

Implementing multifrontal sparse solvers for multicore architectures with Sequential Task Flow runtime systems

Emmanuel Agullo¹, Alfredo Buttari², Abdou Guermouche³, and Florent Lopez⁴

¹Inria-LaBRI

²CNRS-IRIT

³Université de Bordeaux

⁴UPS-IRIT

To face the advent of multicore processors and the ever increasing complexity of hardware architectures, programming models based on DAG parallelism regained popularity in the high performance, scientific computing community. Modern runtime systems offer a programming interface that complies with this paradigm and powerful engines for scheduling the tasks into which the application is decomposed. These tools have already proved their effectiveness on a number of dense linear algebra applications. This paper evaluates the usability and effectiveness of runtime systems based on the Sequential Task Flow model for complex applications, namely, sparse matrix multifrontal factorizations which feature extremely irregular workloads, with tasks of different granularities and characteristics and with a variable memory consumption. Most importantly, it shows how this parallel programming model eases the development of complex features that benefit the performance of sparse, direct solvers as well as their memory consumption. We illustrate our discussion with the multifrontal QR factorization running on top of the StarPU runtime system.

This work is supported by the Agence Nationale de la Recherche, under grant ANR-13-MONU-0007.

This work was performed using HPC resources from GENCI-[CCRT/TGCC/CINES/IDRIS] (Grant 2014-i2014065063)

1 Introduction

The increasing degree of parallelism and complexity of hardware architectures requires the High Performance Computing (HPC) community to develop more and more complex software. To achieve high levels of optimization and fully benefit from their potential, not only the related codes are heavily tuned for the considered architecture, but the software is furthermore often designed as a single whole that aims to cope with both the algorithmic and architectural needs. If this approach may indeed lead to extremely high performance, it is at the price of a tremendous development effort and a decreased maintainability.

Alternatively, a modular approach can be employed. First, the numerical algorithm is written at a high level independently of the hardware architecture as a Directed Acyclic Graph (DAG) of tasks where a vertex represents a task and an edge represents a dependency between tasks. A second layer is in charge of the scheduling. Based on the scheduling decisions, a runtime engine, in the third layer, takes care of performing the actual execution of the tasks, both ensuring that dependencies are satisfied at execution time and maintaining data consistency. The fourth layer consists of the tasks code optimized for the underlying architectures. In most cases, the last three layers need not be written by the application developer. Indeed, it usually exists a very competitive state-of-the-art generic scheduling algorithm (such as work-stealing Arora et al. [2001], Minimum Completion Time Topcuoglu et al. [2002]) matching the algorithmic needs to efficiently exploit the targeted architecture. Otherwise, if needed, as we do in this study for handling efficiently tasks of small granularity, a new scheduling algorithm may be designed which is likely to apply to a whole class of algorithms. The runtime engine only needs to be extended once for each new architecture. Finally, most of the time, the high-level

algorithm can be cast in terms of standard operations (such as BLAS in dense linear algebra) for which vendors provide optimized codes. All in all, with such a modular approach, only the high-level algorithm has to be specifically designed, which ensures a high productivity. The maintainability is also guaranteed since the use of new hardware only requires (in principle) third party effort. Modern tools exist that implement the two middle layers (the scheduler and the runtime engine) and provide a programming model that allows the programmer to conveniently express his workload in the form of a DAG of tasks; we refer to these tools as *runtime systems* or, simply, *runtimes*.

The dense linear algebra community has strongly adopted such a modular approach over the past few years Buttari et al. [2009], Quintana-Ortí et al. [2009], Agullo et al. [2009], Bosilca et al. [2013] and delivered production-quality software relying on it. However, beyond this community, only few research efforts have been conducted to handle large scale codes. The main reason is that irregular problems are complex to design with a clear separation of the software layers without inducing performance loss. On the other hand, the runtime system community has strongly progressed, delivering very reliable and effective tools Augonnet et al. [2011], Bosilca et al. [2012], Badia et al. [2009], Hermann et al. [2010] up to the point that the OpenMP board has included similar features in the latest OpenMP standard 4.0 2013.

In the past Buttari [2013], we have assessed the effectiveness of a DAG based parallelization on the QR factorization of sparse matrices and later on Agullo et al. [2013] we proposed an approach for implementing this method using a runtime system.

This work extends our previous efforts and other related works in three different ways:

1. It proposes a parallelization purely based on a Sequential Task Flow (STF) model: our previous work Agullo et al. [2013] proposed a rather complex strategy for submitting tasks to the runtime system which renders the code difficult to maintain and to update with new features and algorithms; here, instead, we propose an approach that relies on the simplicity and effectiveness of the STF programming model.
2. It introduces the use of 2D algorithms for the factorization of frontal matrices: in the very common case where frontal matrices are (strongly) overdetermined, the block-column partitioning used in our previous work represents a major scalability problem. Here we propose a 2D partitioning of frontal matrices into *tiles* (i.e., blocks) and the use of factorization algorithms by tiles, often referred to as communication-avoiding algorithms. Experimental results show that the higher degree of concurrency provided by these methods can considerably improve the scalability of a sparse matrix QR factorization.
3. It proposes a memory-aware approach for controlling the memory consumption of the parallel multifrontal method, which allows the user to achieve the highest possible performance within a prescribed memory footprint. We show that the resulting parallel STF code is extremely robust as it can consistently achieve a memory consumption as little as the one of a sequential execution. We furthermore show that, on our thirty-two core platform, this memory reduction does not lead to performance degradation with respect to non constrained executions.

Importantly, the design of a complex 2D numerical scheme (item (2)) and of the memory-aware mechanism (item (3)) rely on the properties provided by the STF parallelization (item (1)). Indeed, if 2D algorithms are now commonly provided by dense linear algebra libraries, scheduling such kernels concurrently is original, up to our knowledge, in a sparse multifrontal code. Similarly, limiting the memory consumption of a multifrontal code is often a complex process; we show that STF allows for a very elegant, compact and efficient approach for doing so. Finally, the proposed techniques are not tied to the multifrontal QR algorithm and can be readily applied to any sparse matrix factorization method. In fact, their scope can be even larger than sparse factorizations as they may be of interest for other algorithms with complex workloads and irregular data access patterns.

The implementation approach proposed here is related to previous work on the PaStiX solver Lacoste et al. [2014] where the StarPU runtime system is used to implement supernodal Cholesky and LU factorizations for multicore architectures equipped with GPUs. In the proposed solution, however, dependencies among tasks are explicitly declared and therefore the approach does not take full advantage of the Sequential Task Flow programming model provided by StarPU (see Section 3 for the details). They also proposed a variant based on the PaRSEC (formerly DAGuE) runtime system (which features a different programming model) and presented a comparison between the two implementations. More recently, an article by Kim *et al.* 2014 presented a DAG-based approach for sparse

LDL^T factorizations. In this work OpenMP tasks are used where inter-node dependencies (by node, here, we refer to the elimination tree nodes, as described in Section 2) are implicitly handled through a recursive submission of tasks, whereas intra-node dependencies are essentially handled manually.

The rest of this paper is organized as follows. In sections 2 and 3 we describe the context of our work giving a brief description of the multifrontal QR method and of modern runtime systems along with the STF parallel programming model. In Section 4.1 we show how the STF model can be used to implement the block-column parallelization approach described in our previous work Buttari [2013] with a very compact and elegant code. In the subsequent Section 4.2 we explain how performance can be improved by using communication-avoiding, tiled factorizations for frontal matrices. In Section 4.3 we propose a technique for achieving the multifrontal factorization within a prescribed memory consumption. In Section 5 we provide and analyze experimental results showing the effectiveness of the proposed techniques.

