

Evolving Sorting Networks using Genetic Programming and the Rapidly Reconfigurable Xilinx 6216 Field-Programmable Gate Array

John R. Koza
Computer Science Dept.
Stanford University
Stanford, California 94305-9020
koza@cs.stanford.edu
<http://www-cs-faculty.stanford.edu/~koza/>

Forrest H Bennett III
Visiting Scholar
Computer Science Dept.
Stanford University
Stanford, California 94305
forrest@evolute.com

Jeffrey L. Hutchings
Convergent Design, L.L.C.
3221 E. Hollyhock Hill
Salt Lake City, UT 84121
hutch@Convergent-Design.com

Stephen L. Bade
Convergent Design, L.L.C.
379 North, 900 East
Orem, UT, 84097
bade@Convergent-Design.com

Martin A. Keane
Martin Keane Inc.
5733 West Grover
Chicago, Illinois 60630
makeane@ix.netcom.com

David Andre
Computer Science Division
University of California
Berkeley, California
dandre@cs.berkeley.edu

ABSTRACT

This paper describes how the massive parallelism of the rapidly reconfigurable Xilinx XC6216 FPGA (in conjunction with Virtual Computing Corporation's H.O.T. Works board) can be exploited to accelerate the computationally burdensome fitness measurement task of genetic algorithms and genetic programming. This acceleration is accomplished by *embodying* each individual of the evolving population *into hardware* in order to perform this time-consuming fitness measurement task. A 16-step sorting network for seven items was evolved that has two fewer steps than the sorting network described in the 1962 O'Connor and Nelson patent on sorting networks (and the same number of steps as a 7-sorter that was devised by Floyd and Knuth subsequent to the patent and that is now known to be minimal).

1. Introduction

Field-programmable gate arrays (FPGAs) are often used to facilitate rapid prototyping of new electronic products – particularly those for which time-to-market considerations preclude the design and fabrication of a custom application-specific integrated circuit.

Genetic programming (GP) is an extension to the genetic algorithm (Holland 1975) that automatically creates a computer program to solve a problem using a simulated evolutionary process (Koza 1992, 1994a, 1994b; Koza and Rice 1992). Genetic programming successively transforms a *population* of individual computer programs, each with an associated value of *fitness*, into a new population of *individuals* (i.e., a new *generation*), using the Darwinian principle of survival and *reproduction* of

the fittest and analogs of naturally occurring genetic operations such as *crossover* (sexual recombination) and *mutation*. Additional information about genetic programming can be found in books (Banzhaf, Nordin, Keller, and Francone 1997), edited collections of papers (Kinnear 1994, Angeline and Kinnear 1996), conference proceedings (Koza et al. 1996, 1997), and www.genetic-programming.org.

The dominant component of the computational burden of solving non-trivial problems with the genetic algorithm or genetic programming is the task of measuring the fitness of each individual in each generation of the evolving population. (Relatively little computer time is expended on other tasks of the algorithm, such as the creation of the initial random population at the beginning of the run and the execution of the genetic operations during the run). In a run of the genetic algorithm or genetic programming, the population may contain thousands or even millions of individuals and the algorithm may be run for dozens, hundreds, or thousands of generations. Moreover, the measurement of fitness for just one individual in just one generation typically involves exposing the individual program to hundreds or thousands of different combinations of inputs (called *fitness cases*). Executing one individual program for just one fitness case may, in turn, entail hundreds or thousands of steps.

Field-programmable gate arrays are massively parallel computational devices. Once an FPGA is configured, its thousands of logical function units operate in parallel at the chip's clock rate. The advent of rapidly reconfigurable field-programmable gate arrays (FPGAs) and the idea of evolvable hardware (Higuchi et al. 1993; Sanchez and Tomassini 1996; Higuchi 1997; Thompson 1996) opens the possibility of *embodying* each individual of an evolving population *into hardware*. Since the fitness measurement task residing in the inner loop of genetic algorithm or

genetic programming constitutes the main component of the computational burden of a run, the question arises as to whether the massive parallelism of FPGAs can be used to accelerate this time-consuming task.

This alluring possibility cannot, in practice, be realized with previously available FPGAs for four reasons.

