High performance extendable instruction set computing

Lee, H; Beckett, Paul; Appelbe, William

https://researchrepository.rmit.edu.au/esploro/outputs/9921857680301341/filesAndLinks?institution=61RMIT_INST&index=null

Lee, H., Beckett, P., & Appelbe, W. (2001). High performance extendable instruction set computing. 6th Australasian Computer Systems Architecture Conference, 89–94.
<https://doi.org/10.1109/ACAC.2001.903365>

Published Version: <https://doi.org/10.1109/ACAC.2001.903365>

Repository homepage: <https://researchrepository.rmit.edu.au>

© 2001 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Downloaded On 2023/04/26 18:42:17 +1000

High-Performance Extendable Instruction Set Computing

Heui Lee
Asia Design Corporation
hlee@adc.co.kr

Paul Beckett
RMIT University
pbeckett@rmit.edu.au

Bill Appelbe
RMIT University
bill@cs.rmit.edu.au

Abstract

In this paper, a new architecture called the extendable instruction set computer (EISC) is introduced that addresses the issues of memory size and performance in embedded microprocessor systems. The architecture exhibits an efficient fixed length 16-bit instruction set with short length offset and immediate operands. The offset and immediate operands can be extended to 32 bits via the operation of an extension flag.

The code density of the EISC instruction set and its memory transfer performance is shown to be significantly higher than current architectures making it a suitable candidate for the next generation of embedded computer systems.

The compact EISC instruction set introduces data dependencies that seemingly limit deep pipeline and superscalar implementations. This paper suggests a mechanism by which these dependencies might be removed in hardware.

1. Introduction

Since its development in the 1970s, the microprocessor had been applied in a range of embedded applications, in fields as diverse as home automation and industrial control to PDAs and network computers [1]. The introduction of RISC architectures in the 1980s [2] allowed microprocessor architectures to move into the domain previously occupied by mini-computers. Further, advances in semiconductor technology have seen the development of super scalar architectures [3] as well as significant increases in processor operating speeds [4-7].

However despite all of these architectural advances, program execution still involves accessing program and data in memory. And regardless of either the rapid increase in available memory capacity or the accompanying decrease in access times, the basic performance of memory still does match that of the microprocessor core. For example, in 1980 the access time of DRAM was in the order of 250nsec, and by 1988 its operating speed had increased to 300MHz - 70 times faster. However, in the same period the performance of the microprocessor core increased from 8MHz (e.g. 8086)

to 500MHz (e.g. Pentium-2). Moreover, if the superscalar nature of the Pentium-2 is taken into account, its performance is actually in the order of 1GHz - 120 times faster. Clearly, the performance of microprocessors is now limited by not only the speed difference between memory and the CPU, but also by the physical properties of the bus available to connect the memory and the CPU of the microprocessor (in terms of width and throughput; i.e. its bandwidth) [8].

Further, in embedded microcontroller systems where memory, CPU and I/O circuits tend to be integrated into a single chip, the price of that chip is largely dependent upon its size. As memory circuits (RAM and ROM) take up significant area, the overall cost is especially sensitive to the size of the memory.

In this paper, the architecture of the Extendable Instruction Set Computer (EISC) is presented. The EISC architecture offers a viable solution to the memory size/bandwidth problem by exhibiting an efficient fixed length 16-bit instruction set with short length offset and small immediate operands. The offset and immediate operands can be extended to 32 bits via the operation of an extension flag (the e-flag) resulting in high code density for the 32-bit microprocessor. A solution for avoiding e-flag dependencies in deep pipeline and superscalar configurations is proposed. Comparisons are made with the performance of the MIPS-R3000 to demonstrate the development and operation of the extendable instruction set concept.

2. Embedded Microprocessors

There has been much recent work looking at the issue of bus bandwidth and memory size in embedded processors. One approach has been to add code compression to a variety of architectural styles ([9],[10],[11]). For example, [12] looks at the effect of replacing frequently used sequences of instructions with a code word serving as an index into a list of instruction sequences, while [13] proposes a software method of doing the same thing.

Another approach has been the adaptation of 32-bit RISC architectures to use a compressed 16-bit instruction set. The ARM-7TDMI [14] represents a 16 bit

compressed instruction version of the ARM-7 and the TR4101 is the 16-bit compressed instruction architecture form of the MIPS-R3000 ([1], [14]). These 16-bit compressed-instruction RISCs exhibit complex architectures because of the requirement for compatibility with existing machines. Moreover, the 16-bit instruction set version of these architectures can address only eight registers, further compromising their performance.

