

Improving Offset Assignment through Simultaneous Variable Coalescing

Desiree Ottoni¹, Guilherme Ottoni², Guido Araujo¹, and Rainer Leupers³

¹ IC-UNICAMP - Brazil

² Princeton University, Department of Computer Science - USA

³ Aachen University of Technology, Integrated Signal Processing Systems - Germany

Abstract. Efficient address code optimization is a central problem in code generation for processors with restricted addressing modes, like Digital Signal Processors (DSPs). This paper proposes a new heuristic to solve the Simple Offset Assignment (SOA) problem, the problem of allocating scalar variables to memory so as to minimize addressing code. This new approach, called Coalescing SOA (CSOA), performs variable memory slot coalescing simultaneously to offset assignment computation. Experimental results, based on compiling MediaBench benchmark programs with LANCE compiler, reveal a very significant improvement over the previous solutions to SOA. In fact, CSOA produces, on average, 37.3% fewer update instructions when comparing with the prior solution that perform memory slot coalescing before applying SOA, and 66.2% fewer update instructions when comparing with the best traditional SOA solution.

1 Introduction

The growth of the DSP market and the increasing demand for new and complex applications running on these processors have brought a strong interest to compilers capable of generating efficient DSP code. However, as DSPs have very irregular architectures, traditional compiling techniques designed for general-purpose processors [1, 21] are not capable of generating efficient code for DSPs [14]. As a result, new techniques tailored for these processors have been proposed and intensively studied. Due to their instruction size and performance constraints, DSPs traditionally have no offset addressing mode, containing only indirect addressing, and a few general-purpose registers. In addition, DSPs have specialized *Address Generation Units (AGU)*, that provide address computation in parallel to datapath computation. AGUs perform *auto-increment (decrement)* in address registers (AR) by some fixed values⁴. For different values, the program is required to have an explicit *update instruction* (prior to the memory access) that uses datapath resources to compute the memory address. Therefore, in order to produce efficient code for such DSPs, it is important to use auto-increment (decrement) addressing modes effectively.

⁴ Generally, the values of auto-increment (decrement) are one, but in some architectures these values can sometimes be larger.

The optimization that tries to maximize the use of instructions with auto-increment (decrement) for local scalar variables is called *Offset Assignment* (OA). This optimization finds a stack layout for these variables in such a manner that auto-increment (decrement) addressing modes are used whenever possible. The variation of the OA problem when there is only one address register and auto-increment (decrement) by 1 is called *Simple Offset Assignment (SOA)* [20] and is the focus of this paper.

In this paper we describe a new approach to the SOA optimization problem, called Coalescing SOA (CSOA). It uses liveness information [1, 21] to simultaneously coalesce variable memory slots while solving SOA optimization. The interference graph [21] is used to identify which pairs of variables can be coalesced. Only variables that do not interfere⁵ can be coalesced during CSOA. We show that variable coalescing can lead to a large improvement in code quality (66.2% fewer update instructions) when comparing to the best algorithm in OffsetStone [15, 22], and 37.3% fewer update instructions when comparing with the other coalescing approach described in [25]. This result dismisses the first assumptions to this problem, as in Liao [19], that seemed to indicate the opposite. Moreover, CSOA reduces both the code and the data segment, resulting in 92.3% of the number of memory slots when comparing with the results obtained through SOA-Liao.

The remainder of this paper is organized as follows. Section 2 exhibits an example that illustrates how the use of coalescing can affect the number of update instructions. Section 3 lists the previous work on SOA. Section 4 describes our technique (CSOA) and Section 5 shows the time-complexity of our method. Section 6 shows a small example that demonstrates the workings of the algorithm. In addition, Section 7 evaluates the results of CSOA, while Section 8 summarizes the main results.

2 Motivation

This section shows an example that illustrates how coalescing variables can decrease the number of update instructions. Consider that only a single address register is available in the processor, which can be auto-incremented (decremented) only by one.

