



A Sparse symmetric indefinite direct solver for GPU architectures

J Hogg, E Ovtchinnikov, J Scott

June 2014

Submitted for publication in ACM Transactions on Mathematical Software

RAL Library
STFC Rutherford Appleton Laboratory
R61
Harwell Oxford
Didcot
OX11 0QX

Tel: +44(0)1235 445384
Fax: +44(0)1235 446403
email: libraryral@stfc.ac.uk

Science and Technology Facilities Council preprints are available online
at: <http://epubs.stfc.ac.uk>

ISSN 1361- 4762

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

A sparse symmetric indefinite direct solver for GPU architectures

Jonathan Hogg, Evgueni Ovtchinnikov and Jennifer Scott¹

ABSTRACT

In recent years, there has been considerable interest in the potential for graphics processing units (GPUs) to speed up the performance of sparse direct linear solvers. Efforts have focused on symmetric positive-definite systems for which no pivoting is required while little progress has been reported for the much harder indefinite case. We address this challenge by designing and developing a sparse symmetric indefinite solver SSIDS. This new library-quality solver is designed for use on GPU architectures and incorporates threshold partial pivoting within a multifrontal approach. Both the factorize and the solve phases are performed using the GPU. Another important feature is that the solver produces bit-compatible results. Numerical results for indefinite problems arising from a range of practical applications demonstrate that, for large problems, SSIDS achieves performance improvements of up to a factor of $7\times$ compared with a state-of-the-art multifrontal solver on a multicore CPU.

Keywords: sparse linear systems, indefinite symmetric systems, direct solver, multifrontal, GPU, bit compatibility.

AMS(MOS) subject classifications: 65F05, 65F50

¹ Scientific Computing Department, Rutherford Appleton Laboratory, Harwell Oxford, Oxfordshire, OX11 0QX, UK.

Correspondence to: jonathan.hogg@stfc.ac.uk

This work was supported by EPSRC grant EP/J010553/1.

1 Introduction

The solution of large sparse linear systems of equations, $Ax = b$, lies at the heart of many scientific and engineering problems. Solving such systems frequently represents a computational bottleneck and this has resulted in significant effort being invested over the past fifty years in the development of both iterative and direct solution methods. While the former have the advantage of requiring limited memory and are able to achieve good performance on modern machines through the efficient implementation of matrix-vector products, the latter are popular because of their robustness, allowing them to be frequently used as “black-box” codes within industrial applications. Over the years, much effort has gone into exploiting novel computing architectures to improve the performance of sparse direct algorithms. In this paper, we seek to accelerate the factorization and subsequent solution of large sparse indefinite matrices using graphics processing units (GPUs).

Modern GPUs offer a significant advantage over mass-market CPUs in terms of both performance per pound and performance per watt. Comparing systems utilizing similar amounts of power, GPUs can offer roughly $5\times$ more floating-point performance and memory bandwidth. They also provide a test bed for the alternative programming models likely to be required as the core count in future CPUs increases. Current GPUs can be modelled as vector machines, with each set of cores sharing a common program counter and other resources.

In this paper, we consider only NVIDIA GPUs, as these are by far the most predominant ones used for high-performance computing at present. Resources on such a GPU are arranged into a number of pools called streaming multiprocessors (SMs). Each SM has up to 192 floating-point cores that share (to varying degrees) hardware resources. All cores in an SM are connected to the same pool of fast cache, that also doubles as a shared memory area, allowing efficient exchange of information between them. The hardware supports running multiple threads per core (up to 32) and uses fast context switching to hide both instruction and memory latency.

Symmetric indefinite sparse linear systems arise in a wide variety of important technical and scientific applications. These include mixed finite-element methods in engineering fields such as fluid and solid mechanics, and interior point algorithms in both linear and nonlinear optimization. Given a sparse symmetric indefinite matrix A , an LDL^T factorization algorithm computes a unit lower triangular matrix L and a block diagonal matrix D , with 1×1 and 2×2 blocks on the diagonal, such that $A = LDL^T$. The solution process is completed by solving $Ly = b$ for y , then $Dz = y$ for z and, finally, $L^Tx = z$ for x . In practice, the system matrix is prescaled to help with numerical stability and preordered to minimise the amount of fill in the L factor, so that $\bar{A} = PSASP^T$ is factorized, where S is a diagonal scaling matrix and P is a permutation matrix. For simplicity of notation, throughout this paper, we work only with A , but employ scaling and ordering when performing numerical experiments.

The multifrontal method [6, 8] is often used to perform an indefinite factorization. This method is built upon the decomposition of the factorization into a large number of small dense submatrices (frontal matrices) that are able to exploit high-level BLAS operations, leading to high flop rates and efficient performance. The computational workflow is structured as a tree, with a frontal matrix at each node of the tree. The edges represent data movement in which the results from a child node are assembled into the frontal matrix of its parent. As each child of a parent can be computed independently, parallelism can be achieved both through coarse task parallelism across the tree and fine-grained parallelism within the linear algebra at each node. Modern sparse solvers that run on CPUs and use this parallel multifrontal approach include MUMPS [29], WSMP [12] and HSL_MA97 [16, 18, 20]. A number of recent papers and reports have considered accelerating the multifrontal algorithm using a GPU [10, 24, 25, 26, 28, 32]. In the symmetric case, these have concentrated on positive-definite systems that do not require pivoting for stability; little attention has been paid to the problem of efficiently incorporating robust pivoting strategies that are needed for the numerical stability of indefinite systems. Furthermore, the proposed implementations port only the computationally intensive part of the algorithm to the GPU, that is, they aim to run the factorization of the largest frontal matrices that are close to the root of the tree on the GPU while the rest of the factorization (and the solve phase that uses the computed factors) is performed on the host CPU, thus significantly increasing the data movement involved and limiting the overall gains achieved by employing GPUs. Thus the main contributions

of our work are:

- The design and development of an efficient pivoting strategy within a sparse indefinite multifrontal code for running on GPU architectures. We employ threshold partial pivoting: the novelty of our approach lies in the organisation of the computation to exploit the GPU. Our data structures are not static since we have to accommodate pivots that are delayed (that is, pivots that fail the threshold stability test and have to be delayed until later in the factorization when, after further update operations, they can be used stably). This not only increases the complexity of the coding but can also adversely effect performance so we employ scaling and ordering strategies to minimise delayed pivots [19].
- An implementation that allows all the frontal matrices, including the small ones at the leaf nodes of the tree, to be efficiently factorized on the GPU.
- A solve phase that is implemented on the GPU. As far as we are aware, this issue has not been discussed before as all attention has been on the factorization phase that, traditionally, has been the most time-consuming part of the solution process. However, if the factorization is speeded up sufficiently and if multiple solves in sequence are required following the factorization, a serial solve phase can become a bottleneck. The solve phase is memory bound (see [15] for a discussion); implementing it on a GPU allows exploitation of the increased memory bandwidth.
- A multifrontal code that produces bit-compatible (reproducible) results in the sense that two runs of the solver on the same machine with identical input data produces identical output. Not only is this an important aid in debugging and correctness checking, some industries (e.g. nuclear, finance) require reproducible results to satisfy regulatory requirements.

Most of the past studies into multifrontal solvers on GPUs have produced software prototypes, rather than software that has been made generally available. By contrast, our new solver, which is called **SSIDS** (**S**parse **S**ymmetric **I**ndefinite **D**irect **S**olver), is an open-source package and is part of the **SPRAL** library; it can be downloaded from <http://www.numerical.rl.ac.uk/spral>. The code is written in a mixture of Fortran 2003 and CUDA, using double precision arithmetic throughout. It is fully documented, tested and maintained by the Numerical Analysis Group at the STFC Rutherford Appleton Laboratory.

The remainder of this paper is organised as follows. Section 2 provides a brief introduction to the multifrontal method for indefinite systems. In Section 3, we discuss the design of **SSIDS** and, in particular, we discuss the assembly of the child nodes into the parent, the stable partial factorization of the frontal matrices, the subsequent update operations, and the implementation of the solve phase. In Section 4, we present numerical results for a range of indefinite problems arising from practical applications. The performance of our GPU-accelerated solver is compared with that of our state-of-the-art OpenMP multifrontal code **HSL_MA97**. Results are given for both the factorize and solve phases of **SSIDS**; both demonstrate performance improvements of up to 7x compared to **HSL_MA97**. Finally, we summarize our findings in Section 5.

