

Data Layout Optimization for the LCC Compiler

Viktor Shamparov*, Murad Neiman-zade†

AO "MCST", 24 Vavilova str., Moscow, 119334, Russia

MIPT, 9 Institutskiy per., Dolgoprudny, Moscow Region, 141701, Russia

Email: *Victor.E.Shamparov@mcst.ru, †Murad.I.Neiman-zade@mcst.ru

Abstract—In this research-in-progress report, we propose a novel approach to unified cache usage analysis for implementing data layout optimizations in the LCC compiler for the Elbrus and SPARC architectures. The approach consists of three parts. The first part is generalizing two methods of estimating cache miss amount and choosing more applicable one in the compiler. The second part is finding an applicable solution for the problem of cache miss amount minimization. The third part is implementing this analysis in the compiler and using analysis results for data layout transformations.

Index Terms—Compilers, Compiler Optimization, Cache Analysis, Data Layout Transformation.

I. INTRODUCTION

Improving computer resources usage efficiency by a program is one of the main tasks for optimizing compilers. Particularly, improving memory usage is especially important because hardware developers have introduced multi-level intermediate memory, called cache memory, due to the growing performance difference between memory and CPU. Cache memory capabilities must be used efficiently.

Cache memory is structured for using the following program properties effectively: *temporal locality* and *spatial locality*. *Temporal locality* means that the program often works with the same data in memory. *Spatial locality* means that the program is likely to work with adjacent data. Thus, to make compiled program use cache memory efficiently, the compiler must improve these two programs' properties.

Nowadays, compilers optimize the programs' temporal locality well by loop optimizations, but optimizing spatial locality is more complicated since it requires choosing the correct data structures for the program. Therefore, optimizing spatial locality is often entrusted to the programmer, although data location optimizations are implemented for some relatively simple cases.

In this article, we describe the ongoing research on cache memory usage for the further development of a high-quality automatic cache usage analysis in the compiler for applying an optimal set of data layout optimizations.

The article is organized as follows. In section 2, we substantiate the potential effect of optimizing data layout. In section 3, we state the problem. In section 4, we analyze papers on this topic and related ones. In section 5, we propose further research approach. In section 6, we describe current progress. Finally, in section 7, we provide a conclusion.

II. MOTIVATION

It is known that part of program execution time is spent waiting for data from memory. This is especially evident for processors with in-order execution. They have fewer opportunities to mask this wasted time by executing other instructions than processors with out-of-order execution.

To illustrate this problem and determine the potential effect of optimization, we measured the percentage of test execution time from SPEC CPU benchmark packages that the processor spends waiting for data from memory. This data is shown in Table I. We used a computer with an Elbrus-4C processor for measurement. It has VLIW ISA, in-order execution and two-level cache memory. Benchmarks were compiled with peak options.

The table shows that more than 10% of the execution time is spent waiting for data from memory in 92 from 172 launches, which is more than a half.

Some of this spent time is due to inefficient use of cache memory. Mainly, these inefficiencies are:

- 1) loading unnecessary for further work data into the cache, which fact is a violation of spatial locality;
- 2) conflicts between different data chunks due to hitting the same cache set.

For example, it was found during our previous work that it is possible to reduce the number of cache misses with the help of optimization called Structure Splitting [1]. This optimization improves the spatial locality of the program in some cases. Such CPU pipeline stalls number decrease and consequent execution speeding up are shown in the Table II.

From this example, it can be seen that at least some of the losses due to waiting for data can be removed by data layout transformations improving spatial locality. These transformations require unified analysis for an effective combination.

III. PROBLEM STATEMENT

Thus, we need to:

- 1) Theoretically analyze cache memory usage by programs and develop a method of solving the problem of minimizing time losses based on this theoretical analysis.
- 2) Based on theoretical results, make applicable automatic analysis in the LCC compiler for the Elbrus and SPARC ISA.

TABLE I
NUMBER OF BENCHMARK LAUNCHES FROM SPEC CPU PACKAGES THAT
USE MORE THAN 10% OF TIME TO WAIT FOR DATA

Set	Part of time	Number launches
1995	10...15%	12
1995	15...20%	4
1995	20...25%	0
1995	25...30%	1
1995	≥ 30%	0
1995	Total in set	37
2000	10...15%	6
2000	15...20%	6
2000	20...25%	6
2000	25...30%	1
2000	≥ 30%	6
2000	Total in set	44
f2006	10...15%	4
f2006	15...20%	1
f2006	20...25%	1
f2006	25...30%	1
f2006	≥ 30%	2
f2006	Total in set	20
i2006	10...15%	3
i2006	15...20%	1
i2006	20...25%	3
i2006	25...30%	3
i2006	≥ 30%	12
i2006	Total in set	35
f2017	10...15%	1
f2017	15...20%	2
f2017	20...25%	3
f2017	25...30%	0
f2017	≥ 30%	1
f2017	Total in set	16
i2017	10...15%	1
i2017	15...20%	2
i2017	20...25%	0
i2017	25...30%	3
i2017	≥ 30%	6
i2017	Total in set	20
All	10...15%	27
All	15...20%	16
All	20...25%	13
All	25...30%	9
All	≥ 30%	27
All	Total	172