First, the encoding schemes for the configuration bits of almost all commercially available FPGAs are complex and kept confidential by the FPGA manufacturers.

Second, the tasks of technology mapping, placement, routing, and creation of the configuration bits, consume so much time as to preclude practical use of an FPGA in the inner loop of the genetic algorithm or genetic programming. Even if these four tasks could be reduced from the usual hours or minutes to as little as 10 seconds for each individual in the population, these four tasks would consume 10^6 seconds (278 hours) in a run of the genetic algorithm or genetic programming involving a population as minuscule as 1,000 for as short as 100 generations.

Third, the 500 milliseconds typically required for the task of downloading the configuration bits to the FPGA (brief and insignificant for an engineer who has spent hours, days, or months on a single prototype design) would consume 14 hours for even a minuscule population of 1,000 that was run for as few as 100 generations. Again, a run involving a population of 1,000,000 individuals would multiply this already unacceptably long time (14 hours) by 1,000. What's worse – both of these unacceptably long times (278 hours and 14 hours) are merely *preliminary* to the time required by the FPGA for the actual problem-specific fitness measurement. Thus, there is a discrepancy of *numerous orders of magnitude* between the time required for the technology mapping, placement, routing, bit creation, and downloading tasks and the time available for these preliminaries in the inner loop of a practical run of the genetic algorithm or genetic programming. Reconfigurability is not enough for practical work with genetic algorithms and genetic programming. *Rapid reconfigurability* is what is needed – where "rapid" means times ranging between microseconds to milliseconds for *all five* preliminary tasks (technology mapping, placement, routing, bit creation, and downloading).

Fourth, the genetic algorithm starts with an initial population of randomly created individuals and uses probabilistic operations to breed new candidate individuals. These randomly created individuals typically do not conform to the design principles unconsciously employed by humans and are often quite bizarre. Most commercially available FPGAs are vulnerable to damage caused by combinations of configuration bits that connect contending digital signals to the same line. The process of verifying the acceptability of genetically created combinations of configuration bits is complex and would be prohibitively slow in the inner loop of the genetic algorithm or genetic programming. Invulnerability (or

near invulnerability) to damage is needed in order to make FPGAs practical.

3. The Xilinx XC6216 FPGA

The new Xilinx XC6200 series of rapidly reconfigurable field-programmable gate arrays addresses the four issues of (1) openness, (2) rapid technology mapping, placement, routing, and creation of the configuration bits, (3) rapid downloading of configuration bits, and (4) invulnerability to damage. This chip opens the possibility of exploiting the massive parallelism of FPGAs in the inner loop of the genetic algorithm and genetic programming.

The Xilinx XC6216 chip contains a 64×64 two-dimensional array of identical cells (Xilinx 1997). Each of the chip's 4,096 cells contains several multiplexers and a flip-flop and is capable of implementing all two-argument Boolean functions. The functionality and local routing of each cell is controlled by 24 configuration bits whose meaning is simple, straightforward, and public. The configuration bits of the XC6216 can be randomly accessed, and the memory containing the configuration bits is directly memory-mapped onto the address space of the host processor. It is not necessary to download 100% of the configuration bits in order to change only one bit.

Also, the Xilinx XC6216 FPGA is designed so that no combination of configuration bits for cells can cause internal contention (i.e., conflicting 1 and 0 signals simultaneously driving a destination) and potential damage of the chip. This feature is especially important when the configuration bits are being created by an evolutionary process such as genetic programming. Specifically, it is not possible for two or more signal sources to ever simultaneously drive a routing line or input node of a cell. This is accomplished by obtaining the driving signal for each routing line and each input node from a single multiplexer. Thus, only a single driving signal can be selected regardless of the choice of configuration bits. In contrast, in most other FPGAs, the driving signal is selected by multiple independently programmable interface points (pips).

A H.O.T. Works (Hardware Object Technology) expansion board for PC type computers is available from Virtual Computer Corporation (www.vcc.com). The board contains the Xilinx XC6216, SRAM memory, a programmable oscillator that establishes a suitable clock rate for operating the XC6216, and a PCI interface.