2 The multifrontal QR method and qr_mumps

The multifrontal method, introduced by Duff and Reid 1983 as a method for the factorization of sparse, symmetric linear systems, can be adapted to the QR factorization of a sparse matrix thanks to the fact that the R factor of a matrix A and the Cholesky factor of the normal equation matrix $A^T A$ share the same structure under the hypothesis that the matrix A is *Strong Hall*. As in the Cholesky case, the multifrontal QR factorization is based on the concept of *elimination tree* introduced by Schreiber 1982. This graph, which has a number of nodes that is typically one order of magnitude or more smaller than the number of columns in the original matrix, expresses the dependencies among the computational tasks in the factorization: each node i of the tree is associated with k_i unknowns of A and represents an elimination step of the factorization. The coefficients of the corresponding k_i columns and all the other coefficients affected by their elimination are assembled together into a relatively small dense matrix, called *frontal matrix* or, simply, *front*, associated with the tree node. The multifrontal QR factorization consists in a tree traversal in a **topological order** (i.e., bottom-up) such that, at each node, two operations are performed. First, the frontal matrix is **assembled** by stacking the matrix rows associated with the k_i unknowns with uneliminated rows resulting from the processing of child nodes. Second, the k_i unknowns are eliminated through a **complete QR factorization** of the front. This produces k_i rows of the global R factor, a number of Householder reflectors that implicitly represent the global Q factor and a *contribution block* formed by the remaining rows. Those rows will be assembled into the parent front together with the contribution blocks from all the sibling fronts. The main differences with the better known multifrontal Cholesky factorization are:

1. Frontal matrices are generally rectangular (either over or under-determined) and not square;
2. Frontal matrices are completely and not partially factorized;
3. Assemblies are not extend-add operations but simple memory copies, which means that they can be executed in an embarrassingly parallel fashion;
4. Frontal matrices are not full but have a staircase structure: once a frontal matrix is assembled, its rows are sorted in order of increasing index of the leftmost nonzero. The number of operations can thus be reduced, as well as the fill-in in the Q matrix, by ignoring the zeroes in the bottom-left part of the frontal matrix.

A detailed presentation of the multifrontal QR method, including the optimization techniques described above, can be found in the work by Amestoy *et al.* 1996. The classical approach to the parallelization of the multifrontal QR factorization Amestoy *et al.* [1996], Davis [2011] consists in exploiting separately two distinct sources of concurrency: **tree** and **node parallelism**. The first one stems from the fact that fronts in separate branches are independent and can thus be processed concurrently; the second one from the fact that, if a front is big enough, multiple processes can be used to assemble and factorize it. The baseline of this work, instead, is the parallelization model proposed in the `qr_mumps` software Buttari [2013] which is based on the approach presented earlier in related work on dense matrix factorizations Buttari *et al.* [2009] and extended to the supernodal Cholesky factorization of sparse matrices Hogg *et al.* [2010]. In this approach, frontal matrices are partitioned into block-columns, which allows for a decomposition of the workload into fine-grained tasks.

Figure 2 shows on the left-side a simplified sequential pseudocode that achieves the multifrontal QR factorization in the case where a 1D partitioning into block-columns is applied to the frontal matrices. It makes use of five kernels:

1. **activate**: this routine allocates and initializes the front data structure;
2. **assemble**: for a block-column in the child node c , assembles the corresponding part of the contribution block into f . Note that, in practice, only a subset of the block-columns of f are concerned by this operation which renders most of the **assemble** tasks independent; the actual parallel implementations presented in our previous works and here, in Section 4, take advantage of this property but for the sake of readability we do not present this level of details in the pseudo-code;
3. **panel**: computes the QR factorization of a block-column;
4. **update**: updates a block-column with respect to a previous panel operation.
5. **deactivate**: stores the factors aside and frees the memory containing the contribution block;

Assuming that each function call in this pseudo-code defines a task, the whole multifrontal factorization can be expressed in the form of a Directed Acyclic Graph (DAG) of tasks. This allows for a seamless exploitation of both tree and node parallelism. Additionally it permits to pipeline the processing of a front with those of its children which provides an additional source of concurrency. The benefits of this approach over more traditional techniques have been assessed in our previous work Buttari [2013].

3 Task-based runtime systems and the Sequential Task Flow model

Whereas task-based runtime systems were mainly research tools in the past years, their recent progress makes them now solid candidates for designing advanced scientific software as they provide programming paradigms that allow the programmer to express concurrency in a simple yet effective way and relieve him from the burden of dealing with low-level architectural details. However, no consensus has still been reached on a specific paradigm. For example, Parametrized Task Graph approaches Bosilca et al. [2012], Budimlić et al. [2010], Cosnard and Jeannot [2001] consist of explicitly describing tasks (vertices of the DAG) and their mutual dependencies (edges) by informing the runtime system with a set of dependency rules. In such a way, the DAG is never explicitly built but can be progressively unrolled and traversed in a very effective and flexible way. This approach can achieve a great scalability on a very large number of processors but explicitly expressing the dependencies may be a hard task, especially when designing complex schemes.

On the other hand, the STF model simply consists of submitting a sequence of tasks through a non blocking function call that delegates the execution of the task to the runtime system. Upon submission, the runtime system adds the task to the current DAG along with its dependencies which are automatically computed through data dependency analysis Allen and Kennedy [2002]. The actual execution of the task is then postponed to the moment when its dependencies are satisfied. This paradigm is also sometimes referred to as *Superscalar* since it mimics the functioning of superscalar processors where instructions are issued sequentially from a single stream but can actually be executed in a different order and, possibly, in parallel depending on their mutual dependencies. Figure 1 shows a dummy sequential algorithm and its corresponding STF version. Instead of making three function calls (f , g , h), the equivalent STF submits the three corresponding tasks. The data onto which these functions operate as well as their access mode (Read, Write or Read/Write) are also specified. Because task g accesses data x after task f has accessed it in Write mode, the runtime infers a dependency between tasks f and g . Similarly a dependency is inferred between tasks f and h due to data y . Figure 1 (*right*) shows the DAG corresponding to this STF dummy code.

Many runtime systems Badia et al. [2009], Kurzak and Dongarra [2009], Augonnet et al. [2011] support this paradigm and a complete review of them is out of the scope of this paper. The OpenMP 4.0 standard, for example, supports the STF model through the `task` construct with the recently introduced `depend` clause. Although this OpenMP feature is already available on some compilers (`gcc`

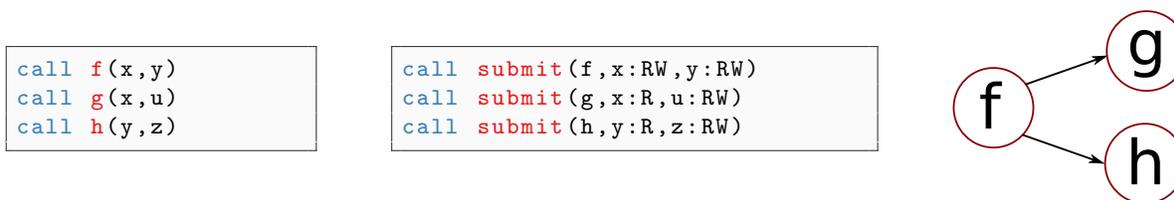


Figure 1: Pseudo-code for a dummy sequential algorithm (*left*), corresponding STF version (*center*) and subsequent DAG (*right*).

and `gfortran`, for instance) we chose to rely on StarPU as it provides a very wide set of features that allows, most importantly, for a better control of the tasks scheduling policy and because it supports accelerators and distributed memory parallelism which will be addressed in future work.

