

Code Generator Amelioration Using Genetic Algorithm Techniques

P. Cockshott and Y.Gdura

Abstract—Genetic algorithms (GAs) are based on techniques inspired by some aspects of natural science such as inheritance, reproduction and mutation, and they are used as optimization technique for searching large solution spaces. In computer science, for example, they could be used in data sorting and searching, circuit design and to improve application performance the quality of designed tools such as code generation. This paper looks at the possibility of using genetic algorithms to ameliorate the automatic construction of code generators. Experimental evidence is provided that the use of such algorithms can improve the quality of automatically constructed code generators.

Index Terms—Genetic algorithms, compilers, code-selection, permutation problem.

I. INTRODUCTION

Mass produced human artifacts undergo evolution in a manner analogous to biological lineages. A mass produced article like a car or a processor chip is the materialisation of a pre-existing data structure produced in the design offices of its manufacturing company. Firms retain elements of previous designs in new products, with the design being modified with successive product releases. This process generates product lineages that, just as with biological ones, allow the reconstruction of an ancestry. The theoretical foundations of genetic algorithms (GAs) are based on techniques inspired by some aspects of natural science such as inheritance, reproduction and mutation.

An implementation of a standard GA starts by encoding a feasible solution to a given problem on a simple data structure; called chromosome, and apply genetic operators on a set of potential solutions (population) in order to select better solutions and reproduce new individuals (offspring). Consequently, the new population is expected to be better than the old population.

Genetic algorithms could be used as tools to evolve a solution for several types of problems [1]-[7]. In computer science, GAs could be used to predict performance of a given application under certain circumstances or to improve the quality of designed tools such as code generation. For example, a genetic algorithm approach was used to model the performance of memory-bound computations on different architecture [5]. GAs has been also used to improve generated code's execution time or its size [1], [2].

In the case of computer processors the most fundamental

part of the inherited design is the instruction set, and it is this that we can view as the inherited genome. The process of evolution of the Intel x86 processor families, for example, is a history of genetic accretion from the genetic code Ur-microprocessor the 4004, through the 8008, 8080, 8086 etc, down to modern chips like the SandyBridge [8]. We even see processes analogous to the formation of the Eukaryota when formerly free living organisms were ingested and incorporated as organelles: mitochondria and chloroplasts. The incorporation of these organelles led to the Eukaryota having dual genomes - nuclear and mitochondrial DNA for animals. In the Intel/AMD/Transmeta/VIA etc lineage the equivalent to the Eukaryotic Revolution was the ingestion of the formerly free living x87 floating point processor into the 486. Since then the processor lineage has incorporated two distinct genomes for floating point and integer code. Further events analogous to chromosome duplication and subsequent specialisation led to the generation of sub-families inheriting the MMX, 3D Now, SSE and AVX instruction sets [5]-[11]. Younger processor lineages like the IBM Power architecture have also evolved, perhaps less ornately, but show a similar process.

Each of these evolutionary events replicated certain basic functions and structures: register sets, register load and store instructions and arithmetic operations. A consequence is that current processors allow you to perform a given calculation in many different ways using instructions that evolved at different times.

Consider the simple operation that we might write in a high level language as $a = a+1$ and how this might be performed on an x86 lineage machine.

If we used the most primitive style instructions descended from the 8080 we might code it as:

```
mov eax,[a]
add eax,1
mov [a],eax
```

or as

```
mov eax,[a]
inc eax
mov [a],eax
```

Adding the instructions from the 8086 vintage to the options we might try

```
mov eax,[a]
lea eax,[eax+1]
mov [a],eax
or simply
inc dword[a]
```

But then what about using the 'Eukaryotic' FPU instructions, which, after the 486 were always there:

```
fild dword[a]
fldl
```

Manuscript received June 10, 2013; revised August 17, 2013.

P. Cockshott is with the School of Computer Science, University of Glasgow, USA (e-mail: wpc@dcs.gla.ac.uk).

Y. Gdura was with Computer Engineering Department, University of Tripoli, USA (e-mail: ygadura@tripoliuniv.edu.ly).

```
faddp
fistp dword[a]
```

With the invention of the MMX and SSE instructions there are even more possibilities, for example we could use the xmm registers and generate the following sequence:

```
movd xmm0,[a]
movd xmm1,[one]
padd xmm0,xmm1
movd [a],xmm0
section .data
one: dd 1
```

Given such a plethora of mechanisms by which even such a simple calculation could be performed, how is the code generator of a compiler to select between them?

