레지스터 프로모션을 이용한 내장형 소프트웨어의 성능 향상

이 종 열[†]

요 약

이 논문에서는 내장형 소프트웨어의 성능 향상을 위하여 사용될 수 있는 레지스터 프로모션의 새로운 기법을 제안한다. 레지스터 프로모션은 프로그램 내의 메모리 접근 연산(memory access)을 레지스터 접근 연산(register access)으로 바꾸어서 프로그램의 성능 향상을 꾀하는 최적화 방법 중의 하나이다. 제안된 방법에서는 프로파일링(profiling)을 통하여 주어진 소스 코드 내에서의 메모리 접근 연산에 대한 트레이스(trace)를 얻는다. 그리고 각 함수의 수행 횟수에 대한 프로파일링 결과로부터 높은 동적 호출 횟수를 가지는 대상 함수를 선정하여 제안된 레지스터 프로모션 기법을 적용한다. 이와 같이 최적화의 대상이 되는 함수의 수를 줄임으로써 컴파일 시간을 줄일 수 있다. 최적화 대상 함수의 메모리 트레이스를 탐색하여 레지스터 접근 연산으로 변경될 경우 수행 사이클을 줄일 수 있는 메모리 접근 연산을 찾는다. 찾아진 메모리 접근 연산에 대해서는 컴파일러의 중간단계 코드를 수정하여 프로모션 레지스터를 할당한다. 이와 같은 과정을 거쳐 메모리 접근 연산이 프로모션 레지스터에 대한 접근 연산으로 대체되고 이로부터 성능향상을 얻을 수 있다. 제안된 레지스터 프로모션 기법을 ARM과 MCORE 프로세서용 컴파일러에 적용한 후 MediaBench와 DSPStone 벤치마크을 이용하여 실험한 결과 ARM과 MCORE 프로세서에 대하여 각각 평균 14%와 18%의 성능향상을 얻을 수 있었다.

Performance Enhancement of Embedded Software Using Register Promotion

Jong-Yeol Lee[†]

ABSTRACT

In this paper, a register promotion technique that translates memory accesses to register accesses is presented to enhance embedded software performance. In the proposed method, a source code is profiled to generate a memory trace. From the profiling results, target functions with high dynamic call counts are selected, and the proposed register promotion technique is applied only to the target functions to save the compilation time. The memory trace of the target functions is searched for the memory accesses that result in cycle count reduction when replaced by register accesses, and they are translated to register accesses by modifying the intermediate code and allocating promotion registers. The experiments on MediaBench and DSPstone benchmark programs show that the proposed method increases the performance by 14% and 18% on the average for ARM and MCORE, respectively.

1. Introduction

In embedded systems based on programmable processors, high-level language compilers play an important role in the system design process. While assembly-level programming is still important to achieve optimized codes, high-level programming gains more and more acceptance for embedded processors to permit shorter design cycles, higher productivity and dependability, and higher oppor-

tunities for reuse. However, the code generated by compilers usually implies an overhead in code size and performance as compared to the hand-optimized assembly code. Although this overhead is acceptable in general-purpose computing, it is often not allowed for embedded systems.

In order to minimize the overhead, embedded compilers have to pay higher attention to code optimization for both small code size and short execution time. As a consequence, a number of code optimization techniques have been developed for embedded processors. Most of these are low-level optimizations that exploit the detailed knowledge of the processor architecture to optimize machine code. For ex-

ample, many low-level techniques were developed for code selection, register allocation, and scheduling [1-3], memory access optimization [4], and optimization of address computations [5, 6].

Since register allocation is one of the most important functions required in an optimizing compiler, many previous works such as [7] have dealt with this important problem. Most of the works mainly focused on the register allocation for scalar variables. Furthermore, they treat non-scalar variables in a particularly naive fashion, making it impossible to determine when a specific element might be reused. Due to the incapability, conventional compilers allocate memory locations for global variables, structure, array and union variables. This is true even with highest optimizations. For example, when compiling a source code using GNU C Compiler (GCC) [11], such variables are allocated to memory even with "-O3" that is the highest optimization option. The structure and array variables are allocated to memory since they are treated as a whole. Compiler optimizations do not usually treat the fields of structures and the elements of arrays as separate variables. In case of global variables, they are allocated to memory because they can be accessed in two or more functions and registers cannot be allocated beyond function boundary. For array variables, the previous works have focused on the dependency analysis of array variables used in loops. In [8], a representation method of array access patterns and a dependency analysis were presented without considering a register allocation. In other works, [9] and [10], pointer analysis was exploited to treat pointer-valued variables and arrays in C. As the pointer analysis used in the works was simple, their results were not as good as we expected. Complementary to previous register allocation techniques, in this paper, we present a code optimization that employs *register promotion* to achieve higher performance. Register promotion is to translate frequently accessed memory locations to register accesses, and can be regarded as an optimization because instructions generated to access data objects in registers are more efficient than the case that the objects are in memory. We need to determine which variables can be safely kept in registers and to rewrite the code to keep the found variables in registers. Based on the results of profiling, the proposed register promotion technique identifies code sections in which it is safe to place the value of a data object in a register, and then promotes memory loca-

