

GPU-resident sparse direct linear solvers for alternating current optimal power flow analysis

Kasia Świrydowicz^a, Nicholson Koukpaizan^b, Tobias Ribizel^c, Fritz Göbel^c, Shrirang Abhyankar^a, Hartwig Anzt^{d,c}, Slaven Peleš^{b,*}

^a Pacific Northwest National Laboratory, Richland, 99352, WA, USA

^b Oak Ridge National Laboratory, 1 Bethel Valley Rd, Oak Ridge, 37830, TN, USA

^c Karlsruhe Institute of Technology, Kaiserstraße 12, Karlsruhe, 76131, BW, Germany

^d University of Tennessee, 203 Claxton Complex, Knoxville, 37996, TN, USA

ARTICLE INFO

MSC:

65F05

65F10

65F50

65K10

65Y05

65Y10

90C51

Keywords:

ACOPF

Economic dispatch

Optimization

Linear solver

GPU

ABSTRACT

Integrating renewable resources within the transmission grid at a wide scale poses significant challenges for economic dispatch as it requires analysis with more optimization parameters, constraints, and sources of uncertainty. This motivates the investigation of more efficient computational methods, especially those for solving the underlying linear systems, which typically take more than half of the overall computation time. In this paper, we present our work on sparse linear solvers that take advantage of hardware accelerators, such as graphical processing units (GPUs), and improve the overall performance when used within economic dispatch computations. We treat the problems as sparse, which allows for faster execution but also makes the implementation of numerical methods more challenging. We present the first GPU-native sparse direct solver that can execute on both AMD and NVIDIA GPUs. We demonstrate significant performance improvements when using high-performance linear solvers within alternating current optimal power flow (ACOPF) analysis. Furthermore, we demonstrate the feasibility of getting significant performance improvements by executing the entire computation on GPU-based hardware. Finally, we identify outstanding research issues and opportunities for even better utilization of heterogeneous systems, including those equipped with GPUs.

1. Introduction

Power grid utilities heavily rely on simulation and analysis tools for near- and long-term planning. With increased energy needs, more variable generation added to the grid, less predictable weather patterns, and ever-increasing cyber threats, the electric grid planning and operation have become more complex and compute-intensive. In particular, the ACOPF [1,2] analysis, which is at the core of economic dispatch, has become computationally more demanding, pushing the limits of existing tools. New heterogeneous computing technologies utilizing GPUs, provide affordable high-density computational power that can address emerging needs in the energy industry. However, mathematical and computational methods typically used for economic dispatch are not designed for and perform poorly on GPUs. This is particularly true when solving underlying linear systems [3], which typically makes up more than half of the overall computational cost.

GPUs deliver high performance by executing groups of threads (warps) within single instruction sequence control. In NVIDIA GPUs, 32 threads form a warp. In AMD GPUs, 64 threads form a wavefront —

which is equivalent to a warp. Standard CPUs have separate instruction sequence control for each thread. This means that the computational power of a GPU can be harnessed only if the computation is expressed in terms of single-instruction multiple-data (SIMD) operations. Hardware accelerators also have smaller level-1 cache memory compared to CPUs, which means that poor data coalescence in device memory can severely degrade computational performance. This also means that moving data in hardware accelerator memory is more expensive than moving data in conventional CPU memory. These constraints need to be factored in when developing mathematical algorithms for execution on heterogeneous hardware.

There have been several efforts to develop GPU-accelerated sparse linear solvers that are effective for computations in the power systems domain, including electromagnetic transient [4,5] and power flow [6] simulations. There is far less reporting on linear solvers suitable for ACOPF analysis due to the special properties of linear systems there. Existing results pertain mainly to dense formulations of the underlying linear systems [7,8]. However, the computational complexity of dense

* Corresponding author.

E-mail address: peless@ornl.gov (S. Peleš).

linear solvers is $O(N^3)$ and their memory complexity is $O(N^2)$, where N is the number of linear equations. For larger grid models, all performance gains on GPU are offset by a cubic increase in computational cost and a quadratic increase in memory as shown in, e.g., [9].

In this paper, we discuss the solution of ACOPF on GPUs using sparse linear solvers. The main contributions of this paper are:

- Two refactorization-based sparse linear solvers — one developed using CUDA libraries [10] and the other developed within open-source GINKGO library [11]. Both outperform existing GPU-enabled solvers by a wide margin when used within ACOPF analysis.
- Demonstration of meaningful speedup when our GPU solvers are used in ACOPF compared to the state-of-the-art baseline on CPU hardware. We observed matrix factorization speedup of 2.4–3.4× and triangular solve speedup of 3.2–5.8×, which projects to 1.6–1.9× overall speedup for the entire ACOPF analysis compared to the CPU baseline.
- Performance profiles of ACOPF analysis on NVIDIA and AMD GPUs, as well as on CPU platforms. We identify key performance bottlenecks and estimate additional performance improvement attainable with currently existing technologies.

The paper is organized as follows: In Section 2, we describe properties of linear systems arising in ACOPF analysis. In Section 3, we describe the environment and use cases for testing the performance of linear solvers, and establish the state-of-the-art baseline. We discuss challenges when solving those linear systems on hardware accelerators in Section 4, and present our results in Section 5. Finally, in Section 6, we summarize our findings and discuss the follow-on research that will lead to further performance improvements on heterogeneous hardware.

2. Problem definition

We test our linear solvers within an ACOPF analysis, which is representative of optimization computations used in economic dispatch. Mathematically, an ACOPF is posed as a nonlinear non-convex optimization problem

$$\min_x F(x) \quad (1a)$$

$$\text{s.t. } g(x) = 0, \quad (1b)$$

$$h^- \leq h(x) \leq h^+, \quad (1c)$$

$$x^- \leq x \leq x^+, \quad (1d)$$

where x is a vector of optimization variables, such as bus voltages and generator power outputs; vectors x^-, x^+ define variable bounds, e.g. resource capacity limits; $F(x)$ is a scalar objective function accounting for generator fuel, power imbalance, wind curtailment, and other power generation costs; $g(x)$ is a vector function that defines equality constraints, such as power balance; vector function $h(x)$ defines security constraints and vectors h^-, h^+ specify security limits. For a more comprehensive discussion of the model aspects of the AC optimal power flow formulation, we refer the reader to [1,12].

We use interior method [13] to solve (1). To keep the presentation streamlined, we recast (1), without loss of generality, in a compact form

$$\min_{y \in \mathbb{R}^n} f(y) \quad (2a)$$

$$\text{s.t. } c(y) = 0, \quad (2b)$$

$$y \geq 0, \quad (2c)$$

where inequality constraints (1c) are expressed as equality constraints $h(x) - h^\pm \pm s^\pm = 0$ using slack variables $s^-, s^+ \geq 0$. Similarly, bounds (1d) are expressed in form (2c) by introducing variables $u^-, u^+ \geq 0$ such that $u^\pm = \mp(x - x^\pm)$. The vector of optimization primal variables is

$y = (u^-, u^+, s^-, s^+) \in \mathbb{R}^n$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a possibly non-convex objective function. All constraints are formulated in $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$, where m is the number of equality and inequality constraints.

