

State-of-the-Art Sparse Direct Solvers



Matthias Bollhöfer, Olaf Schenk, Radim Janalik, Steve Hamm,
and Kiran Gullapalli

1 Introduction

Solving large sparse linear systems is at the heart of many application problems arising from computational science and engineering applications. Advances in combinatorial methods in combination with modern computer architectures have massively influenced the design of the state-of-the-art direct solvers that are feasible for solving larger systems efficiently in a computational environment with rapidly increasing memory resources and cores. Among these advances are novel combinatorial algorithms for improving diagonal dominance which pave the way to a static pivoting approach, thus improving the efficiency of the factorization phase dramatically. Besides, partitioning and reordering the system such that a high level of concurrency is achieved, the objective is to simultaneously achieve the reduction of fill-in and the parallel concurrency. While these achievements already significantly improve the factorization phase, modern computer architectures require one to compute as many operations as possible in the cache of the CPU. This in turn can be achieved when dense subblocks that show up during the factorization can be

M. Bollhöfer

Institute for Numerical Analysis, TU Braunschweig, Braunschweig, Germany
e-mail: m.bollhoefer@tu-bs.de

O. Schenk (✉) · R. Janalik

Institute of Computational Science, Faculty of Informatics, Università della Svizzera Italiana,
Lugano, Switzerland
e-mail: olaf.schenk@usi.ch; radim.janalik@usi.ch

S. Hamm · K. Gullapalli

NXP, Austin, TX, USA
e-mail: steve.hamm@nxp.com; kiran.gullapalli@nxp.com

© Springer Nature Switzerland AG 2020

A. Grama, A. H. Sameh (eds.), *Parallel Algorithms in Computational Science and Engineering*, Modeling and Simulation in Science, Engineering and Technology,
https://doi.org/10.1007/978-3-030-43736-7_1

grouped together into dense submatrices which are handled by multithreaded and cache-optimized dense matrix kernels using level-3 BLAS and LAPACK [3].

This chapter will review some of the basic technologies together with the latest developments for sparse direct solution methods that have led to the state-of-the-art LU decomposition methods. The paper is organized as follows. In Sect. 2 we will start with maximum weighted matchings which is one of the key tools in combinatorial optimization to dramatically improve the diagonal dominance of the underlying system. Next, Sect. 3 will review multilevel nested dissection as a combinatorial method to reorder a system symmetrically such that fill-in and parallelization can be improved simultaneously, once pivoting can be more or less ignored. After that, we will review established graph-theoretical approaches in Sect. 4, in particular the elimination tree, from which most of the properties of the LU factorization can be concluded. Among these properties is the prediction of dense submatrices in the factorization. In this way several subsequent columns of the factors L and U^T are collected in a single dense block. This is the basis for the use of dense matrix kernels using optimized level-3 BLAS as well to exploit fast computation using the cache hierarchy which is discussed in Sect. 5. Finally, we show in Sect. 6 how the ongoing developments in parallel sparse direct solution methods have advanced integrated circuit simulations. We assume that the reader is familiar with some elementary knowledge from graph theory, see, e.g., [15, 21] and some simple computational algorithms based on graphs [1].

2 Maximum Weight Matching

In modern sparse elimination methods the key to success is ability to work with efficient data structures and their underlying numerical templates. If we can increase the size of the diagonal entries as much as possible in advance, pivoting during Gaussian elimination can often be bypassed and we may work with static data structures and the numerical method will be significantly accelerated. A popular method to achieve this goal is the maximum weight matching method [16, 37] which permutes, e.g., the rows of a given nonsingular matrix $A \in \mathbb{R}^{n,n}$ by a permutation matrix $\Pi \in \mathbb{R}^{n,n}$ such that $\Pi^T A$ has a *nonzero diagonal*. Moreover, it maximizes the product of the absolute diagonal values and yields diagonal scaling matrices $D_r, D_c \in \mathbb{R}^{n,n}$ such that $\tilde{A} = \Pi^T D_r A D_c$ satisfies $|\tilde{a}_{ij}| \leq 1$ and $|\tilde{a}_{ii}| = 1$ for all $i, j = 1, \dots, n$. The original idea on which these nonsymmetric permutations and scalings are based is to find a *maximum weighted matching* of a *bipartite graphs*. Finding a maximum weighted matching is a well known assignment problem in operation research and combinatorial analysis.

Definition 1 A graph $G = (V, E)$ with vertices V and edges $E \subset V^2$ is called *bipartite* if V can be partitioned into two sets V_r and V_c , such that no edge $e = (v_1, v_2) \in E$ has both ends v_1, v_2 in V_r or both ends v_1, v_2 in V_c . In this case we denote G by $G_b = (V_r, V_c, E)$.

Definition 2 Given a matrix A , then we can associate with it a canonical bipartite graph $G_b(A) = (V_r, V_c, E)$ by assigning the labels of $V_r = \{r_1, \dots, r_n\}$ with the row indices of A and $V_c = \{c_1, \dots, c_n\}$ being labeled by the column indices. In this case E is defined via $E = \{(r_i, c_j) \mid a_{ij} \neq 0\}$.

For the bipartite graph $G_b(A)$ we see immediately that if $a_{ij} \neq 0$, then we have that $r_i \in V_r$ from the row set is connected by an edge $(r_i, c_j) \in E$ to the column $c_j \in V_c$, but neither rows are connected with each other nor do the columns have interconnections.

Definition 3 A *matching* \mathcal{M} of a given graph $G = (V, E)$ is a subset of edges $e \in E$ such that no two of which share the same vertex.

If \mathcal{M} is a matching of a bipartite graph $G_b(A)$, then each edge $e = (r_i, c_j) \in \mathcal{M}$ corresponds to a row i and a column j and there exists no other edge $\hat{e} = (r_k, c_l) \in \mathcal{M}$ that has the same vertices, neither $r_k = r_i$ nor $c_l = c_j$.

Definition 4 A matching \mathcal{M} of $G = (V, E)$ is called *maximal*, if no other edge from E can be added to \mathcal{M} .

If for an $n \times n$ matrix A a *matching* \mathcal{M} of $G_b(A)$ with maximum cardinality n is found, then by definition the edges must be $(i_1, 1), \dots, (i_n, n)$ with i_1, \dots, i_n being the numbers $1, \dots, n$ in a suitable order and therefore we obtain $a_{i_1,1} \neq 0, \dots, a_{i_n,n} \neq 0$. In this case we have established that the matrix A is at least structurally nonsingular and we can use a row permutation matrix Π^T associated with row ordering i_1, \dots, i_n to place a nonzero entry on each diagonal location of $\Pi^T A$.

Definition 5 A *perfect matching* is a maximal matching with cardinality n .

It can be shown that for a structurally nonsingular matrix A there always exists a perfect matching \mathcal{M} .

Perfect Matching

In Fig. 1, the set of edges $\mathcal{M} = \{(1, 2), (2, 4), (3, 5), (4, 1), (5, 3), (6, 6)\}$ represents a perfect maximum matching of the bipartite graph $G_b(A)$.

The most efficient combinatorial methods for finding maximum matchings in bipartite graphs make use of an *augmenting path*. We will introduce some graph terminology for the construction of perfect matchings.

Definition 6 If an edge $e = (u, v)$ in a graph $G = (V, E)$ joins a vertices $u, v \in V$, then we denote it as uv . A path then consists of edges $u_1u_2, u_2u_3, u_3u_4 \dots, u_{k-1}u_k$, where each $(u_i, u_{i+1}) \in E, i = 1, \dots, k - 1$.




Fig. 1 Perfect matching. Left side: original matrix A . Middle: bipartite representation $G_b(A) = (V_r, V_c, E)$ of the matrix A and perfect matching M . Right side: permuted matrix $\Pi^T A$

If $G_b = (V_r, V_c, E)$ is a bipartite graph, then by definition of a path, any path is alternating between the vertices of V_r and V_c , e.g., paths in G_b could be such as $r_1c_2, c_2r_3, r_3c_4, \dots$

Definition 7 Given a graph $G = (V, E)$, a vertex is called *free* if it is not incident to any other edge in a matching M of G . An *alternating path* relative to a matching M is a path $P = u_1u_2, u_2u_3, \dots, u_{s-1}u_s$ where its edges are alternating between $E \setminus M$ and M . An *augmenting path* relative to a matching M is an alternating path of odd length and both of its vertex endpoints are free.

