



UNIVERSITÀ DEGLI STUDI DI PADOVA
FACOLTÀ DI INGEGNERIA
CORSO DI LAUREA IN
INGEGNERIA DELL'INFORMAZIONE

TESI DI LAUREA

PSORT:
AUTOMATED CODE TUNING

ADVISOR: PROF. ENOCH PESERICO STECCHINI NEGRI DE SALVI

CO-ADVISOR: DOTT. MARCO BRESSAN

CANDIDATE: GIOVANNI DI LIBERTO

ANNO ACCADEMICO 2010 - 2011

ABSTRACT

This thesis describes the design and implementation of an automated code tuner for *psort*, a fast sorting library for large datasets. Our work, motivated by the necessity of guaranteeing a high performance while keeping a low cost on the end user, provides a reusable and portable framework that can be easily extended to automatically tune virtually every portion of the source code, including code that has not yet been written. Experiments show that our system produces code which is significantly faster than original code, suggesting that *psort* should include it among its tools.

SOMMARIO

Questa tesi descrive la progettazione e la realizzazione di un ottimizzatore di codice automatico per *psort*, una libreria di ordinamento veloce per grandi moli di dati. Il nostro lavoro, motivato dalla necessità di garantire alte prestazioni mantenendo un basso costo sull'utente finale, fornisce una infrastruttura riutilizzabile e portabile che può essere facilmente estesa per ottimizzare in maniera automatica virtualmente ogni porzione di codice sorgente, incluso codice che ancora non è stato scritto. Gli esperimenti mostrano che il nostro sistema produce codice che è significativamente più veloce del codice originale, suggerendo che *psort* dovrebbe includerlo tra i suoi strumenti.

Contents

1	Introduction	1
1.1	External sorting	1
1.2	<i>psort</i>	2
1.2.1	Stage one	2
1.2.2	Stage two	5
1.2.3	Performance limitations	5
2	Automated code tuning	9
2.1	Overview	9
2.2	<i>ATLAS</i>	10
2.3	<i>FFTW</i>	12
2.4	<i>PHiPAC</i>	12
3	Implementation	15
3.1	Code Generator	16
3.1.1	Grammars and generation rules	17
3.1.2	Tags	17
3.1.3	Creation of tunable code	19
3.2	Critical optimizations	21
3.3	Local search	22
4	Results	25
5	Conclusions	31
A	Source code	33
A.1	cache_sorters.tun	33
A.2	inlines.tun	50
A.3	tuner.sh	50
A.4	cache_sorters_tuner.gen	51

Chapter 1

Introduction

This thesis describes the design and implementation of an automatic code generator for *psort*, a fast, stable *external sorting* software. This section provides a gentle introduction to the problem of *external sorting* (Section 1.1) before describing the *psort* library (Section 1.2) and its limiting factors (Section 1.3) which are partially addressed by this work.

1.1 External sorting

Sorting is one of the best known and most important problems in computer science, arising in virtually every application. Indeed, some estimates suggest that few decades ago the global amount of time spent in sorting amounted to approximately 25% of the global CPU time (see [4]); and, while this amount has probably decreased today, sorting still plays a crucial role. Due to its primary importance, it has also been adopted as the core task by benchmarks such as the Sort Benchmark (see [8]) to measure the evolution of computing power over time.

Perhaps the best known scenario for the sorting problem is the RAM model – where accessing a datum has always the same cost, independently of the address where it is stored. However, this is not generally true, and in some cases is very inaccurate. This is the case of *external sorting*, where the size of the *main memory* (which is assumed to be unlimited in the above mentioned RAM model) is not sufficient to hold the data at every step of the algorithm, which must then resort to a larger *external memory* of higher access time cost. This happens, for example, when sorting large amounts of data (which is becoming more the rule than the exception) from disk to disk, and in fact this is a main requirement of the Sort Benchmark competition [8]. In all these cases, the *hierarchical memory model* [5] captures the basic essence of the typical structure of modern machines, which consists of several layers of progressively faster (but more expensive and thus smaller) memory. This *memory hierarchy* is exploited by the typical spatial and temporal locality exhibited by real programs, which ideally perform as many accesses as possible at the highest levels and as few as possible at the lowest ones.

In many practical cases, and especially in the case of desktop-class machines, the memory hierarchy often boils down to three main levels: the CPU (L2)

cache, the main memory, and the external storage. Given the typical large gap between the main memory and the external storage (usually represented by mechanical hard disks) in terms of data transfer rate, it is crucial to provide an efficient access to the latter. However, depending on the structure of the input (more precisely, on the ratio between the length of record payloads to the length of record keys) efficient access to the main memory may become at least as crucial as efficient access to external storage. It is thus important to optimize access at every memory level in order to squeeze out a consistent fraction of the computing power of the machine.

1.2 *psort*

psort is a sorting library designed for large amount of data. According to the PennySort benchmark, *psort* was the fastest desktop-class external sorting library from 2008 to 2011[7, 8]. The input files are viewed as sequences of records of arbitrary size; these are sorted, according to an arbitrary infix (the key), by first obtaining individual "runs" of sorted data approximately sized as the main memory and then merging those runs into a single sorted file.

Towards the end of 2010 *psort* had dozens of parameters that needed to be chosen carefully in order to achieve high performance, and this amount is still increasing with the newer optimization for *multi-core* and *multi-disks*. So a manual configuration takes too long, and a user may not be inclined to afford this cost. Even a good programmer might not have in-depth knowledge of the hardware architecture underlying his programs, and knowing how a software will be executed is probably increasingly difficult in modern systems, especially those which are highly parallel [6]. Perhaps an *automated code tuning* of the source code could resolve this problem.

psort uses the methodology described in Section 1.1, in fact it divides the work into two stages. The first creates many sorted runs and the second merges those in a single output. There are many important aspects for each stage. In particular this thesis is focused on the low-level sorting routines in the file `cache_sorters.cpp`, that are the core of the first phase. The reasons are that *stage one* is the most complex and that it represents a serious bottleneck at the highest levels of the memory hierarchy.

1.2.1 Stage one

The stage one of *psort*, like shown in Figure 1.1, essentially involves reading from an input file sequences of data (runs) whose size can be chosen by the user; a data run can't be larger than the main memory. Then *psort* execute the sorting of the runs, saving these in the secondary memory. Finally stage two merges the sorted sequences into the output file, in some cases executing more additional merge steps. For I/O efficiency, *psort* allows the use of direct asynchronous I/O to minimise the cost of transferring data between the disks and a set of userspace buffers. A simple double-buffering technique ensures that the disks never fall idle: while the CPU reads/writes data from/to a buffer, the I/O subsystem reads/writes data from/to another buffer, thus completely overlapping I/O with computation. Unsurprisingly then, the number and size of the I/O buffers are crucial parameters that need to be set carefully.

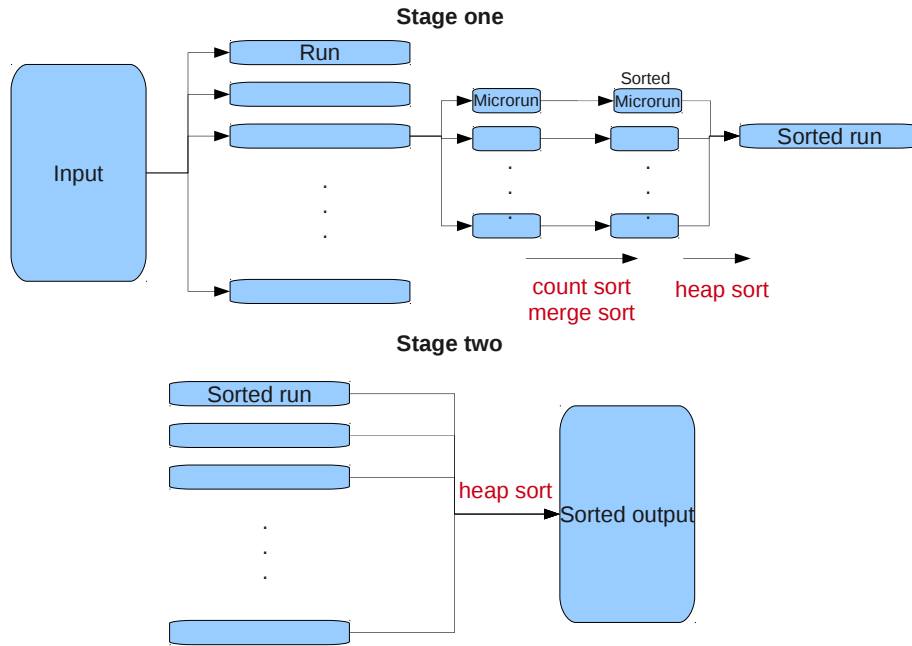


Figure 1.1: Data-flow representing the high-level execution of *psort* (data on cyan background, routines in red)

As described in Section 1.1, *psort* exploits data locality (and thus the hierarchical structure of the underlying machine) by splitting each run in smaller data blocks, called “*microruns*”, of size roughly equivalent to the L2 cache; so the first operation is to sort the microruns and, after that, these will be processed by an heap merge implementation called `kMerger` (*stage two* uses a similar solution to merge the sorted runs). This appears to be a method that could reduce the running time using the higher speed of the cache. Unfortunately there are hardware-related problems, for example caused by the associativity of the CPU cache [3], that can be overcome by a careful *hand-tuning*. At the base case, *psort* employs a small, efficient counting sort to sort short sequences of records whose length is determined by the parameter `chunk_size`. The size of these input sequences is determined by the compilation parameter `chunk_size`. These sorted chunks are then sorted by a carefully designed merge sort routine which consumes only $5/4$ of the total input size, rather than the “common” factor 2. There is also a variant of this algorithm, called `sorter_quasi_in_place_wave` (see Section 3.2), with the same space complexity but in some cases faster (although it is not clear in which cases). With a proper hand-tuning in compilation it is possible to choose the best implementation.

When *keys* are sufficiently small compared to the records’ remaining *pay-*

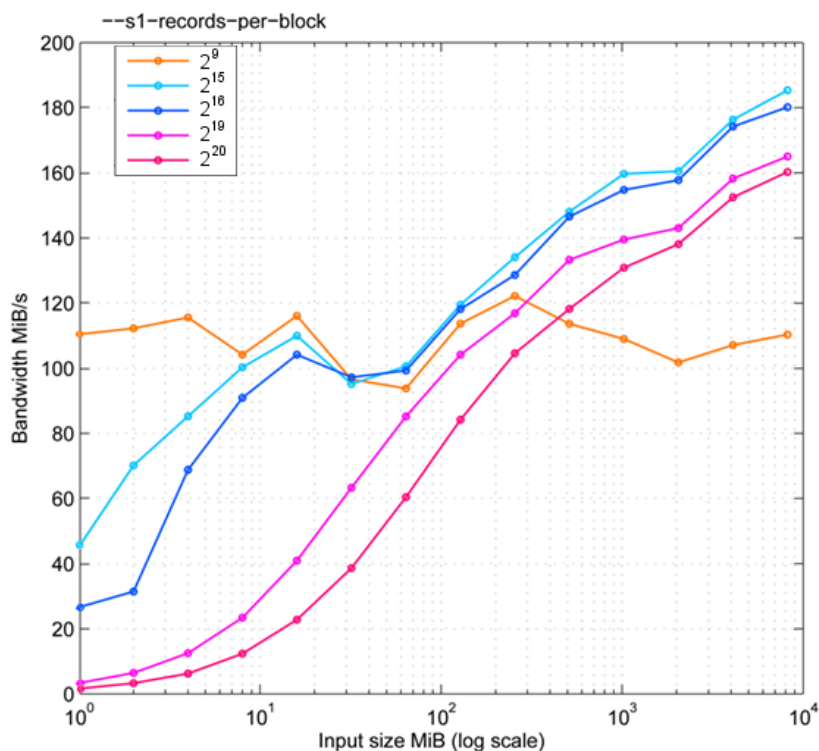


Figure 1.2: *psort* stage one: bandwidth as a function of the input size, for some values of `-s1-records-per-block`. Input consisted of 128-byte records with 8-byte keys, sorted using 3GB of allocated memory, direct I/O access and an I/O total buffer size of 300MB

loads, *psort* detaches keys from payloads, attaching to each key a pointer to the corresponding payload. The performance improvement is evident when payloads are longer than keys, that is a scenario arising in many applications (and in the Sort Benchmark competition). Thus parameter tuning must take into account that *psort* usually works with blocks of *extended keys*. For example, consider the parameter `-s1-records-per-block`, which indicates how many extended keys (or records if the payloads have not been separated by the keys) compose a microrun in the stage one. In order to improve the performance, the size of a microrun should be less than the L2 cache capacity, so the value of this parameter must satisfy the following inequality (up to constants):

$$size(microrun) \leq \frac{size(cache\ L2)}{size(extended\ key)}$$

Figure 1.2 shows that, for input sizes greater than about 50MiB, the optimal choice of `-s1-records-per-block` is the value that makes the microruns size around the size of L2 cache. In fact in this case the latter is 1MB and the optimal values appears to be 2^{15} and 2^{16} records; with an extended keys of $8 + 8 = 16$ bytes, this means respectively 0.5MB and 1MB for each microrun (other results shows

that for larger inputs the trend remains the same).

Starting from the considerations above, it can be understood why the allocation of the memory for I/O buffers is not so simple. About this problem, the following parameters give a high flexibility:

- *-s1-read-buffer-size*: size of each read buffer,
- *-s1-read-buffers*: number of read buffers,
- *-s1-write-buffer-size*: size of each write buffer,
- *-s1-write-buffers*: number of write buffers,
- *-s1-io-space*: fraction of memory allocated for stage one I/O buffers.

In order to tune well these variables it has to be considered that more buffers don't mean necessarily increase the performance; furthermore using more memory for the buffers the runs will be more and of smaller size, thus the second phase could be slower.

The output of the final pass is directly written to the I/O buffers, potentially recombining keys and payloads. Note that disk writes can be almost entirely overlapped with the merge pass(es) described above.

1.2.2 Stage two

The second phase (which only takes place if the data-set does not fit in main memory) is much simpler: w sorted runs at a time are streamed from disk and merged (with the same code that merges microruns in the first phase), and the output is streamed back to the disk. It is reasonable to use a single merge pass if the number of data runs (approximately the ratio between data and memory size) does not exceed the ratio between the size of the memory and that of one "efficient" read from disk. This means that a single merge pass is possible only if the memory is large enough to contain, for each run, an amount of data that is large enough to ensure an efficient transfer from/to the I/O subsystem – which depends on the hardware.

Data are read with direct asynchronous I/O into (userspace) dynamically sized buffers, one per run. When the amount of data in a buffer falls below a threshold (*-s2-read-threshold*, a parameter to be tuned) the buffer is "refilled" from the appropriate run. This can be highly ineffective, however, if data are not consumed uniformly, and in particular if they are consumed more rapidly from recently refilled buffers. For this reason, *psort* allows the user to specify, via the *-s2-geometry-factor* option, a "geometry factor" that determines the dimension of a run's buffer in comparison to the previous run's one, choosing a trade-off between safety and optimization. Figure 1.3 illustrates how this parameter influences the performance of stage two.

1.2.3 Performance limitations

The goal of this work is to design and implement an automatic system to optimize the source code of *psort*. This is actually part of a more general process, consisting of three separate tasks: *automated hardware detection*, *tuning at runtime*, and *automated code tuning*. The first task has the goal to detect some




Figure 1.3: *psort* stage two: bandwidth as a function of the parameter `--s2-rel-read-threshold` (run queue size threshold for read request, relative to a minimum default value of queue size), for various values of the parameter `--s2-geometry` (the geometry factor used in the stage two). The input, consisting in 128-byte records with an 8-byte keys, was sorted into 128 runs, each of 256MiB (overall input size: 32GiB).

hardware parameters; these values have an impact on (and are taken into account by) the other tasks. Tuning at runtime pertains the choice of optimal execution parameters, like the size of each run, the number of records in a single block, or the buffers size. Automated code tuning concerns instead lower level optimizations like *loop-unrolling* and the choice between *bitwise* and *logical compares*.

The first task faced was tuning at runtime. The initial step was to find optimal combinations of the parameters on different machines, with the goal of understanding the relation between the optimal value of a parameter and the hardware architecture. To this end, we have written a Bash script that allows to automatically execute many unsupervised tests and a *Matlab* script that plots automatically the output data. This allowed us to rapidly produce large amounts of custom plots, which were of substantial help to our research. Thus the most important parameters were identified; which means that the running time of *psort* is very sensible to the variation of these variables.

In the stage one, some parameters are assigned at the execution but some

others needs the re-compilation of the source code; the reason is that the routines in which they are used are repeated many times, so adding code to handle conditions at runtime could mean to reduce the performance visibly, and compilers produce significantly faster code if the values of these parameters are known at compile time. Thus the tuning of these parameters cannot be done at runtime. The performance optimization done using variables of this type is called *compiler-based auto-tuning* (see Section 2.1).

Important potential hurdles at this stage are associativity misses (in early experiments these reduced performance by as much as 20% [7]). However *psort* employs a careful data layout to minimize this problem, and the tuning at runtime tries to assign a reasonable value to each parameter using the results of automated hardware detection.

For all these reasons, it is crucial to provide *psort* with an automatic code generation systems, allowing for an efficient and effective generation of code which substitutes the costly (and boring!) manual tuning. The next section is devoted to this.

Chapter 2

Automated code tuning

Automated code tuning is the process with the purpose of finding an efficient solution (which typically means reducing the running time of the resulting code), evaluating different implementations of the same algorithms or data structures; each optimization is generally related to one or more parameters. This process can be (and often is) carried out by a developer, but with the significant disadvantages of take a very long time, of being error-prone and is clearly (except at an extremely high cost) neither portable nor reusable. The values of the working parameters can be calculated using an appropriate model or with an empirical search. Therefore, automated code tuning is the natural choice for high-performance software like *psort*. This chapter presents an overview of some existing solutions (Section 2.1) and a more detailed description of three of them (Sections 2.2-2.4).

2.1 Overview

Modern microprocessors can achieve high performance but this sometimes requires extensive machine-specific hand-tuning (e.g. [6]). There are many factors to be considered in the complex computer architectures: cache (multiple levels, capacity, associativity), memory latency, parallelism. Another point is that the compiler optimizations could be improved with some particular features. In fact, more sophisticated compiler optimizations, including loop unrolling and *software pipelining*, are either not performed or not very effective at producing the highest quality code [6]. Many of these improvements can often be converted in some coding guidelines that permits to write fast codes [12]. There are other low-level choice hard to be predicted, like the faster comparison between logical and bitwise. For these it is sometimes easier try and look which implementation is better. The optimizations at issue cannot be tuned with some execution parameters, like does tuning at runtime. In these cases for each value of such a parameter the code has to be modified and compiled. The automated code tuning has the goal to produce parametrized code generators rather than code by hand, to execute some tests and to obtain an optimal version of the software for the particular machine.

Some automated code tuning systems have already been developed, mainly in high-performance and scientific computing libraries. They usually tune once

when installed to determine the best settings and optimizations for the hardware they are running on. This approach is appropriate when the library will be used in different architectures and should provide good performance for all of them. These systems are clearly very application specific.

One approach is to use *analytical models* to determine an optimal combination of values. Unfortunately if there are too many parameters it could be difficult to build a suitable model that covers the main hardware behaviours [11]. Moreover there is the risk to obtain a non optimal solution, if the model is too simple; nevertheless the probability to obtain a bad model increases with more complex approximations. In a nutshell, when the optimization space is too large a model driven methodology could be intractable.

Another approach is to perform a *global empirical search* over the space of parameters values. For example, the linear algebra library *ATLAS* [10] provides portable performance across a range of CPUs architectures by using a database of precomputed optimizations for known architectures, and by execute a *global search auto-tuning* for new architectures. *PHiPAC* [12] and *OSKI* [16] use similar techniques.

An alternative manner is to use both a simple model and an empirical search. This methodology is called *local search* [11] because the parameters values are searched around the model answer. In this case is also needed a *search engine*, that permits to reduce the work executing only the tests that could give useful results.

Compiler-based auto-tuning is another important part of the problem. It concerns the improvement of some code optimizations, such as loop unrolling, to find the best for a particular program. For example, the compiler uses software pipelining more probably if the code is written following some particular rules. These indication could be dependent from the architecture or the compiler, so having an automated translation of the source code to one more suitable to the hardware used could be help the compiler in the optimization. This sort of operations are often join in a bigger search engine, that is responsible for identifying optimal values also for other parameters

Other examples of auto-tuned software are *FFTW* [13, 14] and *NukadaFFT* provide FFT (Fast Fourier Transform) implementations which use *automated code tuning* to improve performance. *NukadaFFT* runs on GPUs using *CUDA* and uses auto-tuning to choose various GPU optimizations. *FFTW* is a CPU implementation which uses auto-tuning to construct good execution plans. This library follows the idea of problem-by-problem optimization. An execution plan is created based on the hardware being used and the memory layout of the problem being solved. This plan can then be used when solving any problem of the same ‘shape’. *SPIRAL* is another auto-tuning based digital signal processing library.

2.2 ATLAS

The *ATLAS* (Automatically Tuned Linear Algebra Software) project is an ongoing research effort focusing on applying empirical techniques in order to provide portable performance [10, 11]. The original version uses a global search approach; instead the complete *ATLAS* system includes a search engine, a code generator and a test code that measure the execution time of some critical parts

of the source code. It also includes hand-written code for various routines, which may be used to produce the library on some machine if it is found to outperform the code produced by the code generator. This situation could happen using well-known architecture for which precomputed optimizations has been developed very well.




Figure 2.1: Empirical Optimization Architecture

For large parameter ranges, a full global search auto-tuning is very slow, so it is necessary to bound the search space. At the installation *ATLAS* executes a set of micro-benchmarks to measure certain hardware parameters such as the L1 data cache capacity, the number of registers, etc. The search engine bounds the search space using these values. In Figure 2.1 it is schematized the sequence of operations performed by a software like *ATLAS*. The first block, automated hardware detection, indicates the measure of hardware parameters. After that there is a loop in which the search engine decide if the final source code could be produced and compiled, or if more tests are needed. The *Compile-Execute-Measure* block summarize the compilation of the little piece of code, the execution, and the performance measures. Thanks to these data the search engine can decide when the optimal values have been found.