Interior methods enforce bound constraints (2c) by adding barrier functions to the objective (2a)

$$\min_{y \in \mathbb{R}^n} f(y) - \mu \sum_{j=1}^n \ln y_j,$$

where $\mu > 0$ is the barrier parameter. Interior methods are most effective when exact first and second derivatives are available, as we assume for $f(y)$ and $c(y)$. We define $J(y) = \nabla c(y)$ as the sparse $m \times n$ Jacobian matrix for the constraints. The solution of a barrier subproblem satisfies the nonlinear equations

$$\nabla f(y) + J(y)\lambda - z = 0, \quad (3a)$$

$$c(y) = 0, \quad (3b)$$

$$Yz = \mu e, \quad (3c)$$

where $\lambda \in \mathbb{R}^m$ is a vector of Lagrange multipliers (dual variables) for constraints (2b), z are Lagrange multipliers for the bounds (2c), the matrix $Y \equiv \text{diag}(y)$, and e is n -dimensional vector with all elements equal to one.

The barrier subproblem (3) is solved using a variant of a Newton method, and the solution to the original system (1) is recovered by a continuation method driving $\mu \rightarrow 0$. As it searches for the optimal solution, the interior method solves a sequence of the linearized barrier subproblems $K_k \Delta x_k = r_k$, $k = 1, 2, \dots$, where index k denotes optimization solver iteration (including continuation step in μ and Newton iterations). Note that for $\mu = 0$, K_k is singular. A typical implementation of the interior method eliminates the linearized version of (3c) by substituting it in the linearized version of (3a) to obtain a symmetric linear system of Karush-Kuhn–Tucker (KKT) type [13]

$$\begin{bmatrix} \overbrace{H + D_y}^{K_k} & J \\ J^T & 0 \end{bmatrix} \begin{bmatrix} \Delta y \\ \Delta \lambda \end{bmatrix} = \begin{bmatrix} r_y \\ r_\lambda \end{bmatrix}, \quad (4)$$

where Δx_k is a search direction and r_k is residual of (3) evaluated at current values of primal and dual variables. The Hessian

$$H \equiv \nabla^2 f(y) + \sum_{i=1}^m \lambda_i \nabla^2 c_i(y),$$

is a sparse symmetric $n \times n$ matrix and $D_y \equiv \mu Y^{-2}$ is a diagonal $n \times n$ matrix.

The KKT linear system (4) is sparse symmetric indefinite and typically *ill-conditioned*. In ACOPF and other analyses for economic dispatch, K_k matrices are *extremely sparse* due to the connectivity structure of power grids. The interior method needs to return a solution after μ is small enough so the solution of (3) approximates the solution of (2) well, but before K_k in (4) becomes too ill-conditioned. Having a solver that can compute an accurate solution to ill-conditioned linear systems without significant performance penalty is critical for accurate and efficient ACOPF analysis, especially since solving (4) is the major part of the overall computational cost.

We note that all K_k in (4) have the same sparsity pattern. An efficient linear solver shall take advantage of constant sparsity patterns and reuse parts of computations over the sequence of K_k matrices when possible.

3. Test networks and optimization setup

As our test cases, we use synthetic grid models: Northeastern, Eastern, and combined Western and Eastern U.S. grid [14]. The details of our test cases are provided in Table 1.

To establish a baseline for evaluating our new solvers, we use MA57 linear solver [15], which is commonly used in commercial and open-source tools when solving ACOPF and similar optimization problems.

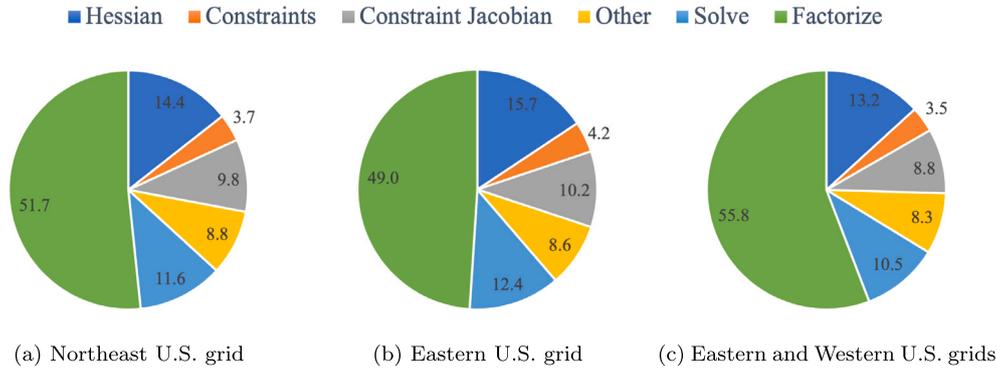


Fig. 1. Computational cost of ACOPF broken down by analysis functions. Linear solver functions (matrix factorization and triangular solve) contribute to 60% of the overall ACOPF compute time.

Table 1

Characteristics of the three test networks specifying the number of buses, generators, and lines. The specifics of the linear system (4) for each of these networks are given in terms of the matrix size (N) and number of non-zeros (nnz) for the matrix K_k . Numbers are rounded to 3 digits. K and M denote 10^3 and 10^6 , respectively.

Grid	Buses	Generators	Lines	$N(K_k)$	$nnz(K_k)$
Northeastern US	25 K	4.8 K	32.3 K	108 K	1.19 M
Eastern US	70 K	10.4 K	88.2 K	296 K	3.20 M
Western and Eastern US	82 K	13.4 K	104.1 K	340 K	3.73 M

When using MA57, symbolic factorization is performed only once for all systems with the same sparsity pattern. The profiling results for the solution of the KKT system for these three networks with MA57 linear solver show that solving the linear system takes up more than 60% of the overall ACOPF cost (see Fig. 1). These results were obtained on an IBM Power 9 CPU; the computing environment is described in more detail in Section 5.3. Starting from this baseline, any meaningful speedup of ACOPF requires accelerating the linear solver.

For the ACOPF analysis we use ExaGOTM [16], an open-source package for solving large-scale optimization problems on CPU and GPU-based platforms involving security, stochastic, and multi-period constraints. ExaGOTM models the ACOPF problem in a power-balance network formulation with a minimum generation cost objective function and enforcing voltage, capacity, and line flow constraints. It uses the open-source HiOp optimization engine [17] to solve the resultant optimization problem. To our knowledge, HiOp is the only open-source optimization solver capable of solving nonlinear optimization problems on GPUs. The developed linear solvers described in this paper are interfaced to and called from HiOp. The ACOPF application provided with ExaGOTM is used to configure and run our tests. ExaGOTM parses input data files for power grids, generates model objects H , J , D_y , r_y , and r_λ , sets variable limits, and passes them to HiOp. The optimization engine assembles (4), updates it at each iteration, and passes it to the linear solver. The solution Δx_k to (4) is used by HiOp to update the vector of primal variables in ExaGOTM for the next iteration. In our computations we set relative tolerance for the interior method in HiOp to 10^{-6} . This leads to barrier parameter values reduced to $\mu \sim 10^{-7}$ before converged solution is obtained.