To an extent this question that can be avoided, since, like bees and flowers, compilers have co-evolved with processors. They provided techniques for generating code for integer operations using the older instruction set before the new alternatives came along, and are likely to retain these primitive code patterns even after more alternatives became available. If one is writing an entirely new code generation system though, the problem of selection between a vast range of semantically equivalent code sequences strikes you afresh.

However, if the processor manufacturers provided detailed timings for each instruction, as used to be done on early generations of microprocessors [9], this would be easy but more recent processor manuals [10] no longer provide these timings. One can infer a number of reasons for this:

- The instructions are common to several chips with different internal structure which may have different numbers of clocks per instruction.
- The timings will vary with the degree of super scalar execution.
- They will vary with the degree of contention for execution units imposed by other instructions.

In the absence of reliable instruction times, one can attempt to improve code selection based on other criteria, for example the number of instructions used to achieve a semantic effect, or the number of memory transfers scheduled by the instructions. But in the simple example given above, memory transfer counts are no help in distinguishing between the several options, nor do instruction counts give an unambiguous answer. Even if we were to favour the memory increment instruction on the grounds that it was the shortest, there is no guarantee that it would run faster than the first alternative since complex instructions like `inc dword [a]` will be broken down by the instruction decoder into a sequence of simpler micro-operations. The micro operations might well be the same as those executed by the explicit load, add, store sequences.

II. EXPERIMENTAL CONFIGURATION

At the University of Glasgow we have a code generator system that supports a significant spread of processor models, a range of CPUs in the x86 family as well as the chips used on the Playstation 2 and Playstation 3. The code generators are automatically written in Java by a compiler which uses as input machine descriptions in the ILCG's notation ([11] appendix A), [12]. The code generators employ a unification

based technique similar to Prolog [13] and the overall approach is functionally motivated by the system described in [14].

Unification was used in Prolog systems to construct logical proofs, and in our approach to code generation, the instructions available on machine M are like the axioms of a formal axiomatic system. The generation of machine code for a programme segment S is the construction of a proof from these axioms that S is derivable from the axiomatic system of machine M . A precondition for this to work is that the abstract source programme and the axioms are represented in the same notation. In our case we translate the source programmes into ILCG syntax trees. The 'axioms' of a given machine are its instructions and addressing modes. These are specified as patterns which are unified with the abstract syntax tree of the programme being translated, with successful unification resulting in the output of the corresponding instruction to the assembler file. The unification based pattern matching often involves the recursive elaboration of patterns. As in Prolog, the order in which patterns are matched can affect which of several possible matches will succeed. Nevertheless, the unification algorithm always outputs the first pattern/instruction whose matching succeeds.

Where multiple alternative patterns are possible the algorithm will attempt to match them left to right:

```
pattern riscaddr
means [offset|baseplusoffset|regindir];
```

So in the above pattern which defines a selection of addressing modes, the unification algorithm will attempt to match a sub expression first against the offset addressing mode, then the base plus offset mode *et al.* The individual elements in the list are themselves patterns such as:

```
pattern baseplusoffset(reg r, offset s )
means [( ^ ( r ) , s )]
assembles [ r '+' s ];
```

Consequently, a processor specification will contain hundreds of patterns describing things like sets of registers, addressing modes or instructions. An example of an instruction pattern is:

```
instruction pattern STORELIT(addrmode rm, type t, int s)
means [ ( ref t ) rm := ( t ) const s ]
assembles [ 'mov ' t ' ' rm ' , ' ' s ];
```

The "means" part is what the unification algorithm matches against an abstract syntax tree and in the process unifies the parameters rm , t , s against the tree. The "assembles" part specifies what assembly code is to be produced [11]. In this context the parameters are replaced by the assembly code that was produced by the matching of the sub-patterns.

The matching or proof process starts by attempting to match an abstract syntax tree against a list of all the axioms/patterns that specify the individual machine instructions. This list gives the individual instruction patterns a definite order. This order is the order of preference in which they will be matched to the abstract syntax. By moving an instruction pattern up this list we can cause the code generator to prefer to use that instruction over other alternatives.