Figure 1(a) shows a fragment of C code with the liveness information annotated at each program point. Figures 1(b) and (c) show two possible memory layouts for the variables, and the sequence in which the variables are accessed in memory. The arrows in Figures 1(b) and (c) indicate that an explicit address calculation instruction (i.e. update instruction) is required to make the address register point to the next variable, because the distance between the variables is greater than one. The layout showed in Figure 1(c) has one slot that is shared between two variables (*b* and *g*) which do not interfere at runtime. By sharing these variables, one less update instruction is required in the program. Clearly,

⁵ Two variables interfere when they are simultaneously live.

coalescing variables increases the closeness between the variables on the stack, thus reducing the number of update instructions.




Fig. 1. (a) A fragment of C code. (b) Memory layout with one slot per variable. (c) Memory layout with more than one variable per slot.

3 Related Work

The Simple Offset Assignment (SOA) problem was first studied by Bartley [5]. Later, Liao et al [20] showed that the graph problem *Maximum Weight Path Cover* (MWPC) (known to be NP-Complete) can be reduced to SOA, thus proving that SOA is NP-Hard. After that, a large number of heuristic techniques have been proposed for SOA [20, 17, 4, 15, 24], making it one of the most studied problems in code generation for DSPs.

Liao et al [20] used a heuristic to solve SOA based on the *Kruskal Minimum Spanning Tree* algorithm [12]. Given a basic block, Liao et al [20] call *access sequence* the sequence used by the program to access variables during execution time. For example, in instruction $a = b \text{ op } c$, the access sequence is bca . Based on the access sequence, Liao et al define an weighted graph $G(V, E)$, called *access graph*, where V is the set of variables in the basic block, and E is the set of edges. An edge $e = (u, v)$, with weight $w(e)$, indicates that there are $w(e)$ consecutive accesses to variables u and v (or v and u) in the access sequence. If two variables u and v are never accessed consecutively, then $(u, v) \notin E$. Once the access graph is constructed, Liao's algorithm tries to find a set of maximum weighted paths, called *assignment* that define the variable layout in memory. The *cost* of an assignment is the addition of the weights of all edges between variables in non-adjacent memory positions, as only auto-increment (decrement) by one is available. To illustrate these concepts consider Figure 2. Figure 2(a) shows a fragment of C code. Figure 2(b) shows the corresponding access sequence, and Figure 2(c) its associated access graph. Liao's heuristic is a greedy algorithm that, at each step, chooses the edge with the greatest weight, taking care not to select an edge that will stay with degree greater than 2 neither a edge that can form a cycle with the already selected edges. By using this heuristic, the

assignment selected in the access graph of Figure 2(c) would be fecadgb, as highlighted in that figure. This choice results in an offset cost of four, i.e. four update instructions are required, corresponding to the non-highlighted edges.




Fig. 2. (a) A fragment of C code. (b) The access sequence of this fragment. (c) The corresponding access graph.

Sudarsanam et al [25] performed graph coloring to coalesce variables before SOA, but their goal was to reduce memory utilization, and they have not shown that this would improve the offset cost. In section 7 the results of this heuristic are showed and compared with the results of CSOA.

Leupers and Marwedel [18] proposed an extension to Liao’s heuristic, called tie-break, that decides what edge to choose when there are edges with the same weight. Rao and Pande [24] described a technique that considers the order of the accesses. This technique optimizes the access sequence through algebraic transformations in the expression tree. In [17], Leupers and David proposed a genetic algorithm to solve SOA. Instead of using the access sequence, they computed the offset assignment directly by a simulation of a natural evolution process.

Many generalizations of the SOA problem have been studied. One important generalization is the *General Offset Assignment (GOA)* problem [20, 18, 17], that is the offset assignment problem when more than one address register is available. In addition, generalizations that exploit the use of modify registers [26, 18, 17], auto-increment (decrement) ranges [25, 11], instruction scheduling coupled with offset assignment [6] and procedure-level offset assignment [9] were also studied.