3. The Extendable Instruction Set

The specification of the EISC instruction set commenced with an analysis of the performance of a number of instruction sets of existing processors (especially the MIPS R3000) using EGCS-1.1 [15], the C library NEWLIB-1.8.1 [16], the C++ library LIBSTDC++-2.8.1 [17] and various benchmark programs. This analysis indicated some general characteristics of embedded microprocessor systems (some of which are common to all processors). For example:

- The availability of sixteen general-purpose registers is close to optimum;
- Load and Store instructions are used often and mostly employ short length offset addressing;
- The frequency of use of small sized constants is high.

To support these characteristics effectively, the EISC was given 16-bit fixed length instruction and the 32-bit instruction set was designed to extend the offset and constant fields. Code density was further increased by the inclusion of register list 'Push and Pop' instructions plus the use of hardware interlocks to resolve pipeline conflicts (thus eliminating the need to insert NOP instructions).

The basic 32-bit EISC system has been implemented as an FPGA, and all of its functions verified at low speed (1.8432MHz). The following sections outline how the analysis of processors such as the R3000 has resulted in the major architectural features of the EISC.

3.1 The EISC Register Set

The MIPS R3000 has thirty-four 32-bit registers. Two of them are the private registers for 'Multiply and Divide', five of them are the special registers for stack, frame pointer and condition codes and the remaining twenty-seven are the general-purpose registers. To study the effect of register availability on code size, the EGCS C/C++ compiler was used to generate code for the C/C++ library and benchmark programs while the number of available registers was varied.

In Table 1, it can be seen that, as the number of the general-purpose registers becomes smaller, the overall program size gets bigger (note that the program size for the case of 27 registers has been normalized to 100). The

frequency and therefore the space taken by Load and Store instructions also grows and, since these instructions use memory and the bus, they directly influence the required data transfer width. There is little change in either the program size or the load and store frequency as the number general-purpose registers reduces from twenty down to sixteen. However, eight registers are clearly too few as, by that stage, the frequency of load and store instructions has almost doubled. Thus the EISC was set up with access to 16 registers.

No. of Regs	Program size	Load/Store	Move
27	100.00	27.90 %	22.58 %
24	100.35	28.21 %	22.31 %
22	100.51	28.34 %	22.27 %
20	100.56	28.38 %	22.24 %
18	100.97	28.85 %	21.93%
16	101.62	30.22 %	20.47 %
14	103.49	31.84 %	19.28 %
12	104.45	34.31 %	16.39 %
10	109.41	41.02 %	10.96%
8	114.76	44.45 %	8.46 %

Table 1. Program size vs. number of registers for MIPS-R3000.

3.2 Load/Store Architecture

Table 2 shows the average instruction frequency in a MIPS-R3000 with sixteen general-purpose registers available to the CPU. The EISC architecture is RISC-like in that all operations use register operands while the only memory access is via load/store instructions.

Instruction	Frequency
Move	20.27 %
lw, sw	28.27 %
Nop	7.26 %
Addiu	7.53 %
Li	2.93 %
Lui	3.76 %
sh, sb, lh, lb, lhu, lbu	1.98%
bnez, bne, beqz, beq, bltz...	6.69%
j, jal	10.36 %
jr	1.79 %
addu, subu, and, or, xor, nor, negu	3.33 %
andi, ori, xori	2.17 %
jalr	0.17 %
slt, sltu, slti, sltiu	1.70 %
sll, srl, sra, sllv, srlv, srav	1.40%
mult, multu, div, divu	0.09 %
break, mfhi, mflo	0.12 %

Table 2 Instruction frequency of MIPS-R3000 with sixteen general-purpose registers.

It is clear from Table 2 that the decision to employ a load/store architecture will very little impact on the performance of the machine because of the low frequency of occurrence of instructions that might otherwise employ a register-memory architecture (e.g. `addu`, `subu`, and etc.).

3.3 16 bit Fixed Length Extendable Instruction

It can also be seen from Table 2 that the `move` instruction is the most frequent, at 20.27% (the combined frequency of load/store is a little larger, but `move` is the most frequent individual instruction). Since the EISC has access to sixteen registers, it requires four bits each to represent the source and destination registers. Instructions such as `move` therefore fit easily within 16 bits. Using a 16 bit fixed length instruction also simplifies the hardware. Whilst most instructions (such as the `move`, above) can be represented within a 16-bit fixed length structure, instructions with constant operands such as load and store are more problematic due to the length of the offset and constant operands. In particular, 93.5% of the total load and store instruction have 32-bit operands.