4 STF parallel multifrontal QR

4.1 1D, block-column parallelization

As described in Section 3, the STF model is such that the parallelization of a code can be achieved by simply replacing each function call in the sequential code with the submission of a task that achieves the actual execution of that function on a specific data set. Figure 2 (*right*) shows this transition for the sequential pseudo-code presented in Section 2. The dependencies among the tasks are automatically inferred by the runtime system based on the order of submission and their data access mode; specifically:

- the activation of a node `f` depends on the activation of its children;
- all the other tasks related to a node `f` depend on the node activation;
- the `assembly` of a block-column `j` of a front `c` into its parent `f` depends on all the `panel` and `update` tasks on `c(j)` and on the activation of `f`;
- the `panel` task on a block-column `p` depends on all the `assembly` and all the previous `update` tasks concerning `p`;
- the `update` task on a block-column `u` with respect to `panel` `p` depends on all the `assembly` and all the previous `update` tasks concerning `u` and the related `panel` task on block-column `p`;
- the deactivation on a front `c` can only be executed once all the related `panel`, `update` and `assemble` tasks are completed.

In the parallel execution, one thread is in charge of submitting the tasks by executing the pseudo-code in Figure 2 (*right*); we refer to this thread as the *master thread*. The execution of tasks is instead achieved by *worker threads*. Note that, in principle, the master thread can also act as a worker; this is not the case in our experimental setting. The function called at the end of the pseudo-code in Figure 2 (*right*) is simply a barrier that prevents the master thread from continuing until all of the submitted tasks are executed. In the case where, for example, multiple factorizations have to be executed, this barrier can be removed in order to obtain a pipelining of the successive operations and thus exploit an additional source of concurrency.

4.2 2D, communication-avoiding parallelization by tiles

The parallel factorization algorithm presented in the previous section as well as in our previous work Agullo et al. [2013], Buttari [2013] is based on a 1D partitioning of frontal matrices into block-columns. In this case the node parallelism is achieved because all the updates related to a panel can be executed concurrently and because panel operations can be executed at the same time as updates related to previous panels (this technique is well known under the name of *lookahead*). It is clear that when frontal matrices are strongly overdetermined (i.e., they have many more rows than columns, which is

<pre> 1 forall fronts f in topological order ! allocate and initialize front 3 call activate(f) 5 forall children c of f forall blockcolumns j=1..n in c 7 ! assemble column j of c into f call assemble(c(j), f) 9 end do ! Deactivate child 11 call deactivate(c) end do 13 forall panels p=1..n in f ! panel reduction of column p 15 call panel(f(p)) forall blockcolumns u=p+1..n in f 17 ! update of column u with panel p call update(f(p), f(u)) 19 end do 21 end do 23 end do </pre>	<pre> forall fronts f in topological order ! allocate and initialize front call submit(activate, f:RW, children(f):R) forall children c of f forall blockcolumns j=1..n in c ! assemble column j of c into f call submit(assemble, c(j):R, f:RW) end do ! Deactivate child call submit(deactivate, c:RW) end do forall panels p=1..n in f ! panel reduction of column p call submit(panel, f(p):RW) forall blockcolumns u=p+1..n in f ! update of column u with panel p call submit(update, f(p):R, f(u):RW) end do end do end do call wait_tasks_completion() </pre>
--	---

Figure 2: Pseudo-code for the sequential (*left*) and STF-parallel (*right*) multifrontal QR factorization with 1D partitioned frontal matrices

the most common case in the multifrontal QR method) this approach does not provide much concurrency. In the multifrontal method this problem is mitigated by the fact that multiple frontal matrices are factorized at the same time. However, considering that in the multifrontal factorization most of the computational weight is related to the topmost nodes where tree parallelism is scarce, a 1D front factorization approach can still seriously limit the scalability.

The node parallelism can be improved by employing 2D dense factorization algorithms. Such methods, which have recently received a considerable attention Dongarra et al. [2013], Demmel et al. [2012], Agullo et al. [2011], Bouwmeester et al. [2011], are based on a 2D decomposition of matrices into tiles (or blocks) as shown in Figure 3 (*left*); this permits to break down the panel factorization and the related updates into smaller tasks which leads to a three-fold advantage:

1. the panel factorization and the related updates can be parallelized;
2. the updates related to a panel stage can be started before the panel is entirely reduced;
3. subsequent panel stages can be started before the panel is entirely reduced.

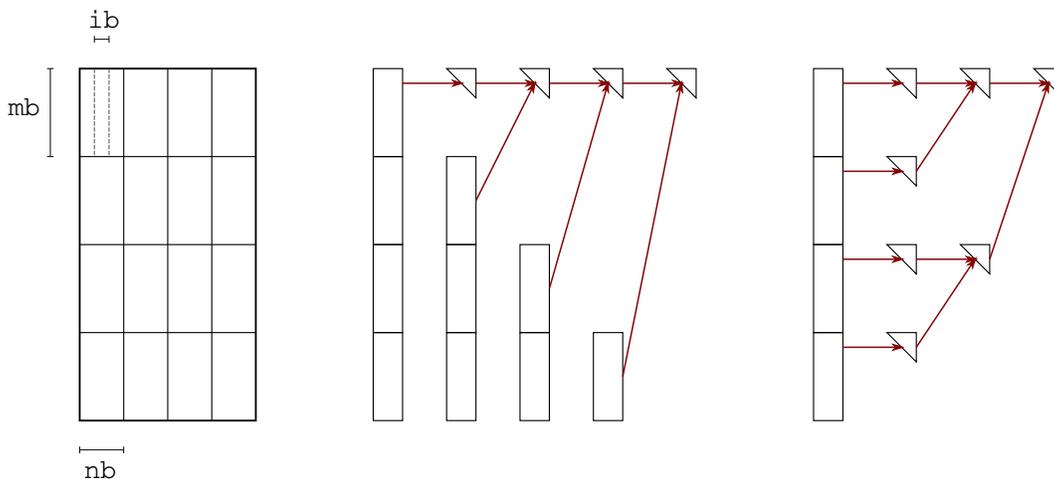


Figure 3: 2D partitioning of a frontal matrix into rectangular tiles of size $mb \times nb$ (*left*). The (*middle*) and (*right*) illustrations show the steps for reducing a panel when a flat ($bh = \infty$) and binary ($bh = 1$) tree is used, respectively.

In these 2D algorithms, the topmost panel tile (the one lying on the diagonal) is used to annihilate the others; this can be achieved in different ways. For example, the topmost tile can be first reduced

into a triangular tile with a general `_GEQRT` LAPACK QR factorization (`GE` stands for general) which can then be used to annihilate the other tiles, one after the other, with `_TPQRT` LAPACK QR factorizations, where `T` and `P` stand for *triangular* and *pentagonal*, respectively. The panel factorization is thus achieved through a flat reduction tree as shown in Figure 3 (*middle*). This approach provides only a moderate additional concurrency with respect to the 1D approach because the tasks within a panel stage cannot be executed concurrently; it allows, however, for a better pipelining of the tasks (points 2 and 3 above). Another possible approach consists in first reducing all the tiles in the panel into triangular tiles using `_GEQRT` operations; this first stage is embarrassingly parallel. Then, through a binary reduction tree, all the triangular tiles except the topmost one are annihilated using `_TPQRT` operations, as shown in Figure 3 (*right*). This second method provides better concurrency but results in a higher number of tasks, some of which are of very small granularity. Also, in this case a worse pipelining of successive panel stages is obtained Dongarra et al. [2013]. In practice a hybrid approach can be used where the panel is split into subsets of size `bh`: each subset is reduced into a triangular tile using a flat reduction tree and then, the remaining tiles are reduced using a binary tree. This is the technique used in the PLASMA library Agullo et al. [2009]. Other reduction trees can also be employed; we refer the reader to the work by Dongarra *et al.* 2013 for an excellent survey of such techniques. Note, also, that all these 2D algorithms induce an extra cost in terms of floating-point operations Buttari et al. [2009]; this can be limited by choosing a relatively small internal block size `ib` for the `_GEQRT`, `_TPQRT` and the corresponding `_GEMQRT` and `_TPMQRT` update operations.