Another problem related to SOA is the problem known as *Array Reference Allocation (ARA)*, which optimizes the access to array variables instead of scalar variables. This problem was originally studied by Araujo et al in [3], and later extended by other researchers [16, 23, 2, 7].

4 Coalescing Simple Offset Assignment

This section describes an optimization for offset assignment that is based on variable liveness information. Our approach, called *Coalescing Simple Offset Assignment (CSOA)*, receives as input the access sequence and the interference graph of the variables. Its output is an offset assignment for the variables in memory.

Our technique is an extension to most of the previous heuristics that solve SOA [5, 20, 18, 4, 15]. For the purpose of testing CSOA, we use the algorithm proposed by Liao et al [20] with the tie-break heuristic in [18] to decide between edges with the same weight. Liao et al try to form a maximum path in the access graph, sorting the edges of the access graph in decreasing order of their weights. After that, their algorithm iterates until all vertices are inserted onto the path or no other edge is available. At each step of the iteration, Liao et al choose the *valid edge*, (i.e. one not already selected, that does not cause a cycle, and does not increase the degree of a vertex on the path to more than two) with maximum weight.

Algorithm 1 presents pseudo-code for CSOA. At each iteration step, instead of always choosing an edge, as in typical SOA solutions, it considers another alternative: coalescing two vertices. Specifically, we do one of two operations: (a) coalesce two vertices u and v in the access graph, if they do not interfere; (b) pick a valid edge of maximum weight from the sorted list of edges, L in the Algorithm 1, as in Liao’s approach.

In Algorithm 1, function *FindCandidatePair* tries to find the two candidates for coalescing. This function returns a quadruple $(coal, u, v, csave)$, where $coal$ is a flag that is set if there are two vertices u and v for coalescing, and $csave$ is the number of update instructions that are saved if u and v are coalesced.

In order to find the two candidates for coalescing, function *FindCandidatePair* (line (7)) searches among all possible combinations of two vertices u and v , in the interference graph, considering only the vertices that satisfy the following conditions:

1. $(u, v) \notin$ the interference graph;
2. Coalescing u and v does not create a cycle, considering only the selected edges;
3. Coalescing u and v , does not cause the coalesced vertex to have degree greater than two, considering only the selected edges.

Then, it picks, among all pairs of vertices that satisfy the above conditions, the pair u and v whose coalescing results in the highest $csave$.

To calculate $csave$, function *FindCandidatePair* computes the following statements, where $Adj_{sel}(y)$ is the set of vertices adjacent to y , considering only the already selected edges:

1. $\forall x \in (Adj_{sel}(u) - Adj_{sel}(v))$, add $w(x, v)$ to $csave$;
2. $\forall x \in (Adj_{sel}(v) - Adj_{sel}(u))$, add $w(x, u)$ to $csave$;
3. Add the weight of the edge between u and v , $w(u, v)$, to $csave$, if the edge was not selected yet.

Algorithm 1 Coalescing-Based SOA

Input: the access sequence L_{AS} ,
the interference graph $G_I(V_I, E_I)$.
Output: the offset assignment.

```
(1)  $G_A(V_A, E_A) \leftarrow \text{BuildAccessGraph}(L_{AS});$ 
(2)  $L =$  sorted list of the  $E_A$ ;
(3)  $coal \leftarrow \text{false};$ 
(4)  $sel \leftarrow \text{false};$ 
(5) repeat
(6)    $rebuild \leftarrow \text{false};$ 
(7)    $(coal, u, v, csave) \leftarrow \text{FindCandidatePair}(G_I, u, v);$ 
(8)    $sel \leftarrow \text{FindEdgeValidNotSel}(L, e);$ 
(9)   if ( $coal \ \&\& \ sel$ )
(10)    if ( $csave \geq w(e)$ )
(11)      $rebuild \leftarrow \text{true};$ 
(12)    else
(13)     mark  $e$  as selected;
(14)   else
(15)    if ( $coal$ )
(16)      $rebuild \leftarrow \text{true};$ 
(17)    else if ( $sel$ )
(18)     mark  $e$  as selected;
(19)   if ( $rebuild$ )
(20)      $\text{RebuildAccessGraph}(G_A, u, v);$ 
(21)      $\text{RebuildInterferenceGraph}(G_I, u, v);$ 
(22)      $\text{RebuildL}(L);$ 
(23) until ( $!(coal \ || \ sel)$ )
(24) return  $\text{BuildOffset}(G_A);$ 
```