The ordering of axioms can thus be crucial both to the

efficiency of the resulting code, and indeed to whether a successful match is obtained at all. Unification of this sort is known to be potentially non-decidable [15]. With certain orderings of the patterns the process of finding a match is potentially non-terminating. We avoid this by running the proof machine as a parallel process and giving it a time quota that is linear in the size of the tree for which it is trying to obtain a proof. In the past a considerably amount of human judgement has had to be used to obtain an order that seems likely both to terminate and to produce efficient code. Such human judgement, whilst certainly much better than nothing, can obviously have no guarantee of producing an optimal instruction ordering, given that the search space over which the selection has to be made is so great. With n instruction axioms, there are $n!$ possible orderings or permutations in which they could be listed. It thus seemed an attractive idea to try and automate the ordering of instruction axioms.

III. RELATED WORK

Genetic algorithms have been used in computer science to seek for better solutions to various types of problems such memory bandwidth, code generator.

A genetic algorithm approach was used to model the performance of memory-bound computations [5]. This approach was used to lean bandwidth and memory performance of many memory-intensive applications on different HPC architectures [5]. The genetic algorithm is based on the results of three main benchmarks; STREAM, Apex-MAPS, and MultiMAPS, to learn bandwidth as a function of cache hit rates on different machines and to predicate the achievable bandwidth from cache hit rates.

Previous work using GAs for improving generated code include the Genetic Algorithm Parallelization System (GAPS) [1] which was used for optimizing parallel loop-based FORTRAN codes. Compiling FORTRAN programs to run on SPMD parallel machines usually involves loop transformations such as loop distribution, loop interchange, loop fusion or statement reordering [1]. GAPS was developed to find out the optimal transformation sequence that results in minimising the SPMD execution time of a given program.

Wu and Li [2] also used GAs for improving machine code for a dual instruction set ARM processor. The ARM is heavily used in embedded applications. It supports, in addition to its standard instruction encoding, a more compactly encoded instruction set, called Thumb with shorter bit-lengths than the original instruction set. A program compiled using only Thumb instructions uses more instructions than the same program compiled using the standard instructions set, and it is consequently slower. Because the dual instruction sets could affect the efficiency of compiled programs in term of performance and space, a GA technique was used here to improve the code generator. The genetic algorithm and other tools helped a code generator to swap between the two instructions sets in order to reduce a program's execution time and its code size.

We differ from [1] and [2] in that we are attempting to ameliorate code generators while the other two approaches

enhance the generated code for a single program at a time. By improving the code generator itself we only have to run the genetic algorithm during compiler development to obtain speed-ups in many programmes subsequently translated by the compiler.

IV. GENETIC ALGORITHM DESIGN

Genetic algorithms are a robust technique for searching large spaces on which some optimality criterion is defined [16]. The basic genetic algorithm procedure encodes solutions as a string and then works with mutation and crossover operations to generate new solutions. The population of solutions is repeatedly expanded by these operations and then shrunk by removing less 'fit' examples. A key issue in the application of genetic algorithms to a domain is designing a solution representation amenable to the mutation and crossover operators.

A. Permutation Operation

Permutation-based GAs is used for scheduling problems such as ordering instruction set or the Travelling Salesman Problem [3], [6]. The classical travelling salesman problem is a similar permutation problem to our own, since in both cases solutions can be represented as lists of elements; either cities or operation codes, in which each element must occur once. Assume that the n instruction patterns or cities are each given an integer in the range $0 \dots n-1$. Any solution must be a permutation of these integers which can be represented as a list with each integer presented only once. Using a simple list representation for the genome gives rise to problems with mutation and crossover operators as the result of applying them to the list is no longer necessarily a permutation.

There has been past study of how to encode travelling salesman problems as Genetic Algorithms, a review of such encodings and modified mutation and crossover operators is given in [3].

In our case, we have chosen a novel three level representation for the genome:

An initial list of length $n=2m$ operation codes filled in with blanks if we have less than n operation codes for our processor. This will typically be our best hand generated operation code ordering.