tions that are frequently accessed in the code segments to registers. As the register promotion is independent of the type of variables, it can be applied to non-scalar variables as well as scalar ones.

The rest of this paper is organized as follows. In Section 2, a motivating example is given and the proposed register promotion is described. After showing the experimental results in Section 3, conclusions are addressed in Section 4.

2. Register Promotion

An example of register promotion is shown in (Figure 1), where memory is allocated for array variables, A and B. In (Figure 1)(b), registers, R0~R3, are reserved for function arguments and are not allocated in register allocation. Registers, R4~R13, are used in register allocation and the remaining two registers, R14 and R15, are used as stack pointer and link register, respectively. We can see that two registers, R4 and R5, are allocated for local variables, i and j. R6~R13 are free registers that are not allocated.
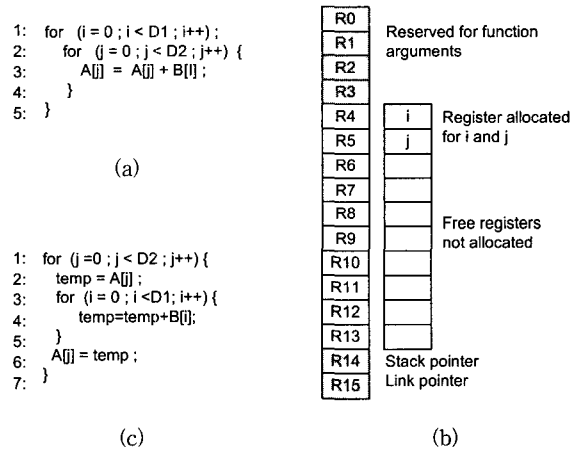
(Figure 1) Example of register promotion (a) C source code (b) Register allocation result (c) Source code representation of a transformation for register promotion

Although there are free registers, memory locations are allocated for the arrays as a result of the inability of compilers stated as above. In (Figure 1)(a), we can see that the element of A accessed multiple times in the inner loop is independent of the index of the inner loop. Therefore, by moving the value of the element into a scalar variable before the inner loop and restoring the variable back to the memory location after the inner loop, we can replace the memory access in the inner loop by a register access. The

source code representation of the result after a transformation for register promotion is shown in (Figure 1)(c), where a new temporary scalar variable that will be allocated to a register is added.

In this paper the performance of the register promotion is determined in terms of *cycle count* and *code size*. The cycle count represents the number of cycles taken when an operation is executed. In some machines the cycle count includes preliminary cycles consumed in preparing a target address for a memory operation. The total execution time of a program can be calculated by multiplying the cycle count and clock cycle time. Therefore, for a given machine where the clock cycle time is fixed, we can compare the execution times by comparing the cycle counts.

The cycle counts, TC_{before}, before the register promotion can be estimated as follows with assuming the cycle counts of memory accesses, C_r and C_w, are two and that of addition, C_{add}, is one.

$$B_{before} = 2 \times C_r + 1 \times C_{add} + 1 \times C_w = 7 \qquad (1)$$

$$TC_{before} = B_{before} \times D_1 \times D_2 = 7 \times D_1 \times D_2 \qquad (2)$$

where B_{before} is the estimated cycle count of the loop body. C_r and C_w are the numbers of cycles consumed when a memory read operation and a memory write operation are executed, respectively. The values include the cycles required to generate the address for a target memory location in ARM for example.