An example of search is about the matrix-multiplication routine: the *naive code* is divided into sub-problems which size is assigned by a parameter. A simple function like that could depends on 4 parameters, for each the search engine have to determine the search space. Thus in a big code dozens of optimization parameters could be found, so it is very important to define small search spaces because the complexity increases quickly. Another issue is if try each combination or if study the trend of each parameter with the others fixed; a good solution could be to test separately some group of related parameters. All these choices depends on the number of parameters, on the search space size, and on how much time the user can spend in this operation. The *brute-force search* of all possible valuations is tractable for little application, otherwise is very slow. The *ATLAS* solution is a global search strategy called *orthogonal line search*. It tries to find the optimal values of a function $y = f(x_1, x_2, \dots, x_n)$ by reducing the *n-dimensional* optimization problem into a sequence of *n 1-dimensional* optimization problems by sorting the parameters x_i , in some order, and optimizing them one at a time using reference values for parameters that have not been optimized yet [15].

2.3 *FFTW*

FFTW (Fastest Fourier Transform in the West) is a *C* subroutine library for computing the discrete Fourier transform (DFT) in one or more dimensions, of arbitrary input size, and of both real and complex data [13, 14]. This package was developed at MIT by Matteo Frigo and Steven G. Johnson. *FFTW*'s performance is portable: the same program will perform well on most architectures without modification. Thanks to this feature the performance is typically superior to that of other publicly available FFT software, and is even competitive with vendor-tuned codes (hand-tuned for a particular hardware)[14]. Like described above, *FFTW* follow the idea of problem-by-problem optimization. This means that an execution plan is based on the hardware being used and the memory layout of the problem being solved. This plan can then be used when solving any problem of the same 'shape'.

In *FFTW*, the computation of the transform is accomplished by an executor that consists of highly optimized, composable blocks of *C* code called *codelets*. A codelet is a specialized piece of code that computes part of the transform. The combination of codelets applied by the *executor* is specified by a special data structure called a plan. The plan is determined at runtime, before the computation begins, by a *planner* which uses a dynamic programming algorithm to find a fast composition of codelets. The planner tries to minimize the actual execution time, and not the number of floating point operations. Consequently, the planner measures the running time of many plans and selects the fastest.

The speed of the executor depends crucially on the efficiency of the codelets, but writing and optimizing them is a tedious and error-prone process. For this reason, *FFTW* generates the codelets automatically with a codelet generator. The main advantages of generating code are that it is simple to experiment with different algorithms or coding strategies, and it is easy to produce many long blocks of unrolled, optimized code.

The user interacts with *FFTW* only through the planner and the executor, as in the following example:

```

1. fftw_plan plan;
2. COMPLEX A[n], B[n];
3. /* plan the computation */
4.   plan = fftw_create_plan(n);
5. /* execute the plan */
6.   fftw(plan, A);
7. /* the plan can be reused for other inputs of size N */
8.   fftw(plan, B);

```

Figure 2.2: Simplified example of *FFTW*'s use. The user must first create a plan, which can be then used at will.

2.4 *PHiPAC*

PHiPAC is a methodology for developing Portable High-Performance linear algebra libraries in *ANSI C*. Its goal is to produce, with minimal effort, libraries for a wide range of systems [12].

The *PHiPAC* methodology has three components. First a generic model of current *C* compilers and microprocessors that provides guidelines for writing high-performance code. Second, rather than hand-code particular routines, it uses parametrized generators that produce optimized code. Third, it automatically tune code for a particular system by varying the generators' parameters and benchmarking the resulting routines.

For example, using the *PHiPAC* methodology in the production of a portable, *BLAS*-compatible matrix multiply generator, the resulting code can achieve over 90% of peak performance on a variety of current workstations, and is sometimes faster than the vendor-optimized libraries [12].

The *PHiPAC code generator* follows some guidelines that can be used directly to improve performance in critical routines. For example, loop unroll explicitly to expose optimization opportunities, or increase locality to improve cache performance. Follows some others coding guidelines (see [12]):

- using local variables, reorder operations to explicitly remove false dependencies;
- exploit multiple integer and floating-point registers;
- minimize pointer updates by striding with constant offsets;
- hide multiple instruction FPU latency with independent operations;
- balance the instruction mix;
- convert integer multiplies to adds;
- minimize branches, avoid magnitude compares.

Chapter 3

Implementation

This thesis has the goal to create a code that generates a machine-optimized version of *psort*. The first choice was to realize a full working source code; so the quantity of tuned code faded into the background, and the focus went on a single critical part of *psort*. In this way the code has achieved a stable version and it uses some well-defined rules. Thus it is possible to expand the optimization to other parts of the software.

The automated code tuning in *psort* uses the local search approach (Section 3.3): the concept is to use a very simple model of the system to determine the values of the parameters. The latter are called start values and the tests should be done around them. The solution determined by this method is optimal, or in any case a good one. This approach generates high-performance code without increasing search time dramatically like in a brute-force code tuning.

First of all it is important to define the main steps of the execution.

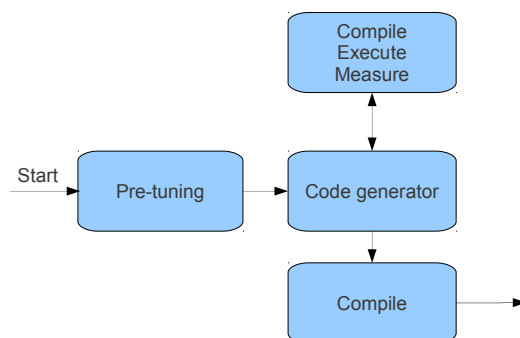


Figure 3.1: Empirical optimization architecture without a search engine

The chosen methodology is similar to the *ATLAS* empirical optimization architecture. In fact there are the automated hardware detection part, a code generator, a test and measure script, and the predisposition for a search engine. The latter isn't necessary because the optimization parameters have a small search space; however the code is thought to be easily adapted if required. The automated hardware detection code provides some hardware information about

cache, hard disk, CPU and RAM. It also makes specific tests that permits a further optimization of the code; for example it tells if bitwise comparisons are faster or slower than logical. All these information are used by the other blocks in Figure 3.1, and also during tuning at runtime. The following block is the code generator: given a file, its parameters, and their search spaces, it generates a big code with the different versions of the interested routines (candidate code). The last block compile this file, executes many tests, and eventually chooses the best parameters. After this phase the code generator produces the final code.

A deep execution of the automated code tuning could take many hours, so there is a parameter that permits to choose the level of detail of the tuning. However this parameter should be set at the maximum value because probably it's better to loose some hours and have a fast software instead of to have a fast tuning but a slow software; in fact the tuning is made only once.

The operations shown in Figure 3.1 are executed by the file `tuner.sh`, a *Bash script* with as input the *tuning-level* (an integer specifying the level of detail of the tests) and some optional parameters. Thus this script has the tuned file as output and, when all the tuning operations have been finished, it compiles the whole *psort*.

3.1 Code Generator

The code generator, like described above, generates from each file that has to be tuned another source code with all the possible versions of the interested routines in the search space. The choice was to use a pre-existent software as support. The most suitable appeared to be *CodeWorker* [9]. It is a versatile Open Source parsing tool and a source code generator. This software permits to expand a code with a transformation source-to-source, following some specific rules described in a separated file in *extended-BNF* syntax, for recognizing the format of the specification to parse - a procedural language for manipulating easily parse trees, string, files, and directories [17].

CodeWorker provides two methods for performing a parsing:

- the reading of tokens is *procedural*;
- the BNF description is *declarative*, and conforms to a kind of BNF (the Backus-Naur Form represents a grammar in a particular syntax) extended with regular expressions.

The first file of *psort* chosen to be optimized is `cache_sorters.cpp`. Like described in Section 1.2 this is a critical file because it contains low level routines used many times in a single execution, so a little optimization could speed up significantly the software. The extension of a file to be tuned has to be change from `.c`, or `.cpp`, to `.tun`; this indicates that the syntax of the file is different from that of a *C* or *C++* source code. In fact, in order to obtain a tunable code it is necessary the uses of particular tags, defined and recognized by the grammar created in the extended-BNF file. For example, these tags permit to the parser to understand when implement a loop unrolling. In this case to the cycle is assigned a parameter with a specific search space and the code generator writes a separate function for each value of the parameter.

3.1.1 Grammars and generation rules

Extended Backus-Naur Form (extended-BNF) is a family of metasyntax notations used for expressing context-free grammars. It is used wherever exact descriptions of languages are needed, so it is suitable for the role of recognizer of programming languages or source-to-source translator. *CodeWorker* uses a variant of extended-BNF.

In order to make a code generator there must be a file that defines the generation rules. This file, with extension *.gen*, describes a grammar. First of all it permits to identify in the code the interested tags; after that it specifies how to substitute this identifier with a valid source code. The extended-BNF, that is used in this kind of files, has the same basic rules of a simple grammar. Thus the programmer writes the production like in Figure 3.2 (where capital letters are used for indicate variables, instead the other characters are used for terminal symbols).

```
S ::= AB | C
A ::= a | aA
B ::= b | bB
C ::= aCb | => { writeText("Hello World");
                  if $exclamation==1$
                    writeText("!");
                  }
```

Figure 3.2: A simple set of productions recognized by *CodeWorker*

The syntax recognized by *CodeWorker* permits to the programmer to write in the grammar definition a *C* like code (see the last production in Figure 3.2). The modified extended-BNF has also many specific keywords that permits to write a more simple code for even complex operations. Some of these are very useful, for example:

1. `implicitCopy` - it puts all scanned characters in the output;
2. `explicitCopy` - the scanned symbols doesn't go automatically in the output;
3. `readCString` - the *C* strings are ignored;
4. `readIdentifier: "STR"` - if it has been scanned the string "STR" this condition is fulfilled.

The file `cache_sorters_tuner.gen` contains an about 500 lines grammar. This has to read a tunable file, to identify the tuning tags (with the last of the keywords described above), to transform properly the code contained in these tags, and to copy the other lines without changes in the output (with the first keyword).

3.1.2 Tags

A tunable code recognizes the parameters and the tags defined in the grammar. Before going deeper it is necessary a short description of the main ones.

Function:

```
< FUN functionName_PARAM1_PARAM2_PARAMN >< _FUN >
```

The code inside these tags will be repeated for each value of the parameters. These must be specified in the function name.

If it is used the parameter < BYTES > (it indicates the size of the extended key) inside FUN tag this code will be repeated for each value of BYTES and the tag is replaced by its value on the specific routine. Furthermore if there are N parameters the code will be repeated for each combination of values.

Assign-Value:

```
< ASSIGN PARAM value1 value2 valuen >
```

This one allows to assign an array of values to the interested parameter until the tag restore, that copies the old values to the variable. These tags were added because it could be useful to give the possibility of make some differences in particular cases. For example, it could be used for the parameter BYTES if we are interested in modify only the 16 bytes versions of a function.

Restore-Value :

```
< RESTORE PARAM >
```

It has been described above.

Only-for-part-1 :

```
< P1 > <_ P1 >
```

The code inside these two tags will be parsed only in the phase one, which generates the code to be tested.

Only-for-part-2 :

```
< P2 > <_ P2 >
```

The code inside these two tags will be parsed only in the phase two, which generates the final code.

Loop:

```
< LOOP _LOOP >
```

This tag is necessary for the loop unrolling. The code inside will be written as many times as the current value of the loop unroll factor specified in the function name. This tag is related to the variable LOOP_UNROLL_FACTOR_CHUNK. There are other elements like this related to other variables.

Select:

```
< SELECT FUNCTION_NAME value > < _SELECT >
```

These tags give a name to the code that they contain. It is useful when TUNING_PART is 2 and only the best sorter function must be written in the final source file. So it has the aim to permit a choice between different implementation of the same function. These tags could also be used when TUNING_PART=1 if the choice has been already given. An example is the selection between bitwise and logical comparison.

The next subsection shows how to use the tags above with some practical examples.

3.1.3 Creation of tunable code

In order to transform a source in a tunable code the programmer has to simply modify the starting file, adding and modifying some lines. Let see this operation with an example.

```

9. inline void chunk_sorter_16b( ureg_t *buffer, const ureg_t size ) {
10.     ureg_t b12, b13, b14, b15, b16, b17, b18, b23, b24, b25, b26, b27;
11.     ureg_t b28, b34, b35, b36, b37, b38, b45, b46, b47, b48, b56, b57;
12.     ureg_t b58, b67, b68, b78;
13.     ureg_t data1[ 2 ], data2[ 2 ], data3[ 2 ], data4[ 2 ], data5[ 2 ];
14.     ureg_t data6[ 2 ], data7[ 2 ], data8[ 2 ];
15.     for ( ureg_t i = 0; i < size * 2; i += 32 ) {
16.         MICRO_SORT_8B_8( i );
17.         MICRO_SORT_8B_8( i + 16 );
18.     }
19. }

```

Figure 3.3: A non-tunable version of `chunk_sorter_16b`

```

20. < ASSIGN BYTES 8 16 24 32 >
21. < P1 >
22. < FUN inline void chunk_sorter_BYTESb_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK >
    ( ureg_t *buffer, const ureg_t size ) {
23. < _P1 >
24. < P2 >
25. < FUN inline void chunk_sorter_BYTESb >( ureg_t *buffer, const ureg_t size ) {
26. < _P2 >
27.     ureg_t < COMPARE_VARIABLES >
28.     ureg_t < DATA_VARIABLES >
29.     for ( ureg_t i = 0; i < size < SIZE_SCALING >; i += < STEP > ) {
30.         < LOOP MICRO_SORT_8B_CHUNK_SIZE( i + INCREMENT ) >; _LOOP >
31.     }
32. }
33. < _FUN >

```

Figure 3.4: A tunable routine `chunk_sorter (8b, 16b, 24b, 32b)`

Looking at Figure 3.4, the code might seem very difficult to understand. Surely after the transformation the it become more complex, but the transition from Figure 3.3 to Figure 3.4 isn't so intricate.

First of all it is added at line 20 an `ASSIGN` tag, indicating that the code has to be repeated for each value of `BYTES` (instead in Figure 3.3 is reported only the 16 *key-bytes* version). After that there is the difficult part: the *phase one* has to produce all the function to be tested, so they must have different names; instead the *phase two* returns only the optimal code with the final function name. At lines 27 and 28 there are particular tags created because for each value of `BYTES` this code needs a different number of variables. Finally there is the core of the optimization: a loop unrolling tag. `SIZE_SCALING` multiplies size by a number from 1 to 4 dependent by the value of `BYTES`, because larger is the *extended key* larger become the microrun size (the variable size is normalized to the 8 *key-bytes* version, so with an *extended key* of 16 bytes the value of size has to be doubled). The tag `STEP` depends on the value of the `loop_unroll_factor`. At line 30 there is a tag `LOOP` that repeats that macro like in Figure 3.3.

Figures 3.5-3.6 show a similar transformation applied on the function `merge`. The parameters and the considerations are similar to those explain above.

```

34. inline void merge_16b
    ( const ureg_t *s1, const ureg_t *s2, ureg_t *dest, const ureg_t size ) {
35.     ureg_t i = 0, j = 0, k = 0;
36.     ureg_t a[ 2 * 2 ];
37.     ureg_t b;
38.     while ( ( i < size * 2 ) & ( j < size * 2 ) ) {
39.         KEY_COPY_16B( &a[ 0 ], &s1[ i ] );
40.         KEY_COPY_16B( &a[ 2 ], &s2[ j ] );
41.         b = KEY_COMPARE_16B( &a[ 0 ], &a[ 2 ] );
42.         KEY_COPY_16B( &dest[ k ], &a[ b * 2 ] );
43.         k += 2;
44.         i += ( 1 - b ) * 2; j += b * 2;
45.     }
46.     for ( ; i < size * 2; i += 2 ) {
47.         KEY_COPY_16B( &dest[ k ], &s1[ i ] ); k += 2;
48.     }
49.     for ( ; j < size * 2; j += 2 ) {
50.         KEY_COPY_16B( &dest[ k ], &s2[ j ] ); k += 2;
51.     }
52. }

```

Figure 3.5: A non-tunable version of merge_16b

```

53. < ASSIGN BYTES 8 16 24 32 >
54. < P1 >
55. < FUN inline void merge_BYTESb_LOOP_UNROLL_FACTOR_MERGE >
56. ( const ureg_t *s1, const ureg_t *s2, ureg_t *dest, const ureg_t size ) {
57.     < INFO >
58. < _P1 >
59. < P2 >
60. < FUN inline void merge_BYTESb >
    ( const ureg_t *s1, const ureg_t *s2, ureg_t *dest, const ureg_t size ) {
61. < _P2 >
62.     ureg_t i = 0, j = 0, k = 0;
63.     ureg_t a[ 2 < SIZE_SCALING > ];
64.     ureg_t b;
65.     const ureg_t iMax =
    (size/< LOOP_UNROLL_FACTOR_MERGE >)*< LOOP_UNROLL_FACTOR_MERGE > < SIZE_SCALING >;
66.     while ( ( i < iMax ) & ( j < iMax ) ) {
67.         < MERGELoop KEY_COPY_< BYTES >B( &a[ 0 ], &s1[ i ] );
68.         KEY_COPY_< BYTES >B( &a[ BYTES_DIVIDED_BY_8 ], &s2[ j ] );
69.         b = KEY_COMPARE_< BYTES >B( &a[ 0 ], &a[ BYTES_DIVIDED_BY_8 ] );
70.         KEY_COPY_< BYTES >B( &dest[ k ], &a[ b * BYTES_DIVIDED_BY_8 ] );
71.         k += BYTES_DIVIDED_BY_8; i += ( 1 - b ) * BYTES_DIVIDED_BY_8;
72.         j += b * BYTES_DIVIDED_BY_8;
73.     _LOOP >
74.     }
75.     while ( ( i < size < SIZE_SCALING > ) & ( j < size < SIZE_SCALING > ) ) {
76.         KEY_COPY_< BYTES >B( &a[ 0 ], &s1[ i ] );
77.         KEY_COPY_< BYTES >B( &a[ < BYTES_DIVIDED_BY_8 > ], &s2[ j ] );
78.         b = KEY_COMPARE_< BYTES >B( &a[ 0 ], &a[ < BYTES_DIVIDED_BY_8 > ] );
79.         KEY_COPY_< BYTES >B( &dest[ k ], &a[ b * < BYTES_DIVIDED_BY_8 > ] );
80.         k += < BYTES_DIVIDED_BY_8 >;
81.         i += ( 1 - b ) * < BYTES_DIVIDED_BY_8 >;
82.         j += b * < BYTES_DIVIDED_BY_8 >;
83.     }
84.     for ( ; i < size < SIZE_SCALING >; i += < BYTES_DIVIDED_BY_8 > ) {
85.         KEY_COPY_< BYTES >B( &dest[ k ], &s1[ i ] );
86.         k += < BYTES_DIVIDED_BY_8 >;
87.     }
88.     for ( ; j < size < SIZE_SCALING >; j += < BYTES_DIVIDED_BY_8 > ) {
89.         KEY_COPY_< BYTES >B( &dest[ k ], &s2[ j ] );
90.         k += < BYTES_DIVIDED_BY_8 >;
91.     }
92. }
93. < _FUN >

```

Figure 3.6: A tunable routine merge (8b, 16b, 24b, 32b)

```

< SELECT COMPARE16 1 >
  #define OR16 |
  #define AND16 &
< _SELECT >
< SELECT COMPARE16 2 >
  #define OR16 ||
  #define AND16 &&
< _SELECT >

```

Figure 3.7: A choice with the tag `SELECT` - only one of the two blocks can be written in the final source code. Which one depends on the value of the parameter `COMPARE16`.

Another example of transformation is from `inlines.h` to `inlines.tun`. The source code in Figure 3.7 shows a simple tag `SELECT` that permits to write on the final code only the best choice between bitwise and logical compares. Without this methodology, in order to make the optimal choice, the programmers should probably compile two times the code and measure the performance. With many parameters like this the hand tuning could become very tedious and error-prone.

The file resulting by the first phase of the tuning, that is the source code with all the variants that has to be tested, is very big. In fact, from the about 3000 lines of the tunable file the obtained source code could have more than 9000 lines; this value depends on what optimization has been enabled and on the size of the search space.

3.2 Critical optimizations

The file `cache_sorters.cpp` has been chosen to be automatically tuned. This source code contains low level routines divided by the key size, many macros used for the hand-tuning, and many versions of the same functions with different values of the parameters `chunk_size` and `loop_unroll_factor`. Thus the source code appeared to be often repeated with minor changes from a function to another. These little variations, together, could speed up the execution over the 10%.

There were 3 main functions that are now tuned: `chunk_sorters`, `merge`, and `sorter`. The first sorts a chunk of data (microruns) of 2, 4, or 8 elements with a count-sort; its size depends on the parameter `chunk_size`. The second merges two sorted *chunks* of data. The third uses `chunk_sorters` and the merge functions to sort a block of data whose size is defined by the execution parameters `-s1-records-per-block` and `-s2-records-per-block`, in one or more steps (it is a *merge-sort* that starts from an array of elements sorted in groups which sizes are the value of `chunk_size`).

The example in Figure 3.8 shows the sorting of a simple input sequence with the function `sorter_8b`. First of all it calls the routine `chunk_sorters` having `chunk_size = 2`. Usually the optimal value of the latter is higher, in fact the advantage is when it permits to avoid more than one merge executions. In this simple case after the first operation the array will be composed by sorted chunks, each of two elements. Thus the execution of the merge needs only two merge passes. These are the steps of *psort* in the sorting of a microrun.

There are also many versions of the functions `merge` and `sorter`. For example, the routine `sorter` has three possible implementations:




Figure 3.8: The two iterations performed by the `sorter()` function on an input containing 8 keys. In this case, `chunk_size` equals 4. With a larger input the merge is repeated many times.