TABLE II
CPU PIPELINE STALLS NUMBER DECREASE AND FOLLOWING PROGRAM
EXECUTION SPEEDING UP

Benchmark	SPEC CPU package	CPU pipeline stalls number decrease	Speed-up
181.mcf	2000	27%	26%
429.mcf	2006	19%	13%

- 3) Implement in the same compiler a set of data layout transformations, which transform data layout of a program based on the analysis results.

In this case, it is necessary to take into account some restrictions arising from the fact that the implementation is planned in the form of compiler optimizations:

- 1) Various data structures need to be handled correctly. Particularly, they are:
 - a) Arrays, structures and their combinations.

- b) Various data structures that use pointers to other elements internally and allocate memory for new elements via `malloc` and similar memory allocation functions. For example, lists and trees.
- 2) We need to handle data structures altogether, as their transformations may conflict with each other. Therefore, it is necessary to analytically process not only regular access to memory but also random access.
 - 3) Analysis and transformations must be static (in the compiler) but can be supported with runtime libraries and special profiling, but not memory access trace.
 - 4) Developed analysis and transformations must correctly work in modular build mode.

IV. RELATED WORK

Several works on related topics have already been written, but each of them does not solve assigned tasks entirely due to different reasons.

Chris Lattner proposed automatic Data Structure Analysis to detect data structures whose elements are allocated on the heap in his thesis "Macroscopic Data Structure Analysis and Optimization" [2]. Using the results of this analysis, he proposed a compiler optimization called Automatic Pool Allocation with runtime support, designed to group the elements of such data structures in specific regions of the heap, which improves the spatial and, in some cases, temporal locality of the program. In addition, he offered several optimizations for code already optimized in this way.

Unfortunately, there is no explicit cache memory usage analysis in Lattner's work.

Christopher Haine in his thesis "Kernel optimization by layout restructuring" [3] offered an analyzer, which detects accessing memory regularly simple data structures like structures and arrays and proposes layout transformations using heuristics data. This analysis is separated from the compiler. In addition, this analyzer provides user with information about the complexities of code vectorization. For our purposes, this work is not suitable since there is no explicit cache memory usage analysis.

Mostafa Hagou and Caroline Tice in their article "Cache Aware Data Layout Reorganization Optimization in GCC" [4] proposed several improving spatial locality optimizations of structures and arrays of structures: Structure Peeling, Structure Splitting, and Field Reordering. These optimizations were later implemented in the GCC compiler. Although the authors limited themselves to working with structures, they implemented an analysis handling every structure access, not just regular access. During optimization, particular Field Reference Graphs are built for each analyzed structure for each procedure. Field Reference Graph (FRG) is an analogue of a control-flow graph, where nodes contain operations accessing fields of the analyzed structure and arcs contain information about the amount of data loaded into the cache between nodes. In fact, this is an implicit analysis of cache memory usage. Further, after processing, this information is used in heuristics to apply the specified optimizations and reduce the computational complexity of further algorithms.

This approach can potentially be used for explicit cache memory usage analysis, provided it is generalized for working on all program data in all procedures.

Ghosh et al. [5] and Fraguera et al. [6] suggested more explicit techniques for cache memory usage analysis for regular access cases.

Ghosh et al. [5] proposed to compose and solve systems of linear Diophantine equations to estimate the number of cache misses for each cycle. They implemented this algorithm in the SUIF compiler and implemented the choice of padding size in the Array Padding optimization as an example. However, they did not implement an automatic solution of systems in parametric form - only a particular solution for Array Padding. In addition, this approach was created only for regular memory access.