After the register promotion, since the memory access, the access to A[j], is replaced by a register access, the access to temp, the total cycle count, TC_{after}, is calculated as follows.

$$B_{after} = 1 \times C_r + 1 \times C_{add} = 3 \qquad (3)$$

$$TC_{after} = D1 \times [1 \times C_r + D2 \times B_{after} + 1 \times C_w]$$

$$= 4 \times D_1 + 3 \times D1 \times D2 \qquad (4)$$

$$TC_{before} - TC_{after} = 4 \times D1 \times (D2 - 1) \qquad (5)$$

By calculating the difference of two total cycle counts, TC_{before} and TC_{after}, as shown in equation (5), we can know whether the register promotion is effective or not. It should be noted that we describe for clear understanding as if the source code is modified directly. In the real implementation, the modification is applied to the intermediate code of compilers such as the RTL of GCC. Therefore, the source code is never changed but the intermediate code generated after

parsing the source code is modified to account for the register promotion to be accommodated.

(Figure 2) Overview of the proposed register allocation method

The technique proposed in this paper performs register promotion for non-scalar variables in a systematic way. (Figure 2) shows the overall flow of the proposed method. In the proposed method, the information needed for the optimizations is obtained by profiling the source code. The advantage of profiling is that it can produce more accurate result than the static analysis such as pointer analysis. For register promotion, memory trace is generated and analyzed to find memory locations accessed frequently. They can be replaced with register accesses to reduce cycle count. After finding such memory locations, the intermediate code is modified to move the data in the memory locations to *promotion registers* reserved for register promotion. Note that, in the proposed method, pointers can be optimized without using pointer analysis because memory trace is used.

2.1 Source Code Profiling

As the first step of register promotion, the source code is profiled to find target functions to which register promotion is applied. To maximize the effect of register promotion and reduce the compilation time, the only functions with a high dynamic call count are considered.

For the profiling, instruction-set simulators are modified to report dynamic calls and generate memory trace. The address and program counter (PC) values are dumped for each load or store instruction with a field indicating the type of memory operation (load or store).

An advantage of the profiling is that the number of functions to be considered in the optimization routine is reduced by selecting only the functions called frequently. As can be seen in (Figure 3), most of dynamic function calls are dedicated to only a small portion of the whole functions.

As a consequence, significant improvement can be achieved with applying optimizations to the heavily-called functions. In our experience, only one or two functions take 95% of total dynamic calls in all the benchmarks used in the experiments. Another advantage is that better results than static pointer analysis can be obtained in the presence of pointers because the aliasing problem can be eliminated by analyzing addresses in the memory trace.

Cumulative fraction of dynamic calls

(Figure 3) Distribution of dynamic function calls. The Y axis shows the cumulative fraction of dynamic function calls and 10% of functions take 90% of total dynamic function calls

2.2 Iteration Tree Construction

The memory trace of heavily executed functions is parsed to find whether register promotion can increase performance. The first step of the implemented memory trace parser reads the trace into memory and the following steps are performed on the trace in memory because file I/O is the most time-consuming task.

In memory trace parsing, for each memory access in the intermediate code, a corresponding memory address is found. This is achieved by inserting debugging information into the executable when compiling the source code. A source-level debugger uses the debugging information to find source lines corresponding to a given PC value. In the proposed method the debugging information is used to map a PC value to a memory operation in the intermediate code.

In addition, loops can be found by examining PC values in trace and debugging information. For the found loops, an iteration tree is constructed where nodes represent instances of loop iterations. Each node has a pointer to a loop and the index value of the corresponding iteration. (Figure 4) shows an iteration tree example. In (Figure 4)(b), the

nodes, i0, i1, and i2, represent the iterations of the outer-most loop whose index variable is i. For a loop without an index variable a pseudo index variable is inserted. For example, the inner-most loop in (Figure 4) is a while loop that has no index variable. The iterations of the while loop with a pseudo index variable, k, are shown at the leaf nodes of (Figure 4)(b).
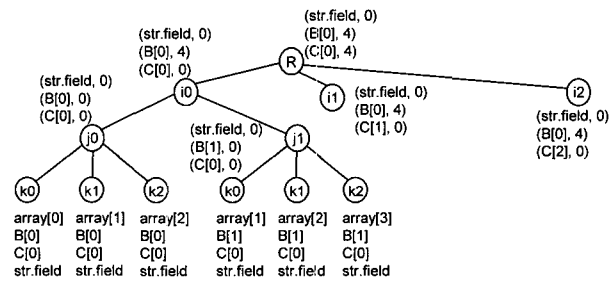
```
1:    for (i = 0 ; i < 2 ; i++) {
2:        for (j = 0 ; j < 1; j++)  {
3:            ptr = array [ j ] ;
4:            while (*ptr) {
5:                *ptr = *ptr + B[ j ] + C[ i ] + str.field++ ;
6:                ptr++ ;
7:            }
8:        }
9:    }
```

(a)

(b)