Augmenting Path

Consider Fig. 1. To better distinguish between row and column vertices we use $\boxed{1}, \boxed{2}, \dots, \boxed{6}$ for the rows and $\textcircled{1}, \textcircled{2}, \dots, \textcircled{6}$ for the columns. A non-perfect but maximal matching is given by $M = \{(\boxed{4}, \textcircled{5}), (\boxed{1}, \textcircled{1}), (\boxed{6}, \textcircled{2}), (\boxed{2}, \textcircled{6}), (\boxed{5}, \textcircled{4})\}$. We can easily see that an augmenting path alternating between rows and columns is given by $\boxed{3} \textcircled{5}, \textcircled{5} \boxed{4}, \boxed{4} \textcircled{1}, \textcircled{1} \boxed{1}, \boxed{1} \textcircled{2}, \textcircled{2} \boxed{6}, \boxed{6} \textcircled{6}, \textcircled{6} \boxed{2}, \boxed{2} \textcircled{4}, \textcircled{4} \boxed{5}, \boxed{5} \textcircled{3}$. Both endpoints $\boxed{3}$ and $\textcircled{3}$ of this augmenting path are free.

In a bipartite graph $G_b = (V_r, V_c, E)$ one vertex endpoint of any augmenting path must be in V_r , whereas the other one must be in V_c . The symmetric difference, $A \oplus B$ of two edge sets A, B , is defined to be $(A \setminus B) \cup (B \setminus A)$.

Using these definitions and notations, the following theorem [5] gives a constructive algorithm for finding perfect matchings in bipartite graphs.

Theorem 1 *If \mathcal{M} is non-maximum matching of a bipartite graph $G_b = (V_r, V_c, E)$, then there exists an augmenting path P relative to \mathcal{M} such that $P = \mathcal{M} \oplus \mathcal{M}$ and \mathcal{M} is a matching with cardinality $|\mathcal{M}| + 1$.*

According to this theorem, a combinatorial method of finding perfect matching in a bipartite graph is to seek augmenting paths.

The perfect matching as discussed so far only takes the nonzero structure of the matrix into account. For their use as static pivoting methods prior to the LU decomposition one requires in addition to maximize the absolute value of the product of the diagonal entries. This is referred to as *maximum weighted matching*. In this case a permutation π has to be found, which maximizes

$$\prod_{i=1}^n |a_{\pi(i)i}|. \quad (1)$$

The maximization of this product is transferred into a minimization of a sum as follows. We define a matrix $C = (c_{ij})$ via

$$c_{ij} = \begin{cases} \log a_i - \log |a_{ij}| & a_{ij} \neq 0 \\ \infty & \text{otherwise,} \end{cases}$$

where $a_i = \max_j |a_{ij}|$ is the maximum element in row i of matrix A . A permutation π which minimizes the sum

$$\sum_{i=1}^n c_{\pi(i)i}$$

also maximizes the product (1). The minimization problem is known as linear-sum assignment problem or bipartite weighted matching problem in combinatorial optimization. The problem is solved by a sparse variant of the Hungarian method. The complexity is $\mathcal{O}(n\tau \log n)$ for sparse matrices with τ entries. For matrices, whose associated graph fulfill special requirements, this bound can be reduced further to $\mathcal{O}(n^\alpha(\tau + n \log n))$ with $\alpha < 1$. All graphs arising from finite-difference or finite element discretizations meet the conditions [24]. As before, we finally get a perfect matching which in turn defines a nonsymmetric permutation.

When solving the assignment problem, two dual vectors $u = (u_i)$ and $v = (v_j)$ are computed which satisfy

$$u_i + v_j = c_{ij} \quad (i, j) \in \mathcal{M}, \quad (2)$$

$$u_i + v_j \leq c_{ij} \quad \text{otherwise.} \quad (3)$$

Using the exponential function these vectors can be used to scale the initial matrix. To do so define two diagonal matrices D_r and D_c through

$$D_r = \text{diag}(d_1^r, d_2^r, \dots, d_n^r), \quad d_i^r = \exp(u_i), \quad (4)$$

$$D_c = \text{diag}(d_1^c, d_2^c, \dots, d_n^c), \quad d_j^c = \exp(v_j)/a_j. \quad (5)$$

Using Eqs. (2) and (3) and the definition of C , it immediately follows that $\tilde{A} = \Pi^T D_r A D_c$ satisfies

$$|\tilde{a}_{ii}| = 1, \quad (6)$$

$$|\tilde{a}_{ij}| \leq 1. \quad (7)$$

The permuted and scaled system \tilde{A} has been observed to have significantly better numerical properties when being used for direct methods or for preconditioned iterative methods, cf., e.g., [4, 16]. Olschowka and Neumaier [37] introduced these scalings and permutation for reducing pivoting in Gaussian elimination of full matrices. The first implementation for sparse matrix problems was introduced by Duff and Koster [16]. For symmetric matrices $|A|$, these nonsymmetric matchings can be converted to a symmetric permutation P and a symmetric scaling $D_s = (D_r D_c)^{1/2}$ such that $P^T D_s A D_s P$ consists mostly of diagonal blocks of size 1×1 and 2×2 satisfying a similar condition as (6) and (7), where in practice it rarely happens that 1×1 blocks are identical to 0 [17]. Recently, successful parallel approaches to compute maximum weighted matchings have been proposed [28, 29].

Example 1: Maximum Weight Matching

To conclude this section we demonstrate the effectiveness of maximum weight matchings using a simple sample matrix “west0479” from the SuiteSparse Matrix Collection. The matrix can also directly be loaded in MATLAB using `load west0479`. In Fig. 2 we display the matrix before and after applying maximum weighted matchings. To illustrate the improved diagonal dominance we further compute $r_i = |a_{ii}| / \sum_{j=1}^n |a_{ij}|$ for each row of A and $\tilde{A} = \Pi^T D_r A D_s$, $i = 1, \dots, n$. r_i can be read as relative diagonal dominance of row i and yields a number between 0 and 1. Moreover, whenever $r_i > \frac{1}{2}$, the row is strictly diagonal dominant, i.e., $|a_{ii}| > \sum_{j:j \neq i} |a_{ij}|$. In Fig. 3 we display for both matrices r_i by sorting its values in increasing order and taking $\frac{1}{2}$ as reference line. We can see the dramatic impact of maximum weighted

(continued)




Fig. 2 Maximum weight matching. Left side: original matrix A . Right side: permuted and rescaled matrix $\tilde{A} = \Pi^T D_r A D_c$




Fig. 3 Diagonal dominance. Left side: r_i for A . Right side: r_i $\tilde{A} = \Pi^T D_r A D_c$

matchings in improving the diagonal dominance of the given matrix and thus paving the way to a static pivoting approach in incomplete or complete LU decomposition methods.

3 Symbolic Symmetric Reordering Techniques

When dealing with large sparse matrices a crucial factor that determines the computation time is the amount of fill that is produced during the factorization of the underlying matrix. To reduce the complexity there exist many mainly symmetric reordering techniques that attempt to reduce the fill-in heuristically. Here we will demonstrate only one of these methods, the so-called nested dissection method.

The main reason for selecting this method is that it can be easily used for parallel computations.

3.1 Multilevel Nested Dissection

Recursive multilevel nested dissection methods for direct decomposition methods were first introduced in the context of multiprocessing. If parallel direct methods are used to solve a sparse system of equations, then a graph partitioning algorithm can be used to compute a fill-reducing ordering that leads to a high degree of concurrency in the factorization phase.

Definition 8 For a matrix $A \in \mathbb{R}^{n,n}$ we define the associated (directed) graph $G_d(A) = (V, E)$, where $V = \{1, \dots, n\}$ and the set of edges $E = \{(i, j) \mid a_{ij} \neq 0\}$. The (undirected) graph is given by $G_d(|A| + |A|^T)$ and is denoted simply by $G(A)$.

In graph terminology for a sparse matrix A we simply have a directed edge (i, j) for any nonzero entry a_{ij} in $G_d(A)$, whereas the orientation of the edge is ignored in $G(A)$.

The research on graph partitioning methods in the mid-nineties has resulted in high-quality software packages, e.g., METIS [25]. These methods often compute orderings that on the one hand lead to small fill-in for (incomplete) factorization methods, while on the other hand they provide a high level of concurrency. We will briefly review the main idea of multilevel nested dissection in terms of graph partitioning.

Definition 9 Let $A \in \mathbb{R}^{n,n}$ and consider its graph $G(A) = (V, E)$. A k -way graph partitioning consists of partitioning V into k disjoint subsets V_1, V_2, \dots, V_k such that $V_i \cap V_j = \emptyset$ for $i \neq j$ and $\cup_i V_i = V$. The subset $E_s = E \cap \cup_{i \neq j} (V_i \times V_j)$ is called *edge separator*.

Typically we want a k -way partitioning to be balanced, i.e., each V_i should satisfy $|V_i| \approx n/k$. The edge separator E_s refers to the edges that have to be taken away from the graph in order to have k separate subgraphs associated with V_1, \dots, V_k and the number of elements of E_s is usually referred to as edge-cut.