A permutation array p of length n made up of integers in the range $0 \dots n-1$ in which each integer is presented exactly once.

A bit string s of length $n-1$ that is the genome used for selection and breeding.

The genome as a bit string does not itself encode the permutation but is instead a programme for a permutation machine. There is some similarity in this approach in to Chaitin's presentation of mutations as Turing Machine programmes that transform the genome [17].

Since in the process of doing this, programmes may be bred that are non-terminating, and thus we used timeouts to detect potentially uncomputable unifications. This approach also resembles the procedure discussed in [17] where the goal is to use evolutionary programming to breed Turing Machine programmes to evaluate the BusyBeaver function. Chaitin's approach resorts to linear time bounded approximation from

below.

Our permutation machine is less general than the Turing complete mutation machine used by Chaitin. Our permutation machine f reads in the bit string s and an initial valid permutation p and then performs a sequence of valid permutations on p such that $q = f(s, p)$; the output q is another valid permutation. If we start with the identity permutation $I = 0, 1, 2, 3, \dots$, then each permutation programme s labels a permutation ps produced by the application $ps = f(s, I)$.

The semantics of the permutation machine f also ensure that binary crossover between and mutation of permutation programmes will again yield a valid permutation programme. Now, if we assume that the permutation machine f code is as follows:

b00 b10 b11 b20 b21 b22 b23 b30 b31 b32 b33 b34 b35
b36 b37 b40...

The permutation machine then proceeds as follows:

if b00 swap the 1st half of p with the 2nd half of p
if b10 swap the 1st quarter of p with the 2nd quarter of p
if b11 swap the 3rd quarter of p with the 4th quarter of p
if b20 swap the 1st eighth of p with the 2nd eighth of p
.....
etc.

Since the permutation of a permutation is still a permutation, and since swapping is a permutation operator, it is clear that the machine f will always produce a valid permutation of p if p is itself a valid permutation. Note that the first bit has more effect than the second and third, and that these have more effect than succeeding ones. In this aspect the encoding has some similarity to structured genetic algorithms [4] which are known to converge faster than unstructured ones.

It is also evident that the space that can be searched using this representation is of the order 2^n which is less than $n!$ for all $n > 3$. The space searched by the GA based on this genome will only be a fraction of the possible permutation space. Since $n!$ is bounded above by $2(n \log n)$, we could construct a permutation programme of length $n \log_2 n$ bits that would be capable of producing any permutation. How can we do this using our existing permutation function f ?

We need to apply $\log n$ independent permutations on the genome in sequence, and in order to do that we need a new function $g(i, s, p)$ which given an integer $i : 0..(\log_2 n) - 1$ a bit string s as before and a permutation p will generate the permutation $\text{rot}(f(s, p), 2^i)$ that is to say it applies the permutation programme s to p as before and then cyclically rotates the permutation list by 2^i places. Suppose we have a genome of length $n \log n$ with $\log n$ bit strings each of length n .

We will denote the i th of these component bit strings as si . The complete permutation space can be scanned by composing g with itself $\log n$ times as follows:

for $i:=0$ to $(\log_2 n) - 1$
do $p := g(i, si, p)$

However, for realistic numbers of operation codes (of the order of 100..200), 2^n already represents a huge optimisation search space and gives the GA plenty of scope to find improvements, so our initial experiments used a genome of length $2[\log_2 n]$.

B. Crossover

The crossover operation allows a new genome to be constructed that inherits information from two parent genomes. Our algorithm is based on the one-point crossover approach which basically picks a random point on both parent genome bit strings and then swaps the portions of the genomes above and below this point to produce new children as shown in Fig. 1.




Fig. 1. Crossover operation.

C. Mutation

The copying of DNA in organisms is imperfect and introduces random code errors or mutations. In Genetic Algorithms one can emulate this by simply flipping m bits at random in a genome of length N . The ratio m / N should be set low to ensure that the algorithm will not degenerate to a random search. Our ameliorater inverts 2 arbitrary bits from the new child genome.