A more detailed analysis of the characteristics of the `lw` (load word) and `sw` (store word) instructions (Table 3) revealed that about 61% of these instructions referenced the Stack Pointer and 40% use the Index Register. In the latter case, 77% of those instructions could be represented by 3-bit offset.

Offset length	Stack pointer (60.9%)	Index register (39.1%)
3 bit	43.2 %	77.0 %
4 bit	72.6 %	81.6 %
5 bit	88.1 %	89.6 %
6 bit	90.5 %	91.5 %
7 bit	95.7 %	93.5 %

Table 3 Characteristics of 'lw' and 'sw' instructions

Instructions with constant operands (e.g. `li` - load immediate) do not occur very often: just 2.9% of all instructions in this analysis. In addition, Table 4 illustrates that using an 8-bit constant will cover 93.6 % of these instructions.

Constant range	Frequency
-32 -- +31	72.4 %
-64 -- +63	86.9 %
-128 -- +127	93.6 %
-256 -- +255	95.2 %
> ±256	100 %

Table 4 Operand Size of 'li' instruction

Thus the great majority of instructions use short offset or constant operands. This situation applies equally to instructions such as `lw` (load word) and `sw` (store word) as well as arithmetic instructions such as `addiu`, `slti`, and `sltiu`.

3.4 The Extension Register and Flag

The EISC architecture synthesizes long operands from adjacent instructions using an Extension Register and Flag (E). The E-flag is set when an operand has just been transferred into the 32-bit Extension Register (i.e. `%ER`). The `leri` instruction (Load Extension Register Immediate) performs a conditional shift (controlled by the E-flag) and load to synthesize long immediate operands in the Extension Register, as shown in Figure 1.

Instruction Mnemonics : LERI
Instruction Format : LERI constant
Instruction Representation :
 bit 15-14 = 01
 bit 13-0 = constant data bit 13-0
Operation :
 If (E flag is 0) Load `%ER` with sign extended constant
 ELSE `%ER` = `%ER` << 14 + Constant
 Set E flag

Figure 1. Operation of the LERI instruction

Instruction Function : Load /Store
Instruction Representation :
 bit 15-14 = 00
 bit 13-12, 7 = Operation

000 :	sign extend 8 bit load	LDB	SRC	DST
001 :	sign extend 16 bit load	LDS	SRC	DST
010 :	32 bit load	LD	SRC	DST
011 :	Zero extend 8 bit load	LDBU	SRC	DST
100 :	8 bit store	STB	SRC	DST
101 :	16 bit store	STS	SRC	DST
110 :	32 bit store	ST	SRC	DST
111 :	Zero extend 16 bit load	LDSU	SRC	DST

bit 11-8 = Source /Destination register. `%R0` to `%R15`.
 bit 6-4 = offset bit 2-0 if 8 bit load/store
 = offset bit 3-1 if 16 bit load/store
 = offset bit 4-2 if 32 bit load/store
 bit 3-0 = Index register. `%R0` thru `%R15`.
 Effective operand address : EA
Operation:
 If (E flag is 0)
 EA = Zero extend offset + Index register
 If (E flag is 1)
 if (32/16 bit load/store)
 EA = `%ER` << 4 + Offset + Index register
 if (8 bit load/store)
 EA = `%ER` << 3 + Offset + Index register

Figure 2. Operation of the LD instruction

Load and store instructions use the E-flag to create the effective address or 32-bit constant. For example, Figure 2, illustrates that if the E-flag is set to 1, the effective address is formed from `%ER` << 4 + Offset + Index register.

As the total number of 32-bit load/store instructions is typically small, this mechanism has little effect on overall performance of the machine.

3.5 Stack pointer

As shown in Table 3, load/store instructions that reference the Stack Pointer and the Index Register tend to exhibit very different operand lengths. For Stack Pointer use, the offset needs to be more than 5 bits, while the majority (77%) of the Index Register load/store operations can utilize a 3-bit operand. For this reason, the 32-bit EISC instruction set includes separate load/store stack pointer instructions with 7-bit offset.