4. Problems Suitable for Genetic Programming and FPGAs

The new Xilinx XC6216 rapidly reconfigurable field-programmable gate array addresses several of the obstacles to using FPGAs for the fitness measurement task of genetic algorithms. First, the XC6216 accelerates the downloading task because the configuration bits are in the address space of the host processor. Second, the encoding

scheme for the configuration bits is public. Third, the encoding scheme for the configuration bits is simple in comparison to most other FPGAs thereby potentially significantly accelerating the technology mapping, placement, routing, and bit creation tasks. This simplicity is critical because these tasks are so time-consuming as to preclude use of conventional CAD tools to create the configuration bits in the inner loop of a genetic algorithm.

The above positive features of the XC6216 must be considered in light of several important negative factors affecting all FPGAs. First, the clock rate (established by a programmable oscillator) at which an FPGA actually operates is often much slower (typically around ten-fold) than that of contemporary microprocessor chips. Second, the operations that can be performed by the logical function units of an FPGA are extremely primitive in comparison to the 32-bit arithmetic operations that can be performed by contemporary microprocessor chips. However, the above negative factors may, in turn, be counterbalanced by the fact that the FPGA's logical function units operate in parallel. The existing XC6216 chip has 4,096 cells.

5. Minimal Sorting Networks

A sorting network is an algorithm for sorting items consisting of a sequence of comparison-exchange operations that are executed in a fixed order. Figure 1 shows a sorting network for four items.



Figure 1 Minimal sorting network for 4 items.

The to-be-sorted items, A_1, A_2, A_3, A_4 , start at the left on the horizontal lines. A vertical line connecting horizontal line i and j indicates that items i and j are to be compared and exchanged, if necessary, so that the larger of the two is on the bottom. In this figure, the first step causes A_1 and A_2 to be exchanged if $A_2 < A_1$. This step and the next three steps cause the largest and smallest items to be routed down and up, respectively. The fifth step ensures that the remaining two items end up in the correct order. The correctly sorted output appears at the right. A five-step network is known to be minimal for four items.

There is considerable interest in sorting networks with a minimum number of comparison-exchange operations. There has been a lively search over the years for smaller sorting networks (Knuth 1973). In U. S. patent 3,029,413, O'Connor and Nelson (1962) described sorting networks for 4, 5, 6, 7, and 8 items using 5, 9, 12, 18, and 19 comparison-exchange operations, respectively.

During the 1960s, Floyd and Knuth devised a 16-step seven-sorter and proved it to be a minimal seven-sorter. They also proved that the four other sorting networks in the 1962 O'Connor and Nelson patent were minimal.

The 16-sorter has received considerable attention. In 1962, Bose and Nelson devised a 65-step 16-sorter. In 1964, Batchner and Knuth presented a 63-step 16-sorter. In 1969, Shapiro discovered a 62-step 16-sorter and, in the same year, Green discovered one with 60 steps.

Hillis (1990, 1992) used a genetic algorithm to evolve 16-sorters with 65 and 61 steps. In this work, Hillis incorporated the first 32 steps of Green's 60-step 16-sorter as a fixed beginning for all sorters (Juille 1995).

Juille (1995) used an evolutionary algorithm to evolve a 13-sorter with 45 steps thereby improving on the 13-sorter with 46 steps presented in Knuth (1973). Juille (1997) has also evolved networks for sorting 14, 15, and 16 items having the same number of steps (i.e., 51, 56, and 60, respectively) as reported in Knuth (1973).

As the number of items to be sorted increases, construction of a minimal sorting network becomes increasingly difficult. In addition, verification of the validity of a network (through analysis, instead of exhaustive enumeration) grows in difficulty as the number of items to be sorted increases. A sorting network can be exhaustively tested for validity by testing all $n!$ permutations of n distinct numbers. However, thanks to the "zero-one principle" (Knuth 1973, page 224), if a sorting network for n items correctly sorts n bits into non-decreasing order (i.e., all the 0's ahead of all the 1's) for all 2^n sequences of n bits, it necessarily will correctly sort any set of n distinct numbers into non-decreasing order. Thus, it is sufficient to test a putative 16-sorter against only $2^{16} = 65,536$ combinations of binary inputs, instead of all $16! \sim 2 \times 10^{13}$ inputs. Nonetheless, in spite of this "zero-one principle," testing a putative 16-sorter consisting of around 60 steps on 65,536 different 16-bit input vectors is a formidable amount of computation when it appears in the inner loop of a genetic algorithm.