(c)

(d)

(Figure 4) Iteration tree (a) An example source code (b) An iteration tree (c) CAS calculation (d) An iteration tree after loop interchanging

2.3 Code Generation

The next step is to find the memory accesses to be promoted to register accesses by examining the iteration tree. The candidate memory accesses for register promotion are

the accesses whose addresses are constant over all the iterations of a loop. We assume that each memory address is associated with a PC value. To find candidate memory accesses, the nodes are levelized by their distance to the root node. For the iteration tree of (Figure 4)(b), the first-level nodes are i0, i1 and i2. A *canonical address set* (*CAS*) is calculated for a node. If memory addresses with the same PC value have a uniform stride, the addresses are called canonical in this paper and can be represented as a form of (starting address, stride). The CAS of a leaf node contains all of the memory addresses accessed in the iteration corresponding to the leaf node, as there is one memory access for a PC value. As in this case, it is difficult to define the stride, but we assume every memory address has a uniform stride initially. Therefore, a CAS contains only canonical addresses, which means a set of memory addresses that have the same PC but do not have a uniform stride is not considered in register promotion. A CAS of a non-leaf node is calculated as CAS_n e {e CAS_i}, where e is an operator that finds the canonical addresses, CAS_n is the CAS of the non-leaf node and CAS_i is the CAS of an immediate descendant of the non-leaf node. An example of CAS calculation is shown in (Figure 4)(c), where each leaf node has memory accesses in the corresponding iteration and non-leaf nodes have their CASs. The CAS of i0 contains (B[0], 4) that is calculated from (B[0], 0) in the CAS of j0 and (B[1], 0) in the CAS of j1. Since in the descendants of i0, B[0] and B[1] are accessed in sequence, the stride is four, which is the size of elements in B.

How to find candidate memory accesses is explained using an example. Let A1, A2, and A3 be canonical addresses in the CASs of i0, i1 and i2, respectively, in (Figure 4)(b). When the following equation holds

$$A2 - A1 = A3 - A2 = d \tag{6}$$

and the number of traces is equal to the statically determined number of iterations, the addresses in iterations can be represented as an arithmetic series

$$A1 = A1 + d \times 0 \tag{7}$$
$$A2 = A1 + d \times 1 \tag{8}$$
$$A3 = A1 + d \times 2 \tag{9}$$

In general, if all the differences of two consecutive addresses are identical, the addresses form an arithmetic series

and the address in i-th iteration can be represented as a general term,

$$A(i) = A1 + d \times i \tag{10}$$

where A1 is the address in the first iteration and d is a stride.

If the stride of memory addresses associated with the same PC value is zero in equation (10), that is, memory addresses are constant, the corresponding memory accesses might be replaced by register accesses with additional load and store instructions. The additional instructions are placed outside the loop corresponding to the level on which the general term is calculated. If equation (10) holds for the first-level CASs in (Figure 4)(b) with a zero stride, a load instruction and a store instruction from/to the canonical address are placed before and after the outer-most loop, respectively. When d is not zero, the address changes with a equal stride, d. In this case, the address can be efficiently implemented by using auto-increment or auto-decrement addressing mode.

In some cases more cycle count reduction can be achieved by interchanging loops. When two or more nodes have the same name in a level and the CASs of the nodes have common canonical addresses, by interchanging the loops corresponding to the level and its upper level the constant addresses can be used in register promotion resulting in cycle count reduction. If CASs of j0s in (Figure 4)(b) have the same constant address, A1, and j1s also have the same constant address, A2, by interchanging the first and second loops corresponding to Level 1 and Level 2, respectively, the canonical addresses will be found at Level 1. As canonical addresses are found in the upper level after the loops are interchanged, the additional load and store instructions are inserted at outer position, and hence, the amount of cycle count increased by the instructions can be reduced in the modified code. The iteration tree with the loop interchanging is shown in (Figure 4)(d).

After finding candidate memory accesses for register promotion, we can calculate the cycle count difference between the original code and the register promoted code by taking into account the execution counts of the loops and functions. If the difference is positive, register promotion is applied and the memory accesses are replaced by register accesses.