A typical implementation of the interior method requires a linear solver to provide matrix inertia (number of positive, negative, and zero matrix eigenvalues). Matrix inertia can be readily obtained in LDL^T , but not in LU factorization. It is important to note that HiOp offers an inertia-free interior method implementation [18], as an option, which allows us to use LU linear solvers for ACOPF analysis.

4. GPU-enabled linear solvers

There exists a vast body of literature on direct linear solvers for heterogeneous compute architectures [19,20]. Moreover, the recent

decade brought increased interest in developing solvers that are capable of using the GPUs efficiently [21]. These efforts, however, usually target *generic* problems that are well-conditioned and characterized by a well-defined structure (i.e., matrices that contain dense blocks of entries). The available parallel techniques rely on multifrontal or supernodal approaches; in both cases, dense blocks within the matrix are essential for satisfactory performance on the GPUs. In [3] we identified and tested five applicable linear solver packages: SuperLU [22], STRUMPACK [20], SPRAL-SSIDS [23], PaStiX [24] and cuSolverSP (we tested the *black-box* QR and LU-based solve functions). The results were compared to the MA57 [15], which is a single thread and CPU-only code. The testing revealed that none of the GPU-accelerated packages was substantially better than MA57. Moreover, turning on GPU acceleration within each package often resulted in performance losses when compared to the CPU-only code. The timing data published in [3] is summarized in Table 2. The presented timing results are averaged over a few representative matrix systems for each solver.

The development of GPU-resident solvers for power flow problems is challenging for several reasons. First of all, it is the nature of sparse direct solvers that allows only for a limited amount of parallelism. The situation becomes worse if the sparse matrix has no inherent block structure, which is the case of power flow problems. As mentioned earlier in this section, in this case, supervariable agglomeration and multifrontal approaches that accumulate elements in dense blocks and invoke dense linear algebra routines are not suitable (which explains the results from [3]). Instead, we require fine-grained scheduling of individual variable eliminations once their dependencies have completed processing. Finally, the power grid problems require pivoting for numerical stability. Pivoting is a bottleneck on GPUs as it requires synchronization and inter-block communication, further degrading the performance.

From a high-level standpoint, sparse direct solver computations can be split into three phases: (i) symbolic factorization, when the matrix is reordered and the structure of L and U factors are set, together with permutation order; (ii) numeric phase, when numeric factorization is computed; and (iii) solve phase, during which the factors obtained in the previous phase are used to compute the solution.

The symbolic phase of the computation is typically done on the CPU, sometimes with marginal GPU offloading. The symbolic factorization does rely on the non-zero structure of the matrix and not on its values. In the case of the KKT matrices, as mentioned before, the non-zero structure of the matrix systems does not change during the optimization solver run, and hence the symbolic phase (or parts of it) can be executed once and reused for all relevant systems. This idea ties in with the idea of *refactorization*. This term is used somewhat loosely in the literature; it usually means that once the permutation vectors and non-zero patterns of triangular factors were computed in the *numerical* factorization, they are reused for the next system in the sequence. Refactorization is implemented in some of the existing direct solvers, for instance, in

Table 2

Test results with various GPU-accelerated linear solvers for representative systems generated with Ipopt [3]. The test linear systems are available at https://github.com/NREL/opf_matrices. We used the last 3 systems for test case 1 (ACTIVSg2000 in the repository) and the last 5 for the two remaining test cases (ACTIVSg10k and ACTIVSg70k in the repository).

Test case	Size	NNZ		SuperLU	STRUMPACK	cuSolver QR	SSIDS	MA57
Case 1	55.7K	268K	CPU (s)	1.1	1.0	2.6	0.7	0.2
			GPU (s)	1.6	1.6	1.8	5.1	–
Case 2	238K	1.11M	CPU (s)	4.0	2.8	18.2	2.6	0.8
			GPU (s)	5.0	3.7	5.7	4.8	–
Case 3	296K	7.67M	CPU (s)	30	24	614	29	6
			GPU (s)	33	24	–	198	–

KLU [25], NISLU [26], NVIDIA's cuSolverRf, and undocumented but publicly released NVIDIA's cuSolverGLU. Refactorization has not been particularly popular in sparse direct CPU solvers as the cost of pivoting on the CPU is generally not a concern. In contrast, pivoting on the GPU is prohibitively expensive, and avoiding pivoting is paramount for GPU speedups.

We found that refactorization approach is very effective for linear systems arising in ACOF (see Section 5). Similar approaches were successfully used before in modeling power systems. For instance, [4] develops a partial refactorization method (on the CPU, within NISLU package) and compares its efficiency to KLU, SuperLU, and non-altered NISLU. The authors use standalone matrices obtained from dynamic phasor simulations as their test cases. In [5], the authors compare the performance of NISLU and KLU (both with refactorization) and GLU [27], which is a sparse direct solver that performs factorization on the GPU (but does not enable refactorization); the approach was also tested on matrices extracted from dynamic phasor simulations and the results favored NISLU. Six different approaches for power flow simulations are tested in [6], two of them being GPU-only, two CPU-only and two hybrid (GPU+CPU). The tested GPU approaches involve refactorization based on cuSolverGLU. The hybrid solver turns out to be the fastest (i.e., preprocessing and symbolic factorization are performed on the CPU, the refactorization happens on the GPU, the triangular solve is done on the CPU).

5. Approach and results

We present two refactorization approaches and demonstrate their performance in the ACOF analysis context. The first approach we developed uses KLU and proprietary cuSolver libraries. The second approach is developed within the open-source GINKGO library. Both are hybrid CPU-GPU approaches that exploit the property that the matrix non-zero structure in (4) does not change from one system to the next so the same pivot sequence can be reused. Both approaches also take advantage of KKT matrix regularization performed by the interior method [13], which has a similar effect as perturbation techniques used in static pivoting [28], and helps us reuse the same pivot sequence over a larger number of linear systems. In contrast to the hybrid approach in [6], we move linear system data to GPU only once, after solving the first system in (4), and perform all subsequent computations there. In this way, we avoid excessive data movement between CPU and GPU, which adds significant overhead to the overall computational time.