- `sorter`,
- `sorter_quasi_in_place`,
- `sorter_quasi_in_place_wave`.

In each of these functions, where needed, is introduced a distinct parameter for the loop unrolling. In particular the function `chunk_sorters` has parameters like `loop_unroll_factor` and `loop_unroll_factor_chunk`, that are related respectively to an explicit loop unrolling and to the value of `chunk_size`, that comports the uses of different macros. These parameters are strongly machine-dependent. So the automated code tuning has to determine the fastest versions of these functions and the optimal values of their variables.

Another tunable file is `inlines.h`. This file defines the compare functions, so for these has to be chosen logical or bitwise compares; that is a little but sometimes an important tuning operation.

3.3 Local search

The automated code tuning in *psort* follows a local search approach (see Section 2.1). Actually the source code has been written to support this methodology, although the search space of the parameters is quite small. Some parameters has been grouped and, for them, the code generator produces routines in all possible combinations in the search space (for example in `chunk_sorters` the parameters `loop_unroll_factor` and `loop_unroll_chunk` are grouped). Our approach is similar to the *ATLAS* global search strategy: we use an orthogonal line search (see Section 2.2), with the additional possibility of grouping some parameters and execute a full global search between them (it tries all the combinations). The fixed values chosen for the parameters are a good compromise that appears to go adequately with many architectures.

Some parameters have a very small search space, for example `BYTES` and `loop_unroll_chunk`. In particular the first is a particular parameter, in the sense that *psort* needs the implementations for all the values in its search space.

In this case, the tuning has to find the best combination of `loop_unroll_chunk` and `loop_unroll_factor` distinctly for each value of `BYTES`.

Chapter 4

Results

This chapter shows some tests on the performance of *psort* and the automated code tuning and argues some consideration starting from them. First of all, the bandwidth improvement is evident by some tests executed on architectures where hand-tuning hasn't been done yet. An important issue is how the trend of the performance depends on the tuning level and if a fast tuning is useful or not.

Figures 4.1 and 4.2 show the performance of the function `merge` for many values of `loop_unroll_factor` and for different sizes of the records (or extended keys). Each test has also been repeated for each available *tuning level*, that is the level of detail of a test: in this case, considering `merge_8b`, with the first level of tuning the test has been executed 20 times with an input of 1MB; with the third level the result is an arithmetic average of 200 executions of the same function with the same input size. This function is described in Subsection 1.2.1: it implements a simple merge sort between two arrays.

Increasing the value of tuning level to 2 has the effect of expanding the search space: the search is not anymore orthogonal but it takes into account more combinations. Furthermore, with tuning level 3 the number of executions of each test increase from 20 to 200, obtaining more reliable values.

Looking at Figures 4.1 and 4.2 it is important to focus on the relative variation of bandwidth. Certainly with another hardware composition the results will be different, but the situation found in this case is quite common. First of all it is important to specify that a larger value of the *tuning level* should mean a greater reliability of the resulting temporal data. Starting from this consideration different trend with tuning level 1 in the plot of `merge_8b` is not probably caused by noise. Unfortunately this situation is not detected with *tuning-level*=1 and this could lead to a non optimal choice. About the case of `merge_24b` the search space should be enlarged, because it is not clear if the bandwidth will fall or not with the increasing `loop_unroll_factor` (this argument follows the local search approach, starting from the value 1). A very different situation is shown in the plot of `merge_16b`, where it is difficult to understand what happens. In this cases, where the values are very near but the trend is systematic the automated code tuning can't do anything but a global search in the whole search space.

The search of the optimal values of `chunk_size` is simpler, in fact the default search space is {2, 4, 8} and the tests show that there is no need to expand it.

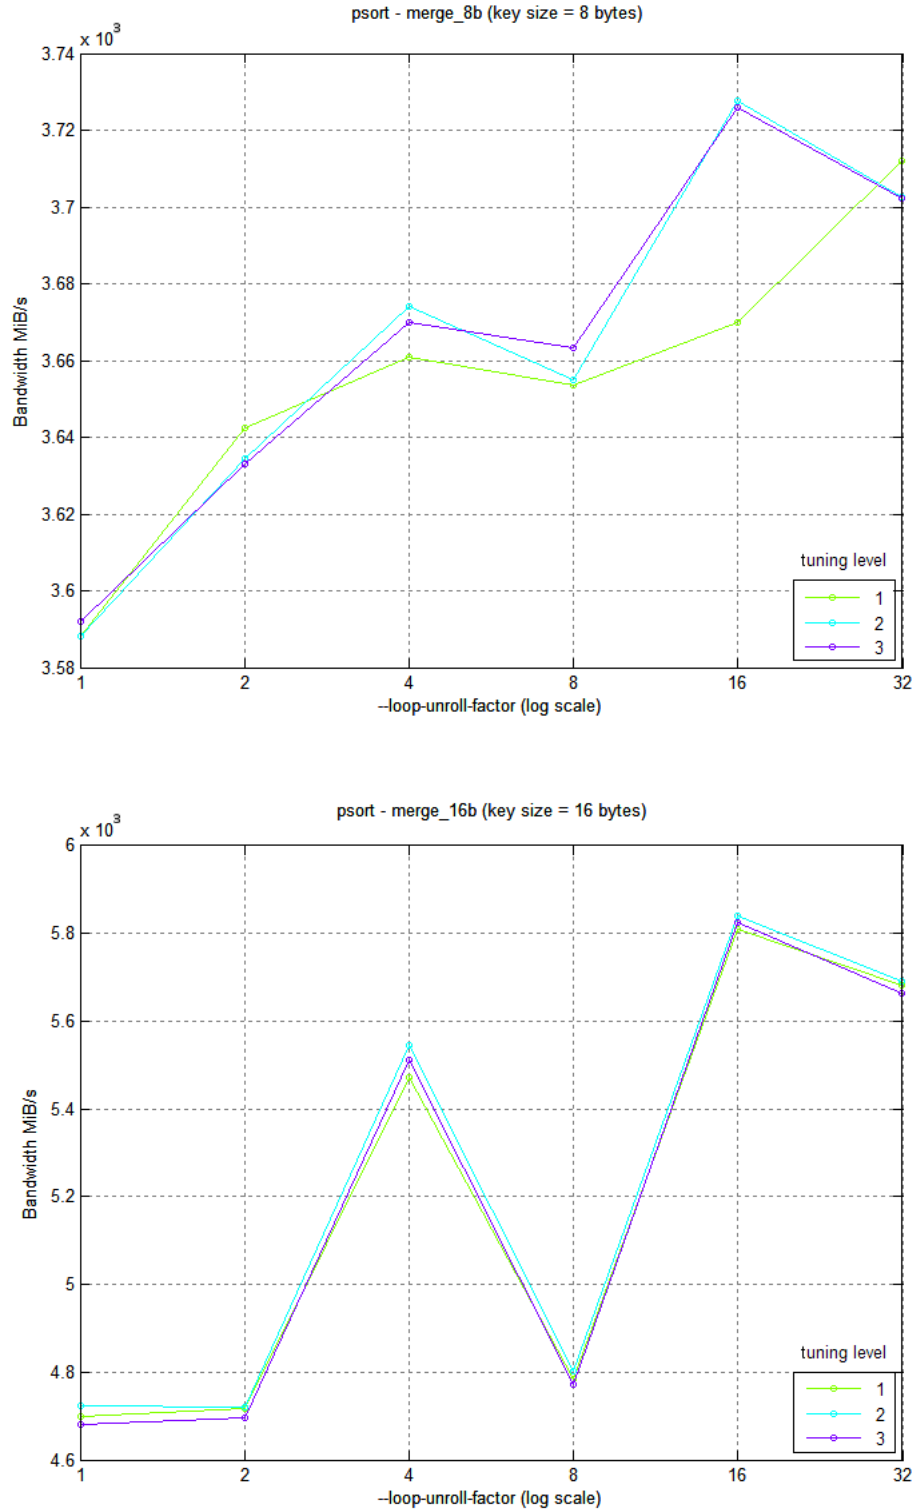


Figure 4.1: Performance of the functions `merge_8b` and `merge_16b` (key size 8 and 16) in `cache_sorters.cpp` for different values of `loop_unroll_factor` and of the tuning level. Hardware used: 3.07 GHz Intel(R) Core(TM) i7-950, 3 banks of 4 GB tri-channel DDR3 running at 1.6 GHz

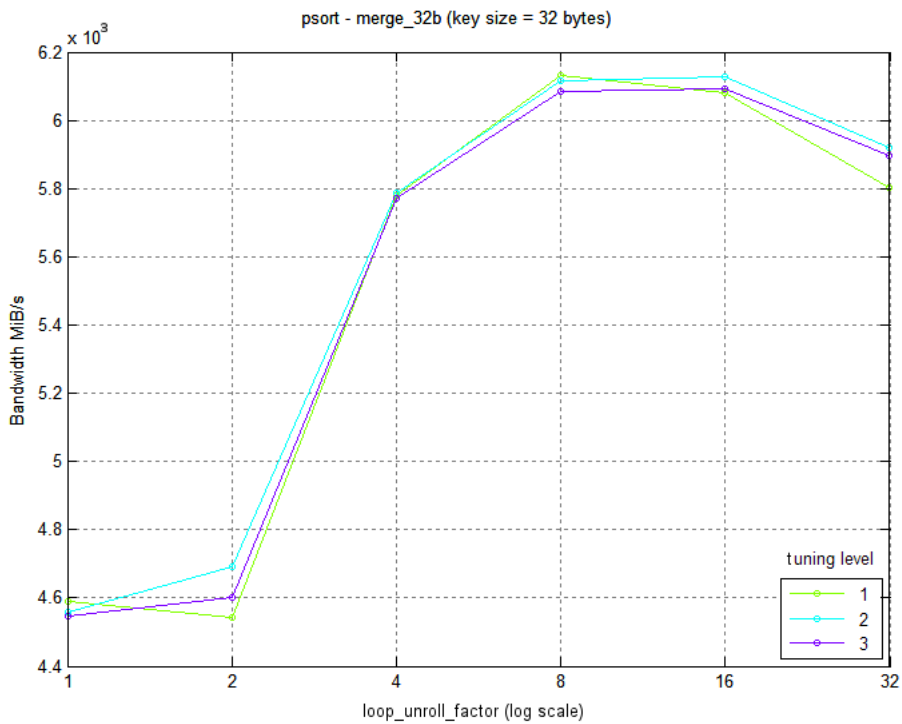



Figure 4.2: Performance of the functions merge_24b and merge_32b (key size 24 and 32) in cache_sorters.cpp for different values of loop_unroll_factor and of the tuning level. Hardware used: 3.07 GHz Intel(R) Core(TM) i7-950, 3 banks of 4 GB tri-channel DDR3 running at 1.6 GHz

tuning level	tuning time [sec]	chunk_size	loop unroll factor	
			chunk_sorter	merge
1	110	8, 4, 8, 8	1, 32, 1, 1	32, 16, 32, 16
2	210	8, 4, 8, 8	1, 2, 1, 8	16, 16, 16, 16
3	695	4, 4, 8, 8	32, 1, 1, 1	16, 16, 32, 16

Table 4.1: Running time and output values of the tuning phase (in these tests the sorter routine used was `sorter_quasi_in_place`), for three different tuning levels. For each tuning level, the cell relative to a routine contains the estimated optimal values for its 8, 16, 24, and 32-byte versions. System setup: 3.07 GHz Intel(R) Core(TM) i7-950, 3 banks of 4 GB tri-channel DDR3 running at 1.6 GHz, RAID 6x1TB

tuning level	input size [GiB]	execution time (stage 1) [sec]	average bandwidth [MB/s]
0	32	113.7	288.2
0	128	450.7	290.8
1	32	112.2	292.0
1	128	446.1	293.8
2	32	110.8	295.7
2	128	444.5	294.9
3	32	109.6	299.0
3	128	440.2	297.8

Table 4.2: Performance of the automatically generated `stage_one()` code for different tuning levels (0 indicates the original version), for 32GiB 128GiB inputs. Hardware used: 3.07 GHz Intel(R) Core(TM) i7-950, 3 banks of 4 GB tri-channel DDR3 running at 1.6 GHz, RAID 6x1TB

Perhaps in some cases it could be useful to add the value 16, but it usually has worse performance.

Table 4.2 shows the performance of `psort`'s stage one using different level of tuning. It is important to remember that these measures of time are only about the tuning of `cache_sorters.cpp` (and `inlines.h`). With more tunable files the following relative considerations become even more justified. From Table 4.1 and 4.2 is simple to verify that with `tuning-level=3` a user saves from 1.3% to 2.3% of the time spent using `tuning-level=1`, in the sorting. This means that the additionally 585 seconds, spent in the tuning phase, will be made up sorting about 7TB of data. Using the `tuning-level=2` a user saves 1.2% of the time respect to the `tuning-level=1` (in the 32GiB version), and the user will be made up the additional 100 seconds after 2TB of data. The machine used for the Sort Benchmark had similar performance, although the hardware used for these tests are probably better for the issue¹. This because the automated code tuning can come closer but hardly exceed the performance of a very well hand-tuned source code. Thus with a more accurate tuning at runtime and an automated code tuning extended to other critical aspects, the performance

¹Sort Benchmark machine: 2.7 Ghz AMD Sempron, 4 GB RAM DDR3 1333, 5x320 GB 7200 RPM Samsung SpinPoint F4 HD332GJ, Linux

Automated Code Tuning test machine: 3.07 GHz Intel(R) Core(TM) i7-950, 3 banks of 4 GB tri-channel DDR3 running at 1.6 GHz, RAID 6x1TB

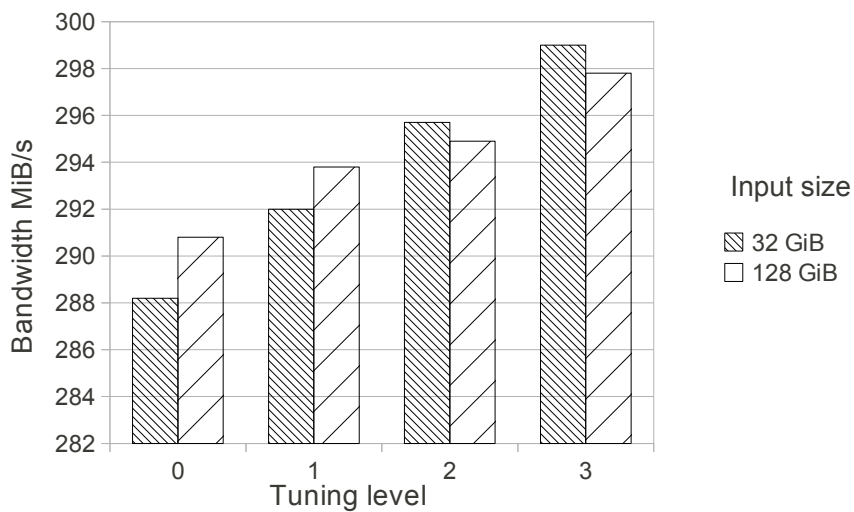


Figure 4.3: Performance of the automatically generated `stage_one()` code for different tuning levels (0 indicates the original version), for 32GiB 128GiB inputs. See Table 4.2. Hardware used: *3.07 GHz Intel(R) Core(TM) i7-950, 3 banks of 4 GB tri-channel DDR3 running at 1.6 GHz, RAID 6x1TB*

could overcome that obtained with the Sort Benchmark machine.

Recently the *psort* team has added the support to multidisks and multicore. These new features increase the software performance, but they also adds many parameters to be tuned. Probably in the future *psort* will become more complex, and thus the hand-tuned approach harder to realize.

Chapter 5

Conclusions

This work describes the design and implementation of an automated code tuning system for the external sorting library *psort*. Our system, which employs a variant of the Backus-Naur Form notation, allows programmers to efficiently and effectively write “tunable code” which is automatically expanded into several different candidate codes. The system then automatically performs code compilation and testing, selects the more efficient candidate code and integrates it into the *psort* library.

Our solution takes into account not only the performance of the resulting code, but also the reusability and modularity of the system itself. Thanks to these guidelines, the system can be extended to support those parts of code that were beyond the (experimental) scope of the present work, including code that has not been written yet. We evaluate the performance of several low-level sorting routines, including *psort*’s merge sort and heap sort, and show that automatically tuned versions perform up to 4% better than their (already highly efficient) original counterparts. Moreover, users can specify the tuning level, exploiting the trade-off between time spent tuning *psort* and time spent running *psort*. Our results show that automated code tuning provides substantial performance improvements at a modest cost in terms of end-user effort (indeed, the whole process aims to be transparent to the end user), and should thus be included in the *psort* library.

We suggest three main future directions of work. First, we propose to further reduce the programmers’ effort by organizing and simplifying the grammar in a strict, well-defined manner. This would result in a simpler, better structured, and more readable code – desirable properties not only for future programmers but also for expert end users who want to hack into the source code. Second, we propose to investigate better strategies to find the (local) optimal point in the space of the automatically generated candidate codes. The current enumerative strategy, far from optimal, could be replaced by smarter search strategies – for example, the well-known taboo search – that would drastically reduce the tuning time, resulting in higher performance. An other issue is the reduction of the tuning time. Now it is acceptable for the user (for example see Table 4.1) but the tuning operations take longer than expected, so it is another direction of work.

Appendix A

Source code

This appendix contains the source code of *psort* used by the automated code tuning. Thus what is reported below are not the entire files.

A.1 cache_sorters.tun

```
1 < P1 >
2 #include <iostream>
3 using namespace std;
4 < _P1 >
5
6 /*****
7 ***** CODE FOR SORTING
8 ***** LEXICOGRAPHICALLY
9 *****/
10
11 /*****
12 * Specific code for 8-byte extended keys
13 *****/
14
15 #define CHUNK_SORT_8B chunk_sorter_8b
16 #define MERGE_8B merge_8b
17
18
19 /* Used for first-level cache sorters; assumes that variables b0 ... bM,
20 data1, ..., dataM exist */
21 #define MICRO_SORT_8B_4( i ) \
22 KEY_COPY_8B( data1, &buffer[ i ] ); \
23 KEY_COPY_8B( data2, &buffer[ i + 1 ] ); \
24 KEY_COPY_8B( data3, &buffer[ i + 2 ] ); \
25 KEY_COPY_8B( data4, &buffer[ i + 3 ] ); \
26 b12 = KEY_COMPARE_8B( data1, data2 ); \
27 b13 = KEY_COMPARE_8B( data1, data3 ); \
28 b14 = KEY_COMPARE_8B( data1, data4 ); \
29 b23 = KEY_COMPARE_8B( data2, data3 ); \
30 b24 = KEY_COMPARE_8B( data2, data4 ); \
31 b34 = KEY_COMPARE_8B( data3, data4 ); \
32 KEY_COPY_8B( &buffer[ i + ( b12 + b13 + b14 ) ], data1 ); \
33 KEY_COPY_8B( &buffer[ i + ( 1 - b12 + b23 + b24 ) ], data2 ); \
34 KEY_COPY_8B( &buffer[ i + ( 2 - b13 - b23 + b34 ) ], data3 ); \
35 KEY_COPY_8B( &buffer[ i + ( 3 - b14 - b24 - b34 ) ], data4 );
36
37 /* Used for first-level cache sorters; assumes that variables b0 ... bM,
38 data1, ..., dataM exist */
39 #define MICRO_SORT_8B_8( i ) \
40 KEY_COPY_8B( data1, &buffer[ i ] ); \
41 KEY_COPY_8B( data2, &buffer[ i + 1 ] ); \
42 KEY_COPY_8B( data3, &buffer[ i + 2 ] ); \
```

```

42 KEY_COPY_8B( data4, &buffer[ i + 3 ] ); \
43 KEY_COPY_8B( data5, &buffer[ i + 4 ] ); \
44 KEY_COPY_8B( data6, &buffer[ i + 5 ] ); \
45 KEY_COPY_8B( data7, &buffer[ i + 6 ] ); \
46 KEY_COPY_8B( data8, &buffer[ i + 7 ] ); \
47 b12 = KEY_COMPARE_8B( data1, data2 ); \
48 b13 = KEY_COMPARE_8B( data1, data3 ); \
49 b14 = KEY_COMPARE_8B( data1, data4 ); \
50 b15 = KEY_COMPARE_8B( data1, data5 ); \
51 b16 = KEY_COMPARE_8B( data1, data6 ); \
52 b17 = KEY_COMPARE_8B( data1, data7 ); \
53 b18 = KEY_COMPARE_8B( data1, data8 ); \
54 b23 = KEY_COMPARE_8B( data2, data3 ); \
55 b24 = KEY_COMPARE_8B( data2, data4 ); \
56 b25 = KEY_COMPARE_8B( data2, data5 ); \
57 b26 = KEY_COMPARE_8B( data2, data6 ); \
58 b27 = KEY_COMPARE_8B( data2, data7 ); \
59 b28 = KEY_COMPARE_8B( data2, data8 ); \
60 b34 = KEY_COMPARE_8B( data3, data4 ); \
61 b35 = KEY_COMPARE_8B( data3, data5 ); \
62 b36 = KEY_COMPARE_8B( data3, data6 ); \
63 b37 = KEY_COMPARE_8B( data3, data7 ); \
64 b38 = KEY_COMPARE_8B( data3, data8 ); \
65 b45 = KEY_COMPARE_8B( data4, data5 ); \
66 b46 = KEY_COMPARE_8B( data4, data6 ); \
67 b47 = KEY_COMPARE_8B( data4, data7 ); \
68 b48 = KEY_COMPARE_8B( data4, data8 ); \
69 b56 = KEY_COMPARE_8B( data5, data6 ); \
70 b57 = KEY_COMPARE_8B( data5, data7 ); \
71 b58 = KEY_COMPARE_8B( data5, data8 ); \
72 b67 = KEY_COMPARE_8B( data6, data7 ); \
73 b68 = KEY_COMPARE_8B( data6, data8 ); \
74 b78 = KEY_COMPARE_8B( data7, data8 ); \
75 KEY_COPY_8B( &buffer[ i + ( b12 + b13 + b14 + b15 + b16 + b17 + b18 ) ],
    data1 ); \
76 KEY_COPY_8B( &buffer[ i + ( 1 - b12 + b23 + b24 + b25 + b26 + b27 + b28 )
    ], data2 ); \
77 KEY_COPY_8B( &buffer[ i + ( 2 - b13 - b23 + b34 + b35 + b36 + b37 + b38 )
    ], data3 ); \
78 KEY_COPY_8B( &buffer[ i + ( 3 - b14 - b24 - b34 + b45 + b46 + b47 + b48 )
    ], data4 ); \
79 KEY_COPY_8B( &buffer[ i + ( 4 - b15 - b25 - b35 - b45 + b56 + b57 + b58 )
    ], data5 ); \
80 KEY_COPY_8B( &buffer[ i + ( 5 - b16 - b26 - b36 - b46 - b56 + b67 + b68 )
    ], data6 ); \
81 KEY_COPY_8B( &buffer[ i + ( 6 - b17 - b27 - b37 - b47 - b57 - b67 + b78 )
    ], data7 ); \
82 KEY_COPY_8B( &buffer[ i + ( 7 - b18 - b28 - b38 - b48 - b58 - b68 - b78 )
    ], data8 );
83
84
85 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */
86 #define MICRO_SORT_8B_2( i ) \
87 KEY_COPY_8B( data1, &buffer[ i ] ); \
88 KEY_COPY_8B( data2, &buffer[ i + 1 ] ); \
89 b12 = KEY_COMPARE_8B( data1, data2 ); \
90 KEY_COPY_8B( &buffer[ i + b12 ], data1 ); \
91 KEY_COPY_8B( &buffer[ i + 1 - b12 ], data2 );
92
93
94 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */
95 #define MICRO_SORT_8B_1( i )
96
97 /* TUNED CODE */
98 /** Sorts quadruples of elements in an array. In particular, independently
99 * sorts every subarray indexed from k * subszie to k * subszie + 1
100 * |param buffer an array of shortrec_t records
101 * |param size the size of the array
102 */
103
104 < P1 >
105 < ASSIGN BYTES 8 16 24 32 >