The first step of the replacement is to modify the intermediate code. In this step, additional variables are defined

and codes that move values from/to the additional variables are inserted. (Figure 5) shows the source modification results of the example code in (Figure 4)(a). At each level of the iteration tree, we search for the same canonical addresses. In Level 1 of (Figure 4)(b), the memory access to str.field is a canonical address. Therefore, a load instruction to a temporary variable, temp1, and a store instruction from the temporary variable are inserted before and after of the outer-most for loop, respectively. Another temporary variable, temp3, is used to promote the accesses to B[j] in (Figure 5)(c) because the addresses of B[j] are constant over the iterations of the inner-most while loop.

```
1 :   temp1 = str.field ;
2 :   for (i = 0 ; i < 2 ; i ++) {
3 :       for (j = 0 ; j < 1 ; j ++) {
4 :           ptr = array [ j ];
5 :           while (*ptr) {
6 :               *ptr = *ptr + B[ j ] + C[ i ] temp1++ ;
7 :               ptr ++ ;
8 :           }
9 :       }
10:   }
11:   str. field = temp1 ;
```

(a)

```
1 :   temp2 = &C[0] ;
2 :   temp1 = str.field ;
3 :   for (i = 0 ; i < 2 ; i ++) {
4 :       temp4 = *temp2 ;
5 :       for (j = 0 ; j < 1 ; j ++) {
6 :           ptr = array [ j ];
7 :           while (*ptr) {
8 :               *ptr = *ptr + B[ j ] + temp4 + temp1++ ;
9 :               ptr ++ ;
10:           }
11:       }
12:       temp2++
13:   }
14:   str. field = temp1 ;
```

(b)

```
1 :   temp2 = &C[0] ;
2 :   temp5 = &B[0] ;
3 :   temp1 = str.field ;
4 :   for (i = 0 ; i < 2 ; i ++) {
5 :       temp4 = *temp2 ;
6 :       for (j = 0 ; j < 1 ; j ++) {
7 :           temp6 = *temp5 ;
8 :           ptr = array [ j ];
9 :           while (*ptr) {
10:               *ptr = *ptr + temp6 + temp4 + temp1++ ;
11:               ptr ++ ;
12:           }
13:           temp5 ;
14:       }
15:       temp2 ;
16:   }
17:   str. field = temp1 ;
```

(c)

```
1 :   temp2 = &C[0] ;
2 :   temp5 = &B[0] ;
3 :   temp1 = str.field ;
4 :   for (j = 0 ; j < 1 ; j ++) {
7 :       temp6 = *temp5 ;
6 :       for (i = 0 ; i < 2 ; i ++) {
5 :           temp4 = *temp2 ;
8 :           ptr = array [ j ];
9 :           while (*ptr) {
10:               *ptr = *ptr + temp6 + temp4 + temp1++ ;
11:               ptr ++ ;
12:           }
15:           temp2++
14:       }
13:       temp5++
16:   }
17:   str. field = temp1 ;
```

(d)

(Figure 5) Code modification. For clear understanding the modification results are shown using source codes zinstead of the intermediate code on which the modification is performed (a) The address of a structure variable, str.field, is constant in the iterations of all loops (b) The addresses of C[i] form an arithmetic series with index variable i (c) Code modification without interchanging the first two for loops. (d) Code modification after interchanging the first two for loops

The addresses of C[i] in (Figure 4)(a) form an arithmetic series with the index variable i. The first address is the address of the first element, C[0]. The difference between the first and second addresses is C[1] − C[0], which is four if C is an array of four-byte integers. Therefore, the address increases by four at each iteration. The code is modified by inserting a statement that loads the address to a

temporary variable, temp2. At the end of the first for loop, the address is incremented by four. In the new statement, temp2++ is used instead of temp2 = temp2 + 4 because temp2++ in C statement is interpreted as the increment of the address by the size of the array elements. The resulting code is shown in (Figure 5)(b).

We can calculate the cycle count before and after loop interchanging shown in (Figure 5)(c) and (d) as follows :

$$C_{before} = 3 \times 2 \times (M + W \times R) \tag{11}$$

$$C_{after} = 2 \times (M + 3 \times W \times R) \tag{12}$$

where M is the cycle count of a memory access, R is that of a register access, and W is the number of iterations of the inner-most while loop. If we assume that the values of M and R are five and three, respectively,

$$C_{before} = 30 + 18W \tag{13}$$

$$C_{after} = 10 + 18W \tag{14}$$

The assumed values are typical values of memory and register accesses in ARM7TDMI. In this calculation, the reduction in cycle count is about 30% when the inner-most while loop iterates two times.

In register allocation step, promotion registers that are special registers dedicated for register promotion are allocated for the compiler-inserted variables, temp1 and temp3. In most cases, the promotion registers need not be saved and restored to preserve their values, since the promotion registers are not used in general register allocations. If there are compiler-inserted variables not mapped to promotion registers because of the restricted number of promotion registers, the compiler may allocate free registers for the variables.