D. How the Ameliorater Works

The Genetic Algorithm ameliorater goes through the following steps:

- 1) Create an initial population of genomes. This list will include the null permutation (all 0) and the all 1 permutation, with the others being selected at random.
- 2) From each genome s produce a permuted list of operation codes by using $f(s, I)$.
- 3) From this list of operation codes produce a modified processor specification file and invoke the compiler to produce a new code generator.
- 4) Using the new code generator compile and run a set of test programmes. Compute the fitness of the genome as $c = t$ where $c = I$ if the programmes all compiled and 0 otherwise and t is the total run time of the programmes.
- 5) Select the fittest $2/3$ of genomes, apply mutation and crossover to produce a new $1/3$ to replace those discarded and repeat steps 2 to 4.

The mutation rate was set to 2 bit flips per offspring, and the working population was set at 90.

E. Results

We tested the optimizer on two architectures: PowerPC and the IA32 with SSE2 instructions. The front end of the compiler was the same in both cases. The PowerPC generator we used was a sub-set generator for the Cell Broadband Engine in which only the common subset of instructions shared by the Cell and the Power PC were included. Tests were run on a Sony Playstation 3 running Linux.

The optimizer was tested on both machines using the same benchmarks in its fitness function: an n-body simulation, a prime number sieve and a vector kernel programme. The n-body problem, which uses single precision floating point, is a simulation of planets or particles evolving under the influence of gravity. The algorithm starts with an initial

position, mass and velocity of a number of bodies at a given time. It then uses that data to compute based on the laws of motion and gravitation the motions of all bodies and to find their positions at later times. It was taken from the Programming Languages Shootout website. The second application, which uses integer data types, is a prime sieve program. The program finds all the prime numbers that are less than or equal a given integer value n using Eratosthenes' method. The last application is a special purpose program that was developed to test the Vector Pascal compilers [5] on various vector operations such as transpose, reduction, dot product operations etc over different data types.

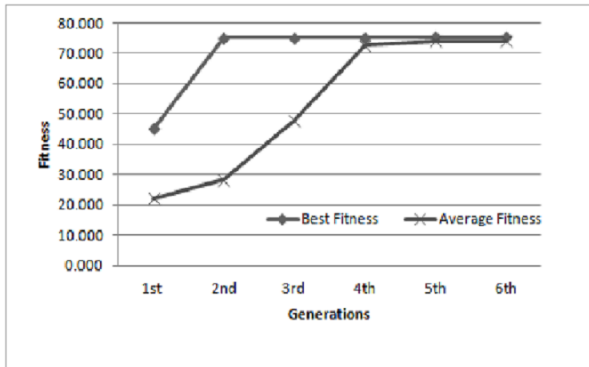


Fig. 2. The results of applying the GA to the hand improved code generator for the PowerPC processor in the Cell-BE.

The performance of a genetic algorithm on optimizing a PowerPC instruction set ordering. The average fitness values reflect the performances of the code generators on three different applications. The code generators are generated using different solutions (IS orderings) offered by different generations. The higher the fitness scores the shorter the average execution time of the selected applications.

To summarise the results, we present the average and the best fitness values of each generation. Note that for genomes which yield successful compilations, the fitness value is proportional to the performance of the final code.

1) *The PowerPC code generator*

The PowerPC machine description includes 184 instructions, and thus the mutation probability here is around 0.01. Fig. 2 shows that the average fitness of the solutions offered by the last generation is 3.4 times better than the solutions provided by the first generation. Fig. 2 shows also that the performances of the PowerPC code generators improved significantly throughout the first three successive generations. The first generation's average fitness value was around 22.2 and at the fourth generation reached approximately 72.6 with an improvement factor of about 3.2. After the fourth generation, the algorithm then got steadier and just a slight improvement were gained in the last two generations. However, the best solution was obtained in the second generation, and it remained the best performance among all generations.

In additional to the three testing programs, which were used in the fitness procedure to evaluate the algorithm, we also tested the optimizer using two other Shootout benchmarks; Spectral-norm and Mandelbrot. Spectral-norm is a program to calculate an eigenvalue using the power method. The final tests were conducted to see if we get the

same performance improvement on applications using the optimized order of the PowerPC instruction set that was produced by our optimizer. Table I compares the performances of these five applications using a default instruction ordering and the optimized ordering.