```



```

106 < FUN inline void chunk_sorter_BYTESb( ureg_t *buffer, const ureg_t size ) {
    } >
107 < _FUN >
108 < FUN inline void merge_BYTESb( const ureg_t *s1, const ureg_t *s2, ureg_t *
    dest, const ureg_t size ) { } >
109 < _FUN >
110 < FUN inline void merge_BYTESb_right( const ureg_t *s1, const ureg_t *s2,
    ureg_t *dest, const ureg_t size ) { } >
111 < _FUN >
112
113 ureg_t log_ceiling( ureg_t x ) {
114     ureg_t depth = 0;
115     ureg_t compare = 1;
116     while( compare < x ) {
117         compare <<= 1;
118         depth++;
119     }
120     return depth;
121 }
122 < _P1 >
123
124 < ASSIGN BYTES 8 16 24 32 >
125 < P1 >
126 < FUN inline void chunk_sorter_BYTESb_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK >(
    ureg_t *buffer, const ureg_t size ) {
127     < INFO >
128     < _P1 >
129     < P2 >
130 < FUN inline void chunk_sorter_BYTESb >( ureg_t *buffer, const ureg_t size )
    {
131 < _P2 >
132     ureg_t < COMPARE_VARIABLES >
133     ureg_t < DATA_VARIABLES >
134     for ( ureg_t i = 0; i < size < SIZE_SCALING >; i += < STEP > ) { // STEP =
        CHUNK_SIZE * LOOP_UNROLL_FACTOR
135     < LOOP MICRO_SORT_8B_CHUNK_SIZE( ( i + INCREMENT ) ); _LOOP >
136     }
137 }< _FUN >
138
139 /** Merges two arrays
140  * Bit-based version
141  */
142 < ASSIGN BYTES 8 16 24 32 >
143 < P1 >
144 < FUN inline void merge_BYTESb_LOOP_UNROLL_FACTOR_MERGE >( const ureg_t *s1,
    const ureg_t *s2, ureg_t *dest, const ureg_t size ) {
145     < INFO >
146     < _P1 >
147     < P2 >
148 < FUN inline void merge_BYTESb >( const ureg_t *s1, const ureg_t *s2, ureg_t
    *dest, const ureg_t size ) {
149 < _P2 >
150     ureg_t i = 0, j = 0, k = 0;
151     ureg_t a[ 2 < SIZE_SCALING >];
152     ureg_t b;
153     const ureg_t iMax = (size/< LOOP_UNROLL_FACTOR_MERGE >)*<
        LOOP_UNROLL_FACTOR_MERGE > < SIZE_SCALING >;
154
155     while ( ( i < iMax ) & ( j < iMax ) ) {
156         < MERGELoop KEY_COPY_< BYTES >B( &a[ 0 ], &s1[ i ] );
157         KEY_COPY_< BYTES >B( &a[ BYTES_DIVIDED_BY_8 ], &s2[ j ] );
158         b = KEY_COMPARE_< BYTES >B( &a[ 0 ], &a[ BYTES_DIVIDED_BY_8 ] );
159         KEY_COPY_< BYTES >B( &dest[ k ], &a[ b * BYTES_DIVIDED_BY_8 ] );
160         k += BYTES_DIVIDED_BY_8;
161         i += ( 1 - b ) * BYTES_DIVIDED_BY_8;
162         j += b * BYTES_DIVIDED_BY_8; _LOOP >
163     }
164     while ( ( i < size < SIZE_SCALING > ) & ( j < size < SIZE_SCALING > ) ) {
165         KEY_COPY_< BYTES >B( &a[ 0 ], &s1[ i ] );
166         KEY_COPY_< BYTES >B( &a[ < BYTES_DIVIDED_BY_8 > ], &s2[ j ] );
167         b = KEY_COMPARE_< BYTES >B( &a[ 0 ], &a[ < BYTES_DIVIDED_BY_8 > ] );
168         KEY_COPY_< BYTES >B( &dest[ k ], &a[ b * < BYTES_DIVIDED_BY_8 > ] );
169         k += < BYTES_DIVIDED_BY_8 >;
170         i += ( 1 - b ) * < BYTES_DIVIDED_BY_8 >;

```

```

171     j += b * < BYTES_DIVIDED_BY_8 >;
172   }
173   for ( ; i < size < SIZE_SCALING >; i += < BYTES_DIVIDED_BY_8 > ) {
174     KEY_COPY_< BYTES >B( &dest[ k ], &s1[ i ] );
175     k += < BYTES_DIVIDED_BY_8 >;
176   }
177   for ( ; j < size < SIZE_SCALING >; j += < BYTES_DIVIDED_BY_8 > ) {
178     KEY_COPY_< BYTES >B( &dest[ k ], &s2[ j ] );
179     k += < BYTES_DIVIDED_BY_8 >;
180   }
181 }< _FUN >
182
183
184 /** Sorts an array using a classical mergesort
185 * |param input the source array
186 * |param output the destination array; must be allocated
187 * |param size number of elements
188 * |return 0 if the sorted data is in input, 1 if in output
189 */
190 < SELECT SORTER 1 > // Ignore this block if the value of sorter is neither 0
    nor 1
191 < ASSIGN BYTES 8 16 24 32 >
192 < P1 >
193 // < ASSIGN LOOP_UNROLL_FACTOR_CHUNK 4 >
194 < ASSIGN LOOP_UNROLL_FACTOR_MERGE 4 >
195 < ASSIGN LOOP_UNROLL_FACTOR_MERGE_RIGHT 4 >
196 < FUN ureg_t sorter_BYTESb_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK >( ureg_t *
    input, ureg_t *output, const ureg_t size ) {
197   < INFO >
198   < _P1 >
199   < P2 >
200 < FUN ureg_t sorter_BYTESb >( ureg_t *input, ureg_t *output, const ureg_t
    size ) {
201   < _P2 >
202     const ureg_t start = log_ceilng( < CHUNK_SIZE > );
203     const ureg_t stop = log_ceilng( size );
204
205   < P1 >   chunk_sorter_< BYTES >b_< CHUNK_SIZE >_< LOOP_UNROLL_FACTOR_CHUNK
    >( input, size );< _P1 >
206   < P2 >   chunk_sorter_< BYTES >b( input, size );< _P2 >
207
208   for ( ureg_t l = start; l < stop; l++ ) {
209     /* merge pairs of subarrays of size 2^l */
210     for ( ureg_t s = 0; s < size; s += 1 << ( l + 1 ) ) {
211   < P1 >   merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >( &input[ s <
    SIZE_SCALING > ], &input[ ( s + ( 1 << l ) ) < SIZE_SCALING > ], &output[ s
    < SIZE_SCALING > ], ( 1 << l ) );< _P1 >
212   < P2 >   merge_< BYTES >b( &input[ s < SIZE_SCALING > ], &input[ ( s + ( 1
    << l ) ) < SIZE_SCALING > ], &output[ s < SIZE_SCALING > ], ( 1 << l ) );<
    _P2 >
213     }
214     TYPE_SWAP( ureg_t*, input, output );
215   }
216
217   return ( ( stop - start ) & 1 );
218 }< _FUN >
219 < P1 >
220 < RESTORE LOOP_UNROLL_FACTOR_MERGE >
221 < RESTORE LOOP_UNROLL_FACTOR_MERGE_RIGHT >
222 // < RESTORE LOOP_UNROLL_FACTOR_CHUNK >
223 < _P1 >
224 < _SELECT >
225
226
227 /** Sorts an array using a classical mergesort, quasi-in-place (uses 1.5
    times the input size)
228 * |param input the source array
229 * |param output the destination array; must be allocated
230 * |param size number of elements
231 * |return 0 if the sorted data is in input, 1 if in output
232 */
233 < SELECT SORTER 2 > // Ignore this block if the value of sorter is neither 0
    nor 2
234 < ASSIGN BYTES 8 16 24 32 >

```

```

235 < P1 >
236 // < ASSIGN LOOP_UNROLL_FACTOR_CHUNK 4 >
237 < ASSIGN LOOP_UNROLL_FACTOR_MERGE 4 >
238 < ASSIGN LOOP_UNROLL_FACTOR_MERGE_RIGHT 4 >
239 < FUN ureg_t sorter_BYTESb_quasi_in_place_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK
    >( ureg_t *input, ureg_t *output, const ureg_t size ) {
240 < INFO >
241 < _P1 >
242 < P2 >
243 < FUN ureg_t sorter_BYTESb_quasi_in_place >( ureg_t *input, ureg_t *output,
    const ureg_t size ) {
244 < _P2 >
245
246     const ureg_t start = log_ceiling( < CHUNK_SIZE > );
247     const ureg_t stop = log_ceiling( size );
248
249 < P1 >
250 chunk_sorter_< BYTES >b_< CHUNK_SIZE >_< LOOP_UNROLL_FACTOR_CHUNK >( input,
    size );< _P1 >
251 < P2 >     chunk_sorter_< BYTES >b( input, size );< _P2 >
252
253     for ( ureg_t l = start; l < stop - 1; l++ ) {
254         /* merge pairs of subarrays of size 2^l */
255         ureg_t s = 0;
256         for ( s = 0; s < size; s += 1 << (l + 1) ) {
257 < P1 >         merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >( &input[ s <
    SIZE_SCALING >], &input[ ( s + (1 << l) ) < SIZE_SCALING >], &input[ ( s
    - (1 << l) ) < SIZE_SCALING >], (1 << l) );< _P1 >
258 < P2 >         merge_< BYTES >b( &input[ s < SIZE_SCALING >], &input[ ( s + (1
    << l) ) < SIZE_SCALING >], &input[ ( s - (1 << l) ) < SIZE_SCALING >],
    (1 << l) );< _P2 >
259         }
260         input -= (1 << l) < SIZE_SCALING >;
261     }
262 < P1 >     merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >( input, &input[ (1 <<
    (stop-1) ) < SIZE_SCALING >], output, (1 << (stop-1) ) );< _P1 >
263 < P2 >     merge_< BYTES >b( input, &input[ (1 << (stop-1) ) < SIZE_SCALING
    >], output, (1 << (stop-1) ) );< _P2 >
264
265     return ( 1 );
266
267 }< _FUN >
268 < P1 >
269 < RESTORE LOOP_UNROLL_FACTOR_MERGE >
270 < RESTORE LOOP_UNROLL_FACTOR_MERGE_RIGHT >
271 // < RESTORE LOOP_UNROLL_FACTOR_CHUNK >
272 < _P1 >
273 < _SELECT >
274
275 #endif
276
277
278 /*****
279  *     Specific code for 16-byte extended keys
280  *****/
281
282 #if ( ENABLE_16B == 0 ) /* If not enabled, use empty functions */
283
284 ureg_t sorter_16b( ureg_t *input, ureg_t *output, const ureg_t size ) {
285     return 0;
286 }
287 ureg_t sorter_16b_quasi_in_place( ureg_t *input, ureg_t *output, const ureg_t
    size ) {
288     return 0;
289 }
290
291 #else
292
293 #define CHUNK_SORT_16B chunk_sorter_16b
294 #define MERGE_16B merge_16b
295 #define MERGE_16B_RIGHT merge_16b_right
296
297

```

```

298 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */
299 #define MICRO_SORT_16B_2( i ) \
300     KEY_COPY_16B( data1, &buffer[ i ] ); \
301     KEY_COPY_16B( data2, &buffer[ i + 2 ] ); \
302     b12 = KEY_COMPARE_16B( data1, data2 ); \
303     KEY_COPY_16B( &buffer[ i + b12 * 2 ], data1 ); \
304     KEY_COPY_16B( &buffer[ i + ( 1 - b12 ) * 2 ], data2 );
305
306
307 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */
308 #define MICRO_SORT_16B_4( i ) \
309     KEY_COPY_16B( data1, &buffer[ i ] ); \
310     KEY_COPY_16B( data2, &buffer[ i + 2 ] ); \
311     KEY_COPY_16B( data3, &buffer[ i + 2 * 2 ] ); \
312     KEY_COPY_16B( data4, &buffer[ i + 3 * 2 ] ); \
313     b12 = KEY_COMPARE_16B( data1, data2 ); \
314     b13 = KEY_COMPARE_16B( data1, data3 ); \
315     b14 = KEY_COMPARE_16B( data1, data4 ); \
316     b23 = KEY_COMPARE_16B( data2, data3 ); \
317     b24 = KEY_COMPARE_16B( data2, data4 ); \
318     b34 = KEY_COMPARE_16B( data3, data4 ); \
319     KEY_COPY_16B( &buffer[ i + ( b12 + b13 + b14 ) * 2 ], data1 ); \
320     KEY_COPY_16B( &buffer[ i + ( 1 - b12 + b23 + b24 ) * 2 ], data2 ); \
321     KEY_COPY_16B( &buffer[ i + ( 2 - b13 - b23 + b34 ) * 2 ], data3 ); \
322     KEY_COPY_16B( &buffer[ i + ( 3 - b14 - b24 - b34 ) * 2 ], data4 );
323
324 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */
325 /* #define MICRO_INSERT_16B( i ) \
326     for ( ureg_t _ii_ = 1; _ii_ < 4; _ii_++ ) { \
327         KEY_COPY_16B( data1, &buffer[ i + _ii_ * 2 ] ); \
328         ureg_t _jj_ = _ii_; \
329         while ( KEY_COMPARE_16B( &buffer[ i + ( _jj_ - 1 ) * 2 ], data1 ) & (
330             _jj_ > 0 ) ) { _jj_--; } \
331         memmove( &buffer[ i + ( _jj_ + 1 ) * 2 ], &buffer[ i + _jj_ * 2 ], ( _ii_
332             - _jj_ ) * 16 ); \
333         KEY_COPY_16B( &buffer[ i + _jj_ * 2 ], data1 ); \
334     } \
335 */
336
337 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */
338 #define MICRO_INSERT_16B( i ) \
339     for ( ureg_t _ii_ = 1; _ii_ < 4; _ii_++ ) { \
340         unsigned int _jj_ = _ii_; \
341         while ( ( _jj_ > 0 ) && ( KEY_COMPARE_16B( &buffer[ i + ( _jj_ - 1 ) * 2
342             ], &buffer[ i + _jj_ * 2 ] ) ) ) { \
343             KEY_COPY_16B( data1, &buffer[ i + _jj_ * 2 ] ); \
344             KEY_COPY_16B( &buffer[ i + _jj_ * 2 ], &buffer[ i + ( _jj_ - 1 ) * 2 ]
345             ); \
346             KEY_COPY_16B( &buffer[ i + ( _jj_ - 1 ) * 2 ], data1 ); \
347             _jj_--; \
348         } \
349     }
350
351 /* Used for first-level cache sorters; assumes that variable data1 exists */
352 #define MICRO_INSERT_16B_bis( i ) \
353     for ( ureg_t _ii_ = 1; _ii_ < 4; _ii_++ ) { \
354         unsigned int _jj_ = _ii_ * 2; \
355         while ( ( _jj_ > 0 ) && KEY_COMPARE_16B( &buffer[ i + _jj_ - 2 ], &buffer
356             [ i + _jj_ ] ) ) { \
357             KEY_COPY_16B( data1, &buffer[ i + _jj_ ] ); \
358             KEY_COPY_16B( &buffer[ i + _jj_ ], &buffer[ i + _jj_ - 2 ] ); \
359             KEY_COPY_16B( &buffer[ i + _jj_ - 2 ], data1 ); \
360             _jj_ -= 2; \
361         } \
362     }
363
364 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
    data1, ..., dataM exist */

```

```

362 #define MICRO_INSERT_16B_4( i )          \
363 for ( ureg_t _ii_ = 1; _ii_ < 8; _ii_++ ) {          \
364     int _jj_ = _ii_;          \
365     while ( KEY_COMPARE_16B( &buffer[ i + ( _jj_ - 1 ) * 2 ], &buffer[ i +
366         _jj_ * 2 ] ) & ( _jj_ > 0 ) ) { \
367         KEY_COPY_16B( data1, &buffer[ i + _jj_ * 2 ] );          \
368         KEY_COPY_16B( &buffer[ i + ( _jj_ - 1 ) * 2 ], &buffer[ i + ( _jj_ - 1 ) * 2 ]
369             ); \
370         _jj_--;          \
371     }          \
372 }
373
374 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
375    data1, ..., dataM exist */
376 #define MICRO_INSERT_16B_4_bis( i )          \
377 for ( ureg_t _ii_ = 1; _ii_ < 8; _ii_ += 2 ) {          \
378     unsigned int _jj_ = _ii_;          \
379     while ( KEY_COMPARE_16B( &buffer[ i + ( _jj_ - 2 ) ], &buffer[ i + _jj_ ]
380         ) & ( _jj_ > 0 ) ) { \
381         KEY_COPY_16B( data1, &buffer[ i + _jj_ ] );          \
382         KEY_COPY_16B( &buffer[ i + _jj_ ], &buffer[ i + ( _jj_ - 2 ) ] ); \
383         KEY_COPY_16B( &buffer[ i + ( _jj_ - 2 ) ], data1 );          \
384         _jj_ -= 2;          \
385     }          \
386 }
387
388
389 /* Used for first-level cache sorters; assumes that variables b0 ... bN,
390    data1, ..., dataM exist */
391 #define MICRO_SORT_16B_8( i )          \
392 KEY_COPY_16B( data1, &buffer[ i ] );          \
393 KEY_COPY_16B( data2, &buffer[ i + 2 ] );          \
394 KEY_COPY_16B( data3, &buffer[ i + 4 ] );          \
395 KEY_COPY_16B( data4, &buffer[ i + 6 ] );          \
396 KEY_COPY_16B( data5, &buffer[ i + 8 ] );          \
397 KEY_COPY_16B( data6, &buffer[ i + 10 ] );          \
398 KEY_COPY_16B( data7, &buffer[ i + 12 ] );          \
399 KEY_COPY_16B( data8, &buffer[ i + 14 ] );          \
400 b12 = KEY_COMPARE_16B( data1, data2 );          \
401 b13 = KEY_COMPARE_16B( data1, data3 );          \
402 b14 = KEY_COMPARE_16B( data1, data4 );          \
403 b15 = KEY_COMPARE_16B( data1, data5 );          \
404 b16 = KEY_COMPARE_16B( data1, data6 );          \
405 b17 = KEY_COMPARE_16B( data1, data7 );          \
406 b18 = KEY_COMPARE_16B( data1, data8 );          \
407 b23 = KEY_COMPARE_16B( data2, data3 );          \
408 b24 = KEY_COMPARE_16B( data2, data4 );          \
409 b25 = KEY_COMPARE_16B( data2, data5 );          \
410 b26 = KEY_COMPARE_16B( data2, data6 );          \
411 b27 = KEY_COMPARE_16B( data2, data7 );          \
412 b28 = KEY_COMPARE_16B( data2, data8 );          \
413 b34 = KEY_COMPARE_16B( data3, data4 );          \
414 b35 = KEY_COMPARE_16B( data3, data5 );          \
415 b36 = KEY_COMPARE_16B( data3, data6 );          \
416 b37 = KEY_COMPARE_16B( data3, data7 );          \
417 b38 = KEY_COMPARE_16B( data3, data8 );          \
418 b45 = KEY_COMPARE_16B( data4, data5 );          \
419 b46 = KEY_COMPARE_16B( data4, data6 );          \
420 b47 = KEY_COMPARE_16B( data4, data7 );          \
421 b48 = KEY_COMPARE_16B( data4, data8 );          \
422 b56 = KEY_COMPARE_16B( data5, data6 );          \
423 b57 = KEY_COMPARE_16B( data5, data7 );          \
424 b58 = KEY_COMPARE_16B( data5, data8 );          \
425 b67 = KEY_COMPARE_16B( data6, data7 );          \
426 b68 = KEY_COMPARE_16B( data6, data8 );          \
427 b78 = KEY_COMPARE_16B( data7, data8 );          \
428 KEY_COPY_16B( &buffer[ i + ( b12 + b13 + b14 + b15 + b16 + b17 + b18 ) * 2
429     ], data1 );          \
430 KEY_COPY_16B( &buffer[ i + ( 1 - b12 + b23 + b24 + b25 + b26 + b27 + b28 )
431     * 2 ], data2 );          \