For the sake of clarity, consider Figure 3. According to the statements above and Figure 3, the value of $csave$ when u and v are coalesced is the weight of edge (x, v) (since x is adjacent to u , edge (x, u) is selected, and edge (x, v) is not selected) plus the weight of the non-selected edge (u, v) . The value of $csave$ becomes 6, 4 from edge (u, v) and 2 from edge (x, v) .

After that, in line (8) of the Algorithm 1, function *FindEdgeValidNotSel* searches for the valid edge e with maximum weight $w(e)$ in the sorted list of edges L , and if it exists, flag sel is set.

Finally, if both $coal$ and sel are true (line (9)), Algorithm 1 chooses (line (10)) the one that makes the best reduction in the number of update instructions.

When two vertices u and v are coalesced, parts of the access and the interference graphs need to be rebuilt in order to reflect the operation. This is performed in lines (19)-(22) of Algorithm 1. In the new access graph, all the old adjacencies of u and v must be redirected to the coalesced vertex (uv) . In the




Fig. 3. (a) One access graph. (b) The access graph after coalescing variables u and v .

new interference graph, the coalesced vertex must interfere with all vertices that were adjacent to either u or v in the old interference graph.

Algorithm 1 uses function *RebuildL*, in line (22), to reconstruct the sorted list of edges (i.e. L) from the new access graph. Algorithm 1 ends when there are no more valid edges that can be chosen and no more vertices to coalesce. This condition is tested by using flags *sel* and *coal* in line (23) of Algorithm 1.

5 Complexity Analysis of CSOA

In this section, it is analyzed the time-complexity of CSOA in the worst case. In this analysis, consider m as the length of the access sequence, and n the number of variables considered for CSOA. In Algorithm 1, the complexity of *BuildAccessGraph* is $O(m + n^2)$. The sorting operation in line (2) takes $O(n^2 \log n)$. After this, the *repeat-until* loop can be executed at most $2(n - 1)$ times (at most $n - 1$ edges are selected and $n - 1$ coalescing operations are performed). The *repeat-until* loop is dominated by the *RebuildL* function, which is $O(n^2 \log n)$. So, this loop has complexity $O(n^3 \log n)$. Finally, the *BuildOffset* function takes $O(n^2)$ time. Therefore, CSOA has time-complexity $O(m + n^3 \log n)$.

It is worth bringing to attention that this is a worst case analysis, and that in practice one can expect a better runtime for CSOA.

6 Example of Coalescing SOA

To better illustrate CSOA, consider the code fragment of Figure 4(a). Each program point in the code shows the set of live variables (assuming that only g is live at the exit of the fragment). When Algorithm 1 is applied to this example, it receives as input the interference graph shown in Figure 4(b) and the access sequence (Figure 4(c)).

As the algorithm proceeds, it produces at each iteration the access graphs shown in Figures 4(d)-(j), after which it reaches the final memory assignment. The edges selected during the assignments are highlighted. The final memory layout is shown in Figure 4(k).