TABLE I: GENETIC ALGORITHM IMPROVEMENTS ON THE CELL-BE (THE FIRST 3 PROGRAMMES WERE IN THE TRAINING SET WHILE THE LAST TWO WERE IN THE TEST SET)

Program	Performance (Sec)		
	Default	GA Optimized	Improvement
Sieve	26.5	20.6	22%
N-Body	39.2	30.5	22%
Vector Op's	30.6	23.2	24%
Spectral-norm	83.2	63.2	24%
Mandelbrot	24.7	21.1	15%

The results in the Table I show that the optimizer generally behaved about the same on most of the applications. Our interpretation for not getting the same improvement on the Mandelbrot benchmark as on the other applications is that all other applications were run many iterations while Mandelbrot was run only once.

2) *The IA32 code generator*

The IA32 machine instruction set description included 252 instructions. This machine's code generator, unlike the PowerPC's one, had been under development for several years, and its instructions ordering had received a considerable manual tuning. The improvement due to the genetic algorithm is not expected to be as good as on the PowerPC. Fig. 3 shows that the average fitness is increasing slightly during the first three generations and declined and started improving after the fourth generation. As the diagram shows, the best solution was not improved during the six generations. This probably indicates that the instruction set was hand tuned to a close to optimal configuration.

Fig. 4 represents a normalised performance gained by optimizing instruction orderings of both machines. The chart shows that the reordering of the PowerPC's instruction set improved the performance of the PowerPC around 3 times while on Pentium the optimizer reports only a slight improvement because its instruction set was well tuned over several years. However, it is possible that more generations of the genetic algorithm might result in improvements, but they are clearly harder to find.

V. CONCLUSIONS

We have shown that it is possible to use genetic algorithms to produce a generalisable improvement in code generator performance. The improvement in code quality extends beyond the training set. However, the automatic genetic algorithm was unable to produce improvements in another code generator whose code selection rules had been subjected to many years of human optimisation.

The experiments are relatively time consuming in computer time since each run of the fitness function on the genome involves the building of a new code generator in Java and its linking to the front end compiler, followed by the compilation of the test set along with the Pascal run time library with the new code generator. This build time is

considerably greater than the run time of the final test programmes. Despite this, the automatic process is:

- Much faster than the process of hand optimisation that we had previously used.
- Allows code generator optimisation to be done by less experienced team members such as final year students.

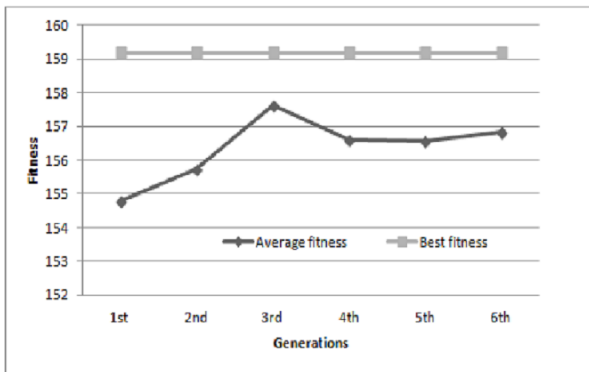


Fig. 3. The results gained by optimizing previously manually optimised code generator for the Pentium using a genetic algorithm.

The average fitness values reflect the performances of the code generators on the three selected applications. The higher the fitness scores the shorter the average execution time of the selected applications.

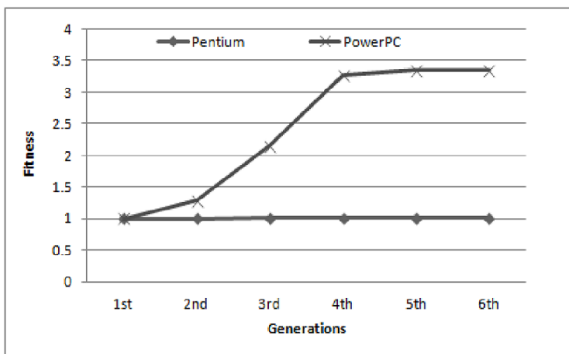


Fig. 4. Normalized performance gained by optimizing instruction set orderings of PowerPC and Pentium 4 architecture using a genetic algorithm.

The average fitness values reflect the performances of the code generators on three selected applications.