2 Overview of the multifrontal method

In common with other sparse direct methods, the multifrontal method comprises the following phases.

- An ordering phase that exploits the sparsity (non-zero) structure of A to determine a pivot sequence (that is, the order in which the Gaussian eliminations will be performed). The choice of pivot sequence significantly influences the memory requirements of the factorization, the fill in the factor L , and the number of floating-point operations required to carry out the factorization.
- An analyse phase that uses the pivot sequence and sparsity pattern of A to establish the work flow and data structures for the factorization.
- A factorization phase that optionally scales the matrix before performing the numerical factorization. Following the analyse phase, more than one matrix with the same sparsity pattern may be factorized.

- A solve phase that performs forward elimination followed by back substitution using the stored factors. Repeated calls to the solve phase with one or more right-hand sides may follow the factorization phase. This is typically used to implement iterative refinement (see, e.g. [11]) to improve the accuracy of the computed solution.

2.1 Selecting an ordering

Considerable effort in the last thirty years has concentrated on the development of algorithms that generate good pivot orders. These include methods based upon the minimum degree algorithm [1, 2, 27, 31] and on nested dissection [9]. The former perform well on many small and medium-sized problems and on highly sparse matrices while the latter has been found to give better orderings for very large problems, particularly those from three-dimensional discretizations. Furthermore, a parallel solver is generally better able to exploit parallelism if a nested dissection ordering is used. Many direct solvers offer a choice of orderings, including either their own implementation of nested dissection or an explicit interface to an established nested-dissection library such as the METIS graph partitioning package [21, 22].

2.2 The analyse phase

Analyse phase is designed to construct the sparsity pattern of the factor L . It does this by building an *elimination tree*. This graph describes the structure of L in terms of the data dependence between pivotal columns. It permits permutations of the pivot order that do not affect the number of entries in L to be identified and allows fast algorithms to be used in determining the exact structure of L .

A *supernode* is a set of contiguous columns of L with the same sparsity structure below a dense (or nearly dense) triangular submatrix. This trapezoidal matrix has zero rows corresponding to variables that are eliminated later in the pivot sequence at supernodes that are not ancestors in the elimination tree. This matrix can be compressed by holding only the nonzero rows, each with an index held in an integer. The condensed version of the elimination tree consisting of supernodes is referred to as the *assembly tree*. Supernodes can be exploited to facilitate the use of efficient dense linear algebra kernels and, in particular, BLAS kernels. These can offer such a large performance increase that it is often advantageous to amalgamate supernodes that have similar (but not exactly the same) nonzero patterns, despite this increasing the fill in L and the operation count. For convenience, throughout the remainder of this paper, we use the term node when referring to a supernode.

2.3 The factorize phase

The factorization of A proceeds using a succession of assembly operations into small dense frontal matrices, interleaved with the partial factorizations of these matrices. The dense frontal matrix consists of those rows of the associated node that are non-zero and the matching set of columns. The aim is to ensure that all columns of the node are *fully summed*, that is all the contributions from A and from any child nodes have been summed (assembled). Those columns of the frontal matrix that do not belong to the node (and are there from symmetry with the non-zero rows) are said to be *non-fully summed* and the assembly aims to sum any contributions from child nodes therein. Once the assembly is complete, a partial factorization of the frontal matrix is performed (that is, pivots are chosen from the fully-summed columns and then eliminations performed). The computed columns of L and D are not needed again until the solve phase (see below) and so can be stored while the rest of the frontal matrix corresponding to the non-fully-summed rows and columns (termed the *generated element* or *contribution block*), together with a list of the variables involved, is stored separately.

At each node of the assembly tree, the $m \times m$ frontal matrix can be expressed in the form

$$\mathcal{F} = \begin{pmatrix} \mathcal{F}_1 & \mathcal{F}_2^T \\ \mathcal{F}_2 & \mathcal{F}_S \end{pmatrix}, \quad (1)$$

where \mathcal{F}_1 and \mathcal{F}_2 are *fully summed* while \mathcal{F}_S is the partially summed block. Pivots are chosen from \mathcal{F}_1 . If \mathcal{F}_1 has order $p \leq m$ and $q \leq p$ pivots are chosen, the partial factorization of \mathcal{F} takes the form

$$\mathcal{F} = P \begin{pmatrix} L_1 & 0 \\ L_2 & I \end{pmatrix} \begin{pmatrix} D_1 & 0 \\ 0 & C \end{pmatrix} \begin{pmatrix} L_1^T & L_2^T \\ 0 & I \end{pmatrix} P^T, \quad (2)$$

where D_1 is a block diagonal matrix of order q , P is a permutation matrix and C is the generated element.

For symmetric positive-definite matrices, the pivot order chosen before the factorization commences can be used without modification. Moreover, the data structures determined by the analyse phase can be static throughout the factorization. For indefinite problems, using the supplied pivot order may be unstable or impossible because of (near) zero diagonal entries and, in general, it must be modified to maintain numerical stability. If symmetry is not to be destroyed, both 1×1 and 2×2 pivots are needed. Ashcraft, Grimes and Lewis [3] showed that bounding the size of the entries of L , together with a backward stable scheme for solving 2×2 linear systems, suffices to give backward stability for the entire solution process. They found that the widely used strategy of Bunch and Kaufman [4] does not have this property whereas the threshold pivoting technique first used by Duff and Reid [8] in their original multifrontal solver does.

Duff and Reid choose the pivots one-by-one, with the aim of limiting the size of the entries $l_{i,j}$ in L :

$$|l_{i,j}| < u^{-1}, \quad (3)$$

where the threshold u lies in the range $0 \leq u \leq 1.0$. At a given stage of the factorization with frontal matrix \mathcal{F} given by (1), let k denote the number of rows and columns of the block diagonal matrix D_1 found so far from \mathcal{F} and let $f_{i,j}$, with $i > k$ and $j > k$, denote an entry of \mathcal{F} after it has been updated by all the permutations and pivot operations so far. For a 1×1 pivot in column $j = k + 1$, the requirement for inequality (3) corresponds to the stability threshold test

$$|f_{k+1,k+1}| > u \max_{i>k+1} |f_{i,k+1}|. \quad (4)$$

In Duff et al. [7], the corresponding stability test for a 2×2 pivot involving columns $k + 1$ and $k + 2$ is

$$\left| \begin{pmatrix} f_{k+1,k+1} & f_{k+1,k+2} \\ f_{k+1,k+2} & f_{k+2,k+2} \end{pmatrix}^{-1} \right| \begin{pmatrix} \max_{i>k+2} |f_{i,k+1}| \\ \max_{i>k+2} |f_{i,k+2}| \end{pmatrix} < \begin{pmatrix} u^{-1} \\ u^{-1} \end{pmatrix}, \quad (5)$$

where the absolute value notation for a matrix refers to the matrix of corresponding absolute values. The choice of u controls the balance between stability and sparsity: in general, the smaller u is, the sparser L will be while the larger u is, the more stable the factorization will be. Typically, a value of 0.01 is recommended.


An alternative pivoting test is used by Kim and Eijkhout [23] that explicitly checks the condition (3). This is slightly weaker than the test (5) but is useful because it does not require that the largest off-diagonal entries in the candidate pivot columns are found before the pivot is applied. Observe, however, that the main drawback of the Kim and Eijkhout approach is that if (3) is not satisfied, some mechanism must be in place to undo the pivot application. We use this *a posteriori pivoting* in our GPU solver; this is discussed further in Section 3.2.

If some of the pivot candidates at a node fail the stability test, they are passed as part of the generated element up the assembly tree to the parent. This means that the data structure at the parent must be modified to accommodate the extra rows and columns. If there are a large number of delayed pivots, this can have serious consequences in terms of time as well as the memory and flops required for the factorization. Thus it is important to choose an ordering and scaling that aims to limit the number of delayed pivots [19].

2.4 The solve phase

The solve phase iterates over the assembly tree, first solving $Ly = b$ (forward) by iterating from the leaf nodes toward the root, then solving backwards $L^T x = y$ (backward) by iterating from the root to the leaf nodes. In the forward solve, at each node a triangular solve is performed with the diagonal block of the factors,

Figure 1: Plot of the front size distribution for matrix GHS_indef/ncvxqp3. The order of the matrix is 75000.



followed by a matrix-vector multiplication with the remainder. In the backward solve the order is reversed and transpose operations applied. Multiple right-hand sides can be solved simultaneously by replacing the vectors x, y, b with matrices X, Y, Z , transforming the BLAS2 operations into more efficient BLAS3 ones.