5.1. Refactorization solver using CUDA libraries

5.1.1. Approach

In this approach, we completely solve the first system on the CPU using KLU with approximate minimum degree (AMD) reordering [29]. Then, we extract the elements of the symbolic factorization (i.e., the permutation vectors and the sparse matrix structures of the factors). Next, we set up an appropriate cuSolver (Rf or GLU) data structure, and we then solve all the remaining systems at each next optimization solver step by calling refactorization function. Typically, hundreds of linear systems need to be solved during ACOF execution, hence the

cost of solving the first system(s) on the CPU is amortized over many optimization solver steps. Both cuSolverRf and cuSolverGLU allow the user to provide permutation vectors and the sparse non-zero structure of L and U matrices obtained by the LU factorization of choice. We tested the native cuSolver factorization methods on CPU, however, we obtained the best results with KLU (this agrees with the results from [6]). KLU also produced the sparsest triangular factors. Since the same pivot sequence is used over and over in refactorization, the computed solution could potentially have a large error. Many (but not all) linear solvers follow the factorization by iterative refinement [30], which is a method of improving the backward error in the solution. For instance, MA57 by default uses up to 10 iterations of iterative refinement; for the test matrices from [3] the average was around 6. How does this work? Let us assume that we solved the linear system (4) using a direct method and computed a solution, $\Delta x_k^{(0)}$. To improve the solution, $\rho^{(0)} = r_k - K_k \Delta x_k^{(0)}$ is computed, and a new system is solved for $\delta^{(0)}$: $K_k \delta^{(0)} = \rho^{(0)}$ using the factors of K_k obtained previously. Then, a new solution is formed $\Delta x_k^{(1)} = \Delta x_k^{(0)} + \delta^{(0)}$. If the solution is still not satisfactory, $\rho^{(1)} = r_k - K_k \Delta x_k^{(1)}$ is computed and the process repeats. This method is often quite effective. An alternative method is to use the triangular factors as a preconditioner inside the linear solver in the same way one would use the incomplete LU (ILU) preconditioner [31]. The literature suggests using flexible GMRES [32] as the iterative method deeming it more stable in this case [33]. Some authors, e.g. [34], mix the two approaches by using an iterative solver to solve the systems $K_k \delta^{(i)} = \rho^{(i)}$. We use the *ILU style approach* because in practice it requires fewer triangular solves and leads to a solution with the same quality.

Another method to address the deterioration in the solution quality is to recompute the symbolic factorization (on the CPU) if the solution quality becomes too poor or if the maximum number of refinement iterations is exceeded. In our test cases, however, we typically needed only one or two refinement iterations to maintain an appropriate level of accuracy. Hence, we never recomputed the symbolic factorization.

5.1.2. Implementation

The refactorization approach is implemented in CUDA/C++ and currently is distributed with HiOp library. The algorithm follows the steps below.

Solve the first system in the sequence. The first system in the sequence (4) is solved completely using KLU (using the usual three-step approach: symbolic factorization, numeric factorization and solve). We use AMD reordering and we allow the factorization to continue even if the matrix is deemed to be singular by the analysis (which is one of the user-defined KLU options).

Set the sparsity pattern of the factors. The non-zero structures of the L and U factors are extracted from KLU together with permutation vectors. The factors are returned in compressed sparse column (CSC) format. If cuSolverGLU is used, the factors are combined into a matrix containing both L and U and the matrix is copied to the GPU. If cuSolverRf is used, both factors are copied to the GPU separately and each factor is converted to compressed sparse row (CSR) format on the GPU.

Refactorization setup. An appropriate `cuSolver` refactorization setup and analysis functions are called (there are different functions for `cuSolverRf` and for `cuSolverGLU`).

Solve subsequent systems. Once the refactorization data is set up, each subsequent system is solved on the GPU. This consists of two functions: resetting the values inside the refactorization data structure and performing the refactorization (factorizing the new matrix using the prior pivot sequence). Each refactorization method has its own functions to execute these operations. `cuSolverRf` follows up with a “solve” function that permutes the solution to the original order and performs two triangular solves. `cuSolverGLU` also adds MA57-style iterative refinement inside its black-box “solve” function. In the case of `cuSolverRf`, we investigated an approach as in [6], in which the triangular solve happens on the CPU. However, for our test cases, the GPU-resident `cuSolverRf` triangular solver was overall faster, as it avoids moving matrix factors and permutation vectors to CPU and solution vector back to GPU each time the triangular solver is called.

Iterative refinement. If desired (or if the error in the solution is too large), `cuSolverRf` solve can be followed by iterative refinement. We implemented it as flexible GMRES with re-orthogonalized classical Gram–Schmidt (CGS2) orthogonalization (to improve numerical stability and GPU performance).

The iterative refinement was implemented to be highly efficient on the GPU. For instance, all the large arrays (such as the ones used to store Krylov vectors) are allocated once and reused for all the systems. The handles (needed for various CUDA libraries, such as `cuBLAS`) and buffers (such as the one used in the matrix–vector product) are also allocated once and reused. For most of the test cases only one or two steps of iterative refinement were needed.

Both `cuSolver` refactorization solvers come with limitations. The faster out of the two, `cuSolverGLU` does not provide access to the triangular factors, which makes it impossible to couple it with alternative (faster) sparse triangular solvers or user-implemented iterative refinement. `cuSolverGLU` also comes with its own iterative refinement (part of the `solve` function) which cannot be turned off and is not parameterized (i.e., the user cannot specify the maximum number of iterations, the algorithm used, or the solver tolerance). Both refactorization solvers are closed source, which makes it difficult to extract pivoting information and impossible to apply problem-specific optimization. Also, both of them can only execute on NVIDIA GPUs, which limits the platform support.

Against this background, we explored the possibility to design and deploy platform-portable open-source direct solver functionality in the GINKGO math library.

5.2. GINKGO refactorization solvers

5.2.1. Approach

For the refactorization approach, we use the open-source software library GINKGO that is developed within the Exascale Computing Project (ECP) [11]. GINKGO focuses on the efficient handling of sparse linear systems on GPUs. Implemented in C++ and featuring backends in the hardware-native languages CUDA (for NVIDIA GPUs), HIP (for AMD GPUs), and DPC++ (for Intel GPUs), the library combines sustainability with performance portability, see Fig. 2.

The implementation using the GINKGO framework consists of three phases: preprocessing, (re)factorization, and triangular solve. The preprocessing step is executed only once, and on the CPU, while the repeated calculations in the factorization happen entirely on the GPU. We preprocess the system matrix to improve numerical stability (pivoting and equilibration using the MC64 algorithm introduced by Duff and Koster [35]) followed by a reordering to reduce fill-in (symmetric AMD on $A + A^T$). The resulting permutation and scaling factors will later be used to map the solution of the transformed system back to the original linear system. Following the preprocessing, we compute a symbolic factorization of the scaled and reordered matrix. The resulting

sparsity pattern of the L and U factors can be reused for all subsequent iterations of the optimization algorithm. Finally, we compute the values of L and U on the GPU by first filling the values of A into the combined storage for L and U and executing the numerical factorization kernel.

5.2.2. Implementation

The components of the algorithm were implemented in the GINKGO library using C++ for the high-level control flow and CUDA/HIP for the factorization kernels. It is included in the Release 1.6.0 of GINKGO available on GitHub [36].

Preprocessing. The reordering and equilibration as well as the symbolic factorization of the resulting reordered matrix happens sequentially on the CPU. The symbolic LU factorization closely follows the `fill1` algorithm by Rose and Tarjan [37].