We integrated the flat/binary hybrid approach described above in our multifrontal STF parallel code. This required modifying all the mentioned LAPACK routines in order to cope with the fronts staircase structure. Lines 14-21 in Figure 2 were replaced by the described 2D factorization algorithm and also lines 6-9 were replaced with a double, nested loop to span all the tiles lying on the contribution block. Each assembly operation reads one tile of a node `c` and assembles its coefficients into a subset of the tiles of its parent `f`. As a consequence, some tiles of a front can be fully assembled and ready to be processed before others and before the child nodes are completely factorized. This finer granularity (with respect to the 1D approach presented in the previous section) leads to more concurrency since a better pipelining between a front and its children is now enabled. In our implementation tiles do not have to be square but can be rectangular with more rows than columns. This is only a minor detail from an algorithmic point of view but, as far as we know, it has never been discussed in the literature and, as described in Section 5, provides considerable performance benefits for our case.

4.3 Memory-aware Sequential Task Flow

The memory needed to achieve the multifrontal factorization (QR as well as LU or other types) is not statically allocated at once before the factorization begins but it is allocated and deallocated as the frontal matrices are activated and deactivated, as described in Section 2. Specifically, each activation task allocates all the memory needed to process a front; this memory can be split into two parts:

1. a persistent memory: once the frontal matrix is factorized, this part contains the factors coefficients and, therefore, once allocated it is never freed, unless in an out-of-core execution (where factors are written on disk)¹;
2. a temporary memory: this part contains the contribution block and it is freed by the `deactivate` task once the coefficients it contains have been assembled into the parent front.

As a result, the memory footprint of the multifrontal method in a sequential execution varies greatly throughout the factorization. Starting at zero, it grows fast at the beginning as the first fronts are activated, it then goes up and down as fronts are activated and deactivated until it reaches a maximum value (we refer to this value as the *sequential peak*) and eventually goes down towards the end of the factorization to the point where the only thing left in memory is the factors. The memory consumption varies depending on the particular topological order followed for traversing the elimination tree and techniques exist to determine the memory minimizing traversal for sequential executions Guermouche et al. [2003].

It has to be noted that parallelism normally increases the memory consumption of the multifrontal method simply because, in order to feed the working processes, more work has to be generated and

¹In some other cases the factors can also be discarded like, for example, when the factorization is done for computing the determinant of the matrix.

thus more fronts have to be activated concurrently. The memory consumption of the STF parallel code discussed so far can be considerably higher than the sequential peak (up to 3 times or more). This is due to the fact that the runtime system tries to execute tasks as soon as they are available and to the fact that activation tasks are extremely fast and only depend upon each other; as a result, all the front in the elimination tree are almost instantly allocated at the beginning of the factorization. This behavior, moreover, is totally unpredictable.

This section proposes a method for limiting the memory consumption of parallel executions of our STF code by forcing it to respect a prescribed memory constraint which has to be equal to or bigger than the sequential peak. This technique fulfills the same objective as the one proposed in the work by Rouet 2012 for distributed memory architectures. In the remainder of this section we assume that the tree traversal order is fixed. Recent theoretical studies Marchal et al. [2013] address the problem of finding tree traversal orders that minimize the execution time within a given memory envelope.

The STF model allows to achieve this objective in a relatively easy way. In essence, the proposed technique amounts to subordinating the submission of tasks to the availability of memory. This is done by suspending the execution of the outer loop in Figure 2 (*right*) if not enough memory is available to activate a new front until the required memory amount is freed by `deactivate` tasks. Special attention has to be devoted to avoiding memory deadlocks, though. A memory deadlock may happen because the execution of a front deactivation task depends (indirectly, through the assembly tasks) on the activation of its parent front; therefore the execution may end up in a situation where no more fronts can be activated due to the unavailability of memory and no more deactivation tasks can be executed because they depend on activation tasks that cannot be submitted. Figure 4 shows this case. The table on the right side of the figure shows the memory consumption for a sequential execution where the tree is traversed in natural order; the corresponding sequential peak is equal to 13 memory units. Assume a parallel execution with a memory constraint equal to sequential peak. If no particular care is taken, nothing prevents the runtime system from activating nodes 1, 2 and 4 at once thus consuming 9 memory units; this would result in a deadlock because no other fronts can be activated nor deactivated without violating the constraint.

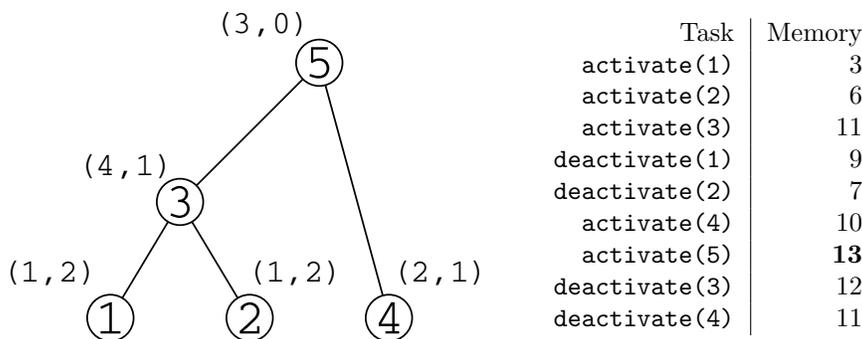


Figure 4: The memory consumption (*right*) for a 5-nodes elimination tree (*left*) assuming a sequential traversal in natural order. Next to each node of the tree the two values corresponding to the factors (permanent) and the contribution block (temporary) sizes in memory units.

This problem can be addressed by ensuring that the fronts are activated in exactly the same order as in a sequential execution: this condition guarantees that, if the tasks submission is suspended due to low memory, it will be possible to execute the deactivation tasks to free the memory required to resume the execution. Note that this only imposes an order in the activation operations and that all the submitted tasks related to activated fronts can still be executed in any order provided that their mutual dependencies are satisfied. This strategy is related to the Banker’s Algorithm proposed by Dijkstra in the early 60’s Dijkstra [1965, 1982].

In our implementation this was achieved as shown in Figure 5. Before performing a front activation (line 4), the master thread, in charge of the tasks submission, checks if enough memory is available to perform the corresponding allocations (line 2). This activation is a very lightweight operation which consists in simple memory bookkeeping (due to the first-touch rule) and therefore does not substantially slow down the tasks submission. The front structure is instead computed in a new `init` task (line 7) submitted to the runtime system which can potentially execute it on any worker thread.

```

1 forall fronts f in topological order
  do while (size(f) > avail_mem) wait
3   ! allocate and initialize front: avail_mem -= size(f)
  call activate(f)
5
  ! initialize the front structure
7  call submit(init, f:RW, children(f):R)
9
  ! front assembly
  forall children c of f
11   ...
   ! Deactivate child: avail_mem += size(cb(f))
13   call submit(deactivate, c:RW)
  end do
15
  ! front factorization
17  ...
19 end do
  call wait_tasks_completion()

```

Figure 5: Pseudo-code showing the implementation of the memory-aware tasks submission.

5 Experimental results

5.1 Experimental setup

We evaluated the above discussed techniques and the code resulting from their implementation on a set of eleven matrices from the UF Sparse Matrix Collection Davis and Hu [2011] plus one (matrix #1) from the HIRLAM² research program. These matrices are listed in Table 1 along with their size, number of nonzeros and operation count obtained when a nested dissection, fill-reducing column permutation is applied and a panel size of 32 is used for the factorization of fronts (this parameter defines how well the fronts staircase structure is exploited and thus affects the global operation count Buttari [2013]); the column permutation is computed with the SCOTCH Pellegrini and Roman [1996] graph partitioning and ordering tool on the graph of $A^T A$. In the case of under-determined systems, the transposed matrix is factorized, as it is commonly done to find the minimum-norm solution of a problem.