6. Preparatory Steps for Genetic Programming

Before applying genetic programming to a problem, the user must perform six major preparatory steps, namely (1) identifying the terminals, (2) identifying the primitive functions, (3) creating the fitness measure, (4) choosing control parameters, (5) setting the termination criterion and method of result designation, and (6) determining the architecture of the program trees in the population.

For the problem of evolving a sorting network for 16 items, the terminal set, \mathcal{T} , is

$$\mathcal{T} = \{D1, \dots, D16, \text{NOOP}\}.$$

Here NOOP is the zero-argument "No Operation" function.

The function set, \mathcal{F} , is

$$\mathcal{F} = \{\text{COMPARE-EXCHANGE}, \text{PROG2}, \text{PROG3}, \text{PROG4}\}.$$

Each individual in the population consists of a constrained syntactic structure composed of primitive functions from the function set, \mathcal{F} , and terminals from the

terminal set, \mathcal{T} such that the root of each program tree is a PROG2, PROG3, or PROG4; each argument to PROG2, PROG3, and PROG4 must be a NOOP or a function from \mathcal{F} ; and both arguments to every COMPARE-EXCHANGE function must be from \mathcal{T} (but not NOOP).

The PROG2, PROG3, and PROG4 functions respectively evaluate each of their two, three, or four arguments sequentially.

The two-argument COMPARE-EXCHANGE function changes the order of the to-be-sorted bits. The result of executing a (COMPARE-EXCHANGE i j) is that the bit currently in position i of the vector is compared with the bit currently in position j of the vector. If the first bit is greater than the second bit, the two bits are exchanged. Table 1 shows the two results R_i and produced by (COMPARE-EXCHANGE i j). Column R_i is the Boolean AND function and column R_j is the OR function.

Table 1 The COMPARE-EXCHANGE function.

Two Arguments		Two Results	
A_i	A_j	R_i	R_j
0	0	0	0
0	1	0	1
1	0	0	1
1	1	1	1

The fitness of each individual program in the population is based on the correctness of its sorting of $2^{16} = 65,536$ fitness cases consisting of all possible vectors of 16 bits. If, after an individual program is executed on a particular fitness case, all the 1's appear below all the 0's, the program has correctly sorted that particular fitness case.

Because our goal is to evolve small (and preferably minimal) sorting networks, we ignore exchanges where $i = j$ and exchanges that are identical to the previous exchange. Moreover, during the depth-first execution of a program tree, only the first $C_{max} = 65$ COMPARE-EXCHANGE functions (i.e., five more steps than in Green's 60-step 16-sorter) in a program are actually executed (thereby relegating the remainder of the program to be unused code).

The fitness measure for this problem is multi-objective in that it involves both the correctness and size of the sorting network. Standardized fitness is defined in a lexical fashion to be the number of fitness cases (0 to 2^{16}) for which the sort is performed incorrectly plus 0.01 times the number (1 to C_{max}) of COMPARE-EXCHANGE functions that are actually executed.

The population size was 1,000. The percentage of genetic operations on each generation was 89% one-offspring crossovers, 10% reproductions, and 1% mutations. The maximum size, H_{rpb} , for the result-producing branch was 300 points. The other parameters

for controlling the runs were the default values specified in Koza 1994a (appendix D).

7. Mapping the Problem onto the Chip

The problem of evolving sorting networks was run on a host PC Pentium type computer with a Virtual Computer Corporation "HOT Works" PCI board containing a Xilinx XC6216 field-programmable gate array. This combination permits the field-programmable gate array to be advantageously used for the computationally burdensome fitness measurement task while permitting the general-purpose host computer to perform all the other tasks.