Numerical Factorization. The LU factorization kernel (Algorithm 1) uses a fine-grained up-looking parallel factorization based on the dependency structure encoded in the lower L factor. The kernel relies on two important components: a sophisticated scheduling approach and an efficient sparse vector addition routine. Each row is mapped to a warp/wavefront on the GPU and gets updated with all dependency rows. The updates add a sparse row from U in-place to the current sparse row, thus eliminating the corresponding variable in the current row. This sparse vector addition is facilitated by a sparsity pattern lookup structure that uses bitmaps or hashtables, depending on the pattern, to compute the nonzero index to each column index in the row. The dependency resolution uses a sync-free scheduling approach based on a `ready` flag for each row, previously used for sparse triangular solvers by Lui et al. [38]. The scheduling takes advantage of GPUs from NVIDIA and AMD scheduling their thread blocks in monotonic order and providing strong forward-progress guarantees inside a single thread block.

```

set ready[i] ← 0 for all i;
parfor row  $i = 0, \dots, n - 1$  do
  for lower nonzero  $l_{id}$  do
    while ready[d] = 0 do wait;
     $\alpha \leftarrow a_{id}/a_{ii}$ ;
     $l_{ii} \leftarrow \alpha$ ;
    parfor upper nonzero  $u_{dj}$  do // warp-parallel
       $a_{ij} \leftarrow a_{ij} - \alpha \cdot u_{dj}$ ;
    end
  end
  split  $a_i$  into  $l_i$  and  $u_i$ ; // conceptually
  ready[i] ← 1;
end

```

Algorithm 1: GINKGO LU factorization algorithm

Refactorization. All of the scaling coefficients and permutation indices from Step 1 and all of the lookup data structures and the symbolic factorization from Step 2 can be reused for subsequent linear systems, leaving only the numerical factorization itself to be recomputed. Finally, we permute and scale the right-hand side and solution vectors into/out of a persistent workspace that needs only to be allocated once.

Triangular Solvers. We solve the resulting triangular systems using the triangular solvers provided by the `cuSPARSE/rocSPARSE` library. Alternatively, the application can be set to use GINKGO’s own triangular solvers based on a configuration option. These triangular solvers use the same sync-free scheduling approach as the factorization (Algorithm 2). We use a thread-per-row mapping, which requires independent thread scheduling available on NVIDIA GPUs since the Volta architecture. On AMD GPUs and older NVIDIA GPUs, we need to modify the control flow to guarantee forward-progress for the entire warp.

5.3. Performance profiling

We evaluated the approach described herein on the Summit super-computer and the Crusher test and development system at the OLCF.

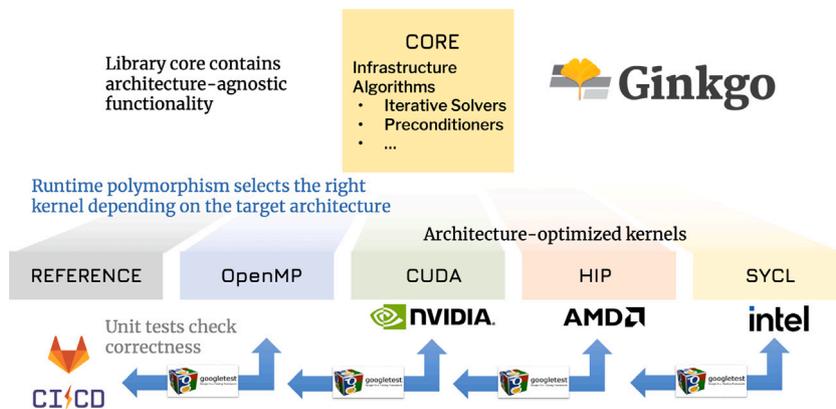


Fig. 2. The overall structure of the Ginkgo math library that was designed for performance portability.

```

set  $x_i \leftarrow \text{NaN}$  for all  $i$ ;
parfor row  $i = 0, \dots, n - 1$  do
     $\tilde{x} \leftarrow y_i$ ;
    for nonzero  $l_{id}$  do
        while  $x_d$  is NaN do wait;
         $\tilde{x} \leftarrow \tilde{x} - l_{id}x_d$ ;
    end
     $x_i \leftarrow \tilde{x}$ ;
end

```

Algorithm 2: GINKGO triangular solver algorithm solving $Lx = y$

Each Summit node is equipped with two 21-core IBM Power 9 CPUs and six NVIDIA V100 GPUs, while Crusher’s node architecture consists of one 64-core AMD EPYC 7A53 CPU and four AMD MI250X GPUs. The AMD GPU contains two graphics complex dies (GCDs) that can be treated as individual GPUs. For the present benchmarks, we consider one CPU core and one GPU/GCD (NVIDIA V100 GPU or AMD MI250X GCD). The code was compiled with CUDA 11.4.2 and the GNU Compiler Collection 10.2.0 on the NVIDIA platform and with the ROCm 5.2.0 software stack on the AMD platform.

We used state-of-the-art MA57 linear solver, which implements LDL^T factorization, as our CPU baseline. The MA57 was configured to perform symbolic factorization only once, for the first linear system in (4), and reuse it for all subsequent systems. Numerical factorization was configured to use default pivoting options and the triangular solve was configured to use no more than one iterative refinement iteration to maximize the performance. These are default MA57 settings in HiOp optimization library. We note it has been reported that performance of MA57 could be further improved if the numerical factorization is configured to use *static pivoting* [28] and triangular solve is followed with 10 or more refinement iterations to ensure sufficient solution accuracy is retained. When tested with our use cases, we found that static pivoting improved linear solver performance by 30%. However, the solution quality deteriorated so much that even with up to 100 iterative refinement iterations, we could not recover solution accuracy needed for the optimization solver to converge.

The MA57 and GINKGO solvers were tested on both systems, as GINKGO supports different GPU backends [39], while the cuSolver approach was only evaluated on NVIDIA GPUs. cuSolverRf was evaluated with a maximum of 20 refinement iterations and with the iterative refinement tolerance set to 10^{-14} (near the double floating point precision).

Tables 3 and 4 report the overall ACOF analysis performance when using different linear solvers for the test cases detailed in Table 1. The reported total runtimes were each measured from a single profiling run. Multiple ACOF runs were subsequently performed to assess the performance variability, resulting in an observed population standard

deviations (normalized by the mean value of the total runtime) lower than 1%.

Figs. 3 and 4 provide additional insight by breaking down the average runtime of an optimization solver (in this case HiOp) step into its components. The averages were obtained by normalizing the total computational time for each component by the number of optimization steps to allow comparison between the approaches, and these averages account for the cost of the first factorization performed on CPU.

We first focus on the evaluation of the benchmark results obtained on the OLCF Summit system equipped with NVIDIA V100 GPU. The results reported in Table 3 reveal that the GPU solvers outperform the MA57 baseline implementation for all test cases. cuSolverRf is 10 – 30% slower overall than cuSolverGLU, though it requires fewer iterations overall to converge. This implies that the iterative refinement improves the solution quality and the rate of convergence for an additional computational cost for the cases evaluated. In terms of overall ACOF compute time, using cuSolverGLU and GINKGO on V100 GPU leads to 1.3–1.4× and 1.05–1.3× faster solution, respectively, compared to the CPU baseline with MA57.