An alternative approach was suggested by Fraguera et al. [6] for regular memory access. It was improved by Andrade in [7] thesis for some cases of irregular memory access: regular access under condition and access to an array, where the indices are read from another array. This approach is based on estimating the probability of cache misses in each analyzable cycle using Probabilistic Miss Equations (PME) generated from regular access characteristics and cache memory characteristics. To do this, for each processed access in the loop, a partial Probabilistic Miss Equation is built, and then they are combined into a complete equation for the loop or loop nest. This complete equation gives an estimation of cache misses amount. In addition, they did not offer any solution to the problem of minimizing cache misses amount and did not handle random memory access. Thus, the PME approach can potentially be applied for explicit cache memory usage analysis, provided the analysis is generalized for working for all irregular memory access.

Data layout transformations were described in many papers. Particularly, a small catalogue of such transformations was created in the article [8]. Following transformations are listed in this article:

- 1) Array Padding - adding padding between arrays to reduce number of conflicts between arrays;
- 2) Array Merging - element-wise arrays merging;
- 3) Array Transpose - changing dimensions order of an array by analogy with transposing a matrix.

In addition to these, in the above-mentioned article [4] and thesis [2] some other transformations were described:

- 1) Structure Peeling - splitting an array of structures element by element into several arrays;
- 2) Structure Splitting - splitting an array of structures element by element into several arrays and addition of links between the elements corresponding to the initial element;
- 3) Field Reordering - changing order of fields inside the structure;
- 4) Automatic Pool Allocation - replacing memory allocation for data structure elements in the heap with memory allocation in a specific pool.

V. PROPOSAL

In this paper, we propose the following approach to research.

Firstly, it is proposed to investigate and compare following methods for cache memory usage analysis:

- 1) the method described in [4] using FRG graphs, generalized for working with all program data in all procedures;
- 2) the method described in [6], [7] using the Probabilistic Miss Equations, generalized for the case of random access.

We propose to choose one method for cache memory usage analysis that is more suitable for implementation in the compiler. The selection criterion is the accuracy of the estimation of cache misses amount. Another selection criterion is analysis time.

Further, we propose to develop an analytical or another compiler-applicable method for solving the problem of minimizing the obtained estimation of the cache misses amount using data layout transformations. This problem is a discrete optimization problem, in which the objective function is the dependence of the cache misses amount on the applied data layout transformations, and a countable set of feasible solutions is the data layout transformations.

Finally, based on the developed analysis method and the method for solving the problem of minimizing the cache misses amount, it is proposed to implement automatic analysis in the compiler that controls a set of data layout transformations. Also, we will need to implement missing transformations.

A. Generalizing FRG analysis

This method should be generalized for working on all program data in all procedures and provide an estimation of cache misses amount. To do this, based on the FRG graph for structures, we need to make a generalized graph for structures, arrays, their combinations and other data structures. Such graphs need to be created for each program object. Let us call such graphs Object Reference Graph - ORG. In addition, we need to build a general RGP (Reference Graph in Procedure) graph consisting of all memory accesses in the procedure and including profile information. So any ORG graph in a procedure contains a subset of RGP nodes; therefore, using RGP, one can estimate the probabilities of transitions through various ORG arcs and cache memory usage characteristics between ORG nodes. In addition, RGP is required to analyze conflicts between different data structures.

It is required to determine the probability of a particular cache line being evicted from the cache memory to estimate the probability of a cache miss in each ORG node. Since the probability of preempting a particular cache line depends on the amount of memory loaded into the cache in the general case in a complex way, it is better to store on the arcs of ORG graphs, not the amount of memory loaded into the cache, but the probability of preempting a particular cache line.

To estimate the probabilities, one must know in which memory regions the memory addressed by each pointer is located and the size of these memory regions. To obtain this

information, we need to use pointer analysis and a particular version of the profile, which collects data on the size of the allocated memory.

B. Generalizing PME analysis

To use this method, we need to generalize it for processing irregular memory access.

For this, we need to:

- 1) Create a way to calculate cache misses probability for random access.
- 2) Generalize PME to those cases of near-regular access where it is possible to estimate cache misses amount more accurately than using a random access model.
- 3) Combine PME for regular access and ones for random access.
- 4) Use the developed techniques for estimating cache misses amount for the entire code, not just for loops.

To estimate the probabilities, one must know in which memory regions the memory addressed by each pointer is located and the size of these memory regions. To obtain this information, we need to use pointer analysis and a particular version of the profile, which collects data on the size of the allocated memory.

VI. CURRENT PROGRESS

In the work [1] we described the particular version of data layout transformation called Structure Splitting, which we had implemented in the LCC compiler for the Elbrus and SPARC architectures. In addition, in this compiler Structure Peeling, Array Transpose, Array Linearization, and Array Padding have already been implemented.