Although not illustrated, the reader should remember that, whenever two vertices in the access graph are coalesced, these vertices are also coalesced in the interference graph. In the first iteration (Figure 4(d)), edge (a, c) is selected, as no pair of vertices can be coalesced to produce saving as high as 2. In the

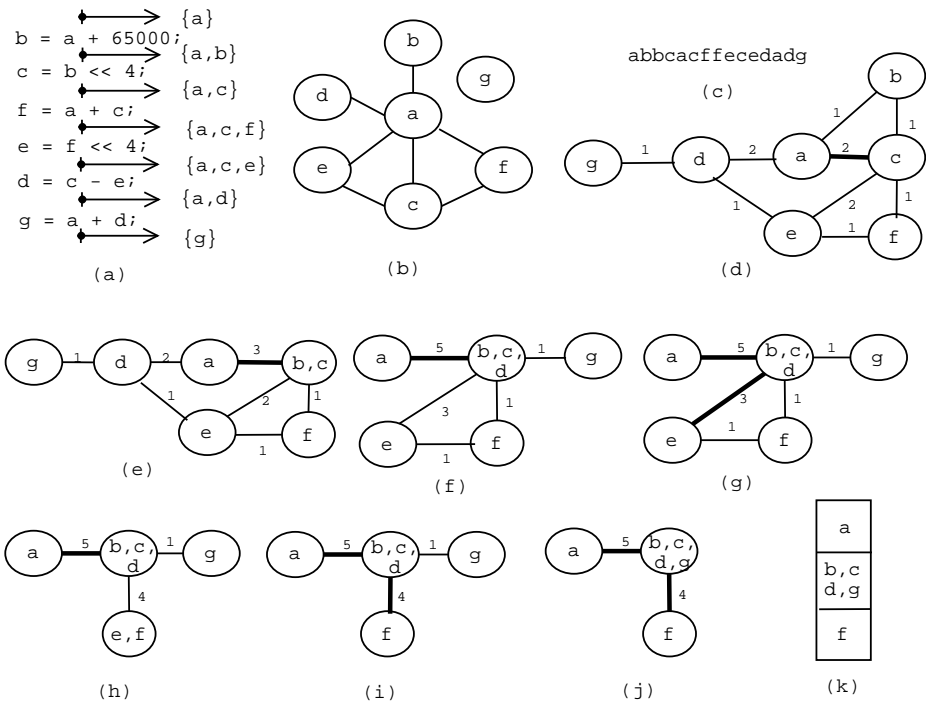


Fig. 4. (a) A fragment of C code with liveness information at each point. (b) The interference graph of the variables. (c) The access sequence of this fragment. (d)-(j) The access graphs resulting after each iteration of the algorithm. (k) The memory layout. Selected edges are highlighted.

next iteration, the best choice is to coalesce vertices b and c , given that this operation results in a saving of 2 (corresponding to the edges (a, b) and (b, c)). The new vertex (bc) becomes adjacent to the vertices that were adjacent to b or c in the previous access graph, that is, a , e and f . Notice that the weight of the edge between a and (bc) becomes 3, the summation of the weights of edges (a, b) and (a, c) in the previous graph. The algorithm proceeds, choosing between coalescing two vertices or selecting an edge, until no more operations are possible, thus resulting in Figure 4(j).

The final cost of applying CSOA to this example is zero, as all edges in the final access graph are selected. Notice that this example is the same as the one in Section 3, for which Liao’s algorithm produces a final cost of four.

7 Experimental Results

In this section, we compare CSOA with four other approaches to SOA. We use the MediaBench benchmark [13] to evaluate the five heuristics.

We implemented our approach using OffsetStone [15, 22], a toolset used to test and evaluate OA algorithms. All benchmark programs were compiled with the Lance [8] compiler front-end, which translates the C source code into three-address code intermediate representation. The code in this intermediate representation was then optimized through a combination of the following optimizations: constant folding, constant propagation, jump optimization, loop invariant code motion, induction variable elimination, global common subexpression elimination, dead code elimination and copy propagation. Access sequences were then extracted from each basic block, and basic block access graphs merged on a function basis. The live ranges of the variables were calculated doing liveness analysis [1] in the intermediate representation after the optimizations described above.