The runtime breakdown in Fig. 3 reveals that the performance advantage of the cuSolverGLU is primarily driven by a faster factorization (which is for the combined Eastern and Western U.S. grid about 3.4× faster than MA57). The faster factorization compensates for the slower triangular solves: The cuSolverGLU triangular solves are about 40% slower than the CPU counterpart. For the GINKGO GPU solver, the story is different: though still faster than the MA57 code, the speedup achieved with the GINKGO factorization is smaller. On the other hand, the triangular solve in GINKGO is faster, mainly because it does not call iterative refinement. We note that it is impossible to combine the cuSolverGLU factorization with the GINKGO sparse triangular solves as the cuSolverGLU does not provide access to the triangular factors.

We now turn to the performance results on the OLCF Crusher system featuring AMD MI250 GPUs. The results in Table 4 reveal that ACOF is overall 1.8 – 2.4× faster when using the GINKGO linear solver on the AMD MI250X GPU than MA57 on the AMD EPYC 7A53 CPU. The runtime breakdown in Fig. 4 reveals that the performance superiority comes from both a faster factorization (3 – 4.8× speedup) and faster triangular solves (1.9 – 3× speedup). Comparing GINKGO’s linear solver (triangular solve and factorization) performance on the MI250X GPU and the NVIDIA V100 GPU, we notice that executing GINKGO on the newer MI250 GPU is 20%–40% faster than on the NVIDIA V100 GPU. MA57 runs slower on the AMD EPYC 7A53 CPU than on the IBM Power 9 CPU.

For completeness, we finally discuss the memory requirements for the sparse linear solvers on GPU. The GPU high bandwidth memory required by the linear solvers for three grids evaluated is less than 3 GB (well below the 16 GB available on NVIDIA V100 and the 64 GB available on a graphics compute die of AMD MI250X). More specifically, on

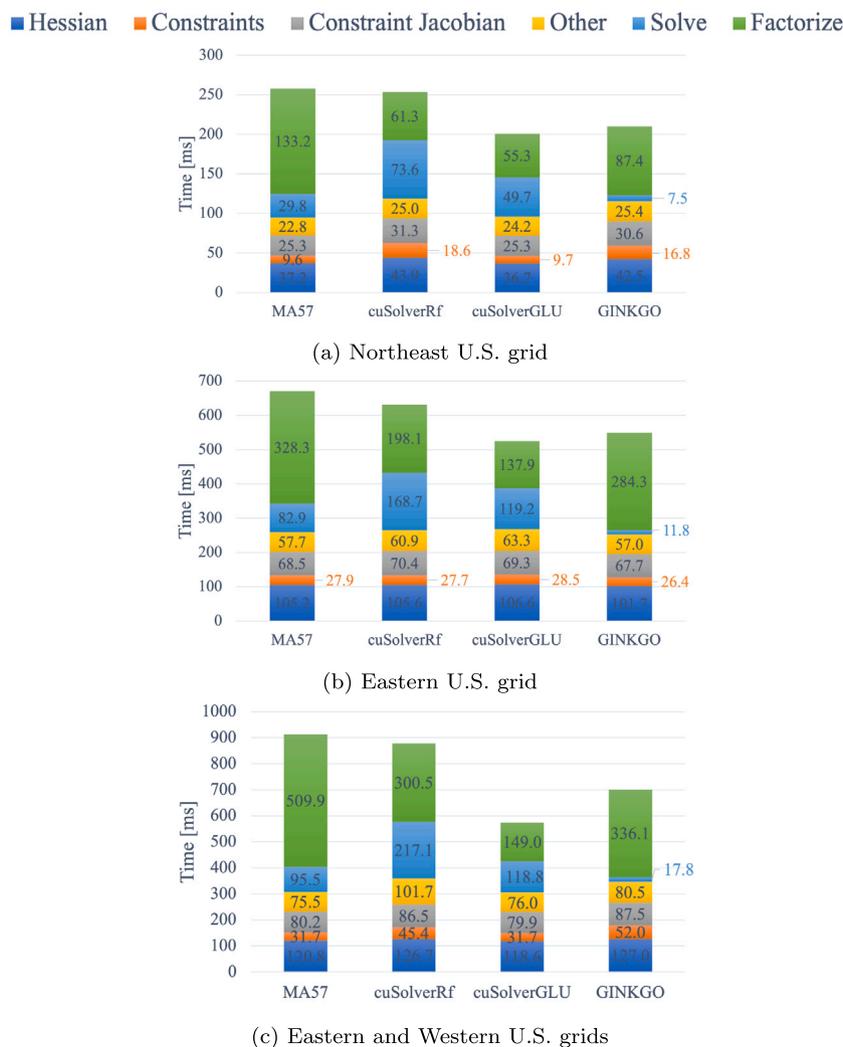


Fig. 3. Comparison of the average computational cost per optimization solver step when different linear solvers are used for ACOPF on OLCF Summit with a breakdown in terms of most expensive operations. The cost of the first step, which is executed on CPU, is accounted for in the averages.

NVIDIA V100, the solver based on cuSolverGLU requires 2383–3065 MiB, the solver based on cuSolverRf requires 673–1125 MiB, and GINKGO requires 725–1109 MiB.

6. Conclusions and next steps

This paper follows up on the analysis from [3] and presents different strategies to overcome the challenge of fast GPU-resident sparse direct solvers. Unlike in [3], we test linear solvers within full optimal power flow applications at interconnection scale, not on standalone test matrices. Our performance analysis reveals that the linear solver dominates the overall performance.

On both the IBM/NVIDIA-based Summit supercomputer and the AMD-based Crusher cluster, the GPU-based sparse linear solvers are faster than the CPU-based solvers. Furthermore, GINKGO’s GPU-resident sparse direct solver functionality brings platform portability to the ACOPF simulation and is, for all considered hardware configurations, faster than the CPU-based counterparts. For our test cases, we observed 1.9 – 2.3× linear solver speedup when comparing the best GPU to the best CPU result (see Figs. 3 and 4). This projects to 1.6–1.9× overall speedup in ACOPF analysis when comparing the best GPU to the best CPU time (see Tables 3 and 4). These and the results in [8] suggest it is feasible to develop methods for efficient ACOPF on GPUs and obtain significant performance gains.

Table 3

Total runtimes for ACOPF when using different linear solvers on OLCF Summit. The number of steps is the total number of optimization solver steps to the converged solution.