In this arrangement, the host PC begins the run by creating the initial random population (with the XC6216 waiting). Then, for generation 0 (and each succeeding generation), the PC creates the necessary configuration bits to enable the XC6216 to measure the fitness of the first individual program in the population (with the XC6216 waiting). Thereafter, the XC6216 measures the fitness of one individual. Note that the PC can simultaneously prepare the configuration bits for the next individual in the population and poll to see if the XC6216 is finished. After the fitness of all individuals in the current generation of the population is measured, the genetic operations (reproduction, crossover, and mutation) are performed (with the XC6216 waiting). This arrangement is beneficial because the computational burden of creating the initial random population and of performing the genetic operations is comparatively small.

The clock rate at which a field-programmable gate array can be run on a problem is considerably slower than that of a contemporary serial microprocessor (e.g., Pentium or PowerPC) that might run a software version of the same problem. Thus, in order to advantageously use the Xilinx XC6216 field-programmable gate array, it is necessary to find a mapping of the fitness measurement task onto the XC6216 that exploits at least some of the massive parallelism of the 4,096 cells of the XC6216.

Figure 2 shows our placement on 32 horizontal rows and 64 vertical columns of the XC6216 chip of eight major computational elements (labeled A through H). Broadly, fitness cases are created in area B, are sorted in areas C, D, and E, and are evaluated in F and G. The figure does not show the ring of input-output blocks on the periphery of the chip that surround the 64×64 area of cells or the physical input-output pins that connect the chip to the outside. The figure does not reflect the fact that two such 32×64 areas operate in parallel on the same chip.

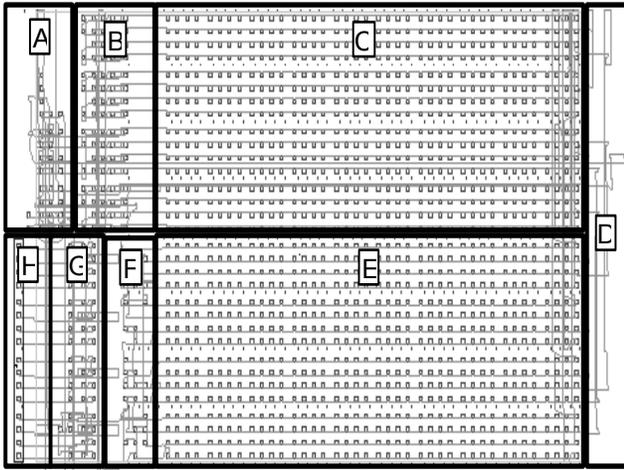


Figure 2 Arrangement of elements A through H on a 32×64 portion of the XC6216 chip.

For a k -sorter ($k \leq 16$), a 16-bit counter B (near the upper left corner of the chip) counts down from $2^k - 2$ to 0 under control of control logic A (upper left corner). The vector of k bits resident in counter B on a given time step represents one fitness case of the sorting network problem. The vector of bits from counter B is fed into the first (leftmost) 16×1 vertical column of cells of the large 16×40 area C. Counter B is an example of a task that is easily performed on a conventional serial microprocessor, but which occupies considerable space (but does not consume not considerable time) on the FPGA.

Each 16×1 vertical column of cells in C (and each cell in similar area E) corresponds to one COMPARE-EXCHANGE operation of an individual candidate sorting network. The vector of 16 bits produced by the 40th (rightmost) sorting step of area C then proceeds to area D.

Area D is a U-turn area that channels the vector of 16 bits from the rightmost column of area C into the first (rightmost) column of the large 16×40 area E.

The final output from area E is checked by answer logic F for whether the individual candidate sorting network has correctly rearranged the original incoming vector of bits so that all the 0's are above all the 1's. The 16-bit accumulator G is incremented by one if the bits are correctly sorted. Note that the 16 bits of accumulator G are sufficient for tallying the number of correctly sorted fitness cases because the host computer starts counter B at $2^k - 2$, thereby skipping the uninteresting fitness case of consisting of all 1's (which cannot be incorrectly sorted by any network). The final value of raw fitness is reported in 16-bit register H after all the $2^k - 2$ fitness cases have been processed.

The logical function units and interconnection resources of areas A, B, D, F, G, and H are permanently

configured to handle the sorting network problem for $k \leq 16$.