#	Mat. name	m	n	nz	op. count (Gflop)
1	hirlam	1385270	452200	2713200	1384160
2	flower_8_4	55081	125361	375266	2851508
3	Rucci1	1977885	109900	7791168	5671282
4	ch8-8-b3	117600	18816	470400	10709211
5	GL7d24	21074	105054	593892	16467844
6	neos2	132568	134128	685087	20170318
7	spal_004	10203	321696	46168124	30335566
8	n4c6-b6	104115	51813	728805	62245957
9	sls	1748122	62729	6804304	65607341
10	TF18	95368	123867	1597545	194472820
11	lp_nug30	95368	123867	1597545	221644546
12	mk13-b5	135135	270270	810810	259751609

Table 1: The set of matrices used for the experiments.

Experiments were done on one node of the ADA supercomputer installed at the IDRIS French

²<http://hirlam.org>

supercomputing center³. This is an IBM x3750-M4 system equipped with four Intel Sandy Bridge E5-4650 (eight cores) processors and 128 GB of memory per node. The cores are clocked at 2.7 GHz and are equipped with Intel AVX SIMD units; the peak performance is of 21.6 Gflop/s per core and thus 691.2 Gflop/s per node for real, double precision computations (which is the arithmetic chosen for all our experiments).

Our code is written in Fortran2003, compiled with the Intel Compilers suite 2015; the code uses Intel MKL v11.2 BLAS and LAPACK routines, the SCOTCH ordering tool v6.0 and the trunk development version of the StarPU runtime system freely available through anonymous SVN access.

For efficiently handling relatively low granularity tasks (especially in the 2D case), we have implemented a scheduling algorithm where each worker thread has its own queue of ready tasks, sorted by priority. Upon completion of a task, the executing worker pushes in its own queue all the newly executable tasks. When a worker runs out of tasks, a work-stealing mechanism is used where victims are selected depending on their distance in the memory hierarchy. This scheduler has been eventually integrated in the StarPU release and is now publicly available under the name *lws* (for *locality work-stealing*).

5.2 Performance and scalability of the 1D and 2D approaches

This section is meant to evaluate the performance of our code. For that we do not impose any memory constraint (the behavior of our memory-aware algorithm is analyzed in Section 5.3). The purpose is to assess the usability and efficiency of runtime systems for complex, irregular workloads such as the QR factorization of a sparse matrix and to evaluate the effectiveness of parallel, 2D factorization algorithms in the multifrontal QR method. The reference sequential execution times are obtained with a purely sequential code (no potential runtime overhead) with no frontal matrix partitioning (no granularity trade-off).

The performance of both parallel 1D and 2D methods depends on the choice of the values for a number of different parameters. For the 1D case these are the block-column **nb** on which depends the amount of concurrency and the internal block size **ib** on which depend the efficiency of elementary BLAS operations and the global amount of flop (this parameter defines how well the staircase structure of each front is exploited). For the parallel 2D STF case these parameters are, the tiles size (**mb**, **nb**), the type of panel reduction algorithm set by the **bh** parameter described in Section 4.2 and the internal block size **ib**. The choice of these values depends on a number of factors, such as the number of working threads, the size and structure of the matrix, the shape of the elimination tree and of frontal matrices and the features of the underlying architecture. It has to be noted that these parameters may be set to different values for each frontal matrix; moreover, it would be possible to let the software automatically choose values for these parameters. Both these tasks are very difficult and challenging and are out of the scope of this work. Therefore, for our experiments we performed a large number of runs with varying values for all these parameters, using the same values for all the fronts in the elimination tree, and selected the best results (shortest running time) among those. For the sequential runs internal block sizes $\text{ib}=\{32, 40, 64, 80, 128\}$ were used for a total of five runs per matrix. For the 1D parallel STF case, the used values were $(\text{nb}, \text{ib})=\{(128, 32), (128, 64), (128, 128), (160, 40), (160, 80)\}$ for a total of five runs per matrix. For the 2D case $(\text{nb}, \text{ib})=\{(160, 32), (160, 40), (192, 32), (192, 64)\}$, $\text{mb}=\{\text{nb}, \text{nb}*2, \text{nb}*3, \text{nb}*4\}$ and $\text{bh}=\{4, 8, 12, 16, 20, 24, \infty\}$ for a total of 112 runs per matrix. This choice of values stems from the following observations. For the 1D case, tasks are of relatively coarse grain and, therefore, it is beneficial to choose smaller values for **nb** with respect to the 2D case in order to improve concurrency. On the other side, in the 2D case concurrency is aplenty and therefore it is beneficial to choose a relatively large **nb** value in order to achieve a better BLAS efficiency and keep the runtime overhead small; the internal block size **ib**, however, has to be relatively small to keep the flop overhead under control.

All of the results presented in this section were produced without storing the factors in order to extend the tests to the largest matrices in our experimental set that could not otherwise be factorized (even in sequential) on the target platform. This was achieved by simply deallocating the tiles containing the factors coefficients at each `deactivate` task rather than keeping them in memory. As confirmed by experiments that we do not report here for the sake of space and readability, this does not have a relevant impact on the following performance analysis.

³<http://www.idris.fr>

Mat.	Sequential reference			Parallel 1D STF				Parallel 2D STF					
	ib	Time (s.)	Gflop/s	nb	ib	Time (s.)	Gflop/s	mb	nb	ib	bh	Time (s.)	Gflop/s
1	40	1.00E+02	14.0	128	64	5.39E+00	269.4	576	192	32	4	4.37E+00	316.1
2	32	1.72E+02	17.1	128	128	9.91E+00	308.8	480	160	40	8	7.39E+00	388.0
3	80	3.48E+02	16.7	128	128	1.99E+01	299.1	480	160	40	12	1.44E+01	394.8
4	128	6.00E+02	18.0	128	128	3.08E+01	356.0	480	160	40	20	2.49E+01	430.6
5	80	9.26E+02	17.9	128	128	4.75E+01	354.0	480	160	32	16	3.78E+01	435.4
6	80	1.18E+03	17.7	128	32	5.16E+01	390.6	640	160	40	4	4.78E+01	424.4
7	128	1.72E+03	17.8	128	128	8.38E+01	366.0	480	160	40	24	6.92E+01	439.0
8	128	3.42E+03	18.5	128	128	1.50E+02	420.6	480	160	32	∞	1.40E+02	441.9
9	128	4.39E+03	15.2	128	64	6.32E+02	104.4	576	192	32	24	1.72E+02	379.5
10	128	1.09E+04	17.9	128	128	4.34E+02	452.6	576	192	32	∞	4.28E+02	453.7
11	128	1.22E+04	18.2	128	128	5.29E+02	423.3	576	192	64	∞	4.81E+02	462.4
12	128	1.51E+04	17.3	128	128	6.24E+02	421.9	576	192	64	20	5.66E+02	460.6

Table 2: Optimum performance on ADA (32 cores).

Table 2 shows for both the 1D and 2D algorithms the parameters values delivering the shortest execution time along with the corresponding attained factorization time and Gflop rate. The Gflop rates reported in the table are related to the operation count achieved with the internal block size `ib` reported therein without including, for the 2D case, the extra flops due to the 2D algorithm; these are difficult to count and can simply be regarded as a price to pay to achieve better concurrency. For the 1D case, the smaller block-column size of 128 always delivers better performance because it offers a better compromise between concurrency and efficiency of BLAS operations, whereas a large internal block size is more desirable because it leads to better BLAS speed despite a worse exploitation of the fronts staircase structure. For the 2D case it is interesting to note that the shortest execution time is never attained with square tiles; this is likely because rectangular tiles lead to a better compromise between granularity of tasks (and, thus, efficiency of operations) and concurrency.