In Table 1, we compare CSOA with four other approaches. We measured the percentage of the number of update instructions inserted by each method, with respect to the number of update instructions inserted by SOA-Liao, the algorithm described in [20]. The four other methods used in the comparison are: SOA-TB, the heuristic described in [18]; SOA-GA, the heuristic described in [17]; SOA-INC-TB [15], the combination of two SOA algorithms, SOA-incremental [4] and SOA-TB [18], and SOA-Color, the optimization described in [25]. The SOA-Color algorithm constructs the interference graph based on the live ranges, and then uses Kempe’s [10] coloring heuristic to coalesce variables that do not interfere. After this, it is applied SOA-Liao heuristic in the coalesced variables.

Notice from Table 1 that CSOA reduces, on average, the number of update instructions to 32.1% of the SOA-Liao cost. This is a significant improvement over the previous algorithms. The best of the other algorithms (SOA-Color) reduced the offset cost, on average, to 51.2% of the SOA-Liao cost. So, the difference between SOA-Color and CSOA, in relation of SOA-Liao, is 19%. This means that, in relation to SOA-Color, CSOA produces 37.3% fewer update instructions than this technique. We believe that this exceptional improvement is

Table 1. Offset costs relative to Liao's algorithm cost.

Benchmarks	TB	GA	INC-TB	SOA-Color	CSOA
adpcm	89.1%	89.1%	89.1%	55.8%	45.6%
epic	96.8%	96.6%	96.6%	74.3%	50.2%
g721	96.2%	96.2%	96.2%	50.6%	27.9%
gsm	96.3%	96.3%	96.3%	26.6%	19.4%
jpeg	96.9%	96.7%	96.7%	52.6%	32.2%
mpeg2	97.3%	97.1%	97.2%	60.2%	34.3%
pegwit	91.1%	90.7%	90.7%	75.2%	38.8%
pgp	94.9%	94.8%	94.8%	55.0%	32.2%
rasta	98.6%	98.5%	98.5%	33.2%	21.1%
Average	95.2%	95.1%	95.1%	51.2%	32.1%

due to the fact that CSOA does not coalesce variables indiscriminately, but tries to make adjacent in memory variables that have many consecutive accesses. This increases the closeness between variables that are accessed consecutively. CSOA, in opposition to other techniques that naively coalesce variable slots [25], wisely takes advantage of coalescing to reduce both the SOA cost and the memory requirement. This is achieved by simultaneously performing variable coalescing while solving SOA.

Table 2 lists, for each benchmark, the following measurements for SOA-Color and CSOA relative to SOA-Liao's algorithm: the percentage of program code size, the percentage of data memory size, and the percentage of total memory size (code plus data). For SOA-Liao's method Table 2 shows the number of memory words used with code, data and code plus data, which is represented by C+D in Table 2. In the percentage of data memory size, only statically allocated variables were considered. To estimate the number of instructions, each three-address IR instruction of the benchmarks was considered to be stored in one memory word. This way, the real effect of the two techniques in the memory size can be better analyzed, since both memory, data and code, are reduced. The data area is reduced by coalescing, and the code by minimizing the number of update instructions.

From Table 2, one can observe that SOA-Color reduces the size of the memory used to store variables to 86.1%, when comparing to the other five methods [20, 18, 17, 15, 4] that do not perform coalescing, while our method reduces to 88.3%. On the other hand, our method reduces the size of the code memory to 89.6%, when comparing with the size of the memory code resulting of SOA-Liao, while SOA-Color reduces to 92.8%. Considering both memory segments, our method reduces memory to 92.3% and SOA-Color to 91.8% when comparing with total size of memory resulting from SOA-Liao. Though CSOA results in 0.5% more memory area than SOA-Color, it produces 37.3% fewer update instructions, thus resulting in a better performance.

Table 2. Number of memory words of code, data and code plus data, using SOA-Liao algorithm, and percentage of memory savings relative to Liao's algorithm when using SOA-Color and CSOA.

