit delivers over 77.86 dB from 36 Hz to 55,000 Hz.

6. Other Examples of Evolutionary Circuit Design Using Genetic Programming

The above techniques have recently been successfully applied to a variety of other problems of evolutionary circuit design.

6.1. Lowpass "Brick Wall" Filter

Genetic programming successfully evolved a design for a lowpass filter with passband below 1,000 Hz and a stopband above 2,000 Hz with requirements equivalent to that of a fifth order elliptic filter (Koza, Bennett, Andre, and Keane 1996a, 1996c). In some runs of this problem, the genetically evolved lowpass filter has a topology that is similar to that employed by human engineers. For example, in one run, a 100% compliant evolved circuit (figure 14) had the recognizable ladder topology of a Butterworth or Chebychev filter (i.e., a composition of series inductors horizontally with capacitors as vertical shunts). In another run, a 100%-compliant recognizable "bridged T" filter was evolved.
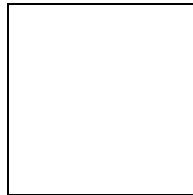

Figure 14 Genetically evolved ladder filter circuit.

6.2. Asymmetric Bandpass Filter

A difficult-to-design asymmetric bandpass filter with requirement equivalent to a tenth-order elliptic filter (Koza, Bennett, Andre, and Keane 1996d) was evolved. Figure 15 shows a 100% compliant evolved asymmetric bandpass filter.
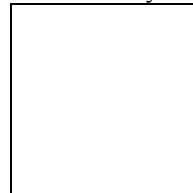

Figure 15 Genetically evolved asymmetric bandpass filter.

6.3. A Crossover (Woofer and Tweeter) Filter

The design for a crossover (woofer and tweeter) filter with a crossover frequency of 2,512 Hz (Koza, Bennett, Andre, and Keane 1996b) was evolved. This problem requires a one-input, two-output embryonic circuit and requires that the fitness be measured at two probe points. The lowpass part of the genetically evolved best-of-run circuit (figure 16) has the Butterworth topology. Except for additional capacitor C36, the highpass part of this circuit also has the Butterworth topology. This circuit is slightly better than the combination of lowpass and highpass Butterworth filters of order 7.
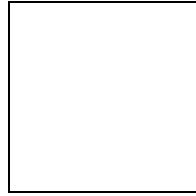
Figure 16 Genetically evolved crossover filter.

7. Conclusion

Genetic programming successfully evolved a 22-transistor amplifier that delivers a DC gain of 60 dB amplification and that has almost no bias or distortion. It generalizes in the frequency domain with a 3 dB bandwidth of 79,433 Hz. Moreover, this genetically evolved 60 dB amplifier generalizes in such a way as to deliver 80 dB of amplification (as measured by the DC sweep) when it is embedded in a test harness that allows 80 dB of amplification.

7.1. Related Paper in this Volume

See also Koza, Bennett, Andre, and Keane (1996e) in this volume.

Acknowledgments

References

Andre, David and Koza, John R. 1996. Parallel genetic programming: A scalable implementation using the transputer architecture. In Angeline, P. J. and Kinnear, K. E. Jr. (editors). 1996. *Advances in Genetic Programming 2*. Cambridge: MIT Press.

Angeline, Peter J. and Kinnear, Kenneth E. Jr. (editors). 1996. *Advances in Genetic Programming 2*. Cambridge, MA: The MIT Press.

Gruau, Frederic. 1996. Artificial cellular development in optimization and compilation. In Sanchez, Eduardo and Tomassini, Marco (editors). 1996. *Towards Evolvable Hardware*. Lecture Notes in Computer Science, Volume 1062. Berlin: Springer-Verlag. Pages 48 – 75.

Hemmi, Hitoshi, Mizoguchi, Jun'ichi, and Shimohara, Katsunori. 1994. Development and evolution of hardware behaviors. In Brooks, R. and Maes, P. (editors). *Artificial Life IV:* Cambridge, MA: MIT Press. Pages 371–376.

Higuchi, Tetsuya, Niwa, Tatsuya, Tanaka, Toshio, Iba, Hitoshi, de Garis, Hugo, and Furuya, Tatsumi. 1993a. In Meyer, Jean-Arcady, Roitblat, Herbert L. and Wilson, Stewart W. (editors). *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: The MIT Press. 1993. Pages 417 – 424.

Higuchi, Tetsuya, Niwa, Tatsuya, Tanaka, Toshio, Iba, Hitoshi, de Garis, Hugo, and Furuya, Tatsumi. 1993b. *Evolvable Hardware – Genetic-Based Generation of Electric Circuitry at Gate and Hardware Description Language (HDL) Levels*. Electrotechnical Laboratory technical report 93-4. Tsukuba, Japan: Electrotechnical Laboratory.

Holland, John H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.