Definition 10 Given $A \in \mathbb{R}^{n,n}$, a *vertex separator* V_s of $G(A) = (V, E)$ is a set of vertices such that there exists a k -way partitioning V_1, V_2, \dots, V_k of $V \setminus V_s$ having no edge $e \in V_i \times V_j$ for $i \neq j$.

A useful vertex separator V_s should not only separate $G(A)$ into k independent subgraphs associated with V_1, \dots, V_k , it is intended that the number of edges $\cup_{i=1}^k \{|e_{is} \in V_i, s \in V_s\}|$ is also small.

Nested dissection recursively splits a graph $G(A) = (V, E)$ into almost equal parts by constructing a vertex separator V_s until the desired number k of partitionings are obtained. If k is a power of 2, then a natural way of obtaining a vertex separator is to first obtain a 2-way partitioning of the graph, a so-called *graph*




Fig. 4 A 2-way partition with vertex separator $V_s = \{1, 4\}$ and the associated reordered matrix placing the two rows and columns associated with V_s to the end

bisection with its associated edge separator E_s . After that a vertex separator V_s is computed from E_s , which gives a 2-way partitioning V_1, V_2 of $V \setminus V_s$. This process is then repeated separately for the subgraphs associated with V_1, V_2 until eventually a $k = 2^l$ -way partitioning is obtained. For the reordering of the underlying matrix A , the vertices associated with V_1 are taken first followed by V_2 and V_s . This reordering is repeated similarly during repeated bisection of each V_i . In general, vertex separators of small size result in low fill-in.

Example 2: Vertex Separators

To illustrate vertex separators, we consider the reordered matrix $\Pi^T A$ from Fig. 1 after a matching is applied. In Fig. 4 we display its graph $G(\Pi^T A)$ ignoring the orientation of the edges. A 2-way partitioning is obtained with $V_1 = \{3, 5\}, V_2 = \{2, 6\}$, and a vertex separator $V_s = \{1, 4\}$. The associated reordering refers to taking the rows and the columns of $\Pi^T A$ in the order 3, 5, 2, 6, 1, 4.

Since a naive approach to compute a recursive graph bisection is typically computationally expensive, combinatorial *multilevel graph bisection* has been used to accelerate the process. The basic structure is simple. The multilevel approach consists of three phases: at first there is a *coarsening phase* which compresses the given graph successively level by level by about half of its size. When the coarsest graph with about a few hundred vertices is reached, the second phase, namely the so-called *bisection*, is applied. This is a high-quality partitioning algorithm. After that, during the *uncoarsening phase*, the given bisection is successively refined as it is prolonged towards the original graph.

3.1.1 Coarsening Phase

The initial graph $G_0 = (V_0, E_0) = G(A)$ of $A \in \mathbb{R}^{n,n}$ is transformed during the coarsening phase into a sequence of graphs G_1, G_2, \dots, G_m of decreasing size

such that $|V_0| \gg |V_1| \gg |V_2| \gg \dots \gg |V_m|$. Given the graph $G_i = (V_i, E_i)$, the next coarser graph G_{i+1} is obtained from G_i by collapsing adjacent vertices. This can be done, e.g., by using a maximal matching \mathcal{M}_i of G_i (cf. Definitions 3 and 4). Using \mathcal{M}_i , the next coarser graph G_{i+1} is constructed from G_i collapsing the vertices being matched into multinodes, i.e., the elements of \mathcal{M}_i together with the unmatched vertices of G_i become the new vertices V_{i+1} of G_{i+1} . The new edges E_{i+1} are the remaining edges from E_i connected with the collapsed vertices. There are various differences in the construction of maximal matchings [9, 25]. One of the most popular and efficient methods is heavy edge matching [25].

3.1.2 Partitioning Phase

At the coarsest level m , a 2-way partitioning $V_{m,1} \dot{\cup} V_{m,2} = V_m$ of $G_m = (V_m, E_m)$ is computed, each of them containing about half of the vertices of G_m . This specific partitioning of G_m can be obtained by using various algorithms such as spectral bisection [19] or combinatorial methods based on Kernighan–Lin variants [18, 27]. It is demonstrated in [25] that for the coarsest graph, combinatorial methods typically compute smaller edge-cut separators compared with spectral bisection methods. However, since the size of the coarsest graph G_m is small (typically $|V_m| < 100$), this step is negligible with respect to the total amount of computation time.

3.1.3 Uncoarsening Phase

Suppose that at the coarsest level m , an edge separator $E_{m,s}$ of G_m associated with the 2-way partitioning has been computed that has led to a sufficient edge-cut of G_m with $V_{m,1}, V_{m,2}$ of almost equal size. Then $E_{m,s}$ is prolonged to G_{m-1} by reversing the process of collapsing matched vertices. This leads to an initial edge separator $E_{m-1,s}$ for G_{m-1} . But since G_{m-1} is finer, $E_{m-1,s}$ is sub-optimal and one usually decreases the edge-cut of the partitioning by local refinement heuristics such as the Kernighan–Lin partitioning algorithm [27] or the Fiduccia–Mattheyses method [18]. Repeating this refinement procedure level-by-level we obtain a sequence of edge separators $E_{m,s}, E_{m-1,s}, \dots, E_{0,s}$ and eventually an edge separator $E_s = E_{0,s}$ of the initial graph $G(A)$ is obtained. If one is seeking for a vertex separator V_s of $G(A)$, then one usually computes V_s from E_s at the end.

There have been a number of methods that are used for graph partitioning, e.g., METIS [25], a parallel MPI version PARMETIS [26], or a recent multithreaded approach MT-METIS [30]. Another example for a parallel partitioning algorithm is SCOTCH [9].




Fig. 5 Application of multilevel nested dissection after the matrix is already rescaled and permuted using maximum weight matching

Multilevel Nested Dissection

We will continue Example 1 using the matrix $\tilde{A} = \Pi^T D_r A D_s$ that has been rescaled and permuted using maximum weight matching. We illustrate in Fig. 5 how multilevel nested dissection changes the pattern $\hat{A} = P^T \tilde{A} P$, where P refers to the permutation matrix associated with the partitioning of $G(\tilde{A})$.

3.2 Other Reordering Methods

One of the first methods to reorder the system was the reverse Cuthill–McKee (RCM) methods [10, 34] which attempts to reduce the bandwidth of a given matrix. Though this algorithm is still attractive for sequential methods and incomplete factorization methods, its use for direct solvers is considered as obsolete. An attractive alternative to nested dissection as reordering method for direct factorization methods is the minimum degree algorithm (MMD) [20, 40] and its recent variants,

in particular the approximate minimum degree algorithm (AMD) [2, 12] with or without constraints. The main objective of the minimum degree algorithm is to simulate the Gaussian elimination process symbolically by investigating the update process $a_{ij} \rightarrow a_{ij} - a_{ik}a_{kk}^{-1}a_{kj}$ by means of graph theory, at least in the case of the undirected graph. The name-giving degree refers to the number of edges connected to a vertex and how the graph and therefore the degrees of its vertices change during the factorization process. Over the years this has led to an evolution of the underlying minimum degree algorithm using the so-called *external degree* for selecting vertices as pivots and further techniques like *incomplete degree update*, *element absorption*, and *multiple elimination* as well as data structures based on cliques. For an overview see [20]. One of the most costly parts in the minimum degree algorithm is to update of the degrees. Instead of computing the exact external degree, in the approximate minimum degree algorithm [2], an approximate external degree is computed that significantly saves time while producing comparable fill-in for the LU decomposition.

We like to conclude this section by mentioning that if nested dissection is computed to produce a vertex separator V_s and a related k -way partitioning $V_1, \dots, V - k$ for the remaining vertices of $V \setminus V_s$ of $G(A) = (V, E)$ which allow for parallel computations, then the entries of each V_i , i, \dots, k could be taken in any order. Certainly, inside V_i one could use nested dissection as well, which is the default choice in multilevel nested dissection methods. However, as soon as the coarsest graph G_m is small enough (typically about 100 vertices), not only the separator is computed, but in addition the remaining entries of G_m are reordered to lead to a fill-reducing ordering. In both cases, for G_m as well as V_1, \dots, V_k one could alternatively use different reordering methods such as variants of the minimum degree algorithm. Indeed, for G_m this is what the METIS software is doing. Furthermore, a reordering method such as the constrained approximate minimum degree algorithm is also suitable as local reordering for V_1, \dots, V_k as alternative to nested dissection, taking into account the edges connected with V_s (also referred to as HALO structure), see, e.g., [38].

4 Sparse LU Decomposition

In this section we will assume that the given matrix $A \in \mathbb{R}^{n,n}$ is nonsingular and that it can be factorized as $A = LU$, where L is a lower triangular matrix with unit diagonal and U is an upper triangular matrix. It is well-known [21], if $A = LU$, where L and U^\top are lower triangular matrices, then in the generic case we will have $G_d(L + U) \supset G_d(A)$, i.e., we will only get additional edges unless some entries cancel by “accident” during the elimination. In the sequel we will ignore cancellations. Throughout this section we will always assume that the diagonal entries of A are nonzero as well. We also assume that $G_d(A)$ is connected.