Bench.	SOA-Liao			SOA-Color			CSOA		
	Code	Data	C+D	Code	Data	C+D	Code	Data	C+D
adpcm	601	29038	29639	89.9%	99.4%	99.2%	87.5%	99.5%	99.3%
epic	14541	137936	152477	92.0%	97.3%	96.8%	84.6%	97.8%	96.6%
g721	2786	2266	5052	90.7%	55.9%	75.1%	86.4%	61.9%	75.4%
gsm	13963	15882	29845	94.3%	71.8%	82.3%	93.7%	76.3%	84.4%
jpeg	65630	53748	119378	94.9%	78.3%	87.4%	92.7%	83.4%	88.5%
mpeg2	31354	148106	179460	92.8%	94.7%	94.4%	88.0%	95.9%	94.6%
pegwit	14685	85217	99902	97.4%	95.6%	95.9%	93.6%	96.9%	96.4%
pgp	618243	148204	766447	99.6%	94.4%	98.6%	99.4%	95.6%	98.7%
rasta	18691	4346642	4365333	84.4%	99.9%	99.8%	81.6%	99.9%	99.9%
Average	17002.0	73550.8	119257.4	92.8%	86.1%	91.8%	89.6%	88.3%	92.3%

Finally, Table 3 shows the number of temporary variables (among those considered for SOA) in each program.⁶ Observe through these numbers that, on average, 64.1% of the variables are temporaries. Memory stored temporaries are very common in DSP architectures, given their reduced number of general-purpose registers. Thus, temporary allocation plays an important role in the final code performance, reinforcing our perception that there are many opportunities for CSOA to coalesce variables in DSP code, as shown by the experimental results.

Table 3. Percentage of temporary variables, considering as temporary a variable that is alive only in one basic block.

Benchmarks	%Temporaries
adpcm	59.6%
epic	48.1%
g721	80.7%
gsm	86.6%
jpeg	65.2%
mpeg2	65.6%
pegwit	72.1%
pgp	67.5%
rasta	43.6%
Average	64.1%

⁶ We consider here as temporaries those variables whose liveness are restricted to a single basic block.

Another result measured is that in 43.9% of the instances of all benchmarks, CSOA resulted in zero cost. So, for at least this percentage of the instances CSOA resulted in the optimal cost, and we believe this percentage to be significantly higher in fact, as many of the other instances may have an optimal cost greater than zero.

8 Conclusions and Future Work

In this paper we proposed a heuristic to solve the Simple Offset Assignment (SOA) problem based on coalescing memory variable slots. The experimental results show that our method (CSOA) eliminates, on average, 37.3% of the update instructions when comparing with the SOA-Color. Another important side effect of our technique is the reduction in the size of the memory layout to 92.3% when comparing with the SOA-Liao approach. The large presence of temporaries in DSP programs and the increased closeness resulting from the coalescing technique seem to explain well these exceptional numbers.

In this paper, we only addressed the SOA problem. We are currently investigating the use of coalescing to partition the access graph in the case of the General Offset Assignment (GOA) problem.

9 Acknowledgments

This work was partially supported by FAPESP (2000/15083-9), and by fellowship grant FAPESP (01/12762-5). We also thank the reviewers for their comments.

References

- [1] A. V. Aho, R. Sethi, and J. D. Ullman. *Addressing Modes for Fast and Optimal Code Generation*. 1987.
- [2] Guido Araujo, Guilherme Ottoni, and Marcelo Cintra. Global array reference allocation. *ACM Trans. on Design Automation of Electronic Systems*, 7(2):336–357, April 2002.
- [3] Guido Araujo, Ashok Sudarsanam, and Sharad Malik. Instruction set design and optimizations for address computation in DSP architectures. In *Proc. of the 9th. ACM/IEEE International Symposium on System Synthesis*, pages 102 – 107, November 1996.
- [4] Sunil Atri, J. Ramanujam, and Mahmut Kandemir. Improving offset assignment for embedded processors. *Lecture Notes in Computer Science*, 2017, 2001.
- [5] David H. Bartley. Optimizing stack frame accesses for processors with restricted addressing modes. *Software - Practice and Experience*, 22(2):101–110, 1992.
- [6] Yoonseo Choi and Taewhan Kim. Address assignment combined with scheduling in DSP code generation. In *Proc. of the 39th Design Automation Conference, DAC 2002*, 2002.
- [7] Marcelo Cintra and Guido Araujo. Array reference allocation using ssaform and live range growth. In *Proc. of the ACM SIGPLAN 2000 LCTES*, pages 26–33, June 2000.