2.5 Tree and node parallelism

Existing works on parallelism in the multifrontal case exploit two distinct sources of parallelism. Tree parallelism exploits the independent nature of different tree branches, which can be worked on simultaneously. Node parallelism refers to exploiting parallelism within the dense linear algebra operation at each node.

Typical problems exhibit a large number of small leaf nodes, reducing to a small number of large nodes near the root of the tree. This leads to a strong synergy between these two types of parallelism. This is illustrated in Figure 1 for the sparse symmetric indefinite matrix GHS_indef/ncvxqp3¹. There is little benefit in exploiting node parallelism on the small leaf nodes, but lots of tree parallelism. Near the root there is limited scope for tree parallelism, but the larger nodes provide greater benefits from exploiting node parallelism.

2.6 Bit compatibility

For sequential solvers, achieving bit compatibility (in the sense that two runs on the same machine with identical input data should produce identical output) is not a problem. But enforcing bit compatibility can limit dynamic parallelism and, when designing a parallel sparse direct solver, the goal of efficiency potentially conflicts with that of bit compatibility.

The critical issue is the way in which N numbers (or, more generally, matrices) are assembled

$$sum = \sum_{j=1}^N S_j,$$

¹All the sparse matrices used in this paper are taken from the University of Florida Sparse Matrix Collection [5].

where the S_j are computed using one or more processors. The assembly is commutative but, because of the potential rounding of the intermediate results, is not associative so that the result *sum* depends on the order in which the S_j are assembled. A straightforward approach to achieving bit compatibility is to enforce a defined order on each assembly operation, independent of the number of processors. As discussed in [17], this is a viable approach for multifrontal codes and is used by the multicore multifrontal solver **HSL_MA97**; the same approach is followed by **SSIDS** (but the two codes are not bit-compatible with respect to each other).

3 Design of SSIDS

Our new GPU sparse indefinite solver **SSIDS** implements a supernodal multifrontal algorithm, as outlined above. In this section, we look at some of details that are designed to optimize performance of the code on a GPU.

Achieving high performance on the GPU requires that sufficient threads are executing to either saturate memory bandwidth (in the case of the memory-bound assembly operations and the entire solve phase) or floating-point capacity (in the case of compute-bound computations during the factorization). The CUDA API provides two ways of launching a significant number of threads. The most straightforward is to launch a large number of threads all running the same kernel code. The second is to run multiple different kernels in parallel. The mechanism for doing so is to create different *streams* of work. The CUDA API guarantees that all kernels launched within the same stream execute in serial (that is, one finishes before the next begins), but allows kernels from different streams to run in parallel. However, at most 16 streams (fewer on older devices) can issue instructions at one time.

Towards the leaves of the assembly tree, with the larger nodes at or near the root(s) of the tree. Given the large number of these very small nodes, merely using streams will be insufficient to ensure that the majority of the GPU is kept busy. Instead, we manually handle the parallelism through the use of an array that describes the parameters of each front and the operations that need to be performed upon it. A kernel with a sufficient number of CUDA thread blocks to cover the available work is launched, and the first action of each thread block is to extract the parameters of its associated task from this data structure. This mechanism allows us to fully exploit both tree- and node-level parallelism.


It would be attractive to have a single kernel that was able to perform any of the tasks required for the factorization. This would facilitate good load balance and minimize the occasions on which the GPU is not fully occupied. However, doing this would have two drawbacks: firstly, the resource requirements of such a multi-kernel would be determined by the worst sub-kernel it contained, severely limiting occupancy; secondly, the code would be so complex that it would overflow the small number of registers available on older Fermi cards, leading to register spill and associated slowdown. For these reasons, each kernel we discuss is implemented and launched independently.

Using our manual approach, synchronization is (normally) achieved through the launch of kernels within the same stream. Where a dependency exists between two tasks, this is enforced by allocating the later operation to a future kernel launch. However, in a limited number of cases, we implement synchronization through GPU global memory and spin-locks (see, for example, the assembly operations described below).

To simplify task generation and memory management, task lists are generated through a level-based approach. Nodes of the assembly tree are allocated to a level depending on their distance from the root (this is illustrated in Figure 2 where the root node is 5). As the generated elements from the child nodes can be discarded once the frontal matrix at the parent node has been assembled, this can be done using only two arrays: one for the even-numbered levels and one for the odd. At each level, one of these arrays holds the generated elements from the child nodes whilst the other holds those generated at this level, swapping roles at the next level.

We split the work to be performed at each node into three distinct operations: assembly into the frontal matrix \mathcal{F} , dense factorization of the fully-summed columns of \mathcal{F} , and formation of the generated element. Each of these operations is itself composed of a number of tasks. In the following sections, we explain how the operations for a single node are handled although, as described above, all the nodes that belong to the same level of the assembly tree are treated simultaneously. To avoid confusion, we refer to *blocks* of CUDA

Figure 2: Allocation of nodes to levels



threads that specify a unit of work that is allocated to a SM, and *tiles* of a matrix (often referred to as blocks in the literature on sparse direct solvers).

We note that the CUBLAS only support the simultaneous execution of multiple small operations of different sizes through the use of multiple streams. However, tests showed that using the specialized kernels we developed is faster and more straightforward, particularly on the older Fermi hardware. CUBLAS are however used when only a single node is being operated on (e.g. at the root).

3.1 Sparse assembly

At each node, entries from A , generated elements from child nodes and any delayed pivots must be assembled. Separate kernels handle each of these.

During the analysis phase, a mapping array from positions in A to positions in L is constructed. This makes the assembly of A into the frontal matrix very straightforward, with the only complication being that, if there are delayed pivots inherited from its child nodes, the dimension of the front is greater than predicted in the analyse phase. In this case, the mapping array is used to calculate the row and column for the insertion, rather than providing the insert location directly.

The kernel that performs the assembly from the child nodes considers one parent-child assembly at a time. Each generated element from a child is divided into a number of tiles, and a thread block launched to assemble each tile into the parent frontal matrix using a simple mapping array to determine the destination row and column of each entry.

To preserve bit-compatibility, we must ensure that the child entries are always assembled in the same order. This could be achieved by using a separate kernel launch to handle each child, however doing so reduces parallel efficiency and introduces additional kernel launch overheads. Instead, we use global memory to signal once a child has completed so the next child may begin (each block atomically increments a counter once to indicate when it has finished).

This approach has two inefficiencies and to discuss these we introduce the concept of potential false dependencies. Suppose a node has nc child nodes. A *potential false dependency* is said to occur when the number of child nodes that contribute to a given entry of the parent frontal matrix is at least one but is fewer than nc . The first inefficiency is that each entry of the parent frontal matrix may be accessed multiple times (once per child). Secondly, false dependencies are often introduced as the number of child nodes that contribute to a particular entry is typically zero or one. An alternative that could overcome these issues would be to use one CUDA block per tile of the parent frontal matrix rather than per tile of the child node generated element. This would have the further advantage of eliminating the need for synchronizations through global memory. However, it may be that substantially more CUDA blocks are required as some entries in the parent frontal matrix may not receive contributions from any of the child nodes.

Table 1 explores the relative cost of the two approaches on a subset of test matrices used in Section 4.

Table 1: A summary of the number of entries of the parent frontal matrices that have contributions from 1 or more child nodes. Statistics in columns 3 to 6 are given as a percentage of the total number of entries summed over all the parent frontal matrices reported in column 2. $nchild = 0$ (respectively, $nchild = 1$) is the percentage of the entries that 0 children (respectively, 1 child) contribute to. Potential false dependencies is the percentage of entries that $0 < nchild < nc$ child nodes contribute to, where nc is the number of children of the parent.