In the preceding sections we have argued that maximum weight matching often leads to a rescaled and reordered matrix such that static pivoting is likely

Fig. 6 Fill-in with respect to $L + U$ is denoted by \times

$$\left(\begin{array}{ccc|ccc} 3 & 0 & 0 & 0 & 0 & 1 \\ 0 & 3 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 3 & 0 & 1 & 0 \\ 0 & 0 & 1 & 2 & \times & 0 \\ \hline 0 & 1 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 & 3 & 4 \end{array} \right)$$

to be enough, i.e., pivoting is restricted to some dense blocks inside the LU factorization. Furthermore, reordering strategies such as multilevel nested dissection have further symmetrically permuted the system such that the fill-in that occurs during Gaussian elimination is acceptable and even parallel approaches could be drawn from this reordering. Thus assuming that A does not need further reordering and a factorization $A = LU$ exists is a realistic scenario in what follows.

4.1 The Elimination Tree

The basis of determining the fill-in in the triangular factors L and U as by-product of the Gaussian elimination can be characterized as follows (see [23] and the references therein).

Theorem 2 *Given $A = LU$ with the aforementioned assumptions, there exists an edge (i, j) in $G_d(L + U)$ if and only if there exists a path*

$$i x_1, x_2 x_3, \dots, x_k j$$

in $G_d(A)$ such that $x_1, \dots, x_k < \min(i, j)$.

In other words, during Gaussian elimination we obtain a fill edge (i, j) for every path from i to j through vertices less than $\min(i, j)$.

Fill-in

We will use the matrix $\Pi^T A$ from Example 2 and sketch the fill-in obtained during Gaussian elimination in Fig. 6.

The fastest known method for predicting the filled graph $G_d(L + U)$ is Gaussian elimination. The situation is simplified if the graph is undirected. In the sequel we ignore the orientation of the edges and simply consider the undirected graph $G(A)$ and $G(L + U)$, respectively.

Fig. 7 Entries of $G(A)$ are denoted by filled circle, fill-in is denoted by times

$$\begin{pmatrix} \bullet & 0 & 0 & 0 & 0 & \bullet \\ 0 & \bullet & 0 & 0 & \bullet & 0 \\ \hline 0 & 0 & \bullet & \bullet & \bullet & 0 \\ 0 & 0 & \bullet & \bullet & \times & \bullet \\ \hline 0 & \bullet & \bullet & \times & \bullet & \bullet \\ \bullet & 0 & \bullet & \bullet & \bullet & \bullet \end{pmatrix}$$

Definition 11 The undirected graph $G(L + U)$ that is derived from the undirected graph $G(A)$ by applying Theorem 2 is called the *filled graph* and it will be denoted by $G_f(A)$.

Fill-in with Respect to the Undirected Graph

When we consider the undirected graph $G(A)$ in Example 4.1, the pattern of $|\Pi^T A| + |\Pi^T A|^T$ and its filled graph $G_f(A)$ now equals $G(A)$ up to positions (5, 4) and (4, 5) (cf. Fig. 7).

The key tool to predict the fill-in easily for the undirected graph is the *elimination tree* [33].

Recall that an undirected and connected graph is called a *tree*, if it does not contain any cycle. Furthermore, one vertex is identified as *root*. As usual we call a vertex j *parent* of i , if there exists an edge (i, j) in the tree such that j is closer to the root. In this case i is called *child* of j . The subtree rooted at vertex j is denoted by $T(j)$ and the vertices of this subtree are called *descendants* of j , whereas j is called their *ancestor*. Initially we will define the elimination tree algorithmically using the depth-first-search algorithm [1]. Later we will state a much simplified algorithm.

Definition 12 Given the filled graph $G_f(A)$ the *elimination tree* $T(A)$ is defined by the following algorithm.

Perform a depth-first-search in $G_f(A)$ starting from vertex n .

When vertex m is visited, choose from its unvisited neighbors i_1, \dots, i_k the index j with the largest number $j = \max\{i_1, \dots, i_k\}$ and continue the search with j .

A leaf of the tree is reached, when all neighbors have already been visited.

We like to point out that the application of the depth-first-search to $G_f(A)$ starting at vertex n behaves significantly different from other graphs. By Theorem 2 it follows that as soon as we visit a vertex m , all its neighbors $j > m$ must have been visited prior to vertex m . Thus the labels of the vertices are strictly decreasing until we reach a leaf node.

Depth-First-Search

We illustrate the depth-first-search using the (filled) graph in Fig. 8 and the pattern from Example 4.1. The extra fill edge is marked by a bold line.

The ongoing depth-first-search visits the vertices in the order $6 \rightarrow 5 \rightarrow 4 \rightarrow 3$. Since at vertex 3, all neighbors of 3 are visited (and indeed have a larger number), the algorithm backtracks to 4 and to 5 and continues the search in the order $5 \rightarrow 2$. Again all neighbors of vertex 2 are visited (and have larger number), thus the algorithm backtracks to 5 and to 6 and continues by $6 \rightarrow 1$. Then the algorithm terminates.




Fig. 8 Filled graph (left) and elimination tree (right)

Remark 1 It follows immediately from the construction of $T(A)$ and Theorem 2 that additional edges of $G_f(A)$ which are not covered by the elimination tree can only show up between a vertex and some of its ancestors (referred to as “back-edges”). In contrast to that, “cross-edges” between unrelated vertices do not exist.

Remark 2 One immediate consequence of Remark 1 is that triangular factors can be computed independently starting from the leaves until the vertices meet a common parent, i.e., column j of L and U^T only depend on those columns s of L and U^T such that s is a descendant of j in the elimination tree $T(A)$.

Elimination Tree

We use the matrix “west0479” from Example 3.1.3, after maximum weight matching and multilevel nested dissection have been applied. We use MATLAB’s `etreeplot` to display its elimination tree (see Fig. 9). The

(continued)




Fig. 9 Elimination tree of “west0479” after maximum weight matching and nested dissection are applied

elimination tree displays the high level of concurrency that is induced by nested dissection, since by Remark 2 the computations can be executed independently at each leaf node towards the root until a common parent vertex is reached.

Further conclusions can be easily derived from the elimination tree, in particular Remark 2 in conjunction with Theorem 2.

Remark 3 Consider some $k \in \{1, \dots, n\}$. Then there exists a (fill) edge (j, k) with $j < k$ if and only if there exists a common descendant i of k, j in $T(A)$ such that $a_{ik} \neq 0$. This follows from the fact that once $a_{ik} \neq 0$, by Theorem 2 this induces (fill) edges (j, k) in the filled graph $G_f(A)$ for all nodes j between i and k in the elimination tree $T(A)$, i.e., for all ancestors of i that are also descendants of k . This way, i propagates fill-edges along the branch from i to k in $T(A)$ and the information $a_{ik} \neq 0$ can be used as path compression to advance from i towards k along the elimination tree.

Path Compression

Consider the graph and the elimination tree from Fig. 8. Since there exists the edge $(3, 5)$ in $G(A)$, therefore another (fill) edge $(4, 5)$ must exist. Similarly, the same conclusion can be drawn from the existence of the edge $(4, 6)$ (here not a fill edge, but a regular edge).

The elimination tree itself can be easily described by a vector p of length n such that for any $i < n$, p_i denotes the parent node, while $p_n = 0$ corresponds to the root. Consider some step k with $a_{ik} \neq 0$, for some $i < k$. By Remark 3, i must be a descendant of k and there could be further ancestors j of i which are also

descendants of k . Possibly not all ancestors of i have been assigned a parent node so far. Thus we can replace i by $j = p_i$ until we end up with $p_j = 0$ or $p_j \geq k$. This way we traverse $T(A)$ from i towards to k until we have found the child node j of k . If the parent of j has not been assigned to j yet, then $p_j = 0$ and k must be the parent of j . If some $l < k$ were the parent of j , then we would have assigned l as parent of j in an earlier step $l < k$. In this case we set $p_j \leftarrow k$. Otherwise, if $p_j \geq k$, then we have already assigned j 's parent in an earlier step $l < k$.

Computation of Parent Nodes

Consider the elimination tree $T(A)$ from Fig. 8. Unless $k = 4$, no parents have been assigned, i.e., $p_i = 0$ for all i .

Now for $k = 4$ we have $a_{34} \neq 0$ and using the fact that $p_3 = 0$ implies that we have to set $a_3 = p_3 \leftarrow 4$.

For $k = 5$, $a_{25} \neq 0$ and again $p_2 = 0$ requires to set $a_2 = p_2 \leftarrow 5$. Next, $a_{35} \neq 0$, path compression enables $a_3 \leftarrow 5$ and after another loop we obtain $a_4 = p_4 \leftarrow 5$.