```

```

429     KEY_COPY_16B( &buffer[ i + ( 2 - b13 - b23 + b34 + b35 + b36 + b37 + b38 )
430                 * 2 ], data3 ); \
431     KEY_COPY_16B( &buffer[ i + ( 3 - b14 - b24 - b34 + b45 + b46 + b47 + b48 )
432                 * 2 ], data4 ); \
433     KEY_COPY_16B( &buffer[ i + ( 4 - b15 - b25 - b35 - b45 + b56 + b57 + b58 )
434                 * 2 ], data5 ); \
435     KEY_COPY_16B( &buffer[ i + ( 5 - b16 - b26 - b36 - b46 - b56 + b67 + b68 )
436                 * 2 ], data6 ); \
437     KEY_COPY_16B( &buffer[ i + ( 6 - b17 - b27 - b37 - b47 - b57 - b67 + b78 )
438                 * 2 ], data7 ); \
439     KEY_COPY_16B( &buffer[ i + ( 7 - b18 - b28 - b38 - b48 - b58 - b68 - b78 )
440                 * 2 ], data8 );
441
442     /** Merges two arrays to the right (starting from the end)
443     * Bit-based version
444     */
445     < ASSIGN BYTES 8 16 24 32 >
446     < P1 >
447     < FUN inline void merge_BYTESb_LOOP_UNROLL_FACTOR_MERGE_RIGHT_right >( const
448     ureg_t *s1, const ureg_t *s2, ureg_t *dest, const ureg_t size ) {
449     < INFO >
450     < _P1 >
451     < P2 >
452     < FUN inline void merge_BYTESb_right >( const ureg_t *s1, const ureg_t *s2,
453     ureg_t *dest, const ureg_t size ) {
454     < _P2 >
455     ureg_t i = size, j = size;
456     ureg_t a[ 2 < SIZE_SCALING > ];
457     ureg_t b;
458
459     s1 -= < BYTES_DIVIDED_BY_8 >;
460     s2 -= < BYTES_DIVIDED_BY_8 >;
461     dest -= < BYTES_DIVIDED_BY_8 >;
462
463     #if ( IF_BASED_MERGER == 1 )
464
465     KEY_COPY_16B( a, s1 ); /* use local vars to avoid load-and-store */
466     KEY_COPY_16B( a + < BYTES_DIVIDED_BY_8 >, s2 ); /* use local vars to
467     avoid load-and-store */
468
469     for ( ureg_t k = size/< LOOP_UNROLL_FACTOR_MERGE_RIGHT >; k > 0; k-- ) {
470     < MERGERIGHTLOOP if ( KEY_COMPARE_< BYTES >B( a, a + BYTES_DIVIDED_BY_8
471     ) ) {
472     KEY_COPY_< BYTES >B( dest, s1 );
473     i--;
474     s1 -= BYTES_DIVIDED_BY_8;
475     KEY_COPY_< BYTES >B( a, s1 );
476     }
477     else {
478     KEY_COPY_< BYTES >B( dest, s2 );
479     j--;
480     s2 -= BYTES_DIVIDED_BY_8;
481     KEY_COPY_< BYTES >B( a + BYTES_DIVIDED_BY_8, s2 );
482     }
483     dest -= BYTES_DIVIDED_BY_8; _LOOP >
484     }
485
486     while ( ( i > 0 ) & ( j > 0 ) ) {
487     if ( KEY_COMPARE_< BYTES >B( a, a + < BYTES_DIVIDED_BY_8 > ) ) {
488     KEY_COPY_< BYTES >B( dest, s1 );
489     i--;
490     s1 -= < BYTES_DIVIDED_BY_8 >;
491     KEY_COPY_< BYTES >B( a, s1 );
492     }
493     else {
494     KEY_COPY_< BYTES >B( dest, s2 );
495     j--;
496     s2 -= < BYTES_DIVIDED_BY_8 >;
497     KEY_COPY_< BYTES >B( a + < BYTES_DIVIDED_BY_8 >, s2 );
498     }
499     dest -= < BYTES_DIVIDED_BY_8 >;
500     }

```

```

493 #else
494     while ( ( i > < LOOP_UNROLL_FACTOR_MERGE_RIGHT >-1 ) & ( j > <
        LOOP_UNROLL_FACTOR_MERGE_RIGHT >-1 ) ) {
495         < MERGERIGHTLOOP KEY_COPY_< BYTES >B( a, s1 );
496         KEY_COPY_< BYTES >B( a + < BYTES_DIVIDED_BY_8 >, s2 );
497         b = 1 - KEY_COMPARE_< BYTES >B( a, a + < BYTES_DIVIDED_BY_8 > );
498         KEY_COPY_< BYTES >B( dest, a + b * < BYTES_DIVIDED_BY_8 > );
499         dest -= < BYTES_DIVIDED_BY_8 >;
500         i -= 1 - b;
501         s1 -= (1 - b) * < BYTES_DIVIDED_BY_8 >;
502         j -= b;
503         s2 -= b * < BYTES_DIVIDED_BY_8 >; _LOOP >
504     }
505 #endif
506
507     while ( i-- > 0 ) {
508         KEY_COPY_< BYTES >B( dest, s1 );
509         dest -= < BYTES_DIVIDED_BY_8 >;
510         s1 -= < BYTES_DIVIDED_BY_8 >;
511     }
512     while ( j-- > 0 ) {
513         KEY_COPY_< BYTES >B( dest, s2 );
514         dest -= < BYTES_DIVIDED_BY_8 >;
515         s2 -= < BYTES_DIVIDED_BY_8 >;
516     }
517 }< _FUN >
518
519 /** Sorts an array using a classical mergesort, quasi-in-place (uses 1.5
    times the input size)
    * |param input the source array
    * |param output the destination array; must be allocated
    * |param size number of elements
    * |return 0 if the sorted data is in input, 1 if in output
    */
525 < SELECT SORTER 3 > // Ignore this block if the value of sorter is neither 0
    nor 3
526 < ASSIGN BYTES 8 16 24 32 >
527 < P1 >
528 // < ASSIGN LOOP_UNROLL_FACTOR_CHUNK 4 >
529 < ASSIGN LOOP_UNROLL_FACTOR_MERGE 4 >
530 < ASSIGN LOOP_UNROLL_FACTOR_MERGE_RIGHT 4 >
531 < FUN ureg_t
    sorter_BYTESb_quasi_in_place_CHUNK_SIZE_wave_LOOP_UNROLL_FACTOR_CHUNK >(
    ureg_t *input, ureg_t *output, const ureg_t size ) {
532     < INFO >
533 < _P1 >
534 < P2 >
535 < FUN ureg_t sorter_BYTESb_quasi_in_place_wave >( ureg_t *input, ureg_t *
    output, const ureg_t size ) {
536 < _P2 >
537     const ureg_t start = log_ceilng( < CHUNK_SIZE > );
538     const ureg_t stop = log_ceilng( size );
539
540     ureg_t count = 0;
541
542 < P1 >     chunk_sorter_< BYTES >b_< CHUNK_SIZE >_< LOOP_UNROLL_FACTOR_CHUNK
    >( input, size );< _P1 >
543 < P2 >     chunk_sorter_< BYTES >b( input, size );< _P2 >
544
545     ureg_t l = start;
546     ureg_t s;
547
548     while ( l < < BYTES_DIVIDED_BY_8 > * ( ( stop - 1 ) / <
        BYTES_DIVIDED_BY_8 > ) ) {
549
550         /* merge pairs of subarrays of size 2^l */
551         for ( s = 0; s < size; ) {
552 < P1 >         merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >( &input[ s * <
            BYTES_DIVIDED_BY_8 > ], &input[ ( s + (1 << l) ) * < BYTES_DIVIDED_BY_8
            > ], &input[ ( s - (1 << l) ) * < BYTES_DIVIDED_BY_8 > ], (1 << l) );<
            _P1 >
553 < P2 >         merge_< BYTES >b( &input[ s * < BYTES_DIVIDED_BY_8 > ], &input[ (
            s + (1 << l) ) * < BYTES_DIVIDED_BY_8 > ], &input[ ( s - (1 << l) ) * <
            BYTES_DIVIDED_BY_8 > ], (1 << l) );< _P2 >

```

```

554
555     s += 1 << ( l + 1 );
556 }
557 input = &input[ - ( 1 << l ) * < BYTES_DIVIDED_BY_8 > + size * <
        BYTES_DIVIDED_BY_8 > ];
558 l++;
559
560     for ( s = 0; s < size; ) {
561 < P1 >     merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >_right( &input[ -s *
        < BYTES_DIVIDED_BY_8 > ], &input[ - ( s + ( 1 << l ) ) * <
        BYTES_DIVIDED_BY_8 > ], &input[ - ( s - ( 1 << l ) ) * <
        BYTES_DIVIDED_BY_8 > ], ( 1 << l ) );< _P1 >
562 < P2 >     merge_< BYTES >b_right( &input[ -s * < BYTES_DIVIDED_BY_8 > ], &
        input[ - ( s + ( 1 << l ) ) * < BYTES_DIVIDED_BY_8 > ], &input[ - ( s - ( 1
        << l ) ) * < BYTES_DIVIDED_BY_8 > ], ( 1 << l ) );< _P2 >
563
564     s += 1 << ( l + 1 );
565 }
566 input = &input[ ( 1 << l ) * < BYTES_DIVIDED_BY_8 > - size * <
        BYTES_DIVIDED_BY_8 > ];
567 l++;
568
569 }
570
571     if ( ( stop - 1 - start ) % 2 == 1 ) {
572         for ( s = 0; s < size; ) {
573 < P1 >     merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >( &input[ s * <
        BYTES_DIVIDED_BY_8 > ], &input[ ( s + ( 1 << l ) ) * < BYTES_DIVIDED_BY_8
        > ], &input[ ( s - ( 1 << l ) ) * < BYTES_DIVIDED_BY_8 > ], ( 1 << l ) );<
        _P1 >
574 < P2 >     merge_< BYTES >b( &input[ s * < BYTES_DIVIDED_BY_8 > ], &input[ (
        s + ( 1 << l ) ) * < BYTES_DIVIDED_BY_8 > ], &input[ ( s - ( 1 << l ) ) * <
        BYTES_DIVIDED_BY_8 > ], ( 1 << l ) );< _P2 >
575
576     s += 1 << ( l + 1 );
577 }
578 input = &input[ - ( 1 << l ) * < BYTES_DIVIDED_BY_8 > ];
579 l++;
580 }
581
582 < P1 >     merge_< BYTES >b_< LOOP_UNROLL_FACTOR_MERGE >( input, &input[ ( 1
        << l ) * < BYTES_DIVIDED_BY_8 > ], output, ( 1 << l ) );< _P1 >
583 < P2 >     merge_< BYTES >b( input, &input[ ( 1 << l ) * < BYTES_DIVIDED_BY_8
        > ], output, ( 1 << l ) );< _P2 >
584
585     return ( 1 );
586
587 }< _FUN >
588 < P1 >
589     < RESTORE LOOP_UNROLL_FACTOR_MERGE >
590     < RESTORE LOOP_UNROLL_FACTOR_MERGE_RIGHT >
591 //     < RESTORE LOOP_UNROLL_FACTOR_CHUNK >
592 < _P1 >
593 < _SELECT >
594
595
596 /*****
597  *      Specific code for 24-byte extended keys
598  *****/
599
600 #if ( ENABLE_24B == 0 ) /* If not enabled, use empty functions */
601
602 ureg_t sorter_24b( ureg_t *input, ureg_t *output, const ureg_t size ) {
603     return 0;
604 }
605 ureg_t sorter_24b_quasi_in_place( ureg_t *input, ureg_t *output, const ureg_t
        size ) {
606     return 0;
607 }
608
609 #else
610
611 #define CHUNK_SORT_24B chunk_sorter_24b
612 #define MERGE_24B merge_24b

```



```

613
614 /* Used for first-level cache sorters; assumes that variables b0 ... bM,
        data1, ..., dataM exist */
615 #define MICRO_SORT_24B_2( i ) \
616     KEY_COPY_24B( data1, &buffer[ i ] ); \
617     KEY_COPY_24B( data2, &buffer[ i + 3 ] ); \
618     b12 = KEY_COMPARE_24B( data1, data2 ); \
619     KEY_COPY_24B( &buffer[ i + b12 * 3 ], data1 ); \
620     KEY_COPY_24B( &buffer[ i + ( 1 - b12 ) * 3 ], data2 );
621
622 /* Used for first-level cache sorters; assumes that variables b0 ... bM,
        data1, ..., dataM exist */
623 #define MICRO_SORT_24B_3( i ) \
624     KEY_COPY_24B( data1, &buffer[ i + 0 * 3 ] ); \
625     KEY_COPY_24B( data2, &buffer[ i + 1 * 3 ] ); \
626     KEY_COPY_24B( data3, &buffer[ i + 2 * 3 ] ); \
627     KEY_COPY_24B( data4, &buffer[ i + 3 * 3 ] ); \
628     b12 = KEY_COMPARE_24B( data1, data2 ); \
629     b13 = KEY_COMPARE_24B( data1, data3 ); \
630     b14 = KEY_COMPARE_24B( data1, data4 ); \
631     b23 = KEY_COMPARE_24B( data2, data3 ); \
632     b24 = KEY_COMPARE_24B( data2, data4 ); \
633     b34 = KEY_COMPARE_24B( data3, data4 ); \
634     KEY_COPY_24B( &buffer[ i + ( b12 + b13 + b14 ) * 3 ], data1 ); \
635     KEY_COPY_24B( &buffer[ i + ( 1 - b12 + b23 + b24 ) * 3 ], data2 ); \
636     KEY_COPY_24B( &buffer[ i + ( 2 - b13 - b23 + b34 ) * 3 ], data3 ); \
637     KEY_COPY_24B( &buffer[ i + ( 3 - b14 - b24 - b34 ) * 3 ], data4 );
638
639
640 /* Used for first-level cache sorters; assumes that variables b0 ... bM,
        data1, ..., dataM exist */
641 #define MICRO_SORT_24B_4( i ) \
642     KEY_COPY_24B( data1, &buffer[ i ] ); \
643     KEY_COPY_24B( data2, &buffer[ i + 3 ] ); \
644     KEY_COPY_24B( data3, &buffer[ i + 6 ] ); \
645     KEY_COPY_24B( data4, &buffer[ i + 9 ] ); \
646     KEY_COPY_24B( data5, &buffer[ i + 12 ] ); \
647     KEY_COPY_24B( data6, &buffer[ i + 15 ] ); \
648     KEY_COPY_24B( data7, &buffer[ i + 18 ] ); \
649     KEY_COPY_24B( data8, &buffer[ i + 21 ] ); \
650     b12 = KEY_COMPARE_24B( data1, data2 ); \
651     b13 = KEY_COMPARE_24B( data1, data3 ); \
652     b14 = KEY_COMPARE_24B( data1, data4 ); \
653     b15 = KEY_COMPARE_24B( data1, data5 ); \
654     b16 = KEY_COMPARE_24B( data1, data6 ); \
655     b17 = KEY_COMPARE_24B( data1, data7 ); \
656     b18 = KEY_COMPARE_24B( data1, data8 ); \
657     b23 = KEY_COMPARE_24B( data2, data3 ); \
658     b24 = KEY_COMPARE_24B( data2, data4 ); \
659     b25 = KEY_COMPARE_24B( data2, data5 ); \
660     b26 = KEY_COMPARE_24B( data2, data6 ); \
661     b27 = KEY_COMPARE_24B( data2, data7 ); \
662     b28 = KEY_COMPARE_24B( data2, data8 ); \
663     b34 = KEY_COMPARE_24B( data3, data4 ); \
664     b35 = KEY_COMPARE_24B( data3, data5 ); \
665     b36 = KEY_COMPARE_24B( data3, data6 ); \
666     b37 = KEY_COMPARE_24B( data3, data7 ); \
667     b38 = KEY_COMPARE_24B( data3, data8 ); \
668     b45 = KEY_COMPARE_24B( data4, data5 ); \
669     b46 = KEY_COMPARE_24B( data4, data6 ); \
670     b47 = KEY_COMPARE_24B( data4, data7 ); \
671     b48 = KEY_COMPARE_24B( data4, data8 ); \
672     b56 = KEY_COMPARE_24B( data5, data6 ); \
673     b57 = KEY_COMPARE_24B( data5, data7 ); \
674     b58 = KEY_COMPARE_24B( data5, data8 ); \
675     b67 = KEY_COMPARE_24B( data6, data7 ); \
676     b68 = KEY_COMPARE_24B( data6, data8 ); \
677     b78 = KEY_COMPARE_24B( data7, data8 ); \
678     KEY_COPY_24B( &buffer[ i + ( b12 + b13 + b14 + b15 + b16 + b17 + b18 ) * 3
        ], data1 ); \
679     KEY_COPY_24B( &buffer[ i + ( 1 - b12 + b23 + b24 + b25 + b26 + b27 + b28 )
        * 3 ], data2 ); \
680     KEY_COPY_24B( &buffer[ i + ( 2 - b13 - b23 + b34 + b35 + b36 + b37 + b38 )
        * 3 ], data3 ); \

```

```

681     KEY_COPY_24B( &buffer[ i + ( 3 - b14 - b24 - b34 + b45 + b46 + b47 + b48 )
        * 3 ], data4 ); \
682     KEY_COPY_24B( &buffer[ i + ( 4 - b15 - b25 - b35 - b45 + b56 + b57 + b58 )
        * 3 ], data5 ); \
683     KEY_COPY_24B( &buffer[ i + ( 5 - b16 - b26 - b36 - b46 - b56 + b67 + b68 )
        * 3 ], data6 ); \
684     KEY_COPY_24B( &buffer[ i + ( 6 - b17 - b27 - b37 - b47 - b57 - b67 + b78 )
        * 3 ], data7 ); \
685     KEY_COPY_24B( &buffer[ i + ( 7 - b18 - b28 - b38 - b48 - b58 - b68 - b78 )
        * 3 ], data8 );

686
687
688
689
690 #endif
691
692
693 /*****
694  *      Specific code for 32-byte extended keys
695  *****/
696
697 #if ( ENABLE_32B == 0 ) /* If not enabled, use empty functions */
698
699 ureg_t sorter_32b( ureg_t *input, ureg_t *output, const ureg_t size ) {
700     return 0;
701 }
702 ureg_t sorter_32b_quasi_in_place( ureg_t *input, ureg_t *output, const ureg_t
        size ) {
703     return 0;
704 }
705
706 #else
707
708 #define CHUNK_SORT_32B chunk_sorter_32b
709 #define MERGE_32B merge_32b
710
711 /* Used for first-level cache sorters; assumes that variables  b0 ... bN,
        data1, ..., dataM exist */
712 #define MICRO_SORT_32B_2( i ) \
713     KEY_COPY_32B( data1, &buffer[ i ] ); \
714     KEY_COPY_32B( data2, &buffer[ i + 4 ] ); \
715     b12 = KEY_COMPARE_32B( data1, data2 ); \
716     KEY_COPY_32B( &buffer[ i + b12 * 4 ], data1 ); \
717     KEY_COPY_32B( &buffer[ i + ( 1 - b12 ) * 4 ], data2 );
718
719 /* Used for first-level cache sorters; assumes that variables  b0 ... bN,
        data1, ..., dataM exist */
720 #define MICRO_SORT_32B( i ) \
721     KEY_COPY_32B( data1, &buffer[ i + 0 * 4 ] ); \
722     KEY_COPY_32B( data2, &buffer[ i + 1 * 4 ] ); \
723     KEY_COPY_32B( data3, &buffer[ i + 2 * 4 ] ); \
724     KEY_COPY_32B( data4, &buffer[ i + 3 * 4 ] ); \
725     b12 = KEY_COMPARE_32B( data1, data2 ); \
726     b13 = KEY_COMPARE_32B( data1, data3 ); \
727     b14 = KEY_COMPARE_32B( data1, data4 ); \
728     b23 = KEY_COMPARE_32B( data2, data3 ); \
729     b24 = KEY_COMPARE_32B( data2, data4 ); \
730     b34 = KEY_COMPARE_32B( data3, data4 ); \
731     KEY_COPY_32B( &buffer[ i + ( b12 + b13 + b14 ) * 4 ], data1 ); \
732     KEY_COPY_32B( &buffer[ i + ( 1 - b12 + b23 + b24 ) * 4 ], data2 ); \
733     KEY_COPY_32B( &buffer[ i + ( 2 - b13 - b23 + b34 ) * 4 ], data3 ); \
734     KEY_COPY_32B( &buffer[ i + ( 3 - b14 - b24 - b34 ) * 4 ], data4 );
735
736
737 /* Used for first-level cache sorters; assumes that variables  b0 ... bN,
        data1, ..., dataM exist */
738 #define MICRO_SORT_32B_4( i ) \
739     KEY_COPY_32B( data1, &buffer[ i + 0 * 4 ] ); \
740     KEY_COPY_32B( data2, &buffer[ i + 1 * 4 ] ); \
741     KEY_COPY_32B( data3, &buffer[ i + 2 * 4 ] ); \
742     KEY_COPY_32B( data4, &buffer[ i + 3 * 4 ] ); \
743     KEY_COPY_32B( data5, &buffer[ i + 4 * 4 ] ); \
744     KEY_COPY_32B( data6, &buffer[ i + 5 * 4 ] ); \
745     KEY_COPY_32B( data7, &buffer[ i + 6 * 4 ] ); \