A. Cache miss probability for random access

To generalize the PME-based analysis, a method was created for calculating the cache misses probability for random access. It is supposed that the memory region is known for this access, but the address of the region beginning is unknown. PME will be merged with this method.

The method is based on determining cache state transformations for each memory access operation. For this, the operations are traversed sequentially in the basic blocks of the procedure, and the transformations on the code blocks are combined according to the probabilities in the profiled control-flow graph. Any operation of the procedure is traversed once for random access case. For any other case number of single operation traversals must be $O(1)$ due to the analysis applicability requirement.

The cache state notation for the general case of regular and random access has not been determined yet, but the following notation has been chosen for the random access model: matrix \mathbf{P} composed of N vectors \mathbf{P}_i corresponding to N memory regions. Each vector has $\mathcal{S} + 1$ size, where \mathcal{S} is the number of cache lines in the cache. The element of the matrix \mathbf{P}_{ij} is the probability that exactly j lines corresponding to the area i are stored in the cache memory at the moment.

An example of the chosen cache state notation for three memory regions called \mathbf{a}_i , where $i = 1..3$, is shown in Table III. In the shown state it is implied that region \mathbf{a}_1 has no lines in cache with 100% probability. Also, probability of \mathbf{a}_2 taking all lines of cache is 90% and probability of \mathbf{a}_3 taking one line and \mathbf{a}_2 taking all other lines is 10%.

TABLE III
CHOSEN CACHE STATE NOTATION EXAMPLE \mathbf{P}_{ij} FOR THREE MEMORY REGIONS CALLED \mathbf{a}_i , $i = 1..3$

j	\mathbf{a}_1	\mathbf{a}_2	\mathbf{a}_3
\mathcal{S}	0%	90%	0%
$\mathcal{S} - 1$	0%	10%	0%
...
1	0%	0%	10%
0	100%	0%	90%

Let us introduce for each operation or code section c an operator for changing the state \mathbf{T}^c . If there was state \mathbf{P}^b before executing c , then state \mathbf{P}^a after executing c is: $\mathbf{P}^a = \mathbf{T}^c \mathbf{P}^b$. We require the following properties for the operator:

- 1) For a code section c , consisting of K consecutive code sections or operations c_1, \dots, c_K , the operator is a composition of operators for parts of the section: $\mathbf{T}^c = \mathbf{T}^{c_K} \dots \mathbf{T}^{c_1}$.
- 2) For a code section consisting of K alternative code sections or operations c_1, \dots, c_K with probabilities of passing through them p_1, \dots, p_K (for example, `if` block and `else` block), with $\sum_{j=1}^K p_j = 1$, the operator is a linear combination of operators for parts of the section: $\mathbf{T}^c = \sum_{j=1}^K p_j \mathbf{T}^{c_j}$.
- 3) Similarly, if during the execution of one operation op one of the K different state changes $\mathbf{T}_1^{op}, \dots, \mathbf{T}_K^{op}$ may occur with probabilities p_1, \dots, p_K , and $\sum_{j=1}^K p_j = 1$, the operator is a linear combination of their operators: $\mathbf{T}^{op} = \sum_{j=1}^K p_j \mathbf{T}_j^{op}$.

For the chosen matrix cache state notation, we also introduce an element-wise product \circ of the operator and coefficients.

Let us consider one memory access operation. It can cause three different outcomes:

- 1) Cache hit. In this case, cache state in the selected notation is not changed.
- 2) Cache miss with a conflict in the memory region. In this case, cache state in the selected notation does not change since it only stores the probabilities of having a certain amount.
- 3) Cache miss with a conflict with another memory region. A new line is loaded into the cache for the memory region the operation is working with. For one of the other memory regions, the line is evicted from the cache.

Thus, change in cache state for a single operation for a specific memory region can consist only in loading a new cache line for memory region, deleting cache line from the cache for memory region, or no changes for memory region. For such changes we introduce operators for the movement of cache state in selected notation:

- 1) \mathbf{M}^+ - moves the matrix values up by 1: if $\mathbf{P}^a = \mathbf{M}^+ \mathbf{P}^b$, then

$$\forall i \in 1 \dots N \mapsto \begin{cases} \mathbf{P}_{ij}^a = \mathbf{P}_{i(j-1)}^b, j = 0 \dots S-1 \\ \mathbf{P}_{iS}^a = \mathbf{P}_{iS}^b + \mathbf{P}_{i(S-1)}^b \\ \mathbf{P}_{i0}^a = 0 \end{cases}$$

- 2) \mathbf{M}^- - moves matrix values down by 1: if $\mathbf{P}^a = \mathbf{M}^- \mathbf{P}^b$, then $\forall i \in 1 \dots N$

$$\forall i \in 1 \dots N \mapsto \begin{cases} \mathbf{P}_{ij}^a = \mathbf{P}_{i(j+1)}^b, j = 1 \dots S, \\ \mathbf{P}_{i0}^a = \mathbf{P}_{i0}^b + \mathbf{P}_{i1}^b \\ \mathbf{P}_{iS}^a = 0 \end{cases}$$

- 3) \mathbf{M}^0 - does not move matrix values.