In the future we intend to extend the work. We will modify the compiler itself so that it includes the genetic algorithm and can run in a training mode to ameliorate its own instruction selection rules. These rules can then be stored in an auxiliary file that is readable at code generator initialisation. We also intend to investigate the effect of more comprehensive genomes capable of searching over the entire $n!$ permutation space of code rule orderings.

REFERENCES

[1] A. Nisbet, "GAPS, Iterative feedback directed parallelisation using genetic algorithms," in *Proc. the Workshop on Profile and Feedback-Directed Compilation*, 1998.

[2] S. Wu and L. S. Kun, "Instruction selection for arm/thumb processors based on a genetic algorithm coupled with critical event tabu search," *Chinese Journal of Computers* 4, pp. 680-685, 2007.

[3] P. Larranaga, C. Kuijpers, R. Murga, I. Inza, and S. Dizdarevic, "Genetic algorithms for the travelling salesman problem: A review of representations and operators," *Artificial Intelligence Review*, vol. 13, no. 2, pp. 129-170, 1999.

[4] D. Dasgupta and D. McGregor, "Nonstationary function optimization using the structured genetic algorithm," *Parallel Problem Solving from Nature*, vol. 2, pp. 145-154, 1992.

[5] M Tikir, L.vCarrington, E.vStrohmaier, and A. Snavely, "A genetic algorithms approach to modeling the performance of memory-bound computations," in *Proc. the 2007 ACM/IEEE Conference on Supercomputing (SC '07)*, 2007, pp. 1-12.

[6] J. Bryant, "Naive and heuristic permutation-coded genetic algorithms for the quadratic knapsack problem," in *Proc. the Fourteenth International Conference on Genetic and Evolutionary Computation Conference Companion*, 2012, pp. 1465-1466.

[7] G. Ahmadreza, A. Kamilia, and F. Nicholas, "Sorting unsigned permutations by reversals using multi-objective evolutionary algorithms with variable size individuals," *Evolutionary Computation (CEC)*, pp. 292-295, 2011.

[8] List of intel microprocessors. [Online]. Available: http://en.wikipedia.org/wiki/List_of_Intel_microprocessors.

[9] MCS-85 Users Manual, Intel, 1977.

[10] Intel 64 and IA-32 Architectures Software Developer's Manual, Intel, 2010.

[11] P. Cockshott and K. Renfrew, *SIMD programming for Windows and Linux*, Springer, 2004.

[12] P. Cockshott, "The Glasgow pascal compiler," *Lulu*, 2012.

[13] D.Warren and A. Center, "An abstract Prolog instruction set," *Technical Note 309*, SRI International, Artificial Intelligence Center, Menlo Park, CA, October 1983.

[14] S. L. Graham, "Table driven code generation," *IEEE Computer*, vol. 13, no. 8, pp. 25-37, 1980.

[15] G. Huet, "The undecidability of unification in third order logic," *Information and Control*, vol. 22, no. 3, pp. 257-267, 1973.

[16] D. Whitley, "A genetic algorithm tutorial," *Statistics and computing*, vol. 4, no. 2, pp. 65-85, 1994.

[17] G. Chaitin. (2010). A mathematical theory of evolution and biological creativity. *Tech. Rep. CDMTCS-397*. University of Auckland, URL. [Online]. Available: <http://www.cs.auckland.ac.nz/CDMTCS//researchreports/397greg.pdf>.



P. Cockshott was trained at McMaster, Manchester, Heriot Watt and Edinburgh Universities. Originally an economist, he later studied computer science. He worked in industry on hardware verification and on the design of database machines. He has been a researcher or academic at the universities of Edinburgh, Glasgow and Strathclyde and is now a reader in computer science at the University of Glasgow. His past research includes persistent programming, database machines, FPGA architectures, data and video compression, and the theory of economic planning. His current research interests are: SIMD and multi-core compilers, econophysics and the physical foundations of computability.



Y. Gdura was born in Tripoli – Libya in 1961. He had his bachelor degree in Computer Engineering from the Higher Institute of Electronics (HIE) in Libya 1982. He got his master degree in computer science from University of Concordia in Canada in 2001, and he got his PhD from University of Glasgow 2012. He worked in government institutions on software development. He also worked as a general manager of IT and Telecoms Company in the private sector in Libya. His is working right now at the University of Tripoli, Libya.