Figure 6, generated with the timing data in Table 2, shows the speedup achieved by the 1D and 2D parallel codes with respect to the sequential one when using all the 32 cores available on the system. This figure shows that the 2D algorithms provide better gains on smaller size matrices or on those where frontal matrices are extremely overdetermined (which is the case for matrix #9) where the 1D method does not provide enough concurrency; on larger size matrices the difference between the two parallelization strategy decreases as the 1D approach achieves very good performance. The average speedup achieved by the 2D code is 24.67 with a standard deviation of 1.05, reaching a maximum of 26.79 for matrix #12. For the 1D case, instead, the average is 19.85 with a standard deviation of 4.84. In conclusion, the 2D code achieves better and more consistent scalability over our heterogeneous set of matrices.

Speedup 1D vs 2D -- 32 cores

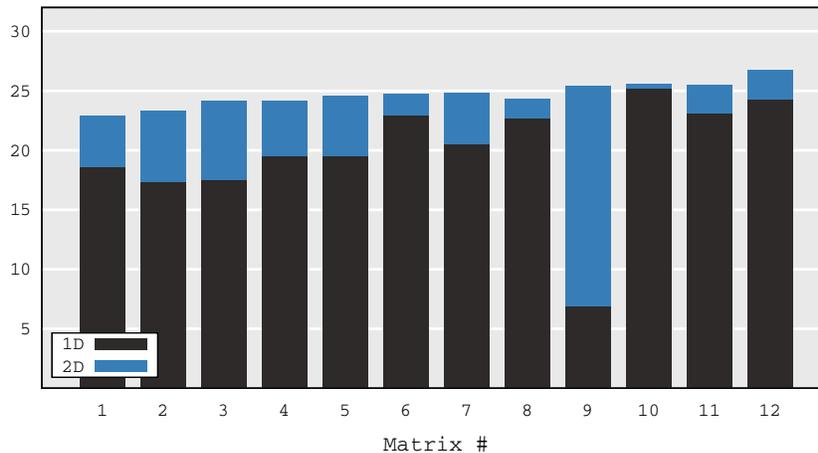


Figure 6: Speedup of the 1D and 2D algorithms with respect to the sequential case on ADA (32 cores).

As in our previous runtime based approach Agullo et al. [2013], we evaluate the performance of our

implementation by means of efficiency measures based on the cumulative time spent inside tasks $t_t(p)$, the cumulative time spent by workers in the runtime system $t_r(p)$ and the cumulative time spent idle $t_i(p)$, p being the number of threads. The efficiency $e(p)$ of the parallelization can then be defined in terms of these cumulative timings as follows:

$$e(p) = \frac{t_t(1)}{t_t(p) + t_r(p) + t_i(p)} = \frac{\overbrace{t_t(1)}^{e_t}}{t_t(p)} \cdot \frac{\overbrace{t_t(p)}^{e_r}}{t_t(p) + t_r(p)} \cdot \frac{\overbrace{t_t(p) + t_r(p)}^{e_s}}{t_t(p) + t_r(p) + t_i(p)}.$$

This expression allows us to decompose the efficiency as the product of three well identified effects:

- e_t : the **tasks efficiency** which measures whether and how the execution of tasks was slowed-down by parallelism. This may be mainly due to two factors. First, our target system is a NUMA architecture therefore the concurrent execution of tasks may suffer from a reduced memory bandwidth and from memory access contention. Second, 2D factorization algorithms induce an extra cost in terms of floating-point operations Buttari et al. [2009]; this overhead may be reduced by choosing an appropriate inner blocking size for the kernels but cannot, in general, be considered negligible.
- e_r : the **runtime efficiency**, which measures the cost of the runtime system with respect to the actual work done.
- e_s : the **scheduling efficiency**, which measures how well the tasks have been pipelined as a result of the scheduling decisions.

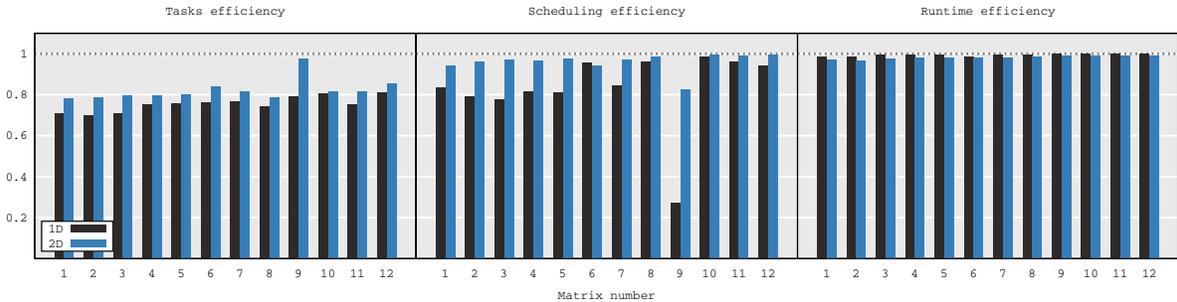


Figure 7: Efficiency measures on ADA (32 cores).

Figure 7 shows a comparison of the three efficiency measures for the 1D and 2D algorithms. The 2D algorithms achieve better tasks efficiency despite the higher number of flop required; this is probably due to a better efficiency of BLAS operations resulting from the 2D partitioning of frontal matrices into tiles which have a more cache-friendly shape than the extremely tall and skinny block-columns used in the 1D algorithm. Note that the sequential code employs a LAPACK-style QR factorization for the frontal matrices (the only difference being in the exploitation of the fronts staircase structure) and may, therefore, suffer from a poor cache behavior in the case of extremely overdetermined frontal matrices; this explains why the tasks efficiency for the 2D code on matrix #9 is much better than for the other matrices and very close to 1. Not surprisingly, the 2D code achieves much better scheduling efficiency (i.e., less idle time) than the 1D code on all matrices: this results from a much higher concurrency, which is the purpose of the 2D code. As for the runtime efficiency, it is in favor of the 1D implementation due to a much smaller number of tasks with bigger granularity and simpler dependencies. However the performance loss induced by the runtime is extremely small in both cases: less than 2% on average and never higher than 4% for the 2D implementation, showing the efficiency of our *lws* scheduling strategy for dealing with tasks of small granularity.

5.3 Performance under a memory constraint

This section describes and analyses experiments that aim at assessing the effectiveness of the memory aware scheduling presented in Section 4.3. Here we are interested in two scenarios:

1. In-Core (IC) execution: this is the most common case where the computed factors are kept in memory. In this case matrices #11 and #12 could not be used because of the excessive memory consumption.
2. Out-Of-Core (OOC) execution: in this scenario the factors are written to disk as they are computed in order to save memory. In this case the memory consumption is more irregular and more considerably increased by parallelism. We simulate this scenario by discarding the factors as we did in the previous section; note that by doing so we are assuming that the overhead of writing data to disk has a negligible effect on the experimental analysis reported here.

These experiments were meant to measure the performance of both the 1D and 2D factorization (with the parameters values in Table 2) within an imposed memory footprint; experiments were performed using 32 cores with memory constraints equal to $\{1.0, 1.2, 1.4, 1.6, 1.8, 2.0\} \times \textit{sequential_peak}$ both when the factors are kept in memory and when they are discarded.

For all 1D and 2D IC tests as well as all 2D OOC tests, a performance as high as the non constrained case (such as discussed in Section 5.2) could be achieved with a memory exactly equal to the sequential peak, which is the lower bound that a parallel execution can achieve. This shows the extreme efficiency of the memory-aware mechanism for achieving high-performance within a limited memory footprint. Combined with the 2D numerical scheme, the memory-aware algorithm is thus extremely robust since it could process all considered matrices at maximum pace with the minimum possible memory consumption. The only cases which requested a higher memory usage for achieving the same performance as the non constrained algorithm were for matrices #1 and #7 in the 1D OOC case.