The two large areas, C and E, together represent the individual candidate sorting network. The configuration of the logical function units and interconnection resources of the 1,280 cells in areas C and E become personalized to the current individual candidate sorting network.

For area C, each cell in a 16×1 vertical column is configured in one of three main ways. First, the logical function unit of exactly one of the 16 cells is configured as a two-argument Boolean AND function (corresponding to result R_i of table 1). Second, the logical function unit of exactly one other cell is configured as a two-argument Boolean OR function (corresponding to result R_j of table 1). Bits i and j become sorted into the correct order by virtue of the fact that the single AND cell in each 16×1 vertical column always appears above the single OR cell. Third, the logical function units of 14 of the 16 cells are configured as "pass through" cells that horizontally pass their input from one vertical column to the next.

For area E, each cell in a 16×1 vertical column is configured in one of three similar main ways.

There are four subtypes each of AND and OR cells and four types of "pass through" cells. Half of these subtypes are required because all the cells in area E differ in chirality (handedness) from those in area C in that they receive their input from their right and deliver output to their left.

If the sorting network has fewer than 80 COMPARE-EXCHANGE operations, 16 "pass through" cells are placed in the last few vertical columns of area E. The genetic operations are constrained so as to not produce networks with more than 80 steps and, as previously mentioned, only the first $C_{max} < 80$ steps are actually executed.

Within each cell of areas C and E, the one-bit output of the cell's logical function unit is stored into a flip-flop. The contents of the 16 flip-flops in one vertical column become the inputs to the next vertical column. The overall arrangement operates as an 87-stage pipeline (the 80 stages of areas C and E, the three stages of answer logic F, and four stages of padding at both ends of C and E).

Figure 3 shows the bottom six cells of an illustrative vertical column from area C whose purpose is to implement a (COMPARE-EXCHANGE 2 5) operation. As can be seen, cell 2 (second from top of the figure) is configured as a two-argument Boolean AND function (*) and cell 5 is configured as a two-argument OR function (+). All the remaining 14 cells of the vertical column (of which only four are shown in this abbreviated figure) are "pass through" cells. These "pass through" cells horizontally convey the bit in the previous vertical column to the next vertical column. Every cell in the Xilinx XC6216 has the additional capacity of being able to convey one signal in each direction as a "fly over" signal that plays no role in the cell's own computation. Thus, the

of performing two disparate operations (AND and OR) in parallel in physically different areas of the FPGA (i.e., different locations in vertical columns of areas C and E).

Fifth, numerous operations are performed in parallel in control logic A, counter B, answer logic F, accumulator G, and reporting register H. Answer logic F of the FPGA is especially advantageous because numerous sequential steps on a conventional serial microprocessor to determine whether k bits are properly sorted. Answer logic F is an example of a multi-step task that is both successfully parallelized and pipelined on the FPGA.

Sixth, most importantly, the 87-step pipeline (80 steps for areas C and E and 7 steps for answer logic F and accumulator G) enables 87 fitness cases to be processed in parallel in the pipeline.

8. Results

A 16-step 7-sorter (figure 4) was evolved that has two fewer steps than the sorting network described in the 1962 O'Connor and Nelson patent on sorting networks. This genetically evolved 7-sorter has the same number of steps as the 7-sorter that was devised by Floyd and Knuth, subsequent to the patent and has been proven to be minimal (Knuth 1973).

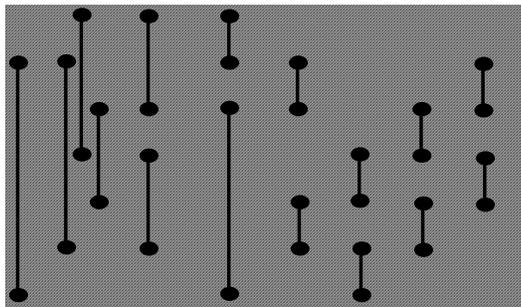


Figure 4 Genetically evolved 7-sorter. Using a population size of 60,000, a 19-step 8-sorter was evolved on generation 58. Using a population size of 100,000, a 25-step 9-sorter was evolved on generation 105. Both of these number of steps are known to be minimal for these two additional problems (Knuth 1973).