Writing down cache state change operator \mathbf{T}^{op} for operation, working with the memory region i , we get:

$$\mathbf{T}^{\text{op}} = \rho_{i+} \circ \mathbf{M}^+ + \rho_{i0} \circ \mathbf{M}^0 + \rho_{i-} \circ \mathbf{M}^-$$

where:

- 1) ρ_{i+} - matrix of coefficients for loading a new line of i into the cache; this matrix consists of a nonzero column for the i -th vector, other coefficients are equal to zero;
- 2) ρ_{i0} - matrix of coefficients for saving cache state as it is;
- 3) ρ_{i-} - matrix of coefficients for evicting a line from the cache when loading a new line of the i area into the cache; this matrix consists of nonzero columns for all vectors except the i -th.

An example of applying operator \mathbf{T}^{op} to cache state example above is shown in Table IV. Operation op accesses memory region \mathbf{a}_1 , so one line of \mathbf{a}_1 is loaded into cache and one line of \mathbf{a}_2 or \mathbf{a}_3 is evicted from the cache.

TABLE IV
RESULT OF APPLYING OPERATOR \mathbf{T}^{op} TO CACHE STATE FROM TABLE III
WHEN op WORKS WITH MEMORY REGION \mathbf{a}_1 ($i = 1$)

j	\mathbf{a}_1	\mathbf{a}_2	\mathbf{a}_3
S	0%	0%	0%
$S-1$	0%	$90\% + 10\% \cdot \frac{1}{S}$	0%
$S-2$	0%	$10\% \cdot \frac{S-1}{S}$	0%
...
1	100%	0%	$10\% \cdot \frac{S-1}{S}$
0	0%	0%	$90\% + 10\% \cdot \frac{1}{S}$

VII. CONCLUSION

Publications analysis showed that there is no unified solution to the problem of improving cache usage of compiled programs. In this paper, we propose a research approach, which can lead to a solution to this problem in compilers.

REFERENCES

- [1] V. E. Shamparov and A. L. Markin, "Structure splitting for elbrus processor compiler," *Programmnaya Ingeneria*, vol. 12, no. 2, p. 82–88, 2021.
- [2] C. Lattner, "Macroscopic Data Structure Analysis and Optimization," Ph.D. dissertation, Computer Science Dept., University of Illinois at Urbana-Champaign, Urbana, IL, May 2005, *See* <http://llvm.cs.uiuc.edu>.
- [3] C. Haine, "Estimation d'efficacité et restructuration automatisées de noyaux de calcul. (kernel optimization by layout restructuring)," Ph.D. dissertation, University of Bordeaux, France, 2017. [Online]. Available: <https://tel.archives-ouvertes.fr/tel-01841485>
- [4] M. Hagog and C. Tice, "Cache aware data layout reorganization optimization in gcc," in *Proceedings of the GCC Developers' Summit*, 2005, pp. 69–92.
- [5] S. Ghosh, M. Martonosi, and S. Malik, "Cache miss equations: A compiler framework for analyzing and tuning memory behavior," *ACM Trans. Program. Lang. Syst.*, vol. 21, no. 4, p. 703–746, Jul. 1999. [Online]. Available: <https://doi.org/10.1145/325478.325479>
- [6] B. B. Fraguera, R. Doallo, and E. L. Zapata, "Probabilistic miss equations: Evaluating memory hierarchy performance," *IEEE Trans. Comput.*, vol. 52, no. 3, p. 321–336, Mar. 2003. [Online]. Available: <https://doi.org/10.1109/TC.2003.1183947>
- [7] D. Andrade, "Systematic analysis of the cache behavior of irregular codes," Ph.D. dissertation, Departamento de Arquitectura e Tecnologías Multimedia, Universidade da Coruña, 2007. [Online]. Available: <http://hdl.handle.net/2183/18378>
- [8] M. Kowarschik and C. Weiß, *An Overview of Cache Optimization Techniques and Cache-Aware Numerical Algorithms*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 213–232. [Online]. Available: https://doi.org/10.1007/3-540-36574-5_10