Problem	total entries	$nchild = 0$ %	$nchild = 1$ %	$nchild > 1$ %	potential false dependencies %
2. Andrianov/mipl	6.27×10^8	3.1	96.5	0.4	0.5
7. GHS_indef/ncvxqp3	8.83×10^8	5.2	93.7	1.0	60.1
8. Oberwolfach/gas_sensor	1.97×10^8	25.6	70.5	3.9	49.2
14. Rothberg/cfd2	2.64×10^8	30.3	65.1	4.6	57.7
15. DNVS/thread	1.70×10^8	28.9	66.0	5.1	58.0
22. AMD/G3_circuit	8.32×10^8	29.5	67.3	3.2	54.2
23. GHS_psdef/bmwcrs_1	5.20×10^8	27.8	66.9	5.3	58.4
29. ND/nd6k	2.01×10^9	4.4	94.3	1.2	8.5
30. Schenk_IBMNA/c-big	3.02×10^9	4.3	94.7	1.0	89.9
32. PARSEC/Si5H12	4.78×10^9	2.7	97.3	0.9	5.7

This shows that only a small number of entries have contributions from more than one child (typically fewer than 5 per cent), so the additional memory bandwidth that is potentially needed by the CUDA block-per-child-tile approach is minimal (it may not need more memory than the the CUDA block-per-parent-tile approach because the latter uses more complex data structures). While the last column headed “potential false dependencies” represents a significant proportion of entries (often more than half), these figures are an overestimate because of the way they are calculated. Under any given assembly order, some of the false dependencies will not occur as the child nodes that do contribute are ordered correctly. Trying to avoid false dependencies is likely to offer only limited speedup as warps only run at the speed of their slowest member. Finally, the $nchild = 0$ column shows the number of excess threads that are launched but have no work to perform under the CUDA block-per-parent-tile approach. Given this analysis, we do not feel that the CUDA block-per-parent-tile approach offers sufficient gains across the range of matrices to offset its increased computational complexity.

A final kernel handles the assembly of entries corresponding to delayed pivots. Assuming the problem has been well-scaled and well ordered, we expect relatively few delays (or at least relatively few children to contribute them) [19]. Thus, a single CUDA block is used to handle all delays contributing to a parent node from any of its child nodes. The same lookup arrays used in the assembly of the generated elements are reused to determine the correct insert locations for entries in delayed columns (delayed rows can be assembled without any permutation).


3.2 Dense factorization of fully assembled columns

We seek to factorize an $m \times p$ submatrix by splitting it into a number of tiles that are each handled by an associated block. Factorization proceeds one column of tiles at a time. We refer to a column of tiles as a *tile column*; it is illustrated in Figure 3 for a tile column comprised of k columns. We note that while F_{11} is always square, the off-diagonal tiles of F_{21} are in fact rectangular. Ignoring pivoting, the factorization proceeds as follows:

1. Factor $F_{11} = L_{11}DL_{11}^T$;
2. Apply factor L_{11} to the off-diagonal part of the tile column, $L_{21} = F_{21}L_{11}^{-T}$; and
3. Update the remaining tile columns to the right $F_s \leftarrow F_s - L_{21}DL_{21}^T$.

Numerical stability requires that we bound entries of L . As discussed in Section 2.3, traditionally this is done by considering the below diagonal entries of largest absolute value in the candidate column(s) (as

Figure 3: $m \times k$ tile column within the fully-summed part of a frontal matrix. Here F_s denotes the fully-summed columns lying to the right of the tile column.



in (4) and (5)). This requires the whole of the tile column to be kept up-to-date as the factorization of F_{11} progresses. On the GPU this is suboptimal as communication between the blocks handling each tile is slow. Further, it may not be possible at all as there is no guarantee that all blocks handling the tile column are run simultaneously (for large matrices there may be insufficient resources). Instead, we use a posteriori pivoting in which F_{11} is factored without reference to F_{21} , and the size of the entries in L_{21} are checked after the application of L_{11} . Suppose an entry in column $r \leq k$ of L_{21} exceeds the threshold tolerance u^{-1} in absolute value. Columns $r, r+1, \dots, k$ in the tile column are marked as failed and are restored to their state before application of the r -th pivot. Our implementation uses an `atomicMin()` operation to determine the number of successful columns (we expect this in general to be equal to k). We emphasize that this strategy is to ensure backward stability and is very different from the approach taken by, for example, the PARDISO solver [30], in which F_{11} is factored without reference to F_{21} but **no** check is made subsequently on the size of the factor entries (and thus no guarantee of stability).

Failed candidate pivot columns can be managed with varying degrees of sophistication. One option is to delay them to the parent node, as is done in the CPU code of Kim and Eijkhout [23]. However, this may lead to significant increases in both the memory usage and the number of flops. Instead, we adopt a simple scheme whereby failed columns are excluded from consideration until they have been updated by at least one other column, at which stage they can be reconsidered. Once all candidate columns in the front have been considered for pivoting, any that have failed are delayed to the parent node.


This scheme can break down at the root node(s), as it only allows partners for 2×2 pivots to be selected from a band of width k around the diagonal. In some (rare) cases this will be insufficient. We cater for these numerically problematic matrices by implementing a special-case kernel that handles any failed pivots that remain at the end of the root node but, as this will not normally be required, we have not invested significant efforts in its efficiency.

3.2.1 Implementation details

We split the dense factorization of fully-summed columns into three different kernels that are run repeatedly until all the fully-summed columns have either been factorized or flagged as delays.

- **factor()**: each block loads its assigned tile(s) $F_{21,j}$ and the diagonal block F_{11} into shared memory. Each block computes $F_{11} = L_{11}D_1L_{11}^T$ and applies L_{11} to its tiles(s) $L_{21,j} = F_{21,j}L_{11}^{-T}$. The threshold condition is checked for each column of $L_{21,j}$ and the number of successful columns is recorded using `atomicMin()`. L_{21} and $(L_{21}D_1)$ are stored in temporary global memory. The first thread block

Figure 4: Frontal matrix data structure after `reorder()`. In order from the left, the columns are (i) those that have been successfully eliminated; (ii) those that have failed (but may be retested later); (iii) the untested candidate pivots, of which the first k will be tested next in this example and the remainder will be tested later; and (iv) the non-fully-summed columns that cannot be eliminated at this node.



associated with each tile column also stores L_{11} , D_1 and the pivot sequence used in the factorization of F_{11} . Once this step is complete, the number of successful pivots is known.

- **reorder()**: a row permutation is applied to any computed columns of L in the frontal matrix that are to the left of the current tile column, and a column permutation is applied so that successful columns are moved to become the leading columns of the frontal matrix. Any failed columns are swapped as necessary with successful columns. This leaves the frontal data structure as shown in Figure 4, with the eliminated columns first, then any failed pivots, followed by any fully-summed columns in tile columns that have not yet been pivoted on. The non-fully-summed part is shown only for illustration purposes, and is not touched by this kernel (it will form the generated element). The temporary array that contained L_{21} can now be discarded. The row and column permutations are applied to the temporary array containing $(L_{21}D_{11})$.
- **update()**: the partial outer product $L_{21}D_{11}L_{21}^T$ that updates the remaining fully-summed columns F_s is formed by performing a matrix multiplication between L_{21} (stored in the factor array) and $(L_{21}D_1)$ (stored in the temporary memory).


When factorizing F_{11} , we say that a column has been *processed* if it has been selected as a pivot column; otherwise it is *non-processed*. A 2×2 pivot (i, j) is chosen so long as

$$\Theta \frac{|a_{ii}a_{jj} - a_{ij}a_{ji}|}{|a_{ii}| + |a_{jj}| + |a_{ij}|} \geq \max(|a_{ii}|, |a_{jj}|) \quad (6)$$

where Θ is the 2×2 pivot bias. A bias towards 2×2 pivots is desirable as it prevents the need for a second pivot test and allows updates to be performed more efficiently. Through experiment, we have chosen a value $\Theta = 100$. The reader might be curious as to the origin of the expression $|a_{ii}| + |a_{jj}| + |a_{ij}|$ used as the denominator: this is a cheap approximation to $\max(|a_{ii}|, |a_{jj}|, |a_{ij}|)$ sufficient to detect scaling issues (it is within a factor 3 of the correct value, which is small compared to Θ). If the test (6) is not met, then a 1×1 pivot is chosen as the larger (in magnitude) of a_{ii} or a_{jj} . The factorization of F_{11} is performed as follows:

1. Initialize the set of processed columns $\mathcal{P} = \phi$.
2. While there are non-processed columns do:

Figure 5: Calculating the outer product for a tile of the contribution block



- Let i be the leftmost non-processed column. Find $j = \max\{|f_{j,i}| : j \notin \mathcal{P}\}$ (break ties in favour of smallest j).
- Check if all entries in column are $< 10^{-20}\alpha$, where α is a bound on the maximum entry in the node matrix before the start of the partial factorization. If so, record zero pivot.
- Select the best pivot of (i, j) , (i) , or (j) (test (6))
- Update \mathcal{P} . Perform the eliminations with the selected pivot column(s).