Kinnear, Kenneth E. Jr. (editor). 1994. *Advances in Genetic Programming*. Cambridge, MA: The MIT Press.

Koza, John R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.

Koza, John R. 1994a. *Genetic Programming II: Automatic Discovery of Reusable Programs*. Cambridge, MA: MIT Press.

Koza, John R. 1994b. *Genetic Programming II Videotape: The Next Generation*. Cambridge, MA: MIT Press.

Koza, John R., Andre, David, Bennett III, Forrest H, and Keane, Martin A. 1996. Use of automatically defined functions and architecture-altering operations in automated circuit synthesis using genetic programming. In Koza, John R., Goldberg, David E., Fogel, David B., and Riolo, Rick L. (editors). 1996. *Genetic Programming 1996: Proceedings of the First Annual Conference, July 28-31, 1996, Stanford University*. Cambridge, MA: The MIT Press.

Koza, John R., Bennett III, Forrest H, Andre, David, and Keane, Martin A. 1996a. Toward evolution of electronic animals using genetic programming. *Artificial*

Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems. Cambridge, MA: The MIT Press.

Koza, John R., Bennett III, Forrest H, Andre, David, and Keane, Martin A. 1996b. Four problems for which a computer program evolved by genetic programming is competitive with human performance. *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation*. IEEE Press. Pages 1–10.

Koza, John R., Bennett III, Forrest H, Andre, David, and Keane, Martin A. 1996c. Automated design of both the topology and sizing of analog electrical circuits using genetic programming. In Gero, John S. and Sudweeks, Fay (editors). *Artificial Intelligence in Design '96*. Dordrecht: Kluwer. Pages 151-170.

Koza, John R., Bennett III, Forrest H, Andre, David, and Keane, Martin A. 1996d. Automated WYWIWYG design of both the topology and component values of analog electrical circuits using genetic programming. In Koza, John R., Goldberg, David E., Fogel, David B., and Riolo, Rick L. (editors). 1996. *Genetic Programming 1996: Proceedings of the First Annual Conference, July 28-31, 1996, Stanford University*. Cambridge, MA: The MIT Press.

Koza, John R., Bennett III, Forrest H, Andre, David, and Keane, Martin A. 1996e. Reuse, parameterized reuse, and hierarchical reuse of substructures in evolving electrical circuits using genetic programming. In this volume.

Koza, John R., Goldberg, David E., Fogel, David B., and Riolo, Rick L. (editors). 1996. *Genetic Programming 1996: Proceedings of the First Annual Conference, July 28-31, 1996, Stanford University*. Cambridge, MA: The MIT Press.

Koza, John R., and Rice, James P. 1992. *Genetic Programming: The Movie*. Cambridge, MA: MIT Press.

Kruiskamp, Marinum Wilhelmus. 1996. *Analog Design Automation using Genetic Algorithms and Polytopes*. Eindhoven, The Netherlands: Data Library Technische Universiteit Eindhoven.

Kruiskamp Marinum Wilhelmus and Leenaerts, Domine. 1995. DARWIN: CMOS opamp synthesis by means of a genetic algorithm. *Proceedings of the 32nd Design Automation Conference*. New York, NY: Association for Computing Machinery. Pages 433–438.

Mizoguchi, Junichi, Hemmi, Hitoshi, and Shimohara, Katsunori. 1994. Production genetic algorithms for automated hardware design through an evolutionary process. *Proceedings of the First IEEE Conference on Evolutionary Computation.* IEEE Press. Vol. I. 661-664.

Quarles, Thomas, Newton, A. R., Pederson, D. O., and Sangiovanni-Vincentelli, A. 1994. *SPICE 3 Version 3F5 User's Manual*. Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA. March 1994.

Rutenbar, R. A. 1993. Analog design automation: Where are we? Where are we going? *Proceedings of the l5th IEEE CICC*. New York: IEEE. 13.1.1-13.1.8.

Sanchez, Eduardo and Tomassini, Marco (editors). 1996.*Towards Evolvable Hardware*. Lecture Notes in Computer Science, Vol 1062. Berlin: Springer-Verlag.

Thompson, Adrian. 1996. Silicon evolution. In Koza, John R., Goldberg, David E., Fogel, David B., and Riolo, Rick L. (editors). 1996. *Genetic Programming*

1996: Proceedings of the First Annual Conference, July 28-31, 1996, Stanford University. Cambridge, MA: MIT Press.

Version 2 – Camera-Ready – Submitted October 28, 1996 for the proceedings of the first International Conference on Evolvable Systems (ICES-96) held in Tsukuba on October 7 – 8, 1996.

Evolution of a 60 Decibel Op Amp Using Genetic Programming

Forrest H Bennett III[1]

John R.Koza[2]

David Andre[3]

Martin A. Keane[4]