Since the case that all the registers are used in register allocation is rare in general compilers, some of the registers can be used for register promotion. If some registers are dedicated for register promotion, the maximum number of registers that can be involved in register allocation is reduced, which may lead to significant effects on the code size and performance. Therefore, the number of promotion registers must be determined by investigating such effects.

3. Experimental Results

We implemented the proposed register promotion method using GCC and performed experiments for ARM and

MCORE. To measure the cycle count variation, ARMulator and MCORE instruction-set simulator (ISS) are used. As listed in <Table 1>, benchmark programs used in the experiments are MediaBench [12] and DSPstone[13]. Some of the programs, DOT_PRODUCT, FIR2DIM, MATRIX_MUL and STR_SUM use pointers heavily and are included to show that the proposed method can handle pointers appropriately.

<Table 1> Summary of Benchmark Programs

Benchmark	Description
RS	Reed-Solomon coder/decoder
STR_SUM	Summation of structures
PEGWIT	Public key encryption
ADPCM	ADPCM coder/decoder
MPEG2	MPEG2 encoder
DOT_PRODUCT	Dot product
FIR2DIM	Two-dimensional FIR filter
MATRIX_MUL	Matrix multiplication
EPIC	Efficient pyramid image coder
G721	CCITT G.721 ADPCM coding algorithm
GSM	An implementation of the final draft GSM 06.10 standard

First, we performed experiments to determine the number of promotion registers. We generated executables by restricting the number of registers. (Figure 6) and (Figure 7) show the effects of register reduction on the code size and cycle count for ARM and MCORE, which are obtained with a modified GCC. In these experiments some selected benchmark programs are used in order to show the effectiveness of the proposed register promotion when the program to be optimized is different from the program used to determine the number of promotion registers. Based on the results in (Figure 6) and (Figure 7), we can determine the number of promotion registers. The amount of code size increase is less than 1% for both processors until the number of registers available for register allocation is reduced by three. However, the register reduction affects cycle count more significantly. In ARM, the cycle count increases more than 30% when the number of registers is reduced by three. In (Figure 6), we can see that reserving two registers for register promotion in ARM does not impose great penalty on code size and cycle count. For MCORE, there is no steep increase in cycle count for three promotion registers. Therefore, two and three promotion registers are assigned for ARM and MCORE, respectively.

We compiled the benchmark programs using the implemented compiler for ARM and MCORE with the determined number of promotion registers. In the experiments, only one or two functions are selected and the register promotion is applied as summarized in <Table 2>. By applying the proposed register promotion to the selected functions the compilation time is reduced as shown in (Figure 8). The reduction is calculated as follows.

$$\text{Compilation time reduction} = \frac{T_A - T_S}{T_A} \qquad (15)$$

where T_A and T_S are compilation times when the proposed method is applied to all the functions and only the selected functions, respectively.
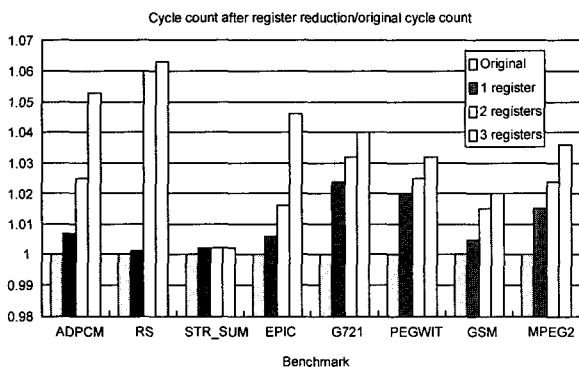

(a)

(b)

(Figure 6) Effects of register reduction in ARM for some selective benchmarks. (a) Code size variable after register reduction Code sizes are compared by calculating "(code size after register reduction/original code size)" Code size increases when the number of available registers is reduced by one, two and three (b) Cycle count variation after register reduction. Cycle counts are compared by calculating "(cycle count after register reduction/original cycle count)" Cycle count increases when the number of available registers is reduced by one, two and three
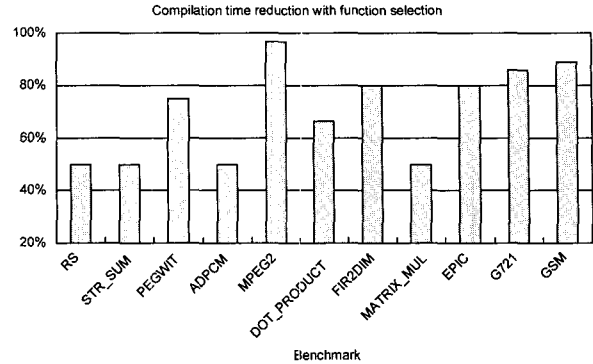
(a)