We note that the processed columns are not permuted to the leading columns of F_{11} . Instead, the array representing \mathcal{P} is used to mask any operations performed.

3.3 Formation of the generated element

The majority of the floating-point operations typically derive from the formation of the generated element $C \leftarrow C - L_2 D_1 L_2^T$ (recall (2)). We observe that this operation cannot be handled by any one standard BLAS operation: instead it is implemented using a custom kernel. The matrices L_2 and $(L_2 D_1)$ are available from the factorization of the fully-summed columns. C is divided into a number of tiles and a thread block launched for each. The entries of C are stored in registers for the duration of the computation. For a given $k_1 \times k_1$ tile of C , the relevant parts of L_2 and $(L_2 D_1)$ are divided into a number of $k_1 \times k_2$ tiles that must be looped over, as shown in Figure 5. Each tile is loaded into shared memory before each thread calculates the portion of matrix-matrix multiplication relevant to the entries of C it is responsible for.

Double buffering is used for performance optimization, allowing the next $k_1 \times k_2$ tiles of L_2 and $(L_2 D_1)$ to load whilst the current values are used for the matrix-matrix multiplication.

3.4 The positive-definite case

In addition to the pivoted kernel for indefinite systems, our code offers an unpivoted Cholesky (LL^T) factorization suitable for positive-definite systems. This is implemented using a drop-in replacement for the dense factorization kernels and a modification to the generated element calculation to allow for the absence of D .

This is provided for convenience, but does not fully exploit properties of the Cholesky factorization that allow considerably more scheduling flexibility and avoid the need to cater for the possibility of pivoting permutations. As such, it is unlikely to be as efficient as a dedicated Cholesky solver.

Table 2: Hardware summary

	CPU		GPU	
	E5620 Westmere-EP	E5-2687W Sandy Bridge-EP	C2050 Fermi	K20 Kepler
Clock speed	2.4 GHz	3.1 GHz	1.15 GHz	0.7 GHz
Cores	4	8	448	2496
Theoretical DP Peak	38.4 GFlop/s	199 GFlop/s	515 GFlop/s	1170 GFlop/s
Achieved dgemm Peak	34.1 GFlop/s	181 GFlop/s	298 GFlop/s	1046 GFlop/s
Max TDP	80 W	150 W	238 W	225 W
Vector length	128 bit	256 bit	n/a	n/a
Memory bandwidth	25.6 GB/s	51.2 GB/s	144 GB/s	208 GB/s
GPU Memory	n/a	n/a	3 GB	5 GB
Launch date	Q1 2010	Q1 2012	Q2 2010	Q4 2012
Launch RRP	\$391	\$1885	\$2499	\$3199

3.5 Implementation of the solve phase

The implementation of the solve phase mirrors that of the factorize phase. Tree parallelism is exploited through the allocation of nodes to levels, and all nodes in a level are addressed by the same kernel launch.

Two different variants of the solve phase have been implemented for different circumstances. The first follows the traditional triangular solve methodology. The second adds additional work to the factorize phase where the explicit inverse of the diagonal blocks is calculated (this can be done stably as they are triangular [13]). Then, in the solve phase, at each node the dense triangular solve is replaced with a matrix-vector multiply. The latter, which we refer to as the *presolve* approach, is more efficient if the additional factorize cost can be amortized across a number of right-hand sides right-hand sides. This is discussed further in Section 4.4.

4 Numerical experiments

Results are given for the hardware summarised in Table 2. We note that CPU results are given for a two-socket system, so the performance statistics for the machines used are double those stated in the table for a single chip. Whilst we compare CPU and GPU systems from the same eras, we note that this may be somewhat misleading due to the disparity in the systems in terms of both cost and power envelope.

Table 3 introduces the problems we use for our experiments. The number of entries in L and number of flops are for a factorization without any numerical pivoting using an ordering returned by the nested dissection ordering package METIS v4 [21]. All problems are drawn from the UFL Sparse Matrix Collection [5]. In each test, the right-hand size was computed so that the exact solution was the vector of all ones.

Unless otherwise indicated, problems are run with the default settings of **HSLMA97** (CPU) (using hyperthreading) and **SSIDS** (GPU). All times are in seconds. Timings on the CPU are taken as the average of 10 runs due to variability (the GPU did not show any measurable variation so a single run was performed).

4.1 Headline results

Figures 6 and 7 show the time and speed of the factorize phase respectively on both CPU and GPU platforms. For these results, all problems are treated as indefinite with a scaling used if it improves the factorization time. In each figure, the first graph shows the performance on the 2010/11 hardware, and the latter on the 2012/13 hardware. Problems 31 and 35 could not be run on the C2050 due to memory constraints, and so are omitted where appropriate. The CPU execution times for problem 35 are not clearly shown on the graphs as they are large in comparison with the times for the other examples: they were 35.5s and 16.2s, respectively, on the Westmere-EP and Sandy Bridge-EP platforms. The results show that significant performance benefits

Table 3: Test problems: n is the dimension of A , $nz(A)$ is the number of non-zero entries in A , $nz(L)$ and Flops give the number of entries in and flops to compute the factors respectively, assuming no pivoting occurs. A \dagger indicates that a problem is positive definite.

	Problem	n	$nz(A)$	$nz(L)$	Flops	Description
1.	Newman/astro-ph	16.7	0.121	3.73	4.08	Collaboration network
2.	Andrianov/mip1	66.5	5.21	11.0	5.24	Optimization
3.	PARSEC/SiNa	5.74	0.102	5.10	6.70	Density functional theory
4.	Schmid/thermal2	1230	4.90	63.0	15.1	Unstructured thermal FEM
5.	McRae/ecology1	1000	3.00	46.9	15.6	Electrical network theory
6.	INPRO/msdoor	416.	10.3	57.1	18.2	Structural problem: medium door
7.	GHS_indef/ncvxqp3	75.0	0.275	19.0	20.7	Nonconvex QP problem
8.	Oberwolfach/gas_sensor	66.9	0.885	24.7	21.3	Thermal model single gas sensor device
9.	ND/nd3k	9.00	1.64	13.0	22.3	3D mesh problem
10.	Boeing/pwtk	218.	5.93	50.8	22.9	Pressurised wind tunnel
11.	GHS_indef/c-71	76.6	0.468	17.2	25.4	Non-linear optimization
12.	BenElechi/BenElechi1	246.	6.70	55.9	27.2	Unknown
13.	GHS_psdef/crankseg_1 \dagger	52.8	5.33	34.0	32.5	Linear static analysis
14.	Rothberg/cfd2 \dagger	123.	1.61	40.0	33.0	CFD pressure matrix
15.	DNVS/thread \dagger	29.7	2.25	24.4	35.0	Threaded connector
16.	DNVS/shipsec1 \dagger	141.	3.98	40.5	38.3	Ship section
17.	DNVS/shipsec8 \dagger	115.	3.38	37.2	38.6	Ship section
18.	Oberwolfach/boneS01 \dagger	127.	3.42	42.1	47.0	Bone micro-FEM
19.	GHS_psdef/crankseg_2 \dagger	63.8	7.11	44.6	47.0	Linear static analysis
20.	Schenk_AFE/af_shell7 \dagger	505.	9.05	99.0	52.9	Sheet metal forming
21.	DNVS/shipsec5 \dagger	180	5.15	55.3	57.7	Ship section
22.	AMD/G3_circuit \dagger	1590	4.62	119.	58.7	Circuit simulation
23.	GHS_psdef/bmwcr1 \dagger	149.	5.40	71.8	61.5	Automotive crankshaft
24.	Schenk_AFE/af_0_k101 \dagger	504.	9.03	104.	61.6	Sheet metal forming
25.	GHS_psdef/lldoor \dagger	952.	23.7	155.	79.9	Structural problem: large door
26.	DNVS/ship_003 \dagger	122.	4.10	62.0	81.9	Ship structure
27.	PARSEC/Si10H16	17.1	0.447	31.3	87.3	Density functional theory
28.	Um/offshore \dagger	260	2.25	88.4	106.	Electromagnetics
29.	ND/nd6k \dagger	18.0	3.46	40.0	111.	3D mesh problem
30.	Schenk_IBMNA/c-big	345.	1.34	52.0	115.	Non-linear optimization
31.	GHS_psdef/inline_1 \dagger	504.	18.7	180.	146.	Inline skate
32.	PARSEC/Si5H12	19.9	0.379	45.0	154.	Density functional theory
33.	GHS_psdef/apache2 \dagger	715.	2.77	149.	176.	3D structural problem
34.	Lin/Lin	256.	1.01	114.	279.	Eigenvalue problem
35.	ND/nd12k \dagger	36.0	7.13	117.	506.	3D mesh problem

Table 4: Factorize phase times when run with the positive definite and the indefinite kernels. A - indicates missing data (insufficient memory).