Northeast U.S. grid				
Linear solver used	MA57	cuSolverRf	cuSolverGLU	GINKGO
Total time (s)	152	127	116	114
Speedup vs. MA57	–	1.2	1.3	1.3
Number of steps	529	425	547	527
Eastern U.S. grid				
Linear solver used	MA57	cuSolverRf	cuSolverGLU	GINKGO
Total time (s)	196	187	147	153
Speedup vs. MA57	–	1.05	1.3	1.3
Number of steps	263	262	263	263
Combined eastern and western U.S. grids				
Linear solver used	MA57	cuSolverRf	cuSolverGLU	GINKGO
Total time (s)	793	674	575	766
Speedup vs. MA57	–	1.2	1.4	1.04
Number of steps	852	735	747	1038

We note that the refactorization approaches we employ in this work reuse the same pivot sequence over a large number of linear systems

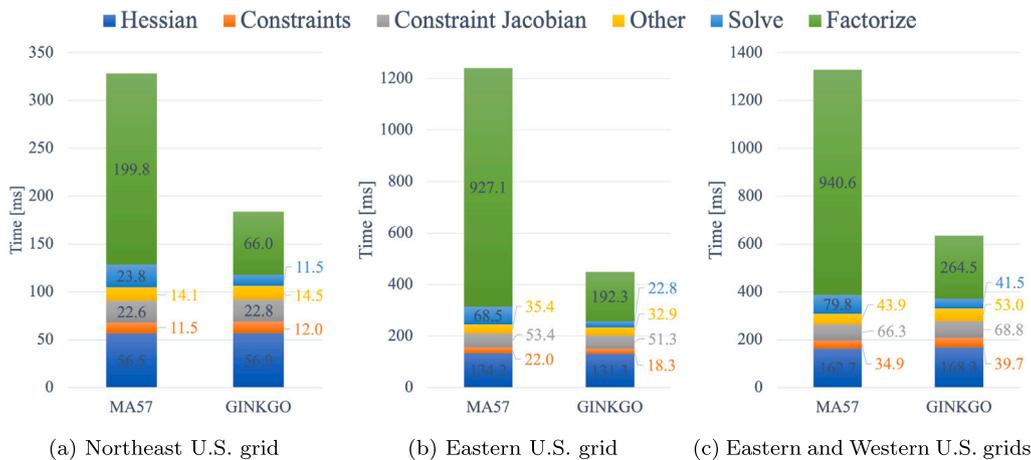


Fig. 4. Comparison of the average computational cost per optimization solver step when different linear solvers are used for ACOPF on OLCF Crusher with a breakdown in terms of most expensive operations. The cost of the first step, which is executed on CPU, is accounted for in the averages.

Table 4

Total runtimes for ACOPF when using MA57 and GINKGO linear solvers on OLCF Crusher. The number of steps is the total number of optimization solver steps to the converged solution.

Northeast U.S. grid		
Linear solver used	MA57	GINKGO
Total time (s)	149	81
Speedup vs. MA57	-	1.8
Number of steps	455	446
Eastern U.S. grid		
Linear solver used	MA57	GINKGO
Total time (s)	293	122
Speedup vs. MA57	-	2.4
Number of steps	234	263
Combined eastern and western U.S. grids		
Linear solver used	MA57	GINKGO
Total time (s)	927	450
Speedup vs. MA57	-	2.1
Number of steps	693	700

(4), and hence come with the risk of deterioration in solution quality. On the other hand, the linear solver error only needs to be small compared to the allowed error in optimization solver iteration. We find that with HiOp’s default tolerance setting of 10^{-8} , refactorization approach without iterative refinement as implemented in GINKGO will perform well in most cases. The exception is the ACOPF analysis for the Eastern and Western U.S. grid, where analysis using GINKGO requires nearly 200 more optimization solver steps than the reference analysis using MA57 linear solver on CPU (see Table 3).

We observed that using high precision iterative refinement with cuSolverRf leads to fewer optimization solver steps to the solution as expected (Table 3). However, iterative refinement can add significant computational overhead, as shown in Fig. 3, as it calls triangular solver repeatedly. Finding the right balance between reducing the number of optimization solver steps and the cost of each step in order to reduce the overall computational cost is a nontrivial problem. We will investigate this further in the follow on research.

We find that numerical (re)factorization in cuSolverGLU is superior to those in cuSolverRf and GINKGO. However, cuSolverGLU does not provide access to the L and U matrix factors, so its factorization function cannot be combined with other solvers. At the same time, its triangular solver function performs poorly compared to MA57

reference, most likely due to built-in iterative refinement, which the user cannot configure or disable. As a next step, we will work towards improving the performance of GINKGO refactorization functions.

Future work will also include algorithmic and technical strategies to speed up triangular solves by using an iterative approximation as proposed in [40]. Furthermore, we will investigate the hybrid strategy developed in [41], which applies both iterative and direct solvers with very promising performance results. Hybrid strategies will likely become more attractive in the future as the hardware manufacturers plan to develop processing units with CPU and GPU cores on the same die.

Finally, the full advantage of deploying ACOPF on heterogeneous hardware will become apparent when the entire analysis is ported to GPU, including the model evaluation and optimization solver. We demonstrated this in [8] where the entire compressed ACOPF model and dense linear solver were run on the GPU. Sparse GPU modules are currently being developed in both, HiOp and ExaGO™ libraries and, when completed, we will use them to deploy the entire analysis on GPU.

CRedit authorship contribution statement

Kasia Świrydowicz: Conceptualization, Investigation, Methodology, Software, Writing – original draft. **Nicholson Koukpaizan:** Formal analysis, Validation, Software, Writing – original draft. **Tobias Ribizel:** Investigation, Methodology, Software, Writing – original draft. **Fritz Göbel:** Investigation, Methodology, Software. **Shrirang Abhyankar:** Software, Writing – original draft. **Hartwig Anzt:** Supervision, Writing – original draft. **Slaven Peleš:** Conceptualization, Software, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All research data and code is available through public repositories referenced in the manuscript.

Acknowledgments

This research has been supported in part by UT-Battelle, LLC, and used resources of the Oak Ridge Leadership Computing Facility under contract DE-AC05-00OR22725 with the U.S. Department of Energy (DOE). This research was also supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the DOE Office of Science and the National Nuclear Security Administration.

The authors thank Cosmin Petra and Nai-Yuan Chiang of Lawrence Livermore National Laboratory for their guidance when using HiOp optimization solver. Warm thanks also go to Phil Roth of Oak Ridge National Laboratory and Christopher Oehmen of Pacific Northwest National Laboratory for their support of this work.

Notice of copyright

This manuscript has been authored in part by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