Acknowledgments

Phillip Freidin of Silicon Spice provided invaluable information concerning FPGAs and helpful comments on this paper. Stefan Ludwig of DEC and Steve Casselman and John Schewel of Virtual Computer Corporation provided assistance concerning operation of the XC6216.

References

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

Banzhaf, Wolfgang, Nordin, Peter, Keller, Robert E., and Francone, Frank D. 1997. *Genetic Programming – An*

Introduction. San Francisco, CA: Morgan Kaufmann and Heidelberg: dpunkt.

Higuchi, Tetsuya, Niwa, Tatsuya, Tanaka, Toshio, Iba, Hitoshi, de Garis, Hugo, and Furuya, Tatsumi. 1993. 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 (editor). 1997. *Proceedings of International Conference on Evolvable Systems: From Biology to Hardware (ICES-96)*. Lecture Notes in Computer Science. Volume 1259. Berlin: Springer-Verlag.

Hillis, W. Daniel. 1990. Co-evolving parasites improve simulated evolution as an optimization procedure. In Forrest, Stephanie (editor). *Emergent Computation: Self-Organizing, Collective, and Cooperative Computing Networks*. Cambridge, MA: The MIT Press.

Hillis, W. Daniel. 1992. Co-evolving parasites improve simulated evolution as an optimization procedure. In Langton, Christopher, Taylor, Charles, Farmer, J. Dooyne, and Rasmussen, Steen (editors). *Artificial Life II, SFI Studies in the Sciences of Complexity*. Volume X. Redwood City, CA: Addison-Wesley. Pages 313-324.

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

Juille, Hugues. 1995. Evolution of non-deterministic incremental algorithms as a new approach for search in state spaces. In Eshelman, L. J. (editor). *Proceedings of the Sixth International Conference on Genetic Algorithms*. San Francisco, CA: Morgan Kaufmann. 351 – 358.

Juille, Hugues. 1997. Personal communication.

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

Knuth, Donald E. 1973. *The Art of Computer Programming*. Volume 3. Reading, MA: Addison-Wesley.

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., Deb, Kalyanmoy, Dorigo, Marco, Fogel, David B., Garzon, Max, Iba, Hitoshi, and Riolo, Rick L. (editors). 1997. *Genetic Programming 1997: Proceedings of the Second Annual Conference, July 13–16, 1997, Stanford University*. San Francisco, CA: Morgan Kaufmann.

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.
- Koza, John R., and Rice, James P. 1992. *Genetic Programming: The Movie.* Cambridge, MA: MIT Press.
- O'Connor, Daniel G. and Nelson, Raymond J. 1962. *Sorting System with N-Line Sorting Switch.* United States Patent number 3,029,413. Issued April 10, 1962.
- Sanchez, Eduardo and Tomassini, Marco (editors). 1996. *Towards Evolvable Hardware.* Lecture Notes in Computer Science, Volume 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.
- Xilinx. 1997. *XC6000 Field Programmable Gate Arrays: Advance Product Information.* January 9, 1997. Version 1.8.

Camera-Ready Copy - Submitted November 3, 1997, to the to be held in Pacific Grove, CA on november 2 – 5 (Sunday – Wednesday) , 1997.

Evolving Sorting Networks using Genetic Programming and the Rapidly Reconfigurable Xilinx 6216 Field-Programmable Gate Array

John R. Koza
Computer Science Dept.
Stanford University
Stanford, California 94305-9020
koza@cs.stanford.edu
<http://www-cs-faculty.stanford.edu/~koza/>

Forrest H Bennett III
Visiting Scholar
Computer Science Dept.
Stanford University
Stanford, California 94305
forrest@evolute.com

Jeffrey L. Hutchings
Convergent Design, L.L.C.
3221 E. Hollyhock Hill
Salt Lake City, UT 84121
hutch@Convergent-Design.com

Stephen L. Bade
Convergent Design, L.L.C.
379 North, 900 East
Orem, UT, 84097
bade@Convergent-Design.com

Martin A. Keane
Martin Keane Inc.
5733 West Grover
Chicago, Illinois 60630
makeane@ix.netcom.com

David Andre
Computer Science Division
University of California
Berkeley, California
dandre@cs.berkeley.edu