Finally, if $k = 6$, we have $a_{16} \neq 0$ and immediately obtain $a_1 = p_1 \leftarrow 6$. Since $a_{46} \neq 0$, a path compression is applied which yields $a_4 \leftarrow 6$ and in the next step we set $a_5 = p_5 \leftarrow 6$. At last $a_{56} \neq 0$ does not cause further changes.

In total we have $p = [6, 5, 4, 5, 6, 0]$ which perfectly reveals the parent properties of the elimination trees in Fig. 8.

By Remark 3 (cf. [12, 43]), we can also make use of path compression. Since our goal is to traverse the branch of the elimination tree from i to k as fast as possible, any ancestor $j = a_i$ of i would be sufficient. With the same argument as before, an ancestor $a_j = 0$ would refer to a vertex that does not have a parent yet. In this case we can again set $p_j \leftarrow k$. Moreover, k is always an ancestor of a_i .

The algorithm including path compression can be summarized as follows (see also [12, 33]).

Computation of the Elimination Tree

Input: $A \in \mathbb{R}^{n,n}$ such that A has the same pattern as $|A| + |A|^T$.

Output: vector $p \in \mathbb{R}^n$ such that p_i is the parent of i , $i = 1, \dots, n-1$, except $p_n = 0$.

- 1: let $a \in \mathbb{R}^n$ be an auxiliary vector used for path compression.
- 2: $p \leftarrow 0$, $a \leftarrow 0$
- 3: **for** $k = 2, \dots, n$ **do**
- 4: **for** all $i < k$ such that $a_{ik} \neq 0$ **do**

(continued)

```

5:   while  $i \neq 0$  and  $i < k$  do
6:      $j \leftarrow a_i$ 
7:      $a_i \leftarrow k$ 
8:     if  $j = 0$  then
9:        $p_i \leftarrow k$ 
10:    end if
11:     $i \leftarrow j$ 
12:  end while
13: end for
14: end for

```

4.2 The Supernodal Approach

We have already seen that the elimination tree reveals information about concurrency. It is further useful to determine the fill-in L and U^T . This information can be computed from the elimination tree $T(A)$ together with $G(A)$. The basis for determining the fill-in in each column is again Remark 3. Suppose we are interested in the nonzero entries of column j of L and U^T . Then for all descendants of j , i.e., the nodes of the subtree $T(j)$ rooted at vertex j , a nonzero entry $a_{ik} \neq 0$ also implies $l_{kj} \neq 0$. Thus, starting at any leaf i , we obtain its fill by all $a_{ik} \neq 0$ such that $k > i$ and when we move forward from i to its parent j , vertex j will inherit the fill from node i for all $k > j$ plus the nonzero entries given by $a_{jk} \neq 0$ such that $k > j$. When we reach a common parent node k with multiple children, the same argument applies using the union of fill-in greater than k from its children together with the nonzero entries $a_{kl} \neq 0$ such that $l > k$. We summarize this result in a very simple algorithm

Computation of Fill-in

Input: $A \in \mathbb{R}^{n,n}$ such that A has the same pattern as $|A| + |A|^T$.

Output: sparse strict lower triangular pattern $P \in \mathbb{R}^{n,n}$ with same pattern as L, U^T .

- 1: compute parent array p of the elimination tree $T(A)$
- 2: **for** $j = 1, \dots, n$ **do**
- 3: supplement nonzeros of column j of P with all $i > j$ such that $a_{ij} \neq 0$
- 4: $k = p_j$
- 5: **if** $k > 0$ **then**

(continued)

```

6:     supplement nonzeros of column  $k$  of  $P$  with nonzeros of column  $j$ 
      of  $P$  greater than  $k$ 
7:   end if
8: end for

```

Algorithm 4.2 only deals with the fill pattern. One additional aspect that allows to raise efficiency and to speed up the numerical factorization significantly is to detect dense submatrices in the factorization. Block structures allow to collect parts of the matrix in dense blocks and to treat them commonly using dense matrix kernels such as level-3 BLAS and LAPACK [13, 14].

Dense blocks can be read off from the elimination tree employing Algorithm 4.2.

Definition 13 Denote by \mathcal{P}_j the nonzero indices of column j of P as computed by Algorithm 4.2. A sequence $k, k + 1, \dots, k + s - 1$ is called *supernode* of size s if the columns of $\mathcal{P}_j = \mathcal{P}_{j+1} \cup \{j + 1\}$ for all $j = k, \dots, k + s - 2$.

In simple words, Definition 13 states that for a supernode s subsequent columns can be grouped together in one dense block with a triangular diagonal block and a dense subdiagonal block since they perfectly match the associated trapezoidal shape. We can thus easily supplement Algorithm 4.2 with a supernode detection.

Computation of Fill-in and Supernodes

Input: $A \in \mathbb{R}^{n,n}$ such that A has the same pattern as $|A| + |A|^T$.

Output: sparse strict lower triangular pattern $P \in \mathbb{R}^{n,n}$ with same pattern as L, U^T as well as column size $s \in \mathbb{R}^m$ of each supernode.


```

1: compute parent array  $p$  of the elimination tree  $T(A)$ 
2:  $m \leftarrow 0$ 
3: for  $j = 1, \dots, n$  do
4:   supplement nonzeros of column  $j$  of  $P$  with all  $i > j$  such that  $a_{ij} \neq 0$ 
5:   denote by  $r$  the number of entries in column  $j$  of  $P$ 
6:   if  $j > 1$  and  $j = p_{j-1}$  and  $s_m + r = l$  then
7:      $s_m \leftarrow s_m + 1$  ▷ continue current supernode
8:   else
9:      $m \leftarrow m + 1, s_m \leftarrow 1, l \leftarrow r$  ▷ start new supernode
10:  end if
11:   $k = p_j$ 

```

(continued)

Fig. 10 Supernodes in the triangular factor



```

12:  if  $k > 0$  then
13:     supplement nonzeros of column  $k$  of  $P$  with nonzeros of column  $j$ 
       of  $P$  greater than  $k$ 
14:  end if
15: end for

```

Supernode Computation

To illustrate the use of supernodes, we consider the matrix pattern from Fig. 7 and illustrate the underlying dense block structure in Fig. 10. Supernodes are the columns 1, 2, 3 as scalar columns as well as columns 4–6 as one single supernode.

Supernodes form the basis of several improvements, e.g., a supernode can be stored as one or two dense matrices. Beside the storage scheme as dense matrices, the nonzero row indices for these blocks need only be stored once. Next the use of dense submatrices allows the usage of dense matrix kernels using level-3 BLAS [13, 14].

Supernodes

We use the matrix “west0479” from Example 3.1.3, after maximum weight matching and multilevel nested dissection have been applied. We use its undirected graph to compute the supernodal structure. Certainly, since the matrix is nonsymmetric, the block structure is only sub-optimal. We display the supernodal structure for the associated Cholesky factor, i.e., for the Cholesky factor of a symmetric positive definite matrix with same undirected graph as our matrix (see left part of Fig. 11). Furthermore, we display the supernodal structure for the factors L and U computed from the nonsymmetric matrix without pivoting (see right part of Fig. 11).

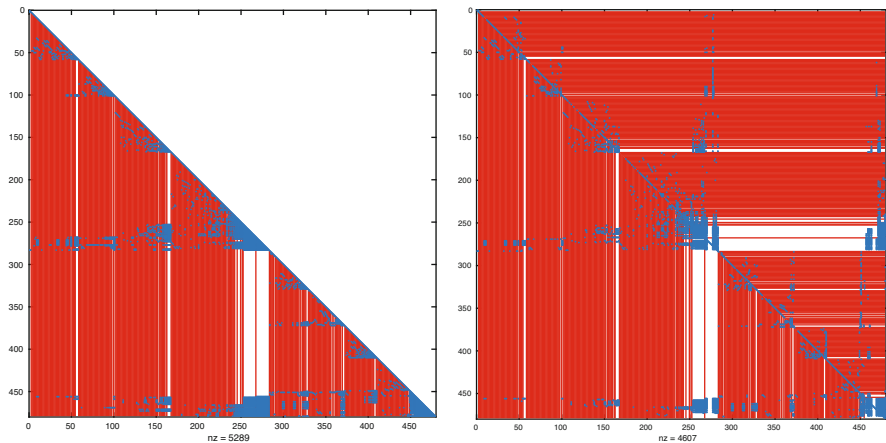


Fig. 11 Supernodal structure. Left: vertical lines display the blocking of the supernodes with respect to the associated Cholesky factor. Right: vertical and horizontal lines display the blocking of the supernodes applied to L and U

While the construction of supernodes is fairly easy in the symmetric case, its generalization for the nonsymmetric case is significantly harder, since one has to deal with pivoting in each step of Gaussian elimination. In this case one uses the column elimination tree [22].