References

- [1] O'Neill RP, Castillo A, Cain MB. The IV formulation and linear approximations of the AC optimal power flow problem (OPF paper 2). FERC Staff Tech Pap 2012;(December):1–18, URL <http://www.ferc.gov/industries/electric/indus-act/market-planning/opf-papers/acopf-2-iv-linearization.pdf>.
- [2] Frank S, Rebennack S, et al. A primer on optimal power flow: Theory, formulation, and practical examples. Tech. rep. 14, Colorado School of Mines; 2012, p. 1–42.
- [3] Świrydowicz K, Darve E, Jones W, Maack J, Regev S, Saunders MA, Thomas SJ, Peleš S. Linear solvers for power grid optimization problems: a review of GPU-accelerated linear solvers. *Parallel Comput* 2022;111:102870.
- [4] Dinkelbach J, Schumacher L, Razik L, Benigni A, Monti A. Factorisation path based refactorisation for high-performance LU decomposition in real-time power system simulation. *Energies* 2021;14(23):7989.
- [5] Razik L, Schumacher L, Monti A, Guironnet A, Bureau G. A comparative analysis of LU decomposition methods for power system simulations. In: 2019 IEEE milan powertech. IEEE; 2019, p. 1–6.
- [6] D'orto M, Sjöblom S, Chien LS, Axner L, Gong J. Comparing different approaches for solving large scale power-flow problems with the Newton-Raphson method. *IEEE Access* 2021;9:56604–15.
- [7] Rakai L, Rosehart W. GPU-accelerated solutions to optimal power flow problems. In: 2014 47th Hawaii international conference on system sciences. IEEE; 2014, p. 2511–6.
- [8] Abhyankar S, Peles S, Rutherford R, Mancinelli A. Evaluation of AC optimal power flow on graphical processing units. In: 2021 IEEE power & energy society general meeting (PESGM). 2021, p. 01–5. <http://dx.doi.org/10.1109/PESGM46819.2021.9638131>.
- [9] Su X, He C, Liu T, Wu L. Full parallel power flow solution: A GPU-CPU-based vectorization parallelization and sparse techniques for Newton–Raphson implementation. *IEEE Trans Smart Grid* 2020;11(3):1833–44.
- [10] NVIDIA. cuSOLVER library, release 11.4. 2021.
- [11] Anzt H, Cojean T, Flegar G, Göbel F, Grützmacher T, Nayak P, Ribizel T, Tsai YM, Quintana-Orti ES. Ginkgo: A Modern Linear Operator Algebra Framework for High Performance Computing. *ACM Trans Math Software* 2022;48(1):2:1–33. <http://dx.doi.org/10.1145/3480935>.
- [12] Zimmerman RD, Murillo-Sánchez CE, Thomas RJ. MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Trans Power Syst* 2011;26(1):12–9. <http://dx.doi.org/10.1109/TPWRS.2010.2051168>.
- [13] Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math Program* 2006;106(1):25–57.
- [14] Birchfield AB, Xu T, Gegner KM, Shetye KS, Overbye TJ. Grid structural characteristics as validation criteria for synthetic networks. *IEEE Trans Power Syst* 2017;32(4):3258–65. <http://dx.doi.org/10.1109/TPWRS.2016.2616385>.
- [15] Duff IS. MA57—a code for the solution of sparse symmetric definite and indefinite systems. *ACM Trans Math Softw* 2004;30(2):118–44.
- [16] Abhyankar S, Peles S, Mancinelli A, Rutherford R, Palmer B. Exascale grid optimization toolkit. 2020, URL <https://https://gitlab.pnnl.gov/exasgd/frameworks/exago>.
- [17] Petra CG, Chiang N, Wang J. HiOp – user guide. Tech. rep. LLNL-SM-743591, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory; 2018.
- [18] Chiang N-Y, Zavala VM. An inertia-free filter line-search algorithm for large-scale nonlinear programming. *Comput Optim Appl* 2016;64(2):327–54.
- [19] Li XS, Demmel JW. SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems. *ACM Trans Math Softw* 2003;29(2):110–40.
- [20] Ghysels P, Li XS, Rouet F-H, Williams S, Napov A. An efficient multicore implementation of a novel HSS-structured multifrontal solver using randomized sampling. *SIAM J Sci Comput* 2016;38(5):S358–84.
- [21] Hogg JD, Ovtchinnikov E, Scott JA. A sparse symmetric indefinite direct solver for GPU architectures. *ACM Trans Math Softw* 2016;42(1):1–25.
- [22] Li XS. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM Trans Math Softw* 2005;31(3):302–25.
- [23] Duff I, Hogg J, Lopez F. A new sparse LDL^T solver using a posteriori threshold pivoting. *SIAM J Sci Comput* 2020;42(2):C23–42. <http://dx.doi.org/10.1137/18M1225963>.
- [24] Hénon P, Ramet P, Roman J. PaStiX: a high-performance parallel direct solver for sparse symmetric positive definite systems. *Parallel Comput* 2002;28(2):301–21.
- [25] Davis TA, Palamadai Natarajan E. Algorithm 907: KLU, a direct sparse solver for circuit simulation problems. *ACM Trans Math Softw* 2010;37(3):1–17.
- [26] Chen X, Wang Y, Yang H. NICSLU: an adaptive sparse matrix solver for parallel circuit simulation. *IEEE Trans Comput-Aided Des Integr Circuits Syst* 2013;32(2):261–74.
- [27] He K, Tan SX-D, Wang H, Shi G. GPU-accelerated parallel sparse LU factorization method for fast circuit analysis. *IEEE Trans Very Large Scale Integr (VLSI) Syst* 2015;24(3):1140–50.
- [28] Duff IS, Pralet S. Towards stable mixed pivoting strategies for the sequential and parallel solution of sparse symmetric indefinite systems. *SIAM J Matrix Anal Appl* 2007;29(3):1007–24. <http://dx.doi.org/10.1137/050629598>.
- [29] Amestoy PR, Davis TA, Duff IS. Algorithm 837: AMD, an approximate minimum degree ordering algorithm. *ACM Trans Math Software* 2004;30(3):381–8. <http://dx.doi.org/10.1145/1024074.1024081>.
- [30] Wilkinson J. Rounding errors in algebraic processes. 1965.
- [31] Saad Y. Iterative methods for sparse linear systems. SIAM; 2003.
- [32] Saad Y. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J Sci Comput* 1993;14(2):461–9.
- [33] Arioli M, Duff IS, Gratton S, Pralet S. A note on GMRES preconditioned by a perturbed LDL^T decomposition with static pivoting. *SIAM J Sci Comput* 2007;29(5):2024–44.
- [34] Carson E, Higham NJ, Pranesh S. Three-precision GMRES-based iterative refinement for least squares problems. *SIAM J Sci Comput* 2020;42(6):A4063–83.
- [35] Duff IS, Koster J. On algorithms for permuting large entries to the diagonal of a sparse matrix. *SIAM J Matrix Anal Appl* 2001;22(4):973–96. <http://dx.doi.org/10.1137/S0895479899358443>.
- [36] Ginkgo Project. Ginkgo release 1.6.0. 2023, URL <https://github.com/ginkgo-project/ginkgo/releases/tag/v1.6.0>.
- [37] Rose DJ, Tarjan RE. Algorithmic aspects of vertex elimination on directed graphs. *SIAM J Appl Math* 1978;34(1):176–97. <http://dx.doi.org/10.1137/0134014>.
- [38] Liu W, Li A, Hogg J, Duff IS, Vinter B. A synchronization-free algorithm for parallel sparse triangular solves. In: Dutot P-F, Trystram D, editors. Euro-Par 2016: parallel processing. Cham: Springer International Publishing; 2016, p. 617–30.
- [39] Cojean T, Tsai Y-HM, Anzt H. Ginkgo - a math library designed for platform portability. *Parallel Comput* 2022;111:102902. <http://dx.doi.org/10.1016/j.parco.2022.102902>.
- [40] Anzt H, Chow E, Dongarra J. Iterative sparse triangular solves for preconditioning. In: European conference on parallel processing. Springer; 2015, p. 650–61.
- [41] Regev S, Chiang N-Y, Darve E, Petra CG, Saunders MA, Świrydowicz K, Peleš S. HyKKT: a hybrid direct-iterative method for solving KKT linear systems. *Optim Methods Softw* 2022;1–24. <http://dx.doi.org/10.1080/10556788.2022.2124990>.