The frequency of push and pop to stack was found to be reasonably high in the experiments reported here (15.8%) and if eight registers were bound together, the average number of registers pushed/popped per instruction was about 4.3. This indicated that it was worth including a "push/pop to list" type of instruction. This is more commonly found in CISC rather than RISC machines and, although it reduces the amount of memory used by push/pop instructions, it can cause problems with superscalar and deep pipeline configurations. A simple solution is to disallow the use of multiple pushes or pops within a single instruction in such configurations. Since binary compatibility is not an objective, this solution suffices.

In Table 2 it can be seen that the frequency of `addiu` (Add Immediate) instructions was 7.53%. Within these instructions, the frequency of stack pointer usage is 35% of which 99.1% employed 7-bit constant operands. The 32-bit EISC instruction set therefore includes the capacity for arithmetic on the stack pointer using 7-bit constant operands.

3.6 Other Instructions

The frequency of conditional branch instructions in Table 2 is 6.69%. In the EISC instruction set, the Carry, Sign, Zero and Overflow flags can be combined such that fourteen kinds of Conditional branch instructions can be formed. The offset of conditional instructions is set at 9 bits and extended to 32 bits via the extension register.

It can also be noted that 48.5% of all ALU instructions used two operands and 51.4% used three. However, three operand instructions can be easily synthesized using a combination of the move instruction and a two-operand instruction and this is what is done in the EISC architecture.

Instructions such as multiply and divide exhibit low frequency of use but they are very useful in applications such as multimedia. As well, the performance of these instruction types depends heavily on their method of

implementation. The EISC includes two 32-bit registers (%ML and %MH) to store the results of multiply operations.

The EISC includes provision for a number of co-processors to perform particular functions. Each co-processor has sixteen general-purpose registers. Co-processor '0' is the system co-processor that manages Cache, Pipeline and Memory control etc. Other co-processors have been defined for Floating point and Multimedia acceleration. Co-processor instructions can extend to 20 or 30 bits by use of the extension register.

4. Performance Evaluation

In order to benchmark the relative performance of the EISC architecture, it was compared to a range of existing microprocessors. The metric used was their Relative Code Density (RCD) defined as follows:

$$RCD = \frac{\text{Code Density of 32bit EISC}}{\text{Code Density of Comparison Microprocessor}} = \frac{\text{Program Size of Comparison Microprocessor}}{\text{Program Size of 32bit EISC}}$$

To perform the evaluation, a cross C/C++ compiler was produced for both the 32-bit EISC and the existing microprocessor and the C/C++ library and benchmark test programs were compiled. The resulting RCD figures are shown in Table 5.

Processor	RCD	Processor	RCD
32 bit EISC	1.00	MC88000	1.59
MIPS-R3000	1.66	MC5200 (Coldfire)	1.43
TR4101 (MIPS-16)	1.07	MC68000	1.35
MIPS-R4000	1.46	MC68332	1.32
MIPSTX-39	1.58	MC68020	1.32
ARM-7	1.64	MN10300	1.17
ARM-7TDMI (THUMB)	1.13	I80386	1.39
Power PC 601	1.92	I80960	1.68
SPARC V8	1.38	ARC	2.19
SPARCLITE	1.78	SH-3	1.38
PA-RISC	2.22	V850	1.21
Alpha-RISC	2.23	M32R	1.44

Table 5. Relative Code Density of 32-bit EISC.

As shown in Table 5, in comparison tests between the EISC and existing architectures, the Relative Code Density of the EISC was found to be higher in all cases (i.e. its code size was smaller). For example, the code density of the EISC is 66% higher than of MIPS R3000 (considered the benchmark 32-bit RISC architecture) for the programs chosen. Compared to the popular ARM 7 processor, the RCD figure was 1.64.

Even compared to CISC processors such as MC68000 and I80386 the RCD figure ranged from 1.2 to 1.4 indicating that the program size of existing CISC machines is 20 to 40% bigger than of the EISC.

The closest figures were that of the ARM-7TDMI and TR4101, which are 16-bit compressed RISC machines. However, these architectures have access to only 8 general-purpose registers, increasing the frequency of load/store instructions (see Table 6), which in turn increases the necessary memory bandwidth. For example, the program size of the TR4101 is 7% bigger than that of the EISC and its load/store frequency is 18 % bigger. Thus the TR4101 would require a 25% higher data transfer rate to match the performance of the EISC. In case of the ARM-7TDMI, the equivalent transfer rate figure is 30 % higher.