Problem	CPU (E5620)			GPU (C2050)		
	indef	posdef	% cost	indef	posdef	% cost
13. GHS_psdef/crankseg_1	1.310	1.175	11.4	0.506	0.450	12.4
14. Rothberg/cfd2	1.192	1.078	10.6	0.504	0.450	11.9
15. DNVs/thread	1.598	1.406	13.7	0.463	0.419	10.4
16. DNVs/shipsec1	1.426	1.294	10.3	0.567	0.502	12.9
17. DNVs/shipsec8	1.571	1.425	10.2	0.557	0.498	11.8
18. Oberwolfach/boneS01	1.569	1.440	9.0	0.614	0.557	10.2
19. GHS_psdef/crankseg_2	1.676	1.492	12.3	0.681	0.630	8.1
20. Schenk_AFE/af_shell7	2.025	1.801	12.4	0.992	0.886	12.0
21. DNVs/shipsec5	2.967	2.701	9.8	0.830	0.754	10.1
22. AMD/G3_circuit	2.627	2.524	4.1	1.662	1.375	20.8
23. GHS_psdef/bmwcrs_1	1.928	1.781	8.2	0.865	0.796	8.7
24. Schenk_AFE/af_0_k101	2.406	2.133	12.8	1.087	0.970	12.0
25. GHS_psdef/lldoor	2.959	2.683	10.3	1.571	1.387	13.3
26. DNVs/ship_003	3.004	2.712	10.7	1.022	0.940	8.7
28. Um/offshore	3.500	3.313	5.6	1.351	1.238	9.1
29. ND/nd6k	6.332	6.443	-1.7	1.328	1.251	6.2
31. GHS_psdef/inline_1	4.426	4.122	7.4	-	-	-
33. GHS_psdef/apache2	5.433	5.194	4.6	2.328	2.097	11.0
35. ND/nd12k	30.089	29.652	1.5	-	-	-

are gained through the use of the GPU hardware, increasing as the problem size becomes larger. The large disparities between the performance behaviour of the CPU and GPU codes on particular problems are down to the significantly different parallel designs of the two codes.

The maximum speedup between the Westmere-EP and Fermi hardware was $6.7\times$ on problem 32, with most problems achieving at least a $2\times$ speedup. The maximum speedup between the Sandy Bridge-EP and Kepler hardware was $7.0\times$ on problem 32, with 13 problems achieving at least a $2\times$ speedup.

4.2 Cost of pivoting

To estimate the cost of incorporating pivoting into our multifrontal algorithm, in Table 4 we compare running the positive-definite test matrices using both the positive-definite and indefinite modes of SSIDS. Here the reported % cost is the overhead of using the indefinite mode. The results only provide an estimate because, as well as the extra cost of pivoting, there are some additional costs associated with using an LDL^T rather than LL^T factorization in the indefinite runs.

The typical overhead on the GPU is around 12%, although a few problems exceed this. This is comparable to the overheads on the CPU and confirms that the pivoting strategy has been implemented efficiently on the GPU.

4.3 Factorization performance

Table 5 gives a detailed breakdown of times and performance metrics between different parts of the factorize phase on the C2050. These results were captured with the profiling tool `nvprof`, and only measure time spent in kernels. The “o” column represents the difference between the sum of these parts and the measured time for the entire factorize phase. This includes such things as memory management, transfer and synchronization. As such it is largely attributable to the dense factorization where information is bounced between the CPU and GPU.

Figure 6: Factorize Time





Figure 7: GFlop/s



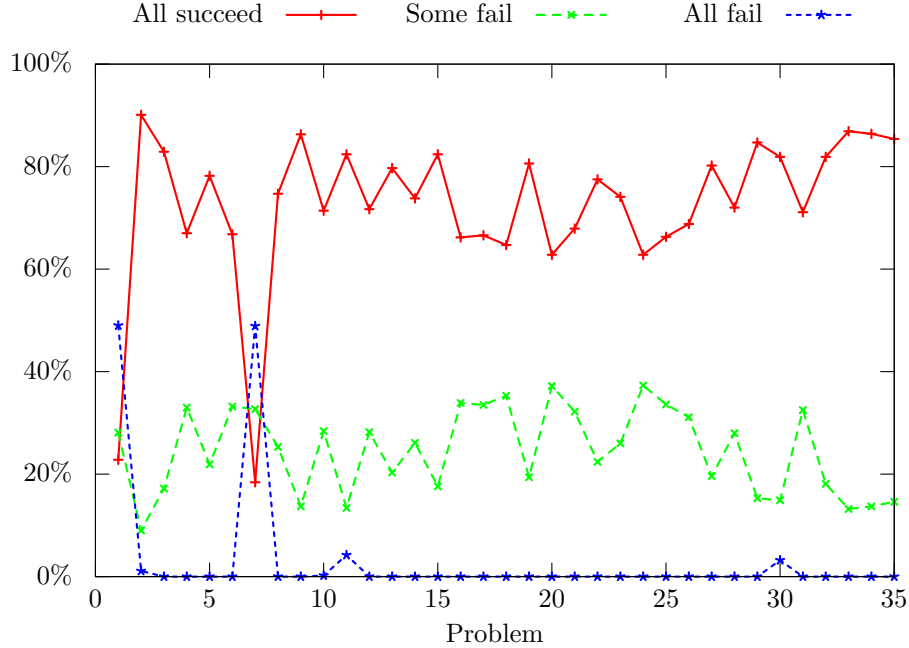


Figure 8: Percentage of candidate pivot columns accepting a given number of pivots (candidate pivot columns can have at most 8 pivots).



The performance of the contribution and assembly kernels is clearly a major factor in dominating the CPU times. By capturing the node sizes involved, we validated the performance of our contribution kernel by comparing it against a simulation using the CUBLAS `dsyrk` routine to perform similar operations, both with and without the use of multiple streams to run multiple matrices in parallel. In all cases, our code was faster, and in the majority of cases, by at least 50%. The assembly phase data rate is calculated based on the number of entries in the child matrices on the assumption that each entry requires two floating-point loads, one integer load and a floating-point store. The achieved data transfer rates are still respectable, especially due to the sparse nature of some loads (where the entire cache line is loaded but only a single value is used).

To assess the efficacy of the *a posteriori* pivoting strategy, Figure 8 shows the percentages of pivots accepted in each column. We see that for the majority of matrices, somewhere between 60% and 80% of pivot columns succeed entirely, with the remainder partially succeeding. For those matrices with a large number of delayed pivots (1, 7, 11, and 30) we see a significant departure from this behaviour, as one might reasonably expect. Only about 20% of columns completely succeed, this is in part skewed by the re-testing of the same pivots until they have been delayed sufficiently that they can be safely eliminated. However, even in these cases the need to fall back on the special-case kernel that handles break down at root nodes is rare: in our test set, only problem 1 required it to resolve the final 31 pivots.

4.4 Solve

Figures 9 and 10 show (for a single right-hand side) the performance of the solve phase with and without using the presolve, compared with the CPU solve phase of `HSL_MA97`. Note that these figures show only the time for the solve phase, and not the overhead on the factorize phase required by the presolve.

These show, as expected, that the GPU is able to deliver a speedup of about $3\times$ on average over the CPU code, and significantly more using the presolve option. Those problems where the presolve significantly outperforms the traditional solve (4, 5, 6, 22, 25, 33) are characterised by a large number of small nodes. As the substantial difference between the standard and presolve algorithms is the replacement of the triangular

Table 5: Split of the operations and time between different routines within the factorize phase on the C2050, with performance metrics. Data captured using `nvprof`. 'f' represents the dense matrix factorization, 'a' the assembly, 'c' the calculation of the contribution block, and 'o' other operations not reported in the profiling data, including data transfer, memory management and synchronization. A - indicates missing data (insufficient memory).