```

```

746 KEY_COPY_32B( data8, &buffer[ i + 7 * 4 ] );          \
747 b12 = KEY_COMPARE_32B( data1, data2 );              \
748 b13 = KEY_COMPARE_32B( data1, data3 );              \
749 b14 = KEY_COMPARE_32B( data1, data4 );              \
750 b15 = KEY_COMPARE_32B( data1, data5 );              \
751 b16 = KEY_COMPARE_32B( data1, data6 );              \
752 b17 = KEY_COMPARE_32B( data1, data7 );              \
753 b18 = KEY_COMPARE_32B( data1, data8 );              \
754 b23 = KEY_COMPARE_32B( data2, data3 );              \
755 b24 = KEY_COMPARE_32B( data2, data4 );              \
756 b25 = KEY_COMPARE_32B( data2, data5 );              \
757 b26 = KEY_COMPARE_32B( data2, data6 );              \
758 b27 = KEY_COMPARE_32B( data2, data7 );              \
759 b28 = KEY_COMPARE_32B( data2, data8 );              \
760 b34 = KEY_COMPARE_32B( data3, data4 );              \
761 b35 = KEY_COMPARE_32B( data3, data5 );              \
762 b36 = KEY_COMPARE_32B( data3, data6 );              \
763 b37 = KEY_COMPARE_32B( data3, data7 );              \
764 b38 = KEY_COMPARE_32B( data3, data8 );              \
765 b45 = KEY_COMPARE_32B( data4, data5 );              \
766 b46 = KEY_COMPARE_32B( data4, data6 );              \
767 b47 = KEY_COMPARE_32B( data4, data7 );              \
768 b48 = KEY_COMPARE_32B( data4, data8 );              \
769 b56 = KEY_COMPARE_32B( data5, data6 );              \
770 b57 = KEY_COMPARE_32B( data5, data7 );              \
771 b58 = KEY_COMPARE_32B( data5, data8 );              \
772 b67 = KEY_COMPARE_32B( data6, data7 );              \
773 b68 = KEY_COMPARE_32B( data6, data8 );              \
774 b78 = KEY_COMPARE_32B( data7, data8 );              \
775 KEY_COPY_32B( &buffer[ i + ( b12 + b13 + b14 + b15 + b16 + b17 + b18 ) * 4
], data1 );
776 KEY_COPY_32B( &buffer[ i + ( 1 - b12 + b23 + b24 + b25 + b26 + b27 + b28 )
* 4 ], data2 ); \
777 KEY_COPY_32B( &buffer[ i + ( 2 - b13 - b23 + b34 + b35 + b36 + b37 + b38 )
* 4 ], data3 ); \
778 KEY_COPY_32B( &buffer[ i + ( 3 - b14 - b24 - b34 + b45 + b46 + b47 + b48 )
* 4 ], data4 ); \
779 KEY_COPY_32B( &buffer[ i + ( 4 - b15 - b25 - b35 - b45 + b56 + b57 + b58 )
* 4 ], data5 ); \
780 KEY_COPY_32B( &buffer[ i + ( 5 - b16 - b26 - b36 - b46 - b56 + b67 + b68 )
* 4 ], data6 ); \
781 KEY_COPY_32B( &buffer[ i + ( 6 - b17 - b27 - b37 - b47 - b57 - b67 + b78 )
* 4 ], data7 ); \
782 KEY_COPY_32B( &buffer[ i + ( 7 - b18 - b28 - b38 - b48 - b58 - b68 - b78 )
* 4 ], data8 );

783
784
785
786 #endif
787
788 < P1 >
789 void randomOperations(ureg_t *buffer, ureg_t num_recs){
790     for ( ureg_t i = 0; i < num_recs; i++ ) {
791         buffer[ i ] = (i % 16 + 1)*rand();
792         buffer[ i ] *= buffer[ num_recs ] + buffer[ num_recs/2 ];
793     }
794 }
795
796 void generateRecords(ureg_t *buffer, ureg_t num_recs){
797     for ( ureg_t i = 0; i < num_recs; i++ ) {
798         buffer[ i ] = rand()%(1024);
799     }
800 }
801
802 int checkSort(ureg_t *buffer, ureg_t num_recs, int bytes) {
803     bytes /= 8;
804     for (ureg_t i=0; i<(num_recs-2)*bytes; i+=bytes){
805         double j = buffer[i+bytes]-buffer[i];
806         if (j<0){ cout << endl << "Indice=" <<i<<"_Valore=" <<buffer[i] <<
endl<< "Indice=" <<(i+bytes)<<"_Valore=" << endl; return 0;}
807     }
808     return 1;
809 }
810

```

```

811 #define RANDOM_OPERATIONS randomOperations(buffer, (input_size)) //
      randomOperations(buffer, (num_recs*BYTES/8)/10)
812
813 int main( int argc, char **argv ) {
814     double elapsed; // variable used to store the execution time
815     ureg_t tuning_level_factor = < TUNING_LEVEL >;
816     if (tuning_level_factor==3) tuning_level_factor=4;
817     ureg_t num_recs, input_size = 1024*1024; //tuning_level_factor
      *4*1024*1024; // 16Mb * BYTES (it must be a power of 2)
818     ureg_t *buffer = new ureg_t[input_size*32]; // 4*1.5
819     ureg_t *bufferOut = new ureg_t[input_size*2];
820     ureg_t num_executions;
821     ureg_t chunk_tuned[4], chunk_current;
822     ureg_t sorter_tuned[4]; // 1 == sorter, 2 == sorter_quasi_in_place
823     ureg_t loopChunk_tuned[4], loopChunk_current;
824     ureg_t loopMerge_tuned[4], loopMerge_current;
825     ureg_t loopMergeRight_tuned[4], loopMergeRight_current;
826     ureg_t bytes_current; // 8, 16, 24, 32
827     double bandBest[4], bandTmp;
828     for (int i=0; i<4; i++){ bandBest[i]=0; chunk_tuned[i]=0; loopChunk_tuned[
      i]=0; loopMerge_tuned[i]=0; loopMergeRight_tuned[i]=0; }
829
830     FILE *fResults;
831     if((fResults=fopen("tunerFinal.sh", "w")) == NULL) {
832         printf("Error: Cannot open file for writing.\n");
833         exit(1);
834     }
835
836     // This first execution could be slower because of the cpu energy saving
837     generateRecords(buffer, input_size*32);
838     sorter_8b_4_1( ureg_t *(buffer), (ureg_t *(bufferOut)), (ureg_t)(
      input_size/16) );
839
840
841     // Valuating the best chunk_size value (LOOP_UNROLL_FACTOR are arbitrarily
      fixed) - first sorter function
842     if (< TUNING_LEVEL > == 1) num_executions = 20;
843     else if (< TUNING_LEVEL > == 2) num_executions = 20;
844     else if (< TUNING_LEVEL > == 3) num_executions = 200;
845     cout << "Tuning Chunk_size using sorter_8b, sorter_16b, ..." << endl;
846     cout << "Bytes\tChunk_size\ttime[s]\t\tband(obj/sec)" << endl;
847     < ASSIGN BYTES 8 16 24 32 >
848     // < ASSIGN LOOP_UNROLL_FACTOR_CHUNK 4 > // arbitrary - these values
      have to be the same setted before the relative sorter function
849     < ASSIGN LOOP_UNROLL_FACTOR_MERGE 4 >
850     < ASSIGN LOOP_UNROLL_FACTOR_MERGE_RIGHT 4 >
851     < FUN print( 1, "sorter_eval" ); num_recs=input_size/(BYTES/8); for (
      int i=0; i<num_executions; i++){ generateRecords(buffer, num_recs*
      BYTES/8); RANDOM_OPERATIONS; prstart(1);
      sorter_BYTESb_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK ( ureg_t *(
      buffer), (ureg_t *(bufferOut)), (ureg_t)(num_recs/2) ); pracc(1);
      /*cout << checkSort((ureg_t *(bufferOut)), (ureg_t)(num_recs),
      BYTES);*/ >
852     prteval( 1, &elapsed );
853     prterm();
854     bandTmp = num_recs/2 * num_executions / elapsed;
855     chunk_current = < CHUNK_SIZE >;
856     bytes_current = < BYTES_DIVIDED_BY_8 > -1;
857     cout << < BYTES > << "\t" << chunk_current;
858     cout << "\t\t" << elapsed << "\t\t" << bandTmp << endl;
859     if (bandTmp > bandBest[bytes_current]) {
860         bandBest[bytes_current] = bandTmp;
861         sorter_tuned[bytes_current] = 1;
862         chunk_tuned[bytes_current] = chunk_current;
863         loopChunk_tuned[bytes_current] = < LOOP_UNROLL_FACTOR_CHUNK >;
864     }
865     < _FUN >
866     // < RESTORE LOOP_UNROLL_FACTOR_CHUNK >
867     < RESTORE LOOP_UNROLL_FACTOR_MERGE >
868     < RESTORE LOOP_UNROLL_FACTOR_MERGE_RIGHT >
869     /* for (int i=0; i<4; i++)
870         cout << endl << "The best value with " << ((i+1)*8) << "b is
      Chunk_size = " << chunk_tuned[i] << endl; */
871     cout << endl;

```

```

872
873 // Valuating the best chunk_size value (LOOP_UNROLL_FACTOR are
      arbitrarily fixed) - second sorter function
874 cout << "Tuning\Chunk_size\using\sorter_8b_quasi_in_place,\u
      sorter_16b_quasi_in_place,\u...\u" << endl;
875 cout << "Bytes\tChunk_size\ttime[s]\t\tband(obj/sec)" << endl;
876 < ASSIGN BYTES 8 16 24 32 >
877 // < ASSIGN LOOP_UNROLL_FACTOR_CHUNK 4 > // arbitrary - these values
      have to be the same setted before the relative sorter function
878 < ASSIGN LOOP_UNROLL_FACTOR_MERGE 4 >
879 < ASSIGN LOOP_UNROLL_FACTOR_MERGE_RIGHT 4 >
880 < FUN print( 1, "sorter_eval" ); num_recs=input_size/(BYTES/8); for (
      int i=0; i<num_executions; i++){ generateRecords((ureg_t *) (buffer)
      , num_recs*BYTES/8); RANDOM_OPERATIONS; prstart(1);
      sorter_BYTESb_quasi_in_place_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK (
      (ureg_t *) (buffer), (ureg_t *) (bufferOut), (ureg_t) (num_recs/2) );
      pracc(1); /*cout << checkSort((ureg_t *) (bufferOut), (ureg_t) (
      num_recs), BYTES); */ >
881 prteval( 1, &elapsed );
882 prterm();
883 bandTmp = num_recs/2 * num_executions / elapsed;
884 chunk_current = < CHUNK_SIZE >;
885 bytes_current = < BYTES_DIVIDED_BY_8 > -1;
886 cout << < BYTES > << "\t" << chunk_current;
887 cout << "\t\t" << elapsed << "\t\t" << bandTmp << endl;
888 if (bandTmp > bandBest[bytes_current]) {
889     bandBest[bytes_current] = bandTmp;
890     sorter_tuned[bytes_current] = 2;
891     chunk_tuned[bytes_current] = chunk_current;
892     loopChunk_tuned[bytes_current] = < LOOP_UNROLL_FACTOR_CHUNK >;
893 }
894 < _FUN >
895 // < RESTORE LOOP_UNROLL_FACTOR_CHUNK >
896 < RESTORE LOOP_UNROLL_FACTOR_MERGE >
897 < RESTORE LOOP_UNROLL_FACTOR_MERGE_RIGHT >
898
899 cout << endl;
900
901 // Valuating the best chunk_size value (LOOP_UNROLL_FACTOR are
      arbitrarily fixed) - third sorter function
902 cout << "Tuning\Chunk_size\using\sorter_8b_quasi_in_place_wave,\u
      sorter_16b_quasi_in_place_wave,\u...\u" << endl;
903 cout << "Bytes\tChunk_size\ttime[s]\t\tband(obj/sec)" << endl;
904 < ASSIGN BYTES 8 16 24 32 >
905 // < ASSIGN LOOP_UNROLL_FACTOR_CHUNK 4 > // arbitrary - these values
      have to be the same setted before the relative sorter function
906 < ASSIGN LOOP_UNROLL_FACTOR_MERGE 4 >
907 < ASSIGN LOOP_UNROLL_FACTOR_MERGE_RIGHT 4 >
908 < FUN print( 1, "sorter_eval" ); num_recs=input_size/(BYTES/8); for (
      int i=0; i<num_executions; i++){ generateRecords((ureg_t *) (buffer)
      , num_recs*BYTES/8); RANDOM_OPERATIONS; prstart(1);
      sorter_BYTESb_quasi_in_place_CHUNK_SIZE_wave_LOOP_UNROLL_FACTOR_CHUNK
      ( (ureg_t *) (buffer), (ureg_t *) (bufferOut), (ureg_t) (num_recs/2)
      ); pracc(1); /*cout << checkSort((ureg_t *) (bufferOut), (ureg_t) (
      num_recs), BYTES); */ >
909 prteval( 1, &elapsed );
910 prterm();
911 bandTmp = num_recs/2 * num_executions / elapsed;
912 chunk_current = < CHUNK_SIZE >;
913 bytes_current = < BYTES_DIVIDED_BY_8 > -1;
914 cout << < BYTES > << "\t" << chunk_current;
915 cout << "\t\t" << elapsed << "\t\t" << bandTmp << endl;
916 if (bandTmp > bandBest[bytes_current]) {
917     bandBest[bytes_current] = bandTmp;
918     sorter_tuned[bytes_current] = 3;
919     chunk_tuned[bytes_current] = chunk_current;
920     loopChunk_tuned[bytes_current] = < LOOP_UNROLL_FACTOR_CHUNK >;
921 }
922 < _FUN >
923 // < RESTORE LOOP_UNROLL_FACTOR_CHUNK >
924 < RESTORE LOOP_UNROLL_FACTOR_MERGE >
925 < RESTORE LOOP_UNROLL_FACTOR_MERGE_RIGHT >
926     for (int i=0; i<4; i++){

```

```

927         cout << endl << "The_best_value_with" << ((i+1)*8) << "b_is_
           Chunk_size=" << chunk_tuned[i];
928         if (sorter_tuned[i]==1) cout << "using_sorter" << endl;
929         else if (sorter_tuned[i]==2) cout << "using_
           sorter_quasi_in_place" << endl;
930         else if (sorter_tuned[i]==3) cout << "using_
           sorter_quasi_in_place_wave" << endl;
931     }
932     cout << endl;
933
934 // Valuating loop_unroll_factorChunk. The best value of chunk_size has
           already been found
935     if (< TUNING_LEVEL > == 1){
936         for (int i=0; i<4; i++) bandBest[i]=0;
937         if (< TUNING_LEVEL > == 1) num_executions = 20;
938         else if (< TUNING_LEVEL > == 2) num_executions = 20;
939         else if (< TUNING_LEVEL > == 3) num_executions = 200;
940         cout << "Tuning_Loop_unroll_factor_for_the_functions_chunk_sorter_8b_
           chunk_sorter_16b_..." << endl;
941         cout << "Bytes\tChunk_size\tLoop_unroll_factorChunk\t\ttime[s]\t\tband(
           obj/sec)" << endl;
942         < ASSIGN BYTES 8 16 24 32 >
943         < FUN bytes_current = BYTES / 8 -1; num_recs=input_size/(BYTES/8);
           chunk_current = CHUNK_SIZE; if (chunk_current==chunk_tuned[
           bytes_current]){ print( 1, "sorter_eval" ); for (int i=0; i<
           num_executions; i++){ generateRecords(buffer, num_recs*BYTES/8);
           RANDOM_OPERATIONS; prstart(1);
           chunk_sorter_BYTESb_CHUNK_SIZE_LOOP_UNROLL_FACTOR_CHUNK ((ureg_t
           *) (buffer), (ureg_t)(num_recs)); pracc(1); } >
944         prteval( 1, &elapsed );
945         prterm();
946         bandTmp = num_recs * num_executions / elapsed;
947         loopChunk_current = < LOOP_UNROLL_FACTOR_CHUNK >;
948         cout << < BYTES > << "\t" << chunk_current << "\t\t" <<
           loopChunk_current;
949         cout << "\t\t\t\t" << elapsed << "\t" << bandTmp << endl;
950         if (bandTmp > bandBest[bytes_current]) {
951             bandBest[bytes_current] = bandTmp;
952             loopChunk_tuned[bytes_current] = loopChunk_current;
953             chunk_tuned[bytes_current] = chunk_current;
954         }
955     }
956     < _FUN >
957     for (int i=0; i<4; i++)
958         cout << endl << "The_best_value_with" << ((i+1)*8) << "b_is_
           Loop_unroll_factor_chunk=" << loopChunk_tuned[i] << endl;
959     cout << endl;
960 }
961 // Valuating loop_unroll_factorMerge and loop_unroll_factorMergeRight.
           The best values of chunk_size and loop_unroll_factor_chunk have
           already been found
962     for (int i=0; i<4; i++) bandBest[i]=0;
963     if (< TUNING_LEVEL > == 1) num_executions = 20;
964     else if (< TUNING_LEVEL > == 2) num_executions = 20;
965     else if (< TUNING_LEVEL > == 3) num_executions = 200;
966     cout << "Tuning_Loop_unroll_factor_for_the_functions_merge_8b_
           merge_16b_..." << endl;
967     cout << "Bytes\tChunk_size\tLoop_unroll_factorMerge\t\ttime[s]\t\tband(
           obj/sec)" << endl;
968     < ASSIGN BYTES 8 16 24 32 >
969     < FUN bytes_current = BYTES / 8 -1; num_recs=input_size/(BYTES/8);
           chunk_current = CHUNK_SIZE; if (chunk_current==chunk_tuned[
           bytes_current]){ print( 1, "sorter_eval" ); for (int i=0; i<
           num_executions; i++){ generateRecords(buffer, num_recs*BYTES/8);
           RANDOM_OPERATIONS; prstart(1);
           merge_BYTESb_LOOP_UNROLL_FACTOR_MERGE ((ureg_t *) (buffer), (
           ureg_t *) (buffer+(num_recs/2)*BYTES_DIVIDED_BY_8), (ureg_t *) (
           bufferOut), (ureg_t)(num_recs/2)); pracc(1); } >
970     prteval( 1, &elapsed );
971     prterm();
972     bandTmp = num_recs * num_executions / elapsed;
973     loopMerge_current = < LOOP_UNROLL_FACTOR_MERGE >;
974     cout << < BYTES > << "\t" << chunk_current << "\t\t" <<
           loopMerge_current;

```

```

975         cout << "\t\t\t\t" << elapsed << "\t" << bandTmp << endl;
976         if (bandTmp > bandBest[bytes_current]) {
977             bandBest[bytes_current] = bandTmp;
978             loopMerge_tuned[bytes_current] = loopMerge_current;
979         }
980     }
981     < _FUN >
982     for (int i=0; i<4; i++){
983         cout << endl << "The_best_value_with" << ((i+1)*8) << " is
          Loop_unroll_factor_merge=" << loopMerge_tuned[i] << endl;
984     }
985     cout << endl;
986
987     for (int i=0; i<4; i++) bandBest[i]=0;
988     cout << "Tuning Loop_unroll_factor for the functions merge_8b_right,
          merge_16b_right.." << endl;
989     cout << "Bytes\tChunk_size\tLoop_unroll_factorMergeRight\t\ttime [s]\t\
          tband(obj/sec)" << endl;
990     < ASSIGN BYTES 8 16 24 32 >
991     < FUN bytes_current = BYTES / 8 -1; num_recs=input_size/(BYTES/8);
          chunk_current = CHUNK_SIZE; if ((chunk_current==chunk_tuned[
          bytes_current])&&(sorter_tuned[bytes_current]==3)){ print( 1, "
          sorter_eval" ); for (int i=0; i<num_executions; i++){
          generateRecords(buffer, num_recs*BYTES/8); RANDOM_OPERATIONS;
          prstart(1); merge_BYTESb_LOOP_UNROLL_FACTOR_MERGE_RIGHT_right ((
          ureg_t *) (buffer+(num_recs/2)*BYTES_DIVIDED_BY_8), (ureg_t *) (
          buffer+(num_recs)*BYTES_DIVIDED_BY_8), (ureg_t *) (bufferOut+(
          num_recs)*BYTES_DIVIDED_BY_8), (ureg_t)(num_recs/2)); pracc(1);
          } >
992     prteval( 1, &elapsed );
993     prterm();
994     bandTmp = num_recs * num_executions / elapsed;
995     loopMergeRight_current = < LOOP_UNROLL_FACTOR_MERGE_RIGHT >;
996     cout << < BYTES > << "\t" << chunk_current << "\t\t" <<
          loopMergeRight_current;
997     cout << "\t\t\t\t" << elapsed << "\t" << bandTmp << endl;
998     if (bandTmp > bandBest[bytes_current]) {
999         bandBest[bytes_current] = bandTmp;
1000         loopMergeRight_tuned[bytes_current] = loopMergeRight_current;
1001     }
1002 }
1003 < _FUN >
1004 for (int i=0; i<4; i++){
1005     if (sorter_tuned[i]==3)
1006         cout << endl << "The_best_value_with" << ((i+1)*8) << " is
          Loop_unroll_factor_merge_right=" << loopMergeRight_tuned
          [i] << endl;
1007 }
1008 cout << endl;
1009
1010 fprintf(fResults, ".codeworker_translate_cache_sorters_tuner.gen
          cache_sorters.tun_cache_sorters.cpp-args_TUNING_PART2");
1011
1012 fprintf(fResults, "LOOP_UNROLL_FACTOR_CHUNK_%u_%u_%u_%u
          LOOP_UNROLL_FACTOR_MERGE_%u_%u_%u_%u_LOOP_UNROLL_FACTOR_MERGE_RIGHT_%
          u_%u_%u_%u_CHUNK_SIZE_%u_%u_%u_%u_SORTER_%u_%u_%u_%u_COMPARE_%u
          TUNING_LEVEL_%u",
1013         loopChunk_tuned[0], loopChunk_tuned[1], loopChunk_tuned
          [2], loopChunk_tuned[3],
1014         loopMerge_tuned[0], loopMerge_tuned[1], loopMerge_tuned
          [2], loopMerge_tuned[3],
1015         loopMergeRight_tuned[0], loopMergeRight_tuned[1],
          loopMergeRight_tuned[2], loopMergeRight_tuned[3],
1016         chunk_tuned[0], chunk_tuned[1], chunk_tuned[2],
          chunk_tuned[3],
1017         sorter_tuned[0], sorter_tuned[1], sorter_tuned[2],
          sorter_tuned[3],
1018         < COMPARE_TYPE >, < TUNING_LEVEL >);
1019 fclose(fResults);
1020 }
1021 < _P1 >