	32 bit EISC	TR4101	ARM- 7TDMI
RCD	1.00	1.07	1.13
Load / Store	30.2 %	48.4 %	46.5 %

Table 6 Comparison between 32-bit EISC and 16-bit compressed RISC

4.1 E-flag Dependencies

Once the instruction set was defined, the performance of the EISC architecture was investigated using the Dhrystone benchmark [18] running on a micro-architecture simulator for a basic five-level pipeline. This was repeated for various levels of superscalar operation. The micro-architecture simulator was adapted from SuperDLX [19] initially developed at McGill University in Canada and further extended at RMIT University.

Of particular interest was the effect of register dependencies in the EISC architecture as these threatened to impact on its performance in deep pipeline and superscalar configurations. In the EISC, dependency on the E-register and, particularly, the E-flag is a special case that is of particular importance as their use is pervasive throughout the instruction set. The vast majority of instructions in the EISC use the E-flag in some manner (either clearing it, or using it during an effective address calculation).

It was identified early in the analysis that the operation of the E-Flag creates artificial pipeline dependencies that threaten to restrict the operating efficiency of the pipeline. While the overall efficiency of the EISC pipeline was still high (>78% in our tests), it was clear that this figure could be further improved by removing E-flag and procedural dependencies. Further, these dependencies had the potential to prevent the architecture from achieving any meaningful gains from superscalar organization.

Fortunately, there is a fairly straightforward solution to this problem as almost all of the necessary information about the E-flag is known at decode time. This solution is outlined in the following section.

In Figure 3, the arrows show the dependencies operating within the example code fragment. Apart from the true data dependency on r1 at lines 1 and 2 (andi & cmpi), it can be seen that this code fragment exhibits an almost continuous serial dependency on the E-flag.

The effect of this is to severely limit the available parallelism in the code. Each instruction has to complete to the commit stage before the next instruction can be decoded – i.e. in the five stage pipeline of the simulator, the next instruction could not commence until the commit stage had been reached by the previous instructions, forcing at least three stall cycles for each instruction. In our benchmark tests it was found that the immediate instructions (e.g. addi, cmpi) were amongst those most heavily dependent on the E-flag and register, and a very low level of overall instruction parallelism was observed (<2.5% best case).

	Instruction	Read ^x	Write
1.	andi %r3, 0x3, %r1	R3	(R1)
2	cmpi %r1, 0x1	R1	(E)
3.	Ld (%sp, 0x10), %r0	E, R0	SP, E
4.	Ld (%sp, 0x14), %r3	E, R3	SP, E
5.	Ldi 0x0, %r2	E	R2, E
6.	Tsti %r3, 0x1	E, R3	R3, E
7.	Jz 0x8b4 <.L2>	E	E

^o E = E-flag, ER = E-register

^x in the case of the E-flag, 'Read' means that the instruction execution depends on the flag, write means the instruction clears the flag.

Figure 3. Instruction Dependencies in EISC

4.2 "Virtualising" the E-flag

Dependency analysis does not distinguish between a "write" and a "clear" of the E-flag. However, there is a real distinction that could be made by the micro-architecture simulator. A "clear" creates a dependence that can be removed if the architecture "virtualises" the e-flag. After an instruction that clears the e-flag is decoded, we can record that the next instruction that reads the e-flag will read a new, or virtual value of zero. This instruction could then be executed in parallel.

Virtualising the e-register/flag in this way is a simple extension in hardware of existing virtual register schemes, and exposes significant parallelism. In typical implementation schemes, virtual registers are assigned at decode time, with a tag that identifies their status. If an instruction unconditionally clears the e-flag, subsequent instructions (that use the e-flag) reference a virtual flag with the value zero. In the instruction fragment of Figure 3, removing the E-flag dependency, in combination with branch prediction would allow instructions pairs to be forwarded to separate execution pipes and executed in parallel. This type of superscalar operation can clearly make a significant impact on the performance of the EISC architecture.

5. Conclusions

This paper has introduced a new architecture called the extendable instruction set computer (EISC) aimed at the embedded systems market. The overall performance and cost of this class of processors are particularly sensitive to the size and access bandwidth of their memory systems. The EISC has been shown to offer significant improvements over existing architectures in these areas. The use of the Extension Register and Extension Flag allows the architecture to use an efficient fixed length 16-bit instruction set with short length offset and immediate operands. The offset and immediate operands can then be extended to 32 bits via the operation of the E-flag.