5 Sparse Direct Solvers—Supernodal Data Structures

High-performance sparse solver libraries have been a very important part of scientific and engineering computing for years, and their importance continues to grow as microprocessor architectures become more complex and software libraries become better designed to integrate easily within applications. Despite the fact that there are various science and engineering applications, the underlying algorithms typically have remarkable similarities, especially those algorithms that are most challenging to implement well in parallel. It is not too strong a statement to say that these software libraries are essential to the broad success of scalable high-performance computing in computational sciences. In this section we demonstrate the benefit of supernodal data structures within the sparse solver package PARDISO [42]. We illustrate it by using the triangular solution process. The forward and backward substitution is performed column wise with respect to the columns of L , starting with the first column, as depicted in Fig. 12. The data dependencies here allow to store vectors y , z , b , and x in only one vector r . When column j is reached, r_j contains the solution for y_j . All other elements of L in this column, i. e. L_{ij} with $i = j + 1, \dots, N$, are used to update the remaining entries in r by

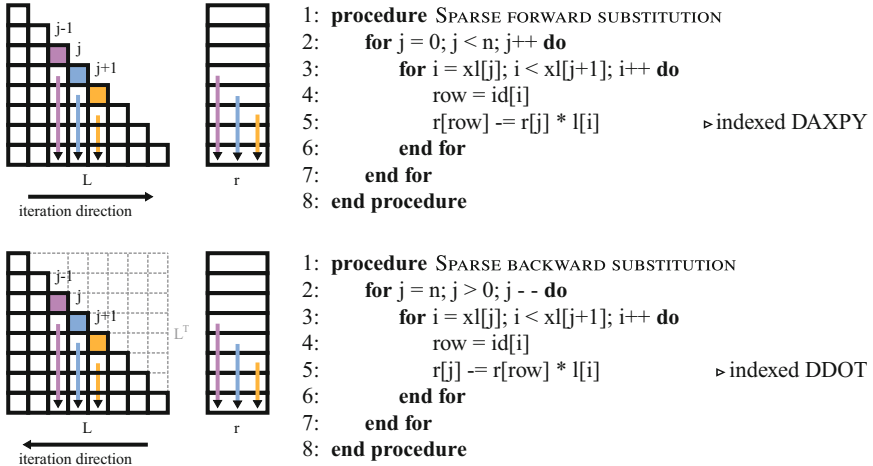


Fig. 12 Sparse triangular substitution in CSC format based on indexed DAXPY/DDOT kernel operations

$$r_i = r_i - r_j L_{ij}. \quad (8)$$


The backward substitution with L^T will take place row wise, since we use L and perform the substitution column wise with respect to L , as shown in the lower part of Fig. 12. In contrast to the forward substitution the iteration over columns starts at the last column N and proceeds to the first one. If column j is reached, then r_j , which contains the j -component of the solution vector x_j , is computed by subtracting the dot-product of the remaining elements in the column L_{ij} and the corresponding elements of r_i with $i = j + 1, \dots, N$ from it:

$$r_j = r_j - r_i L_{ij}. \quad (9)$$

After all columns have been processed r contains the required solution x . It is important to note that line 5 represents in both substitutions an indexed DAXPY and indexed DDOT kernel operations that has to be computed during the streaming operations of the vector r and the column j of the numerical factor L . As we are dealing with sparse matrices it makes no sense to store the lower triangular matrix L as a dense matrix. Hence, PARDISO uses its own data structure to store L , as shown in Fig. 13.

Adjacent columns exhibiting the same row sparsity structure form a *panel*, also known as *supernode*. A panel's column count is called the *panel size* n_p . The columns of a panel are stored consecutively in memory excluding the zero entries. Note that columns of panels are padded in the front with zeros so they get the same length as the first column inside their panel. The padding is of utmost performance for the PARDISO solver to use Level-3 BLAS and LAPACK functionalities [41]. Furthermore panels are stored consecutively in the l array. Row and column

Fig. 13 Sparse matrix data structures in PARDISO. Adjacent columns of L exhibiting the same structure form panels also known as supernodes. Groups of panels which touch independent elements of the right-hand side r are parts. The last part in the lower triangular matrix L is called separator



information is now stored in accompanying arrays. The x_{super} array stores for each panel the index of its first column. Also note that here column indices are the running count of nonzero columns. Column indices are used as indices into x_l array to lookup the start of the column in the l array which contains the numerical values of the factor L . To determine the row index of a column's element an additional array id is used, which holds for each panel the row indices. The start of a panel inside id is found via x_{id} array. The first row index of panel p is $id[x_{id}[p]]$. For serial execution this information is enough. However, during parallel forward/backward substitution concurrent updates to the same entry of r must be avoided. The $parts$ structure contains the start (and end) indices of the panels which can be updated independently as they do not touch the same entries of r . Two parts, colored blue and orange, are shown in Fig. 13. The last part in the bottom right corner of L is special and is called the *separator* and is colored green. Parts which would touch entries of r in the range of the separator perform their updates into separate temporary arrays t . Before the separator is then serially updated, the results of the temporary arrays are gathered back into r . The backward substitution works the same, just reversed and only updates to different temporary arrays are not required. The complete forward substitution and backward substitution is listed in Algorithms 1 and 2.

6 Application—Circuit Simulation

In this section we demonstrate how these developments in sparse direct linear solvers have advanced integrated circuit simulations. Integrated circuits are composed of interconnected transistors. The interconnects are modeled primarily with

Algorithm 1 Forward substitution in PARDISO. Note that in case of serial execution separated updates to temporary arrays in Lines 10–13 are not necessary and can be handled via the loop in Lines 6–9

```

1: procedure FORWARD
2:   for part o in parts do                                     ▷ parallel execution
3:     for panel p in part p do
4:       for column j in panel do                               ▷ unroll
5:         i = xid[p] + offset
6:         for k = xl[j] + offset; k < sep; ++k do
7:           row = id[i++]
8:           r[row] -= r[j] l[k]                               ▷ indexed DAXPY
9:         end for
10:        for k = sep + 1; k < xl[j+1]; ++k do
11:          row = id[i++]
12:          t[row,p] -= r[j] l[k]                             ▷ indexed DAXPY
13:        end for
14:      end for
15:    end for
16:  end for
17:  r[i] = r[i] - sum(t[i,:])                                  ▷ gather temporary arrays
18:  for panel p in separator do                                ▷ serial execution
19:    for column j in panel do                                   ▷ unroll
20:      i = xid[p] + offset
21:      for k = xl[j] + offset; k < xl[j+1]; ++k do
22:        row = id[i++]
23:        r[row] -= r[j] l[k]                               ▷ indexed DAXPY
24:      end for
25:    end for
26:  end for
27: end procedure

```

resistors, capacitors, and inductors. The interconnects route signals through the circuit, and also deliver power. Circuit equations arise out of Kirchhoff's current law, applied at each node, and are generally nonlinear differential-algebraic equations. In transient simulation of the circuit, the differential portion is handled by discretizing the time derivative of the node charge by an implicit integration formula. The associated set of nonlinear equations is handled through use of quasi-Newton methods or continuation methods, which change the nonlinear problem into a series of linear algebraic solutions. Each component in the circuit contributes only to a few equations. Hence, the resulting systems of linear algebraic equations are extremely sparse, and most reliably solved by using direct sparse matrix techniques. Circuit simulation matrices are peculiar in the universe of matrices, having the following characteristics [11]:

- they are nonsymmetric, although often nearly structurally symmetric;
- they have a few dense rows and columns (e.g., power and ground connections);
- they are *very* sparse and the straightforward usage of BLAS routines (as in SuperLU[32]) may be ineffective;

Algorithm 2 Backward substitution in PARDISO. Separator (sep.), parts, and panels are iterated over in reversed (rev.) order

```

1: procedure BACKWARD
2:   for panel  $p$  in sep. rev. do                                     ▷ serial execution
3:     for col.  $j$  in panel  $p$  rev. do                                   ▷ unroll
4:        $i = \text{xl}[p] + \text{offset}$ 
5:       for  $k = \text{xl}[j] + \text{offset}; k < \text{xl}[j+1]; ++k$  do
6:          $\text{row} = \text{id}[i++]$ 
7:          $r[j] -= r[\text{row}] l[k]$                                        ▷ indexed DDOT
8:       end for
9:        $\text{offset} = \text{offset} - 1$ 
10:    end for
11:  end for
12:  for part in parts do                                             ▷ parallel execution
13:    for panel  $p$  in part rev. do
14:      for col.  $j$  in panel  $p$  rev. do                               ▷ unroll
15:         $i = \text{xl}[p] + \text{offset}$ 
16:        for  $k = \text{xl}[j] + \text{offset}; k < \text{xl}[j+1]; ++k$  do
17:           $\text{row} = \text{id}[i++]$ 
18:           $r[j] -= r[\text{row}] l[k]$                                        ▷ indexed DDOT
19:        end for
20:         $\text{offset} = \text{offset} - 1$ 
21:      end for
22:    end for
23:  end for
24: end procedure

```

- their LU factors remain sparse if well-ordered;
- they can have high fill-in if ordered with typical strategies;
- and being unstructured, the highly irregular memory access causes factorization to proceed only at a few percent of the peak flop-rate.