Problem	% flops			%time				GFlop/s		GB/s
	f	a	c	f	a	c	o	f	c	
1. Newman/astro-ph	35.3	1.6	63.1	45.2	22.6	5.9	26.3	10.9	148.5	27.6
2. Andrianov/mip1	6.3	5.5	88.2	22.9	28.6	14.9	33.6	4.5	98.1	89.5
3. PARSEC/SiNa	30.3	2.0	67.8	30.8	20.6	20.8	27.8	37.3	124.0	102.3
4. Schmid/thermal2	35.4	0.6	64.0	49.2	7.3	9.5	34.0	12.3	115.0	39.0
5. McRae/ecology1	30.7	0.5	68.8	49.3	7.2	12.0	31.6	13.9	128.2	39.4
6. INPRO/msdoor	27.1	0.7	72.3	51.5	10.4	17.9	20.1	17.4	133.3	60.0
7. GHS_indef/ncvxqp3	23.3	1.3	75.3	44.5	26.7	10.3	18.6	12.6	177.0	33.6
8. Oberwolfach/gas_sensor	29.2	0.4	70.5	49.8	8.1	26.6	15.5	35.5	160.4	75.2
9. ND/nd3k	24.2	0.9	74.9	33.7	18.0	30.9	17.5	45.0	152.0	90.6
10. Boeing/pwtk	27.6	0.5	71.9	53.6	9.4	22.0	15.0	23.6	150.0	68.5
11. GHS_indef/c-71	9.7	1.3	89.0	27.6	20.6	29.6	22.2	17.1	146.2	84.9
12. BenElechi/BenElechi1	34.2	0.4	65.3	55.3	8.5	20.8	15.5	29.6	150.4	70.3
13. GHS_psdef/crankseg_1	32.3	0.3	67.4	51.5	8.0	26.7	13.8	40.6	163.3	66.6
14. Rothberg/cfd2	30.6	0.3	69.1	52.5	7.8	28.2	11.4	38.4	161.3	69.9
15. DNVs/thread	49.7	0.2	50.1	61.1	5.4	22.9	10.6	61.8	166.4	72.8
16. DNVs/shipsec1	30.5	0.3	69.2	50.5	7.2	28.6	13.7	41.1	164.8	77.9
17. DNVs/shipsec8	32.9	0.3	66.8	51.1	7.5	28.3	13.1	44.9	164.5	79.5
18. Oberwolfach/boneS01	27.4	0.2	72.3	48.0	7.0	33.8	11.2	44.1	164.8	73.2
19. GHS_psdef/crankseg_2	35.2	0.3	64.5	55.3	7.7	27.3	9.7	44.3	164.0	67.0
20. Schenk_AFE/af_shell7	33.5	0.4	66.1	55.1	8.6	23.8	12.6	32.7	149.9	69.0
21. DNVs/shipsec5	36.0	0.3	63.7	54.7	7.4	27.4	10.5	46.0	162.6	80.5
22. AMD/G3_circuit	22.3	0.5	77.2	50.8	9.2	19.0	21.0	15.7	145.3	56.9
23. GHS_psdef/bmwera_1	29.7	0.3	70.0	52.6	8.5	30.9	8.0	40.4	162.0	77.1
24. Schenk_AFE/af_0_k101	33.7	0.4	65.9	55.0	8.3	24.8	11.9	35.1	151.9	69.7
25. GHS_psdef/ldoor	27.8	0.4	71.8	52.0	9.5	24.6	13.8	27.5	149.6	64.7
26. DNVs/ship_003	36.8	0.3	62.8	52.5	8.2	30.5	8.8	56.6	166.0	90.4
27. PARSEC/Si10H16	21.4	1.5	77.1	23.1	24.8	39.5	12.6	63.8	134.5	116.2
28. Um/offshore	27.3	0.3	72.3	47.8	9.1	34.8	8.4	45.3	164.4	82.8
29. ND/nd6k	15.1	0.9	84.0	24.7	19.1	46.1	10.1	51.7	154.4	107.8
30. Schenk_IBMNA/c-big	14.3	1.2	84.5	27.6	21.1	34.5	16.8	29.9	141.7	95.7
31. GHS_psdef/inline_1	26.9	0.3	72.8	-	-	-	-	-	-	-
32. PARSEC/Si5H12	21.6	1.5	76.9	21.1	25.8	42.9	10.1	74.3	130.3	118.4
33. GHS_psdef/apache2	28.4	0.3	71.3	49.2	8.0	32.8	10.0	43.9	165.0	73.6
34. Lin/Lin	18.4	0.4	81.2	33.2	10.7	49.7	6.4	56.4	166.3	104.3
35. ND/nd12k	12.2	0.7	87.1	-	-	-	-	-	-	-

Figure 9: Times for the solve phase in seconds, E5620/C2050.

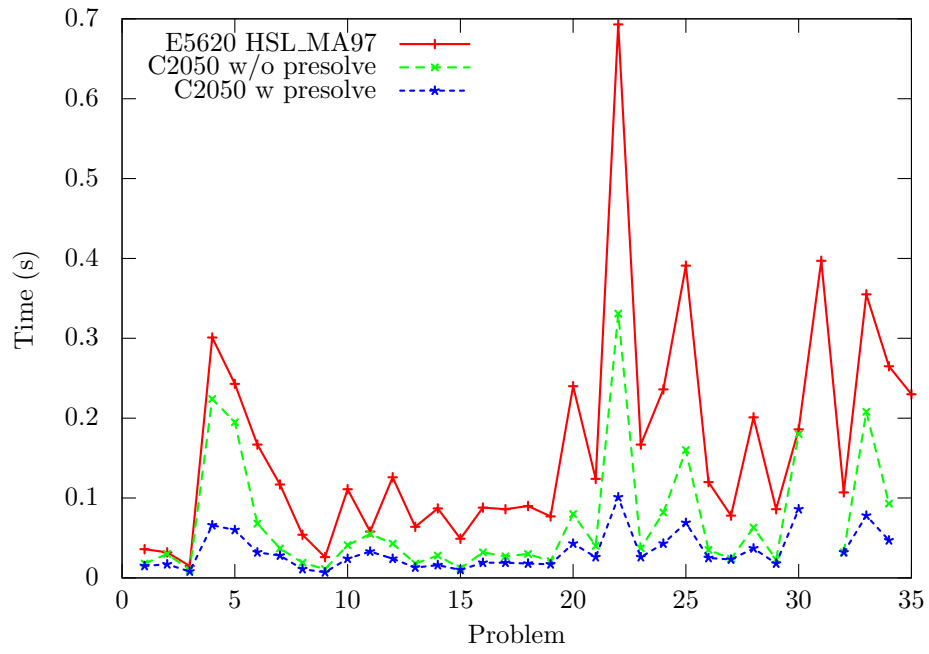
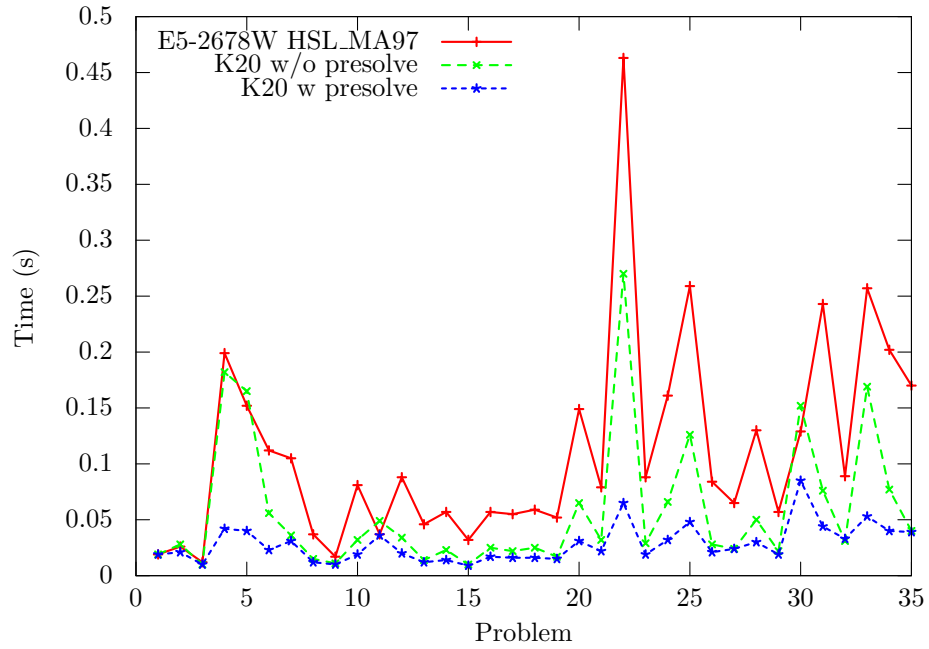


Figure 10: Times for the solve phase in seconds, E5-2687W/K20.



solve (`trsv`) with a matrix-vector multiply (`gemv`), it is unsurprising that profiling shows that the triangular solve kernel is the dominant cost for these problems. Both the `trsv` and `gemv` kernels are memory-bound and capable of achieving similar peak bandwidth on large problems. Hence it is only on problems with a significant number of small operations that we see the benefits of using the communication-avoiding presolve technique.