```

A.2 inlines.tun

```

1  #define OR |
2  #define AND &
3  < SELECT COMPARE16 1 >
4  #define OR16 |
5  #define AND16 &
6  < _SELECT >
7  < SELECT COMPARE16 2 >
8  #define OR16 ||
9  #define AND16 &&
10 < _SELECT >
11 < SELECT COMPARE24 1 >
12 #define OR24 |
13 #define AND24 &
14 < _SELECT >
15 < SELECT COMPARE24 2 >
16 #define OR24 ||
17 #define AND24 &&
18 < _SELECT >
19 < SELECT COMPARE32 1 >
20 #define OR32 |
21 #define AND32 &
22 < _SELECT >
23 < SELECT COMPARE32 2 >
24 #define OR32 ||
25 #define AND32 &&
26 < _SELECT >
27 [...]
```

A.3 tuner.sh

```

1  #!/bin/bash
2  #
3  #-----
4  # Name          : tuner.sh
5  # Author        : Giovanni Di Liberto
6  # Description   : psort Code Tuning. This script executes the operations
7  #               necessary to the code tuning in pre-compilation.
8  #               It tests many critical functions in order to make the
9  #               optimal choice and then it writes the final source code.
10 #               It needed of some parameters that must be indicated as
11 #               ARGS.
12 #-----
13 #
14 #
15 # How to use ./codeworker:
16 # ./codeworker -translate cache_sorters_tuner.gen cache_sorters.cpp
17 #               function_name.c -args loop_unrol_values chunk_size_values
18 #
19 #for i in $(seq $PARINITVALUE $PARINCRVALUE $PARENDVALUE)
20 # do
21 #
22 # SETTING DEFAULT VALUES
23 COMPARE [0]=2;
24 COMPARE [1]=2;
25 COMPARE [2]=2;
26 COMPARE [3]=2;
27 TUNING_LEVEL=2;
28 #
29 # PARSING ARGS
30 while [ "$*" != "" ];
31 do
32     if [ "$1" == "--compare" ] ; then
33         for i in $(seq 0 1 3)
34         do
35             shift
36             test -n "$1" && COMPARE[$i]=$1
37             if [ "${COMPARE[$i]}" != "1" ]&&[ "${COMPARE[$i]}" != "2" ] ;
38                 then
```



```

38         echo "Value${i}_compare_incorrect_it_has_been_assigned_the_
           value_2"
39         COMPARE[$i]=2;
40     fi
41     done
42     elif [ "$1" == "--tuning-level" ]; then
43     shift
44     test -n "$1" && TUNING_LEVEL=$1
45     if [ "$TUNING_LEVEL" != "1" ]&&[ "$TUNING_LEVEL" != "2" ]&&[ "
       $TUNING_LEVEL" != "3" ] ; then
46     echo "Value_tuning_incorrect_it_has_been_assigned_the_value_2"
47     TUNING_LEVEL=2;
48     fi
49     elif [ "$1" == "--help" ]; then
50     echo "Usage: bash_tuner.sh [--compare_value]"
51     echo ""
52     echo "  --compare: it needed 4 parameters, one for each value_
       of_key_length(8, 16, 24, 32);"
53     echo "  bitwisecomparisons, 2==logical_
       comparisons."
54     echo "  --tuning-level: 1==fast, 2==medium, 3==better(slow)."
55     echo ""
56     exit 0
57     fi
58     shift
59     done
60
61
62 # IT SAVES THE OLD BINARY AND THE OLD LOG FILE
63 touch tuningTests_cache_sorters
64 mv tuningTests_cache_sorters tuningTests_cache_sorter_old
65 touch tuningLog_cache_sorters.txt
66 mv tuningLog_cache_sorters.txt tuningLog_cache_sorters.old
67 #touch
68 #tail -f tuningLog_cache_sorters.txt 0
69
70 # MOVING THE OLD STABLE SOURCE FILES (we don't want to delete the last stable
       sources)
71 touch inlines.h
72 mv inlines.h inlines_old.h
73 touch cache_sorters.cpp
74 mv cache_sorters.cpp cache_sorters_old.cpp
75
76 # TUNING
77 echo "***_inlines.tun***"
78 echo "***_Translating_the_tuning_file_to_a_C++_source_code***"
79 ./codeworker -translate cache_sorters_tuner.gen inlines.tun inlines.h -args
       TUNING_PART 2 COMPARE ${COMPARE[0]} ${COMPARE[1]} ${COMPARE[2]} ${
       COMPARE[3]} >> tuningLog_cache_sorters.txt 2>&1
80 echo ""
81 echo "***_cache_sorters.tun***"
82 echo "***_Translating_the_tuning_file_to_a_C++_source_code***"
83 ./codeworker -translate cache_sorters_tuner.gen cache_sorters.tun
       cache_sorters_for_tuning.cpp -args TUNING_PART 1
       LOOP_UNROLL_FACTOR_CHUNK 1 2 4 8 16 LOOP_UNROLL_FACTOR_MERGE 1 2 4 8 16
       LOOP_UNROLL_FACTOR_MERGE_RIGHT 1 2 4 8 16 CHUNK_SIZE 2 4 8 COMPARE ${
       COMPARE[0]} ${COMPARE[1]} ${COMPARE[2]} ${COMPARE[3]} TUNING_LEVEL
       $TUNING_LEVEL >> tuningLog_cache_sorters.txt 2>&1
84 echo "***_Compiling_the_extended_source_code***"
85 make tuning >> tuningLog_cache_sorters.txt 2>&1
86 echo "***_Executing_test_and_evaluating_the_best_options***"
87 ./tuningTests_cache_sorters >> tuningLog_cache_sorters.txt 2>&1
88 echo "***_Generating_the_optimal_source_code***"
89 bash tunerFinal.sh >> tuningLog_cache_sorters.txt 2>&1 # this file is
       generated by tuningTests_cache_sorters
90
91 echo "The_details_has_been_saved_in_tuningLog_cache_sorters.txt"

```

A.4 cache_sorters_tuner.gen

```
1 //=====
```

```

2 // Name      : cache_sorters_tuner.gen
3 // Author    : Giovanni Di Liberto
4 // Description : psort translation script for automated code tuning.
5 //           Translation script: an extended-BNF script that allows
6 //           generating code in the same time.
7 //           It extend the source code when it founds some particular
8 //           tags. Useful when the developer wants to obtain many
9 //           version of the same code with different values of
10 //           a param (loop unroll is a typical example)
11 //=====
12
13 // We want to put all scanned characters in the output
14 #implicitCopy
15
16 // The head of the grammar is the first production rule encountered
17 inlineCodeExpander ::=
18   =>{ global tuningPart=1;           // 1==tuning code, 2==final code
19       (default 1)
20       global loopName; // It permits to use the same code (the
21         production expandMERGELOOPMarkup) for differents loops
22         /* param Loop unroll factor variables */
23         global loop_unroll_factorChunk; // array with the interested
24         values for this param
25         global loopChunkCount=0; // the number of interested
26         values for this param
27         global loopChunkCurrent=0; // the current value
28         global loopChunkActive=0; // !=0 if the current
29         function is related to this param
30         global loop_unroll_factorChunk_stored; // when there is an ASSIGN
31         tag the old values must be copied here
32         /* param Loop unroll factor merge variables */
33         global loop_unroll_factorMerge;
34         global loopMergeCount=0;
35         global loopMergeCurrent=0;
36         global loopMergeActive=0;
37         global loop_unroll_factorMerge_stored;
38         /* param Loop unroll factor merge right variables */
39         global loop_unroll_factorMergeRight;
40         global loopMergeRightCount=0;
41         global loopMergeRightCurrent=0;
42         global loopMergeRightActive=0;
43         global loop_unroll_factorMergeRight_stored;
44         /* param Chunk size variables */
45         global chunk_size;
46         global chunkCount=0;
47         global chunkCurrent=0;
48         global chunkActive=0;
49         global chunk_size_stored;
50
51         global bytes; // psort extended key-length - select the
52           values with, for example < ASSIGN BYTES 8 16 >
53         bytes[0]=8; // default: only 8 bytes
54         global bytesCount=1; // size of the array bytes
55         global bytesCurrent=0;
56
57         global functionName; // the name of the current function (it
58           will be rewrite many times)
59         global functionLength=0;
60         local functionCount=0;
61         global isFunctionStarted=0; // if we are between two tags < FUN ...
62           > and < _FUN >
63         global strFunctionName="";
64
65         global sorter; // array[4] - the IDs of the best sorter
66           functions for each value of bytes
67         global sorterCount=0;
68         global sorterCurrent=0;
69         global tuningLevel=2; // 1 = fast, 2 = medium, 3 = better (slow)
70         global compareType; // 1 = bitwise, 2 = logical
71         compareType[0]=2; compareType[1]=2; compareType[2]=2; compareType
72           [3]=2;
73         global compareCurrent;

```

```

63     compareCurrent[0]=0; compareCurrent[1]=0; compareCurrent[2]=0;
64         compareCurrent[3]=0;
65     global compareCount=0;
66     // sorter has to be init to 0 0 0 0 (the SELECT tag with the SORTER
67     // variable will not ignore)
68     sorter[0]=0; sorter[1]=0; sorter[2]=0; sorter[3]=0;
69
70     global iStartInputPosition=0; // useful when we want to overwrite a
71     // tag in the generated file.
72 }
73
74 // Parsing _ARGS
75 => { local current_variable=0;
76     foreach i in _ARGS{
77         //traceText(i+endl());
78         switch(i){
79             case "LOOP_UNROLL_FACTOR_CHUNK":          set current_variable=1;
80                 break; // it admits an array of values
81             case "LOOP_UNROLL_FACTOR_MERGE":         set current_variable=2;
82                 break; // it admits an array of values
83             case "TUNING_PART":                      set current_variable=3;
84                 break; // it admits only one value
85             case "LOOP_UNROLL_FACTOR_MERGE_RIGHT":  set current_variable=4;
86                 break; // it admits an array of values
87             case "SORTER":                          set current_variable=5;
88                 break; // it admits 4 values (one for each value of bytes)
89             case "LOOP_UNROLL_FACTOR_MERGE_RIGHT": set current_variable=6;
90                 break; // it admits an array of values
91             case "COMPARE":                          set current_variable=7;
92                 break; // it admits 4 values
93             case "TUNING_LEVEL":                     set current_variable=8;
94                 break; // it admits one value
95             default:
96                 switch(current_variable){
97                     case "1": loop_unroll_factorChunk[loopChunkCount]=i;
98                         loopChunkCount = $loopChunkCount+1$;
99                         break;
100                    case "2": chunk_size[chunkCount]=i;
101                        chunkCount = $chunkCount+1$;
102                        break;
103                    case "3": tuningPart=i;
104                        break;
105                    case "4": loop_unroll_factorMerge[loopMergeCount]=i;
106                        loopMergeCount = $loopMergeCount+1$;
107                        break;
108                    case "5": sorter[sorterCount]=i;
109                        sorterCount = $sorterCount+1$;
110                        break;
111                    case "6": loop_unroll_factorMergeRight[loopMergeRightCount
112                        ]=i;
113                        loopMergeRightCount = $loopMergeRightCount+1$;
114                        break;
115                    case "7": compareType[compareCount]=i;
116                        compareCount = $compareCount+1$;
117                        break;
118                    case "8": tuningLevel=i;
119                        break;
120                }
121            }
122        }
123    }
124
125 #ignore(C++) // ignore C++ comments and whitespaces between terminals
126 // and non-terminals.
127
128 [
129     => local iOutputCurrentPosition = getOutputLocation();
130     #readCString // Jump over strings
131     |
132     // Handle a markup
133     '<',
134     // The script keeps the position just before '<'
135     => local iStartPosition = $getOutputLocation() - 1$;
136     // from now, scanned characters aren't put in the output
137     #explicitCopy
138     [

```

```

124 #readIdentifier:"FUN"[ // function start
125 #implicitCopy
126 [
127     => {
128         setOutputLocation($iStartPosition$);
129         /* If it is writing a sorter function and the best
130            sorter function are specified then we skip
131            the non-optimal combination */
132         bytesCurrent=0;
133         if (sorterCurrent>0){
134             while ((bytesCurrent<4)&&(sorter[bytesCurrent]!=
135                 sorterCurrent)&&(sorter[bytesCurrent]!=0))
136                 bytesCurrent=$bytesCurrent+1$;
137         }
138         if (bytesCurrent<4){
139             if (tuningPart==2){ // If it's generating the final
140                 code
141                 chunkActive=1;
142                 loopChunkActive=1;
143                 loopMergeActive=1;
144                 loopMergeRightActive=1;
145             }
146             isFunctionStarted=1;
147             chunkCurrent=0;
148             loopChunkCurrent=0;
149             loopMergeCurrent=0;
150             loopMergeRightCurrent=0;
151
152             while (lookAhead(" ")) readChar(); // ignore the ' '
153             strFunctionName="";
154             while (!lookAhead(" >")){
155                 strFunctionName += readChar();
156             }
157             readChar();readChar(); // jump over the
158                 substring " >"
159             iStartInputPosition=$getInputLocation()$;
160         }
161     }
162 }
163 #check(isFunctionStarted==0)[ // jump to the end of the
164     function
165     #explicitCopy
166     [~"_FUN >"]*
167     #explicitCopy
168     "_FUN >"
169 ]|[
170     writeFunctionName
171 ]
172 ]
173 |
174 #readIdentifier:"_FUN" [ // function end
175 =>{
176     if(isFunctionStarted==1){
177         setOutputLocation(iStartPosition);
178         writeText(endl()+endl());
179         if(tuningPart==1){ // tuning part
180             chunkCurrent=$chunkCurrent+1$;
181             if(($chunkCurrent<chunkCount$)&&(chunkActive!=0)){
182                 setInputLocation(iStartInputPosition); // it turns
183                 back to the start of the block, if it hasn't
184                 finished
185             }else{
186                 if (chunkActive!=0) chunkCurrent=0;
187                 loopChunkCurrent=$loopChunkCurrent+1$;
188                 if(($loopChunkCurrent<loopChunkCount$)&&(
189                     loopChunkActive!=0)){ // loopChunk and loopMerge
190                     won't be used at the same time (they are
191                     mutually exclusive)
192                 }
193                 setInputLocation(iStartInputPosition); // it turns
194                 back to the start of the block, if it hasn't
195             }else{
196                 if (loopChunkActive!=0) loopChunkCurrent=0;
197                 loopMergeCurrent=$loopMergeCurrent+1$;

```

```

186         if(($loopMergeCurrent<loopMergeCount$)&&(
187             loopMergeActive!=0)){
188             setInputLocation(iStartInputPosition); // it
189                 turns back to the start of the block, if it
190                 hasn't
191         }else{
192             if(loopMergeActive!=0) loopMergeCurrent=0;
193             loopMergeRightCurrent=$loopMergeRightCurrent+1$
194             ;
195             if(($loopMergeRightCurrent<loopMergeRightCount$
196                 )&&(loopMergeRightActive!=0)){
197                 setInputLocation(iStartInputPosition); // it
198                     turns back to the start of the block, if it
199                     hasn't
200             }else{
201                 if(loopMergeRightActive!=0)
202                     loopMergeRightCurrent=0;
203                 bytesCurrent=$bytesCurrent+1$;
204                 if (sorterCurrent>0){
205                     while ((bytesCurrent<4)&&(sorter [
206                         bytesCurrent]!=sorterCurrent)&&(
207                             sorter [bytesCurrent]!=0))
208                         bytesCurrent=$bytesCurrent+1$;
209                 }
210                 if(bytesCurrent<$bytesCount$){
211                     setInputLocation(iStartInputPosition); //
212                         it turns back to the start of the
213                         block, if it hasn't
214                     loopChunkCurrent=0;
215                     loopMergeCurrent=0;
216                     loopMergeRightCurrent=0;
217                     chunkCurrent=0;
218                 }else { isFunctionStarted=0; readChar ();
219                     readChar ();}
220             }}}}
221     }}}}
222 }
223 writeFunctionName
224 ]
225 |
226 #readIdentifier:" INFO"
227 => setOutputLocation(iStartPosition);
228 expandINFOMarkup
229 |
230 #readIdentifier:" DATA_VARIABLES "
231 => setOutputLocation(iStartPosition);
232 expandDATA_VARIABLESMarkup
233 |
234 #readIdentifier:" COMPARE_VARIABLES "
235 => setOutputLocation(iStartPosition);
236 expandCOMPARE_VARIABLESMarkup
237 |
238 #readIdentifier:" STEP"
239 => setOutputLocation(iStartPosition);
240 expandSTEPMarkup
241 |
242 #readIdentifier:" SIZE_SCALING"
243 => setOutputLocation(iStartPosition);
244 expandSIZE_SCALINGMarkup

```

```

245 |
246 #readIdentifier:"LOOP_UNROLL_FACTOR_CHUNK"
247 => setOutputLocation(iStartPosition);
248 expandLOOP_UNROLL_FACTOR_CHUNKMarkup
249 |
250 #readIdentifier:"LOOP_UNROLL_FACTOR_MERGE_RIGHT"
251 => setOutputLocation(iStartPosition);
252 expandLOOP_UNROLL_FACTOR_MERGE_RIGHTMarkup
253 |
254 #readIdentifier:"LOOP_UNROLL_FACTOR_MERGE"
255 => setOutputLocation(iStartPosition);
256 expandLOOP_UNROLL_FACTOR_MERGEMarkup
257 |
258 #readIdentifier:"CHUNK_SIZE"
259 => setOutputLocation(iStartPosition);
260 expandCHUNK_SIZEMarkup
261 |
262 #readIdentifier:"COMPARE_TYPE"
263 => setOutputLocation(iStartPosition);
264 expandCOMPARE_TYPEMarkup
265 |
266 #readIdentifier:"LOOP"
267 => setOutputLocation(iStartPosition);
268 expandLOOPMarkup
269 |
270 #readIdentifier:"MERGERIGHTLOOP"
271 => setOutputLocation(iStartPosition);
272 => loopName="MergeRight";
273 expandMERGELOOPMarkup
274 |
275 #readIdentifier:"MERGELOOP"
276 => setOutputLocation(iStartPosition);
277 => loopName="Merge";
278 expandMERGELOOPMarkup
279 |
280 #readIdentifier:"ASSIGN"
281 => setOutputLocation(iStartPosition);
282 expandASSIGNMarkup
283 |
284 #readIdentifier:"RESTORE"
285 => setOutputLocation(iStartPosition);
286 expandRESTOREMarkup
287 |
288 #readIdentifier:"BYTES_DIVIDED_BY_8"
289 => setOutputLocation(iStartPosition);
290 expandBYTES_DIVIDED_BY_8Markup
291 |
292 #readIdentifier:"BYTES"
293 => setOutputLocation(iStartPosition);
294 expandBYTESMarkup
295 |
296 #readIdentifier:"TUNING_LEVEL"
297 => setOutputLocation(iStartPosition);
298 expandTUNING_LEVELMarkup
299 |
300 #readIdentifier:"P1" // only for tuning
301 [ => setOutputLocation(iStartPosition);
302 #check(tuningPart!=1)[
303 #explicitCopy // ignore all the block
304 [~"_P1 >"]*
305 #explicitCopy
306 "_P1 >"
307 => if (lookAhead(endl())) readChar();
308 ]|[#explicitCopy // ignore only the characters until we find
309 '>',
310 [^>']*
311 #explicitCopy
312 '>'
313 ]
314 |
315 #readIdentifier:"_P1" // only for tuning // la ''' è un problema
!! forse è un carattere speciale
=> setOutputLocation(iStartPosition);
316

```

```

317         #explicitCopy
318         [^>']*
319         #explicitCopy
320         '>'
321     |
322     #readIdentifier:"P2" // only for final source code
323     [ => setOutputLocation(iStartPosition);
324       #check(tuningPart!=2)[
325         #explicitCopy
326         [^"_P2 >"]*
327         #explicitCopy
328         "_P2 >"
329         => if (lookAhead(endl())) readChar();
330       ]|#explicitCopy
331       [^>']*
332       #explicitCopy
333       '>'
334     ]
335 ]
336 |
337 #readIdentifier:"_P2" // only for tuning // the '/' can't be
      used
338     => setOutputLocation(iStartPosition);
339     #explicitCopy
340     [^>']*
341     #explicitCopy
342     '>'
343 |
344 #readIdentifier:"SELECT" // it tells which version of the
      indicated function is parsed
345 [ => setOutputLocation(iStartPosition);
346   expandSELECTMarkup
347
348   #check(($ (compareCurrent[0]!=0)&&(compareCurrent[0]!=
      compareType[0]))$)
349       ||($ (compareCurrent[1]!=0)&&(compareCurrent[1]!=
      compareType[1]))$)
350       ||($ (compareCurrent[2]!=0)&&(compareCurrent[2]!=
      compareType[2]))$)
351       ||($ (compareCurrent[3]!=0)&&(compareCurrent[3]!=
      compareType[3]))$)[ // we ignore this section if
      it is relative to a compare value and if that isn't
      the value we have chosen
352     #explicitCopy
353     [^"_SELECT >"]*
354     #explicitCopy
355     "_SELECT >"
356     => if (lookAhead(endl())) readChar();
357   ]|[
358     #explicitCopy
359     [^>']*
360     // ignore the end of the markup - only one: '>'
361     #explicitCopy
362     '>'
363   ]
364 ]
365 |
366 #readIdentifier:"_SELECT"
367     => setOutputLocation(iStartPosition);
368     =>{ sorterCurrent=0; }
369     #explicitCopy
370     [^>']*
371     #explicitCopy
372     '>'
373 ]
374 |
375 // If not a string and not a markup, reading of any
376 // character
377 #readChar
378 ]*;
379
380 // Replaces a LOOP markup with the C++ corresponding code
381 expandLOOPMarkup ::=
382     => local count=0;