(b)

(Figure 7) Effects of register reduction in MCORE for some selective benchmarks (a) Code size variable after register reduction. Code sizes are compared by calculating "(code size after register reduction/original code size)". Code size increases when the number of available registers is reduced by one, two and three (b) Cycle count variation after register reduction. Cycle counts are compared by calculating "(cycle count after register reduction/original cycle count)". Cycle count increases when the number of available registers is reduced by one, two and three

⟨Table 2⟩ Number of functions to which the register promotion is applied

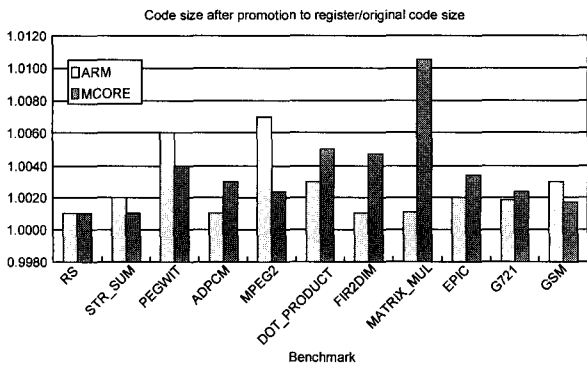| Benchmark	Number of functions
RS	2
STR_SUM	1
. PEGWIT	2
ADPCM	2
MPEG2	1
DOT_PRODUCT	1
FIR2DIM	1
MATRIX_MUL	1
EPIC	1
G721	1
GSM	1

(Figure 8) Compile time reduction with function selection. The compilation time is reduced by selectively applying the proposed register promotion to the heavily executed functions. The reduction is calculated using the equation (15)

By executing the generated executables on ARMulator and MCORE ISS, the cycle counts are measured. (Figure 9) shows the code size and cycle count variation after register promotion. The maximum code size increase is about 1% for MCORE, and the code size remains almost constant in case of ARM. However, the cycle count is reduced for both ARM and MCORE. The average cycle count reductions are about 14% and 18% for ARM and MCORE, respectively. The maximum cycle count reductions are 30% and 35% for ARM and MCORE, respectively. (Figure 9) also indicates that the proposed method can handle pointers effectively because the average cycle count reductions for programs that heavily use pointers are about 15.2% and 20.5% for ARM and MCORE, respectively.
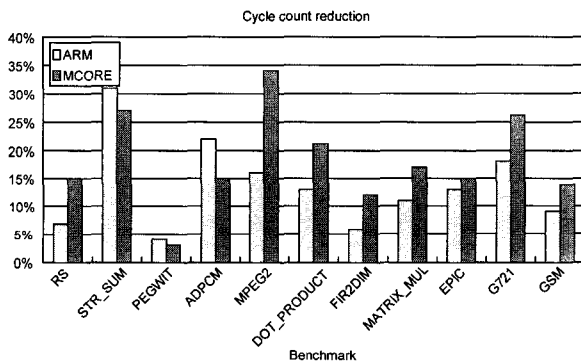
Although the proposed register promotion is applied for two functions in RS, PEGWIT and ADPCM, the amount of the cycle count reduction is not large compared to other benchmark programs where the register promotion is applied for only one function. The proposed register promotion is less effective in RS, PEGWIT and ADPCM because there are fewer variables that can be promoted to registers in those programs after previous code optimization such as register allocation.

In the most cases of (Figure 9), the amount of the cycle count reduction is higher in MCORE than ARM. This is because the GCCs ported to ARM and MCORE, ARM-GCC and MCORE-GCC, have different characteristics and registers are more heavily used in the ARM-GCC generated codes. That is, there are fewer variables that the register promotion can be applied to in ARM-GCC generated codes than in MCORE-GCC generated codes. This

is most evident in MPEG2 where the cycle count reduction of MCORE-GCC generated code is two times larger than that of ARM-GCC generated code. However, in STR_SUM, PEGWIT, and ADPCM, MCORE-GCC generates more optimized code than ARM-GCC and hence, the reduction rate is higher in ARM.