Circuit simulation matrices also vary from being positive definite to being *extremely* ill-conditioned, making pivoting for stability important also. As circuit size increases, and depending on how much of the interconnect is modeled, sparse matrix factorization is the dominant cost in the transient analysis.

To overcome the complexity of matrix factorization a new class of simulators arose in the 1990s, called fast-SPICE [39]. These simulators partition the circuit into subcircuits and use a variety of techniques, including model order reduction and multirate integration, to overcome the matrix bottleneck. However, the resulting simulation methods generally incur unacceptable errors for analog and tightly coupled circuits. As accuracy demands increase, these techniques become much slower than traditional SPICE methods. Even so, since much of the research effort was directed at fast-SPICE simulators, it brought some relief from impossibly slow simulations when some accuracy trade-off was acceptable. Because these simulators partitioned the circuit, and did not require the simultaneous solution of the entire system of linear equations at any given time, they did not push the state of the art in sparse matrix solvers.

Starting in the mid-2000s, increasing demands on accuracy, due to advancing semiconductor technology, brought attention back to traditional SPICE techniques. This was aided by the proliferation of multicore CPUs. Parallel circuit simulation, an area of much research focus in the 1980s and 1990s, but not particularly in practice, received renewed interest as a way to speed up simulation without sacrificing accuracy. Along with improved implementations to avoid cache misses, rearchitecture of code for parallel computing, and better techniques for exploitation of circuit latency, improved sparse matrix solvers, most notably the release of KLU [11], played a crucial role in expanding the utility of SPICE.

Along with the ability to simulate ever larger circuits with full SPICE accuracy came the opportunity to further improve sparse matrix techniques. A sparse matrix package for transient simulation needs to have the following features:

- must be parallel;
- fast matrix reordering;
- incremental update of the L and U factors when only a few nonzeros change;
- fast computation of the diagonal entries of the inverse matrix;
- fast computation of Schur-complements for a submatrix;
- allow for multiple LU factors of the same structure to be stored;
- use the best-in-class method across the spectrum of sparsity;
- use iterative solvers with fast construction of sparse preconditioners;
- run on various hardware platforms (e.g., GPU acceleration).

Some of these features must be available in a single package. Others, such as iterative solvers and construction of preconditioners, can be implemented with a combination of different packages. The PARDISO solver¹ stands out as a package that does most of these very well. Here we touch on a few of these features.

When applied in the simulation of very large circuits, the difference between a “good” and a “bad” matrix ordering can be the difference between seconds and days. PARDISO offers AMD and nested dissection methods for matrix ordering, as well as permitting user-defined ordering. Because the matrix reordering method which has been used most often in circuit simulation is due to Markowitz [35], and because modern sparse matrix packages do not include this ordering method, we briefly describe it here. The Markowitz method is quite well-adapted for circuit simulation. Some desirable aspects of the typical implementation of the Markowitz method, as opposed to the MD variants, are that it works for nonsymmetric matrices and combines pivot choice with numerical decomposition, such that a pivot choice is a numerically “good” pivot which generates in a local sense the least fill-in at that step of the decomposition. Choosing pivots based on the Markowitz score often produces very good results: near-minimal fill-in, unfortunately at the cost of an $O(n^3)$ algorithm (for dense blocks). Even though the Markowitz algorithm has some good properties when applied to circuit matrices, the complexity of the algorithm has become quite burdensome. When SPICE [36] was originally conceived, a hundred-

¹The PARDISO solver is available from <http://www.pardiso-project.org>.




Fig. 14 Performance improvements of PARDISO 6.2 against Intel MKL PARDISO (one thread) for various circuit simulation matrices

node circuit was huge and the Markowitz algorithm was not a problem. Now we routinely see netlists with hundreds of thousands of nodes and post-layout netlists with millions of elements. As matrix order and element counts increase, Markowitz reordering time can become an obstruction. Even as improved implementations of the Markowitz method have extended its reach, AMD and nested dissection have become the mainstay of simulation of large denser-than-usual matrices.

Next we turn our attention to parallel performance. While KLU remains a benchmark for serial solvers, for parallel solvers, MKL-PARDISO is often cited as the benchmark [6, 8]. To give the reader a sense of the progress in parallel sparse matrix methods, in Fig. 14 we compare KLU, PARDISO (Version 6.2) to MKL-PARDISO on up to 16 cores on an Intel Xeon E7-4880 architecture with 2.5 GHz processors.

Some of the matrices here can be obtained from the SuiteSparse Matrix Collection, and arise in transistor level full-chip and memory array simulations. It is clear that implementation of sparse matrix solvers has improved significantly over the years.

Exploiting latency in all parts of the SPICE algorithm is very important in enabling accurate circuit simulation, especially as the circuit size increases. By latency, we mean that only a few entries in the matrix change from one Newton iteration to the next, and from one timepoint to the next. As the matrix depends on the time-step, some simulators hold the time-steps constant as much as feasible to allow increased reuse of matrix factorizations. The nonzero entries of a matrix

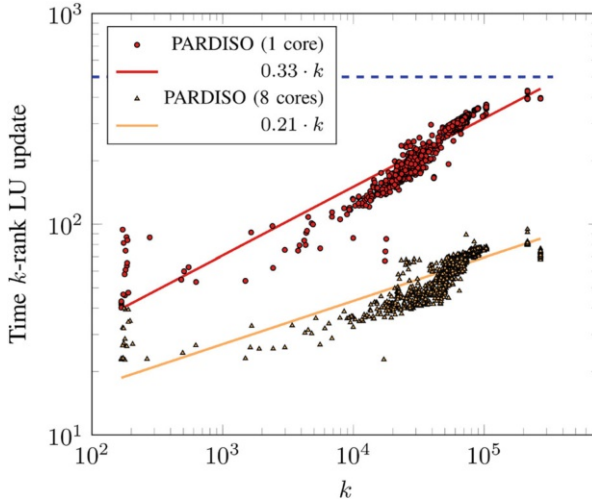


Fig. 15 Regression analysis on the rank- k update LU factorization in PARDISO

change only when the transistors and other nonlinear devices change their operation point. In most circuits, very few devices change state from one iteration to the next and from one time-step to the next. Nonzeros contributed by entirely linear components do not change value during the simulation. This makes incremental LU factorization a very useful feature of any matrix solver used in circuit simulation. As of April 2019 the version PARDISO 6.2 has a very efficient exploitation of incremental LU factorization, both serial and parallel. In Fig. 15 we show that PARDISO scales linearly with number of updated columns, and also scales well with number of cores. Here, the series of matrices were obtained from a full simulation of a post-layout circuit that includes all interconnects, power and ground networks. The factorization time is plotted against the number of columns that changed compared to the previous factorization. The scatter plot shows the number of rank- k update and the corresponding factorization time in milliseconds. The regression analysis clearly demonstrates a linear trend both for the single and the multiple core versions. The dashed line shows the time for the full factorization.

Another recent useful feature in PARDISO is parallel selective inverse matrix computation as demonstrated in Table 1. In circuit simulation, the diagonal of the inverse matrix is the driving point impedance. It is often required to flag nodes in the circuit with very high driving point impedance. Such nodes would indicate failed interfaces between different subcircuits, leading to undefined state and high current leakage and power dissipation. A naive approach to this is to solve for the driving point impedance, the diagonal of the inverse matrix, by N triangular solves. This is sometimes unacceptably expensive even with exploiting the sparsity of the right-hand side, and minimizing the number of entries needed in the diagonal of the inverse. To bypass this complexity, heuristics to compute the impedance of

Table 1 Details of the benchmark matrices

Matrix	N	$\text{nnz}(A)$	$\text{nnz}(\frac{L+U}{A})$	A^{-1}	Selected A^{-1}
circuit5M_DC	3,523,317	19,194,193	2.87	82.3 h	1.3 s
circuit5M	5,558,326	59,524,291	1.04	371.1 h	2.1 s
Freescall	3,428,755	18,920,347	2.94	89.8 h	1.0 s
Freescall2	2,999,349	23,042,677	2.92	8.5 h	1.2 s
FullChip	2,987,012	26,621,990	7.41	162.9 h	11.9 s
memchip	2,707,524	14,810,202	4.40	62.5 h	0.9 s

“ N ” is the number of matrix rows and “ nnz ” is the number of nonzeros. The table shows the fill-in factor related to the number of nonzeros in $\frac{L+U}{A}$, the time for computing all diagonal elements of the inverse A^{-1} using N multiple forward/backward substitution in hours, and using the selected inverse method in PARDISO for computing all diagonal elements of the inverse A^{-1} in seconds.

connected components are used. But this is error prone with many false positives and also false negatives. In the circuit Freescall, PARDISO, e.g., finished the required impedance calculations in 11.9 s compared to the traditional computation that consumed 162.9 h.

