

SuperLU



Supernodal Sparse LU Direct Solver. Flexible, user-friendly interfaces.
Examples show various use scenarios. Testing code for unit-test. BSD license.

Capabilities

- Serial (thread-safe), shared-memory (SuperLU_MT, OpenMP or Pthreads), distributed-memory (SuperLU_DIST, hybrid MPI+ OpenM + CUDA/HIP).
 - Written in C, with Fortran interface
- Sparse LU decomposition (can be nonsymmetric sparsity pattern), triangular solution with multiple right-hand sides
- Incomplete LU (ILUTP) preconditioner in serial SuperLU
- Sparsity-preserving ordering: minimum degree or graph partitioning applied to $A^T A$ or $A^T + A$
- User-controllable pivoting: partial pivoting, threshold pivoting, static pivoting
- Condition number estimation, iterative refinement, componentwise error bounds

Exascale early systems GPU-readiness




- Available: Nvidia GPU (CUDA), AMD GPU (HIP)
- In progress: Intel GPU (DPC++ planned)

Parallel Scalability


- Factorization strong scales to 32,768 cores with 4096 GPUs (IPDPS'18, PARCO'19)
- Triangular solve strong scales to 4000 cores (SIAM CSC'18, SIAM PP'20)

Open source software

- Used in a vast range of applications, can be used through PETSc and Trilinos, available on github



ITER tokamak



quantum mechanics

Widely used in commercial software, including AMD (circuit simulation), Boeing (aircraft design), Chevron, ExxonMobile (geology), Cray's LibSci, FEMLAB, HP's MathLib, IMSL, NAG, SciPy, OptimaNumerics, Walt Disney Animation.



<https://portal.nersc.gov/project/sparse/superlu/>