(Figure 9) Code size variation and cycle count reduction after register promotion (a) Code size variation. Code sizes are compared by calculating "code size after register reduction/Original code size" (b) Cycle count reduction. Cycle count reductions are compared by calculating "(Original cycle count - cycle count after register reduction)/Original cycle count"

Since the proposed register promotion can be applied after conventional optimization developed in the previous works such as register allocation, it can be efficiently used for the optimization of embedded software as a complementary tool to the previous works as shown the experiments with the embedded software benchmarks, Media-Bench and DSPStone. Furthermore, the proposed method uses the profiling and can register-promote pointers effectively. Another advantage of the register promotion with profiling is that the proposed method can effectively handle non-scalar variables such as arrays and structures. In the

profiling method, the trace of memory accesses are generated and analyzed to perform the register promotion.

4. Conclusions

In this paper, we have presented a register promotion technique to enhance performance by moving memory accesses to register accesses. In the proposed method, a given source code is profiled to generate a memory trace. The profiling result is used to find target functions with high dynamic call counts, and the proposed register promotion is applied only to the target functions to save the compilation time. By examining the memory trace of the target functions, we search for memory accesses that can result in cycle count reduction if changed to register accesses. Such a memory access is replaced by a register access by modifying the code and allocating a promotion register. The experimental results show that the proposed method is effective in that average cycle count reduction is about 14% with increasing less than 1% of the code size.

References

[1] C. Liem, T. May and P. Paulin, "Instruction-Set Matching and Selection for DSP and ASIP Code Generation," Proceedings of European Design and Test Conference (ED & TC), pp.31-37, 1994.

[2] G. Araujo and S. Malik, "Optimal Code Generation for Embedded Memory Non-Homogeneous Register Architecture," Proceedings of 8th Int. Symp. On System Synthesis (ISSS), pp.13-15, 1995.

[3] S. Liao, S. Devadas, K. Keutzer, S. Tjiang, G. Araujo and S. Malik, "Instruction Selection Using Binate Covering for Code Size Optimization," Proceedings of Int. Conf. on Computer-Aided Design (ICCAD), pp.5-9, 1995.

[4] A. Sudarsanam and S. Malik, "Memory Bank and Register Allocation in Software Synthesis for ASIPs," Proceedings of Int. Conf. on Computer-Aided Design (ICCAD), pp.393-399, 1995.

[5] S. Liao, S. Devadas, K. Keutzer, S. Tjiang and A. Wang, "Storage Assignment to decrease code size," ACM Trans. Program. Lang. and Sys., Vol.18, No.3, pp.186-195, May, 1996.

[6] R. Leupers and P. Marwedel, "Algorithms for Address Assignment in DSP Code Generation," Proceedings of Int. Conf. on Computer-Aided Design (ICCAD), pp.109-112, 1996.

[7] A. Aho, R. Sethi and J. Ullman, 'Compilers – Principles, Techniques, and Tools,' Reading, MA : Addison-Wesley, 1986.

[8] M. Wolfe, 'High Performance Compilers for Parallel Computing,' Redwood City, CA : Addison-Wesley, 1996.

[9] M. hind, "Pointer Analysis : Haven't We Solved This Problem Yet?," Proceedings of ACM SIGPLAN – SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE), pp.54-61, 2001.

[10] R. P. Wilson and M. S. Lam, "Efficient Context-Sensitive Pointer Analysis for C Programs," Proceedings of SIGPLAN Programming Language Design and Implementation (PLDI), pp.1-12, 1995.

[11] R. M. Stallman. Using and Porting GNU CC. [Online]. Available, http://gcc.gnu.org/onlinedocs/gcc-2.95.3/gcc.html.

[12] Chunho Lee, M. Potkonjak and W. H. Mangione-Smith, MediaBench. [Online]. Available, http://www.cs.ucla.edu/leec/mediabench/.

[13] V. Zivojnovici, J. Martinez Velarde and C. Schlager, "DSPstone : A DSP-oriented Benchmarking Methodology," Proceedings of Int. Conf. On Signal Processing and Technology, pp.715-720, 1994.

이 종 열

e-mail : jong@chonbuk.ac.kr
1993년 한국과학기술원 전자전산학과
(공학사)
1996년 한국과학기술원 전자전산학과
(공학석사)
2002년 한국과학기술원 전자전산학과
(공학박사)
2002년~2003년 하이닉스 반도체 선임연구원
2003년~2004년 한국과학기술원 BK21 초빙교수
2004년~현재 전북대학교 전자정보공학부 전임강사
관심분야 : SOC 설계, 내장형 프로세서 설계, 내장형 소프트웨어
최적화