Using this mechanism, the EISC architecture has achieved in the order of 140 to 220 % better code density than existing RISC processors and a figure of about 120 to 140 % compared to CISC machines. Even compared to 16-bit compressed instruction RISC machines such as the ARM-7TDMI, the program size of the EISC has been found to be 5 to 15% smaller and the frequency of Load and Store instructions about 15% lower. As a result, the EISC architecture can be seen to be well suited to embedded applications where small code size and low memory transfer width are desirable attributes.

The effect of register dependencies in the EISC architecture was investigated thoroughly - especially those caused by the E-flag and E-register. The operation of the E-Flag in particular has the potential to create artificial dependencies that could restrict the operating efficiency of the pipeline and impact on the overall performance of the architecture in deep pipeline and superscalar configurations. "Virtualising" the e-flag, using a similar mechanism to existing virtual register schemes, will expose significant parallelism in the instruction set and allow the architecture to operate efficiently in superscalar and deep pipeline configurations.

Comparing its code density and memory transfer performance against existing processors, the EISC

architecture can be seen to offer significant advantages that will make it highly suited to a range of embedded computer systems applications.

6. References

- [1] Manfred Schlett, "Trends in Embedded-Microprocessor Design", *IEEE Computer*, pp. 44-50, Aug. 1998.
- [2] D. Patterson, "Reduced Instruction Set Computer", *Comm. ACM*, Vol. 28, No. 1, pp. 8-21, Jan. 1985.
- [3] Dezso Sima et al., "Superscalar Instruction Issue", *IEEE Micro*, pp. 28-39, Oct. 1987.
- [4] B. Gieseke et al., "A 600MHz Superscalar RISC Microprocessor with out-of-order execution", *ISSCC Digest Tech. Papers*, pp. 176-177, Feb. 1997.
- [5] C. A. Maier et al., "A 533MHz BiCMOS Superscalar RISC Microprocessor", *IEEE Journal of Solid-State Circuits*, Vol. 32, No. 11, pp. 1625-1634, Nov. 1997.
- [6] Charles F. Webb et al., "A 400MHz S/390 Microprocessor", *IEEE Journal of Solid-State Circuits*, Vol. 32, No. 11, pp. 1665-1675, Nov. 1997.
- [7] Paul E. Gronowski et al., "High-Performance Microprocessor Design", *IEEE Journal of Solid-State Circuits*, Vol. 33, No. 5, pp. 676-686, May 1998.
- [8] Doug Burger, "Limited Bandwidth to Affect Processor Design", *IEEE Micro*, pp. 55-62, Dec. 1997.
- [9] A. Wolfe & A. Chanin, "Executing Compressed Programs on an Embedded RISC Architecture", *Proceedings of the 25th Annual International Symposium on Microarchitecture*, December 1992.
- [10] M. Kozuch and A. Wolfe, "Compression of Embedded System Programs", *IEEE International Conference on Computer Design*, 1994.
- [11] C. W. Fraser, T. A. Proebsting, "Custom Instruction Sets for Code Compression", unpublished, <http://www.cs.arizona.edu/people/todd/papers/pldi2.ps>, October 1995.
- [12] C. Lefurgy and T. Mudge, "Code Compression for DSP", *CASES '98*, Dec. 4-5, 1998, <http://www.eecs.umich.edu/~tnm/compress>
- [13] C. Lefurgy and T. Mudge, "Fast Software-managed Code Decompression", *CASES'99*, October 1-3, 1999, <http://www.eecs.umich.edu/~tnm/compress>
- [14] S. Segars et al., "Embedded Control Problems, Thumb, and the ARM7TDMI", *IEEE Micro*, pp. 22-30, Oct. 1995.
- [15] <ftp://cair-archive.kaist.ac.kr/pub/gnu/egcs/releases/egcs-1.1b/egcs-1.1b.tar.gz>
- [16] <ftp://ftp.cygus.com/pub/newlib/newlib-1.8.1.tar.gz>
- [17] <ftp://cair-archive.kaist.ac.kr/pub/gnu/released/libstdc++-2.8.1.tar.gz>
- [18] Reinhold P. Weicker, "DHRystone Benchmark Program", *CACM* Vol 27, No 10, 10/84 pg. 1013. (Translated from ADA by Rick Richardson).
- [19] Moura, C., "SuperDLX: A Generic Superscalar Simulator", *ACAPS Technical Memo 64*, May 1993.