To explain this extreme efficiency, we performed the following analysis. As explained in Section 4.3, prior to activating a front, the master thread checks whether enough memory is available to achieve this operation. If it is not the case, the master thread is put to sleep and later woken up as soon as one `deactivate` task is executed; at this time the master thread checks again for the availability of memory. The master thread stays in this loop until enough deactivation tasks have been executed to free up the memory needed to proceed with the next front activation. Every time the master thread was suspended or resumed we recorded the time stamp and the number of ready tasks (i.e., those whose dependencies were all satisfied).

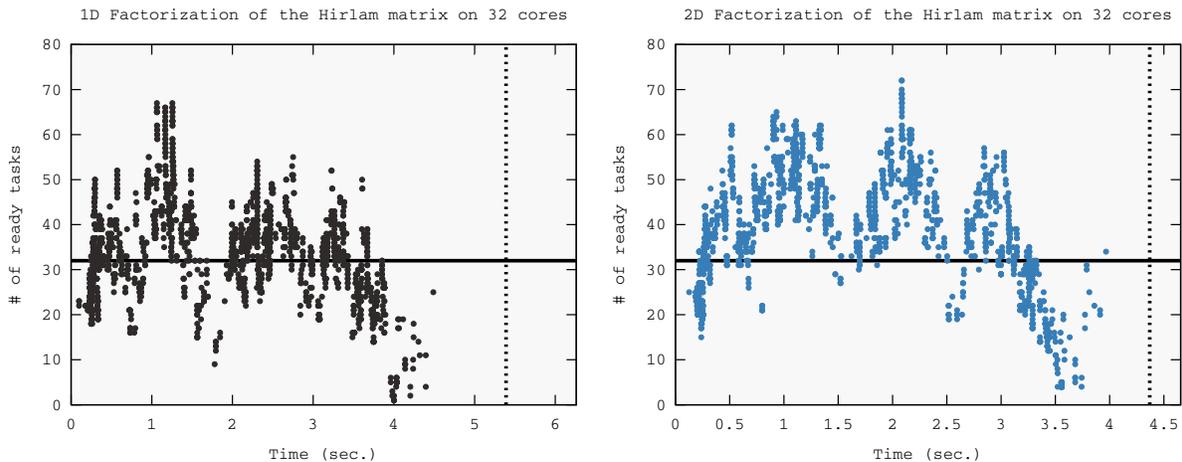


Figure 8: Concurrency under a memory constraint for the Hirlam matrix on ADA (32 cores).

Figure 8 shows the collected data for the Hirlam matrix with an imposed memory consumption equal to the sequential peak, in the OOC case using both the 1D (*left*) and 2D (*right*) methods. In this figure, each (x, y) point means that at time x the master thread was suspended or resumed and that, at that time, y tasks were ready for execution or being executed. The width of each graph shows the execution time of the memory constrained factorization whereas the vertical dashed line shows the execution time when no limit on the memory consumption is imposed. The figure leads to the following observations:

- in both the 1D and 2D factorizations, the number of ready tasks falls, at some point, below the number of available cores (the horizontal, solid line); this lack of tasks is responsible for a longer execution time with respect to the unconstrained case.
- in the 1D factorization this lack of tasks is more important; this can be explained by the fact that the 1D method delivers much lower concurrency than the 2D one and therefore, suspending the submission of tasks may lead more quickly to threads starvation. As a result, the difference in the execution times of the constrained and unconstrained executions is more evident in the 1D factorization.

For the 1D, OOC factorization of matrix #7, the performance increases smoothly when the memory constraint is gradually increased as described above. The corresponding execution times are: 126.7, 118.3, 116.1, 114.9, 94.7, 93.0 seconds.

For all other tests, either the number of tasks is always (much) higher than the number of used cores or the tasks submission is never (or almost never) interrupted due to the lack of memory; as a result, no relevant performance degradation was observed with respect to the case where no memory constraint is imposed. This behavior mainly results from two properties of the multifrontal QR factorization. First, the size of the contribution blocks is normally very small compared to the size of factors, especially in the case where frontal matrices are overdetermined. Second, the size of a front is always greater than or equal to the sum of the sizes of all the contribution blocks associated with its children (because in the assembly operation, contribution blocks are not summed to each other but stacked). As a result, in the sequential multifrontal QR factorization, the memory consumption grows almost monotonically and in most cases the sequential peak is achieved on the root node or very close to it. For this reason, when the tasks submission is interrupted in a memory-constrained execution, a large portion of the elimination tree has already been submitted and the number of available tasks is considerably larger than the number of working threads. We expect that other types of multifrontal factorizations (LU , for instance) are more sensitive to the memory constraint because they do not possess the two properties described above. By the same token, it is reasonable to expect that imposing a memory constraint could more adversely affect performance when larger numbers of threads are used.

6 Conclusions

With this work we aimed at assessing the usability and effectiveness, in a shared-memory environment, of modern runtime systems for handling a complex, irregular workload such as the multifrontal QR factorization of a sparse matrix which is characterized by a very high number of extremely heterogeneous tasks with complex dependencies. We showed how to use the Sequential Task Flow programming model to implement the parallelization approach presented in our previous work Buttari [2013] for such method; this model leads to a simpler, more portable, maintainable, and scalable code. We showed how the resulting code could be improved with a relatively limited effort thanks to the simplicity of the STF programming model. First, we addressed the performance scalability issues of the 1D parallel approach by implementing communication avoiding fronts factorization by tiles. Second, we presented an approach that, by simply controlling the submission of tasks to the runtime system, allows for achieving the factorization within a prescribed memory footprint.

We presented experimental results that assess the effectiveness of our approach by showing that modern runtime systems are very lightweight and efficient, that tile, communication-avoiding parallel algorithms provide better concurrency and scalability to the multifrontal QR factorization and that the memory consumption of this method can be simply yet effectively controlled without introducing a performance penalty in nearly all of the considered test cases. In conclusion, the resulting code achieves a great performance and scalability as well as an excellent robustness when it comes to memory consumption.

7 Acknowledgments

We wish to thank the RUNTIME team of Inria Bordeaux sud-ouest for their help in using the StarPU runtime system and for promptly addressing our requests. We also thank B. Lizé and M. Faverge for their constructive suggestions on a preliminary version of this manuscript.