The productivity gap in simulation continues to grow, and challenges remain. Signoff simulations demand $10\times$ speedup in sparse matrix factorization. Simply using more cores does not help unless the matrices are very large and complex. For a majority of simulations, scaling beyond eight cores is difficult. As a result, some of these simulations can take a few months to complete, making them essentially impossible. Some of the problems in parallelizing sparse matrix operations for circuit simulation are fundamental. Others may be related to implementation. Research on sparse matrix factorization for circuit simulation continues to draw attention, especially in the area of acceleration with Intel’s many integrated core (MIC) architecture [6] and GPUs [7, 31]. Other techniques for acceleration include improved preconditioners for iterative solvers [44]. We are presently addressing the need for runtime selection of optimal strategies for factorization, and also GPU acceleration. Given that circuits present a wide spectrum of matrices, no matter how we categorize them, it is possible to obtain a solver that is $2\text{--}10\times$ better on a given problem. Improvements in parallel sparse matrix factorization targeted at circuit simulation is more necessary today than ever and will continue to drive applicability of traditional SPICE simulation methods. Availability of sparse matrix packages such as PARDISO that completely satisfy the needs of various circuit simulation methods is necessary for continued performance gains.

References

1. A. Aho, J. Hopcroft, and J. Ullman. *Data structures and algorithms*. Addison-Wesley, 1983.
2. P. Amestoy, T. A. Davis, and I. S. Duff. An approximate minimum degree ordering algorithm. *SIAM J. Matrix Anal. Appl.*, 17(4):886–905, 1996.

3. E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide, Second Edition*. SIAM Publications, 1995.
4. M. Benzi, J. Haws, and M. Tuma. Preconditioning highly indefinite and nonsymmetric matrices. *SIAM J. Sci. Comput.*, 22(4):1333–1353, 2000.
5. C. Berge. Two theorems in graph theory. In *Proceedings of National Academy of Science*, pages 842–844, USA, 1957.
6. J. D. Booth, N. D. Ellingwoodb, and S. R. Heidi K. Thornquist. Basker: Parallel sparse LU factorization utilizing hierarchical parallelism and data layouts. *Parallel Computing*, 68:17–31, 2017.
7. X. Chen, L. Ren, Y. Wang, and H. Yang. GPU-accelerated sparse LU factorization for circuit simulation with performance modeling. *IEEE Transactions on Parallel and Distributed Systems*, 26:786–795, 2015.
8. X. Chen, Y. Wang, and H. Yang. NICSLU: An adaptive sparse matrix solver for parallel circuit simulation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 32:261–274, 2013.
9. C. Chevalier and F. Pellegrini. PT-SCOTCH: a tool for efficient parallel graph ordering. *Parallel Comput.*, 34(6–8):318–331, 2008.
10. E. Cuthill and J. McKee. Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 24th national conference of the ACM*. ACM, 1969.
11. T. Davis and K. Stanley. Sparse LU factorization of circuit simulation matrices. In *Numerical Aspects of Circuit and Device Modeling Workshop*, June 2004.
12. T. A. Davis. *Direct Methods for Sparse Linear Systems*. SIAM Publications, 2006.
13. D. Dodson and J. G. Lewis. Issues relating to extension of the basic linear algebra subprograms. *ACM SIGNUM Newslett.*, 20:2–18, 1985.
14. J. J. Dongarra, J. Du Croz, S. Hammarling, and R. J. Hanson. Issues relating to extension of the basic linear algebra subprograms. *ACM SIGNUM Newslett.*, 20:2–18, 1985.
15. I. S. Duff, A. Erisman, and J. Reid. *Direct Methods for Sparse Matrices*. Oxford University Press, 1986.
16. I. S. Duff and J. Koster. The design and use of algorithms for permuting large entries to the diagonal of sparse matrices. *SIAM J. Matrix Anal. Appl.*, 20:889–901, 1999.
17. I. S. Duff and S. Pralet. Strategies for scaling and pivoting for sparse symmetric indefinite problems. Technical Report TR/PA/04/59, CERFACS, Toulouse, France, 2004.
18. C. M. Fiduccia and R. M. Mattheyses. A linear-time heuristic for improving network partitions. In *Proceedings of the 19th Design Automation Conference*, pages 175–181. IEEE, 1997.
19. M. Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory. *Czechoslovak Mathematical Journal*, 25(100):619–633, 1975.
20. A. George and J. W. H. Liu. The evolution of the minimum degree ordering algorithm. *SIAM Review*, 31:1–19, 1989.
21. J. A. George and J. W. Liu. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1981.
22. J. A. George and E. Ng. An implementation of Gaussian elimination with partial pivoting for sparse systems. *SIAM J. Sci. Statist. Comput.*, 6(2):390–409, 1985.
23. J. R. Gilbert. Predicting structure in sparse matrix computations. *SIAM J. Matrix Anal. Appl.*, 15(1):162–79, 1994.
24. A. Gupta and L. Ying. On algorithms for finding maximum matchings in bipartite graphs. Technical Report RC 21576 (97320), IBM T. J. Watson Research Center, Yorktown Heights, NY, October 25, 1999.
25. G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1998.
26. G. Karypis, K. Schloegel, and V. Kumar. *ParMeTis: Parallel Graph Partitioning and Sparse Matrix Ordering Library, Version 2.0*. University of Minnesota, Dept. of Computer Science, September 1999.

27. B. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *The Bell System Technical Journal*, 29(2):291–307, 1970.
28. J. Langguth, A. Azad, and F. Manne. On parallel push-relabel based algorithms for bipartite maximum matching. *Parallel Comput.*, 40(7):289–308, 2014.
29. J. Langguth, M. M. A. Patwary, and F. Manne. Parallel algorithms for bipartite matching problems on distributed memory computers. *Parallel Comput.*, 37(12):820–845, 2011.
30. D. LaSalle and G. Karypis. Multi-threaded graph partitioning. Technical report, Department of Computer Science & Engineering, University of Minnesota, Minneapolis, 2013.
31. W.-K. Lee, R. Achar, and M. S. Nakhla. Dynamic GPU parallel sparse LU factorization for fast circuit simulation. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, pages 1–12, 2018.
32. X. S. Li. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM Trans. Math. Softw.*, 31(3):302–325, 2005.
33. J. W. Liu. The role of elimination trees in sparse factorization. *SIAM J. Matrix Anal. Appl.*, 11(1):134–172, 1990.
34. J. W. Liu and A. Sherman. Comparative analysis of the Cuthill–McKee and the reverse Cuthill–McKee ordering algorithms for sparse matrices. *SIAM J. Numer. Anal.*, 13:198–213, 1976.
35. H. M. Markowitz. The elimination form of the inverse and its application to linear programming. *Management Science*, 3:255–269, April 1957.
36. L. W. Nagel. SPICE2: a computer program to simulate semiconductor circuits. Memorandum No. ERL-M520, University of California, Berkeley, California, May 1975.
37. M. Olschowka and A. Neumaier. A new pivoting strategy for Gaussian elimination. *Linear Algebra and its Applications*, 240:131–151, 1996.
38. F. Pellegrini, J. Roman, and P. Amestoy. Hybridizing nested dissection and halo approximate minimum degree for efficient sparse matrix ordering. *Concurrency: Practice and Experience*, 12:69–84, 2000.
39. M. Rewienski. A perspective on fast-spice simulation technology. Springer, Dordrecht, 2011.
40. D. J. Rose. A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations. In R. C. Read, editor, *Graph Theory and Computing*. Academic Press, 1972.
41. O. Schenk. *Scalable parallel sparse LU factorization methods on shared memory multiprocessors*. PhD thesis, ETH Zurich, 2000. Diss. Technische Wissenschaften ETH Zurich, Nr. 13515, 2000.
42. O. Schenk and K. Gärtner. Solving unsymmetric sparse systems of linear equations with PARDISO. *Journal of Future Generation Computer Systems*, 20(3):475–487, 2004.
43. R. E. Tarjan. Data structures and network algorithms. In *CBMS–NSF Regional Conference Series in Applied Mathematics*, volume 44, 1983.
44. X. Zhao, L. Han, and Z. Feng. A performance-guided graph sparsification approach to scalable and robust spice-accurate integrated circuit simulations. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34:1639–1651, 2015.