- [8] LANCE Retargetable C compiler. <http://ls12-www.cs.uni-dortmund.de/lance/>.
- [9] Erik Eckstein and Andreas Krall. Minimizing cost of local variables access for DSP-processors. In *Proc. of the ACM SIGPLAN 1999 LCTES*, 1999.
- [10] A. Kempe. On the geographical problem of four colors. *Amer. J. Math*, 2, 1879.
- [11] Nakaba Kogure, Nobuhiko Sugino, and Akinori Nishihara. Memory address allocation method with ± 2 update operations in indirect addressing. In *European Conference on Circuit Theory and Design (ECCTD)*, 1997.
- [12] J. B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society*, 7:48–50, 1956.
- [13] Chunho Lee, Miodrag Potkonjak, William H., and Mangione-Smith. Mediabench: A tool for evaluating and synthesizing multimedia and communications systems. In *Proc. of the 30th Annual International Symposium on Microarchitecture (Micro 30)*, December 1997.
- [14] Rainer Leupers. Code generation for embedded processors. In *International System Synthesis Symposium*, 2000.
- [15] Rainer Leupers. Offset assignment showdown: Evaluation of DSP address code optimization algorithms. In *Proceedings of the 12th International Conference on Compiler Construction*, April 2003.
- [16] Rainer Leupers, Anupam Basu, and Peter Marwedel. Optimized array index computation in DSP programs. In *Proc. of the Asia South Pacific Design Automation Conference (ASP-DAC)*. IEEE, February 1998.
- [17] Rainer Leupers and Fabian David. A uniform optimization technique for offset assignment problems. In *Proc. of the International Symposium on System Synthesis (ISSS)*, pages 3–8, 1998.
- [18] Rainer Leupers and Peter Marwedel. Algorithms for address assignment in DSP code generation. In *International Conference on Computer-Aided Design (ICCAD)*, pages 109–112, 1996.
- [19] Stan Liao. *Code Generation and Optimization for Embedded Digital Signal Processors*. PhD thesis, Massachusetts Institute of Technology, 1996.
- [20] Stan Liao, Srinivas Devadas, Kurt Keutzer, Steven Tjiang, and Albert Wang. Storage assignment to decrease code size. *ACM Transactions on Programming Languages and Systems*, 18(3):235–253, May 1996.
- [21] Steven S. Muchnick. *Advanced Compiler Design and Implementation*. Morgan Kaufmann, 1997.
- [22] OffsetStone. <http://www.address-code-optimization.org>.
- [23] Guilherme Ottoni, Sandro Rigo, Guido Araujo, Subramanian Rajagopalan, and Sharad Malik. Optimal live range merge for address register allocation in embedded programs. In *Proceedings of the 10th International Conference on Compiler Construction, CC2001, LNCS 2027*, pages 274–288. Springer, April 2001.
- [24] Amit Rao and Santosh Pande. Storage assignment optimizations to generate compact and efficient code on embedded DSPs. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 128–138, 1999.
- [25] Ashok Sudarsanam, Stan Liao, and Srinivas Devadas. Analysis and evaluation of address arithmetic capabilities in custom DSP architectures. In *Design Automation Conference*, pages 287–292, 1997.
- [26] Bernhard Wess and Martin Gotschlich. Optimal DSP memory layout generations a quadratic assignment problem. In *Int. Symp. on Circuits and Systems (ISCAS)*, 1997.