The additional cost of applying the presolve in the factorize phase can add as much as 40% (DNVS/thread) or as little as 7% (Newman/astro-ph) to the factorization time. On the best problems the extra cost is amortized over only 3 solves, but on average it requires 20-30 solves. On some problems there is no measurable advantage to using the presolve algorithm. Due to these performance overheads, the presolve algorithm must be explicitly enabled by the user.

It may be that a significant redesign of the triangular solve could dispatch large numbers of small operations more efficiently (see, for example, the dedicated small matrix work in [14]). However, the cost of the presolve in the factorize phase could be similarly reduced by limiting the inversion to a smaller band around the diagonal, rather than the entire diagonal block of each node.

5 Concluding remarks

In this paper, we have reported on our work to address the challenging problem of designing and developing an efficient and robust symmetric indefinite sparse direct solver for use on NVIDIA GPUs. The new library-quality open-source solver is called **SSIDS** and is available from <http://www.numerical.rl.ac.uk/spiral/>. **SSIDS** implements a multifrontal algorithm and one of its key features is that it allows all the frontal matrices to be factorized on the GPU. Furthermore, the code produces bit compatible results and incorporates threshold partial pivoting to maintain numerical stability: this has not been done in other GPU sparse solvers. Another novel feature is that **SSIDS** implements the solve phase on the GPU. Both the factorize and solve phases have been shown to achieve performance improvement over our recent multicore CPU code **HSL_MA97**.

Acknowledgements

We wish to thank Mike Giles of the University of Oxford for his useful comments on a draft of this paper. This work was supported by the gift of the C2050 card from NVIDIA, and EPSRC funding through the CCP ASEArch project (EP/J010553/1).

References

- [1] P. R. Amestoy, T. A. Davis, and I. S. Duff. An approximate minimum degree ordering algorithm. *SIAM J. on Matrix Analysis and Applications*, 17:886–905, 1996.
- [2] P. R. Amestoy, T. A. Davis, and I. S. Duff. Algorithm 837: AMD, an approximate minimum degree ordering algorithm. *ACM Transactions on Mathematical Software*, 30(3):381–388, 2004.
- [3] Cleve Ashcraft, Roger G. Grimes, and John G. Lewis. Accurate symmetric indefinite linear equation solvers. *SIAM J. on Matrix Analysis and Applications*, 20(2):513–561, 1999.
- [4] J. R. Bunch and L. Kaufman. Some stable methods for calculating inertia and solving symmetric linear systems. *Mathematics of Computation*, 31:1634–179, 1977.
- [5] T.A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software*, 38(1), 2011.
- [6] I. S. Duff, A. M. Erisman, and J. K. Reid. *Direct Methods for Sparse Matrices*. Oxford University Press, 1989.

- [7] I. S. Duff, N. I. M. Gould, J. K. Reid, J. A. Scott, and K. Turner. Factorization of sparse symmetric indefinite matrices. *IMA Journal of Numerical Analysis*, 11:181–204, 1991.
- [8] I. S. Duff and J. K. Reid. The multifrontal solution of indefinite sparse symmetric linear equations. *ACM Transactions on Mathematical Software*, 9(3):302–325, 1983.
- [9] A. George. Nested dissection of a regular finite-element mesh. *SIAM J. on Numerical Analysis*, 10:345–363, 1973.
- [10] T. George, V. Saxena, A. Gupta, A. Singh, and A. R. Choudhury. Multifrontal factorization of sparse SPD matrices on GPUs. In *Parallel and Distributed Processing Symposium (IPDPS)*, pages 372–383, 2011.
- [11] G. H. Golub and C. F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, third edition, October 1996.
- [12] A. Gupta and H. Avron. WSMP: Watson Sparse Matrix Package. Part I - direct solution of symmetric systems. Version 13.06. Technical Report RC 21886, IBM T. J. Watson Research Center, Yorktown Heights, NY, 2013. <http://www.research.ibm.com/projects/wsmg>.
- [13] N. J. Higham. Stability of parallel triangular system solvers. *SIAM J. on Scientific Computing*, 16(2):400–413, 1995.
- [14] J. D. Hogg. A fast dense triangular solve in cuda. *SIAM Journal on Scientific Computing*, 35(3):C303–C322, 2013.
- [15] J. D. Hogg and J. A. Scott. A note on the solve phase of a multicore solver. Technical Report RAL-TR-2010-007, STFC Rutherford Appleton Laboratory, 2010.
- [16] J. D. Hogg and J. A. Scott. HSL_MA97: a bit-compatible multifrontal code for sparse symmetric systems. Technical Report RAL-TR-2011-024, STFC Rutherford Appleton Laboratory, 2011.
- [17] J. D. Hogg and J. A. Scott. Achieving bit compatibility in sparse direct solvers. Technical Report RAL-P-2012-005, STFC Rutherford Appleton Laboratory, 2012.
- [18] J. D. Hogg and J. A. Scott. New parallel sparse direct solvers for multicore architectures. *Algorithms*, 6:702–725, 2013.
- [19] J. D. Hogg and J. A. Scott. Pivoting strategies for tough sparse indefinite systems. *ACM Transactions on Mathematical Software*, 40, 2014. to appear.
- [20] HSL. A collection of Fortran codes for large-scale scientific computation, 2013. <http://www.hsl.rl.ac.uk>.
- [21] G. Karypis and V. Kumar. METIS: A software package for partitioning unstructured graphs, partitioning meshes and computing fill-reducing orderings of sparse matrices - version 4.0, 1998. <http://www-users.cs.umn.edu/~karypis/metis/>.
- [22] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. on Scientific Computing*, 20:359–392, 1999.
- [23] K. Kim and V. Eijkhout. A parallel sparse direct solver via hierarchical dag scheduling. Technical Report TR-12-04, Texas Advanced Computing Center (TACC), 2012.
- [24] K. Kim and V. Eijkhout. Scheduling a parallel sparse direct solver to multiple GPUs. In *Proceedings 14th IEEE Workshop on Parallel and Distributed Scientific and Engineering Computing*, 2013.

- [25] G. P. Krawezik and G. Poole. Accelerating the ANSYS direct sparse solver with GPUs, 2010. Symposium on Application Accelerators in High Performance Computing (SAAHPC’10).
- [26] X. Lacoste, P. Ramet, M. Faverge, Y. Ichitaro, and J. Dongarra. Sparse direct solvers with accelerators over DAG runtimes. Technical Report No. 7972, INRIA, University of Bordeaux, France, 2013.
- [27] J. W. H. Liu. Modification of the minimum-degree algorithm by multiple elimination. *ACM Transactions on Mathematical Software*, 11(2):141–153, 1985.
- [28] R. F. Lucas, G. Wagenbreth, J. J. Tran, and D. M. Davis. Multifrontal sparse matrix factorization on graphics processing units, 2012. Unpublished manuscript, available at <ftp://ftp.mosis.com/isi-pubs/tr-677.pdf>.
- [29] MUMPS. MUMPS: a multifrontal massively parallel sparse direct solver, 2013. <http://graal.ens-lyon.fr/MUMPS/>.
- [30] O. Schenk and K. Gärtner. On fast factorization pivoting methods for symmetric indefinite systems. *Electronic Transactions on Numerical Analysis*, 23:158–179, 2006.
- [31] W. F. Tinney and J. W. Walker. Direct solutions of sparse network equations by optimally ordered triangular factorization. *Proceedings of the IEEE*, 55:1801–1809, 1967.
- [32] C. D. Yu, W. Wang, and D. L. Pierce. A CPU-GPU hybrid approach for the unsymmetric multifrontal method. *Parallel Computing*, 37(12):759–770, 2011.