```

```

383 => local strRepeat;
384 => local strIndex=0;
385 => local strCount=0;
386
387 // Copy of the query in the output
388 #implicitCopy
389 => {
390   while (!lookAhead("_LOOP >")){           // copying in strRepeat each input
391     line until " >" is found
392     strRepeat[strCount] += readChar();
393     if (lookAhead(endl())){ readChar(); strCount=$strCount+1$; } //
394       ignore endl and increment the counter
395   }
396
397   while $count < loop_unroll_factorChunk[loopChunkCurrent]*chunk_size[
398     chunkCurrent]*(bytes[bytesCurrent]/8){
399     strIndex=0;
400     while (strIndex <= strCount){
401       if ($count!=0$) || ($count==0$ && $strIndex!=0$) writeText(endl());
402       if $strIndex==0$ && $count!=0$ @ @
403         local s=strRepeat[strIndex];
404         if $count==0$ { s=replaceString(" + INCREMENT", "", s); s=
405           replaceString("+ INCREMENT", "", s); } // ugly but it works
406         else
407           s=replaceString("INCREMENT", count, s);
408
409         s=replaceString("CHUNK_SIZE", chunk_size[chunkCurrent], s);
410         s=replaceString("< BYTES_DIVIDED_BY_8 >", $bytes[bytesCurrent]/8$, s
411           );
412         s=replaceString("BYTES_DIVIDED_BY_8", $bytes[bytesCurrent]/8$, s);
413         s=replaceString("< BYTES >", bytes[bytesCurrent], s);
414         s=replaceString("BYTES", bytes[bytesCurrent], s);
415         writeText(s);
416         strIndex = $strIndex+1$;
417     }
418     count = $count + chunk_size[chunkCurrent]*(bytes[bytesCurrent]/8)$;
419   }
420 }
421 // ignore the end of the markup - ne ignora solo 1: '>'
422 #explicitCopy
423 "_LOOP >"
424 ;
425
426 // Replaces a MERGELOOP markup with the C++ corresponding code
427 expandMERGELOOPMarkup ::=
428 => local count=0;
429 => local strRepeat;
430 => local strIndex=0;
431 => local strCount=0;
432
433 // Copy of the query in the output
434 #implicitCopy
435 => {
436   local maxValue;
437   if (loopName=="Merge") maxValue=loop_unroll_factorMerge[loopMergeCurrent
438     ];
439   else if (loopName=="MergeRight") maxValue=loop_unroll_factorMergeRight[
440     loopMergeRightCurrent];
441   while (!lookAhead("_LOOP >")){ // copying in strRepeat each input
442     line until " >" is found
443     strRepeat[strCount] += readChar();
444     if (lookAhead(endl())){readChar(); strCount=$strCount+1$; } // ignore
445       endl and increment the counter
446   }
447   while $count < maxValue${
448     strIndex=0;
449     while ($strIndex <= strCount){
450       if ($count!=0$) || ($count==0$ && $strIndex!=0$) writeText(endl());
451       if $strIndex==0$ && $count!=0$ @ @
452         local s=strRepeat[strIndex];
453         if $count==0$ { s=replaceString(" + INCREMENT", "", s); s=
454           replaceString("+ INCREMENT", "", s); } // ugly but it works
455         else
456           s=replaceString("INCREMENT", count, s);
457         s=replaceString("CHUNK_SIZE", chunk_size[chunkCurrent], s);

```



```

446         s=replaceString("< BYTES_DIVIDED_BY_8 >", $bytes[bytesCurrent]/8$, s
447         );
448         s=replaceString("BYTES_DIVIDED_BY_8", $bytes[bytesCurrent]/8$, s);
449         s=replaceString("< BYTES >", bytes[bytesCurrent], s);
450         s=replaceString("BYTES", bytes[bytesCurrent], s);
451         writeText(s);
452         strIndex = $strIndex+1$;
453     }
454     count = $count + 1$;
455 }}
456 // ignore the end of the markup - ne ignora solo 1: '>'
457 #explicitCopy
458 "_LOOP >"
459 ;
460
461
462 // Replaces a STEP markup with the C++ corresponding code
463 expandSTEPMarkup ::=
464 => {
465     if (loopChunkActive!=0)
466         writeText($loop_unroll_factorChunk [loopChunkCurrent]*chunk_size
467         [chunkCurrent]*(bytes [bytesCurrent]/8$));
468     else if (loopMergeActive)
469         writeText($loop_unroll_factorMerge [loopMergeCurrent]*chunk_size
470         [chunkCurrent]*(bytes [bytesCurrent]/8$));
471     else if (loopMergeRightActive)
472         writeText($loop_unroll_factorMergeRight [loopMergeRightCurrent]*
473         chunk_size [chunkCurrent]*(bytes [bytesCurrent]/8$));
474 }
475 #implicitCopy
476 [^>']*
477 #explicitCopy
478 '>'
479 ;
480
481 // Replaces a SIZE_SCALING markup with the C++ corresponding code
482 expandSIZE_SCALINGMarkup ::=
483 => {
484     if $(bytes [bytesCurrent]/8)>1$
485         writeText("* "+$bytes [bytesCurrent]/8$);
486 }
487 #implicitCopy
488 [^>']*
489 #explicitCopy
490 '>'
491 ;
492
493 // Replaces a LOOP_UNROLL_FACTOR_CHUNK markup with the C++ corresponding code
494 expandLOOP_UNROLL_FACTOR_CHUNKMarkup ::=
495 => writeText(loop_unroll_factorChunk [loopChunkCurrent]);
496 #implicitCopy
497 [^>']*
498 #explicitCopy
499 '>'
500 ;
501
502 // Replaces a LOOP_UNROLL_FACTOR_MERGE markup with the C++ corresponding code
503 expandLOOP_UNROLL_FACTOR_MERGEMarkup ::=
504 => writeText(loop_unroll_factorMerge [loopMergeCurrent]);
505 #implicitCopy
506 [^>']*
507 #explicitCopy
508 '>'
509 ;
510
511 // Replaces a LOOP_UNROLL_FACTOR_MERGE_RIGHT markup with the C++
512 // corresponding code
513 expandLOOP_UNROLL_FACTOR_MERGE_RIGHTMarkup ::=
514 => writeText(loop_unroll_factorMergeRight [loopMergeRightCurrent]);
515 #implicitCopy
516 [^>']*
517 #explicitCopy
518 '>'
519 ;

```

```

515         ;
516
517 // Replaces a CHUNK_SIZE markup with the C++ corresponding code
518 expandCHUNK_SIZEMarkup ::=
519     => writeText(chunk_size[chunkCurrent]);
520     #explicitCopy
521     [^>']*
522     #explicitCopy
523     '>'
524     ;
525
526 // Replaces a COMPARE_TYPE markup with the C++ corresponding code
527 expandCOMPARE_TYPEMarkup ::=
528     => writeText(compareType[0]+","+compareType[1]+","+compareType
529         [2]+","+compareType[3]);
530     #explicitCopy
531     [^>']*
532     #explicitCopy
533     '>'
534     ;
535 // Replaces a BYTES_DIVIDED_BY_8 markup with the C++ corresponding code
536 expandBYTES_DIVIDED_BY_8Markup ::=
537     => writeText($bytes[bytesCurrent]/8$);
538     #explicitCopy
539     [^>']*
540     #explicitCopy
541     '>'
542     ;
543
544 // Replaces a BYTES markup with the C++ corresponding code
545 expandBYTESMarkup ::=
546     => writeText(bytes[bytesCurrent]);
547     #explicitCopy
548     [^>']*
549     #explicitCopy
550     '>'
551     ;
552
553 // Replaces a TUNING_LEVEL markup with the C++ corresponding code
554 expandTUNING_LEVELMarkup ::=
555     => writeText(tuningLevel);
556     #explicitCopy
557     [^>']*
558     #explicitCopy
559     '>'
560     ;
561
562 // Replaces a INFO markup with the C++ corresponding code
563 expandINFOMarkup ::=
564     => {
565         writeText("// ");
566         if (chunkActive!=0) writeText("CHUNK_SIZE == "+chunk_size[
567             chunkCurrent]+" - ");
568         if (loopChunkActive!=0) writeText("LOOP_UNROLL_FACTOR_CHUNK = "+
569             loop_unroll_factorChunk[loopChunkCurrent]+" - ");
570         if (loopMergeActive!=0) writeText("LOOP_UNROLL_FACTOR_MERGE = "+
571             loop_unroll_factorMerge[loopMergeCurrent]);
572         if (loopMergeRightActive!=0) writeText("
573             LOOP_UNROLL_FACTOR_MERGE_RIGHT = "+
574             loop_unroll_factorMergeRight[loopMergeRightCurrent]);
575     }
576     #implicitCopy
577     [^>']*
578     #explicitCopy
579     '>'
580     ;
581
582 // Replaces a DATA_VARIABLES markup with the C++ corresponding code
583 expandDATA_VARIABLESMarkup ::=
584     => {
585         local i=0;
586         while $i<chunk_size[chunkCurrent]$ {
587             if $i>0$ writeText(", ");

```

```

583         writeText("data"+$i+1$+" [ "+$bytes[bytesCurrent]/8$+" ]");
584         i=$i+1$;
585     }
586     writeText(";");
587 }
588 #implicitCopy
589 [^>']*
590 #explicitCopy
591 '>'
592 ;
593
594 // Replaces a COMPARE_VARIABLES markup with the C++ corresponding code
595 expandCOMPARE_VARIABLESMarkup ::=
596 => {
597     local i=1;
598     while $i<chunk_size[chunkCurrent]$ {
599         local j=$i+1$;
600         while $j<=chunk_size[chunkCurrent]$ {
601             if $i>1$ || $j>2$ writeText(", ");
602             writeText("b"+i+j);
603             j=$j+1$;
604         }
605         i=$i+1$;
606     }
607     writeText(";");
608 }
609 #implicitCopy
610 [^>']*
611 #explicitCopy
612 '>'
613 ;
614
615 // Replaces a ASSIGN markup with the C++ corresponding code
616 expandASSIGNMarkup ::=
617 #readIdentifier:"BYTES"
618 => {
619     readChar(); // the ' '
620     local s="";
621     local c=readChar();
622     bytesCount=0;
623     while c>='0' && c<='9' {
624         s=s+c;
625         c=readChar();
626         if (c==' ')||(c=='>'){ // now we'll store the next value
627             bytes[bytesCount]=s;
628             s="";
629             bytesCount=$bytesCount+1$;
630             do{
631                 c=readChar(); // moves to the next symbol
632             }while (c==' ');
633         }
634     }
635 }
636 |
637 #readIdentifier:"CHUNK_SIZE"
638 => {
639     readChar(); // the ' '
640     local s="";
641     local c=readChar();
642     // Storing the old data
643     chunkCount=0;
644     foreach i in chunk_size {
645         chunk_size_stored[chunkCount] = i;
646         chunkCount=$chunkCount+1$;
647     }
648     chunkCount=0;
649     while c>='0' && c<='9' {
650         s=s+c;
651         c=readChar();
652         if (c==' ')||(c=='>'){ // now we'll store the next value
653             chunk_size[chunkCount]=s;
654             s="";
655             chunkCount=$chunkCount+1$;
656         }

```

```

657         c=readChar(); // moves to the next symbol
658     }while (c==' ');
659     }
660 }
661 }
662 |
663 #readIdentifier:"LOOP_UNROLL_FACTOR_CHUNK"
664 => {
665     readChar(); // the ' '
666     local s="";
667     local c=readChar();
668     // Storing the old data
669     loopChunkCount=0;
670     foreach i in loop_unroll_factorChunk {
671         loop_unroll_factorChunk_stored[loopChunkCount] = i;
672         loopChunkCount=$loopChunkCount+1$;
673     }
674     loopChunkCount=0;
675     while c>='0' && c<='9' {
676         s=s+c;
677         c=readChar();
678         if (c==' ')||(c=='>'){ // now we'll store the next value
679             loop_unroll_factorChunk[loopChunkCount]=s;
680             s="";
681             loopChunkCount=$loopChunkCount+1$;
682             do{
683                 c=readChar(); // moves to the next symbol
684             }while (c==' ');
685         }
686     }
687 }
688 |
689 #readIdentifier:"LOOP_UNROLL_FACTOR_MERGE_RIGHT"
690 => {
691     readChar(); // the ' '
692     local s="";
693     local c=readChar();
694     // Storing the old data
695     loopMergeRightCount=0;
696     foreach i in loop_unroll_factorMergeRight {
697         loop_unroll_factorMergeRight_stored[loopMergeRightCount] = i
698         ;
699         loopMergeRightCount=$loopMergeRightCount+1$;
700     }
701     loopMergeRightCount=0;
702     while c>='0' && c<='9' {
703         s=s+c;
704         c=readChar();
705         if (c==' ')||(c=='>'){ // now we'll store the next value
706             loop_unroll_factorMergeRight[loopMergeRightCount]=s;
707             s="";
708             loopMergeRightCount=$loopMergeRightCount+1$;
709             do{
710                 c=readChar(); // moves to the next symbol
711             }while (c==' ');
712         }
713     }
714 |
715 #readIdentifier:"LOOP_UNROLL_FACTOR_MERGE"
716 => {
717     readChar(); // the ' '
718     local s="";
719     local c=readChar();
720     // Storing the old data
721     loopMergeCount=0;
722     foreach i in loop_unroll_factorMerge {
723         loop_unroll_factorMerge_stored[loopMergeCount] = i;
724         loopMergeCount=$loopMergeCount+1$;
725     }
726     loopMergeCount=0;
727     while c>='0' && c<='9' {
728         s=s+c;
729         c=readChar();

```

```

730         if (c==' ')||(c=='>'){ // now we'll store the next value
731             loop_unroll_factorMerge[loopMergeCount]=s;
732             s="";
733             loopMergeCount=$loopMergeCount+1$;
734             do{
735                 c=readChar(); // moves to the next symbol
736             }while (c==' ');
737         }
738     }
739 }
740 ;
741
742 // Replaces a SELECT markup with the C++ corresponding code
743 expandSELECTMarkup ::=
744 => { compareCurrent[0]=0; compareCurrent[1]=0; compareCurrent[2]=0;
745     compareCurrent[3]=0; }
746 #readIdentifier:"SORTER"
747 => {
748     readChar(); // the ' '
749     local s="";
750     local c=readChar();
751     bytesCount=0;
752     while c>='0' && c<='9' {
753         s=s+c;
754         c=readChar();
755     }
756     sorterCurrent=s;
757 }
758 |
759 #readIdentifier:"COMPARE8"
760 => {
761     readChar(); // the ' '
762     local s="";
763     local c=readChar();
764     while c>='0' && c<='9' {
765         s=s+c;
766         c=readChar();
767     }
768     compareCurrent[0]=s;
769 }
770 |
771 #readIdentifier:"COMPARE16"
772 => {
773     readChar(); // the ' '
774     local s="";
775     local c=readChar();
776     while c>='0' && c<='9' {
777         s=s+c;
778         c=readChar();
779     }
780     compareCurrent[1]=s;
781 }
782 |
783 #readIdentifier:"COMPARE24"
784 => {
785     readChar(); // the ' '
786     local s="";
787     local c=readChar();
788     while c>='0' && c<='9' {
789         s=s+c;
790         c=readChar();
791     }
792     compareCurrent[2]=s;
793 }
794 |
795 #readIdentifier:"COMPARE32"
796 => {
797     readChar(); // the ' '
798     local s="";
799     local c=readChar();
800     while c>='0' && c<='9' {
801         s=s+c;
802         c=readChar();

```

```

803         compareCurrent [3]=s;
804     }
805     ;
806
807
808 // Replaces a RESTORE markup with the C++ corresponding code
809 expandRESTOREMarkup ::=
810     [#readIdentifier:"CHUNK_SIZE"
811     => {
812         // Restoring the old data
813         chunkCount=0;
814         foreach i in chunk_size_stored {
815             chunk_size [chunkCount] = i;
816             chunkCount=$chunkCount+1$;
817         }
818     }
819 |
820     [#readIdentifier:"LOOP_UNROLL_FACTOR_MERGE_RIGHT"
821     => {
822         // Restoring the old data
823         loopMergeRightCount=0;
824         foreach i in loop_unroll_factorMergeRight_stored {
825             loop_unroll_factorMergeRight [loopMergeRightCount] = i;
826             loopMergeRightCount=$loopMergeRightCount+1$;
827         }
828     }
829 |
830     [#readIdentifier:"LOOP_UNROLL_FACTOR_MERGE"
831     => {
832         // Restoring the old data
833         loopMergeCount=0;
834         foreach i in loop_unroll_factorMerge_stored {
835             loop_unroll_factorMerge [loopMergeCount] = i;
836             loopMergeCount=$loopMergeCount+1$;
837         }
838     }
839 |
840     [#readIdentifier:"LOOP_UNROLL_FACTOR_CHUNK"
841     => {
842         // Restoring the old data
843         loopChunkCount=0;
844         foreach i in loop_unroll_factorChunk_stored {
845             loop_unroll_factorChunk [loopChunkCount] = i;
846             loopChunkCount=$loopChunkCount+1$;
847         }
848     }
849 ]]
850
851 #implicitCopy
852 [^>'']*
853 #explicitCopy
854 '>'
855 ;
856
857 writeFunctionName ::=
858 => { if (isFunctionStarted==1){
859     local s=strFunctionName;
860     // This flag is necessary to iterate only the declared params
861     s=replaceString("BYTES_DIVIDED_BY_8", $bytes [bytesCurrent]/8$, s)
862     ;
863     s=replaceString("BYTES", bytes [bytesCurrent], s);
864     chunkActive=countStringOccurrences(s, "CHUNK_SIZE");
865     s=replaceString("CHUNK_SIZE", chunk_size [chunkCurrent], s);
866     loopChunkActive=countStringOccurrences(s, "
867         LOOP_UNROLL_FACTOR_CHUNK");
868     s=replaceString("LOOP_UNROLL_FACTOR_CHUNK",
869         loop_unroll_factorChunk [loopChunkCurrent], s);
870     loopMergeRightActive=countStringOccurrences(s, "
871         LOOP_UNROLL_FACTOR_MERGE_RIGHT");
872     s=replaceString("LOOP_UNROLL_FACTOR_MERGE_RIGHT",
873         loop_unroll_factorMergeRight [loopMergeRightCurrent], s);
874     loopMergeActive=countStringOccurrences(s, "
875         LOOP_UNROLL_FACTOR_MERGE");
876     s=replaceString("LOOP_UNROLL_FACTOR_MERGE",
877         loop_unroll_factorMerge [loopMergeCurrent], s);

```

```
870         writeText(s);
871         if (tuningPart==2){
872             chunkActive=1;
873             loopChunkActive=1;
874             loopMergeActive=1;
875             loopMergeRightActive=1;
876         }
877         setInputLocation(iStartInputPosition);
878     }
879 }
880 ;
```


Bibliography

- [1] J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*, 3rd ed., Pearson Education, 2008.
- [2] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein. *Introduction to Algorithms*, 3rd ed., The MIT Press, 2009.
- [3] John L. Hennessy, and David A. Patterson. *Computer Architecture: A Quantitative Approach*, Kaufmann u.a., Amsterdam u.a., 2007.
- [4] D. E. Knuth. *Art of Computer Programming, Volume 3: Sorting and Searching*, 2nd ed., Addison-Wesley Professional, April 1998.
- [5] A. Aggarwal, B. Alpern, A. Chandra, and M. Snir. A model for hierarchical memory. *In Proceedings of the 19th ACM Symposium on Theory of Computing (STOC)*, pp. 305-314, 1987.
- [6] E. Allan and K. Kennedy. *Optimizing Compilers for Modern Architectures*. Morgan Kaufmann Publishers, 2002.
- [7] Paolo Bertasi, Marco Bressan, and Enoch Peserico. *psort*, yet another fast stable external sorting software. *Proceedings of the 8th Symposium on Experimental Algorithms*, 2009.
- [8] Sort Benchmark Home Page. <http://sortbenchmark.org/>.
- [9] CodeWorker Home Page. <http://www.codeworker.org/>.
- [10] ATLAS Home Page. <http://math-atlas.sourceforge.net/>.
- [11] K. Yotov, K. Pingali, and P. Stodghill. *Think Globally, Search Locally*, Cornell University, <http://hdl.handle.net/1813/5669>.
- [12] J. Bilmes and K. Asanovi'c and J. Demmel and D. Lam and C.W. Chin. *PHiPAC: A Portable, High-Performance, ANSI C Coding Methodology and its application to Matrix Multiply, LAPACK Working Note 111*, UTK, <http://www.netlib.org/lapack/lawn>, 1996.
- [13] M. Frigo and S.G. Johnson. *FFTW: an adaptive software architecture for the FFT*, *Proceedings of the IEEE 93(2)*, 2005. <http://www.fftw.org/fftw-paper-icassp.pdf>.

- [14] M. Frigo and S.G. Johnson. The Design and Implementation of FFTW3, *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, <http://www.fftw.org/fftw-paper-ieee.pdf>.
- [15] Karmen Yotov, Xiaoming Li, Gang Ren, Maria Garzaran, David Padue, Keshav Pingali, and Paul Stodghill. Is search really necessary to generate high-performance BLAS? *Proceedings of the IEEE 93(2), 2005. Special issue on "Program Generation, Optimization, and Adaptation"*.
- [16] Vuduc, R., Demmel, J.W., Yelick, K.A.: OSKI: A library of automatically tuned sparse matrix kernels. *J. Phys. Conf. Ser. 16, 521–530*, 200, <http://iopscience.iop.org/>.
- [17] ISO/IEC 14977:1996. *Information technology - Syntactic metalanguage - Extended BNF*, <http://www.iso.org/iso/>.
- [18] Alberto Bedin. Personal communication.