References

- Emmanuel Agullo, Jim Demmel, Jack Dongarra, Bilel Hadri, Jakub Kurzak, Julien Langou, Hatem Ltaief, Piotr Luszczek, and Stanimire Tomov. Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects. *Journal of Physics: Conference Series*, 180(1):012037, 2009. URL <http://stacks.iop.org/1742-6596/180/i=1/a=012037>.
- Emmanuel Agullo, Jack Dongarra, Rajib Nath, and Stanimire Tomov. Fully empirical autotuned QR factorization for multicore architectures. *CoRR*, abs/1102.5328, 2011.
- Emmanuel Agullo, Alfredo Buttari, Abdou Guermouche, and Florent Lopez. Multifrontal qr factorization for multicore architectures over runtime systems. In Felix Wolf, Bernd Mohr, and Dieter Mey, editors, *Euro-Par 2013 Parallel Processing*, volume 8097 of *Lecture Notes in Computer Science*, pages 521–532. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-40046-9. doi: 10.1007/978-3-642-40047-6_53. URL http://dx.doi.org/10.1007/978-3-642-40047-6_53.
- Randy Allen and Ken Kennedy. *Optimizing Compilers for Modern Architectures: A Dependence-Based Approach*. Morgan Kaufmann, 2002.
- Patrick R. Amestoy, Iain S. Duff, and Chiara Puglisi. Multifrontal QR factorization in a multiprocessor environment. *Int. Journal of Num. Linear Alg. and Appl.*, 3(4):275–300, 1996.
- The OpenMP architecture review board. OpenMP 4.0 complete specifications, 2013.
- Nimar S. Arora, Robert D. Blumofe, and C. Greg Plaxton. Thread scheduling for multiprogrammed multiprocessors. *Theory Comput. Syst.*, 34(2):115–144, 2001.
- Cedric Augonnet, Samuel Thibault, Raymond Namyst, and Pierre-André Wacrenier. StarPU: A Unified Platform for Task Scheduling on Heterogeneous Multicore Architectures. *Concurrency and Computation: Practice and Experience, Special Issue: Euro-Par 2009*, 23:187–198, February 2011. doi: 10.1002/cpe.1631. URL <http://hal.inria.fr/inria-00550877>.
- Rosa M. Badia, José R. Herrero, Jesús Labarta, Josep M. Pérez, Enrique S. Quintana-Ortí, and Gregorio Quintana-Ortí. Parallelizing dense and banded linear algebra libraries using SMPSS. *Concurrency and Computation: Practice and Experience*, 21(18):2438–2456, 2009.
- George Bosilca, Aurelien Bouteiller, Anthony Danalis, Thomas Héroult, Pierre Lemarinier, and Jack Dongarra. DAGuE: A generic distributed DAG engine for high performance computing. *Parallel Computing*, 38(1-2):37–51, 2012.
- George Bosilca, Aurelien Bouteiller, Anthony Danalis, Thomas Herault, Piotr Luszczek, and Jack Dongarra. Dense linear algebra on distributed heterogeneous hardware with a symbolic dag approach. *Scalable Computing and Communications: Theory and Practice*, 2013.
- Henricus Bouwmeester, Mathias Jacquelin, Julien Langou, and Yves Robert. Tiled qr factorization algorithms. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '11, pages 7:1–7:11, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0771-0. doi: 10.1145/2063384.2063393. URL <http://doi.acm.org/10.1145/2063384.2063393>.
- Zoran Budimčić, Michael Burke, Vincent Cavé, Kathleen Knobe, Geoff Lowney, Ryan Newton, Jens Palsberg, David Peixotto, Vivek Sarkar, Frank Schlimbach, et al. Concurrent collections. *Scientific Programming*, 18(3):203–217, 2010.
- Alfredo Buttari. Fine-grained multithreading for the multifrontal *qr* factorization of sparse matrices. *SIAM Journal on Scientific Computing*, 35(4):C323–C345, 2013. doi: 10.1137/110846427. URL <http://epubs.siam.org/doi/abs/10.1137/110846427>.
- Alfredo Buttari, Julien Langou, Jakub Kurzak, and Jack Dongarra. A class of parallel tiled linear algebra algorithms for multicore architectures. *Parallel Comput.*, 35:38–53, January 2009. ISSN 0167-8191. doi: 10.1016/j.parco.2008.10.002. URL <http://dl.acm.org/citation.cfm?id=1486274.1486415>.

- Michel Cosnard and Emmanuel Jeannot. Automatic parallelization techniques based on compact DAG extraction and symbolic scheduling. *Parallel Processing Letters*, 11(01):151–168, 2001. doi: 10.1142/S012962640100049X. URL <http://www.worldscientific.com/doi/abs/10.1142/S012962640100049X>.
- Timothy A. Davis. Algorithm 915, SuiteSparseQR: Multifrontal multithreaded rank-revealing sparse QR factorization. *ACM Trans. Math. Softw.*, 38(1):8:1–8:22, December 2011. ISSN 0098-3500. doi: 10.1145/2049662.2049670. URL <http://doi.acm.org/10.1145/2049662.2049670>.
- Timothy A. Davis and Yifan Hu. The university of Florida sparse matrix collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, December 2011. ISSN 0098-3500. doi: 10.1145/2049662.2049663. URL <http://doi.acm.org/10.1145/2049662.2049663>.
- James Demmel, Laura Grigori, Mark Hoemmen, and Julien Langou. Communication-optimal parallel and sequential qr and lu factorizations. *SIAM J. Sci. Comput.*, 34(1):206–239, February 2012. ISSN 1064-8275. doi: 10.1137/080731992. URL <http://dx.doi.org/10.1137/080731992>.
- Edsger W. Dijkstra. Een algoritme ter voorkoming van de dodelijke omarming. circulated privately, 1965. URL <http://www.cs.utexas.edu/users/EWD/ewd01xx/EWD108.PDF>.
- Edsger W. Dijkstra. The mathematics behind the banker’s algorithm. In *Selected Writings on Computing: A personal Perspective*, Texts and Monographs in Computer Science, pages 308–312. Springer New York, 1982. ISBN 978-1-4612-5697-7. doi: 10.1007/978-1-4612-5695-3_54. URL http://dx.doi.org/10.1007/978-1-4612-5695-3_54.
- Jack Dongarra, Mathieu Faverge, Thomas Hérault, Mathias Jacquelin, Julien Langou, and Yves Robert. Hierarchical QR factorization algorithms for multi-core clusters. *Parallel Computing*, 39(4-5):212–232, 2013. doi: 10.1016/j.parco.2013.01.003. URL <http://hal.inria.fr/hal-00809770>.
- Iain S. Duff and John K. Reid. The multifrontal solution of indefinite sparse symmetric linear systems. *ACM Transactions On Mathematical Software*, 9:302–325, 1983.
- Abdou Guermouche, Jean-Yves L’Excellent, and Gilles Utard. Impact of reordering on the memory of a multifrontal solver. *Parallel Computing*, 29(9):1191–1218, 2003.
- Everton Hermann, Bruno Raffin, François Faure, Thierry Gautier, and Jérémie Allard. Multi-GPU and multi-CPU parallelization for interactive physics simulations. In *Euro-Par (2)*, pages 235–246, 2010.
- Jonathan Hogg, John K. Reid, and Jennifer A. Scott. Design of a multicore sparse Cholesky factorization using DAGs. *SIAM J. Scientific Computing*, 32(6):3627–3649, 2010.
- Kyungjoo Kim and Victor Eijkhout. A parallel sparse direct solver via hierarchical DAG scheduling. *ACM Trans. Math. Softw.*, 41(1):3:1–3:27, October 2014. ISSN 0098-3500.
- Jakub Kurzak and Jack Dongarra. Fully dynamic scheduler for numerical computing on multicore processors. *LAPACK working note*, lawn220, 2009.
- Xavier Lacoste, Mathieu Faverge, Pierre Ramet, Samuel Thibault, and George Bosilca. Taking advantage of hybrid systems for sparse direct solvers via task-based runtimes, 05/2014 2014.
- Loris Marchal, Oliver Sinnen, and Frederic Vivien. Scheduling tree-shaped task graphs to minimize memory and makespan. volume 0, pages 839–850, Los Alamitos, CA, USA, 2013. IEEE Computer Society. doi: <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2013.55>.
- François Pellegrini and Jean Roman. SCOTCH: A Software Package for Static Mapping by Dual Recursive Bipartitioning of Process and Architecture Graphs. In *Proceedings of HPCN’96, Brussels*, LNCS 1067, pages 493–498, April 1996.
- Gregorio Quintana-Ortí, Enrique S. Quintana-Ortí, Robert A. Van De Geijn, Field G. Van Zee, and Ernie Chan. Programming matrix algorithms-by-blocks for thread-level parallelism. *ACM Trans. Math. Softw.*, 36(3), 2009.

François-Henry Rouet. *Memory and performance issues in parallel multifrontal factorizations and triangular solutions with sparse right-hand sides*. Thèse de doctorat, Institut National Polytechnique de Toulouse, Toulouse, France, octobre 2012. URL <http://tel.archives-ouvertes.fr/tel-00785748>.

Robert Schreiber. A new implementation of sparse Gaussian elimination. *ACM Transactions On Mathematical Software*, 8:256–276, 1982.

Haluk Topcuoglu, Salim Hariri, and Min-you Wu. Performance-effective and low-complexity task scheduling for heterogeneous computing. *IEEE Transactions on Parallel and Distributed Systems*, 13(3):260–274, Mar 2002. ISSN 1045-9219. doi: 10.1109/71.993206.