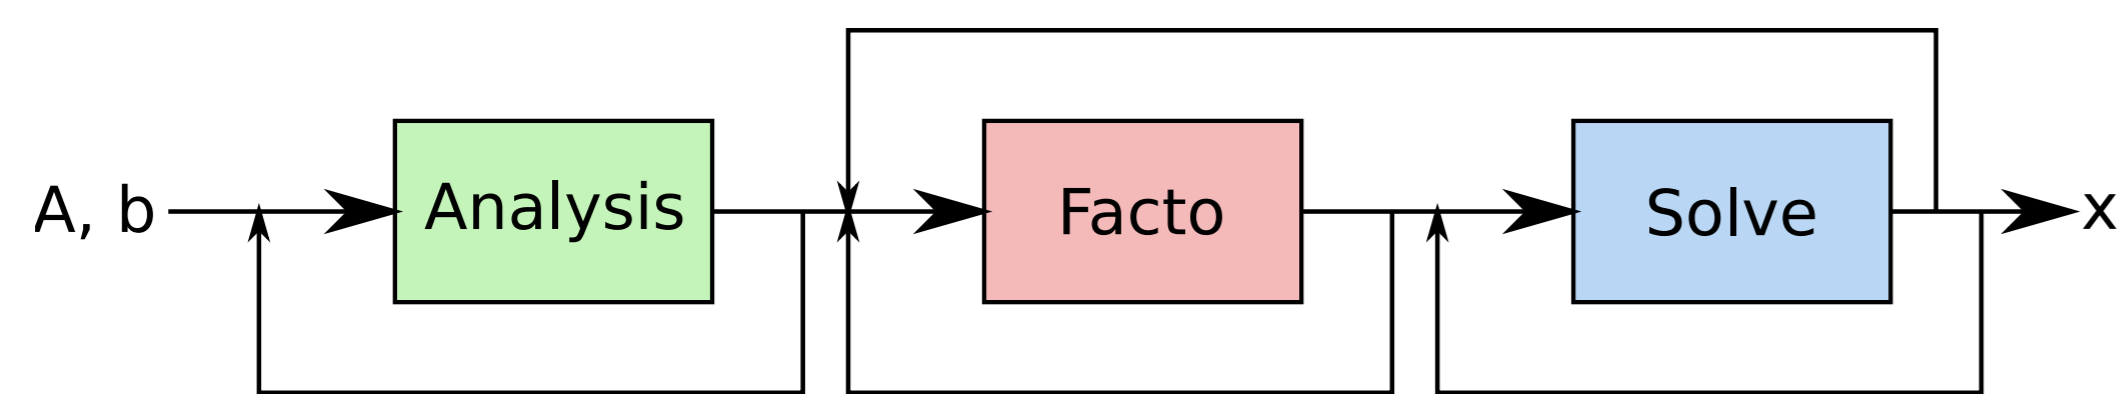


## Sparse direct solver on top of task-based runtime systems

One of the first attempts to build a complex and irregular application on top of task-based runtime systems

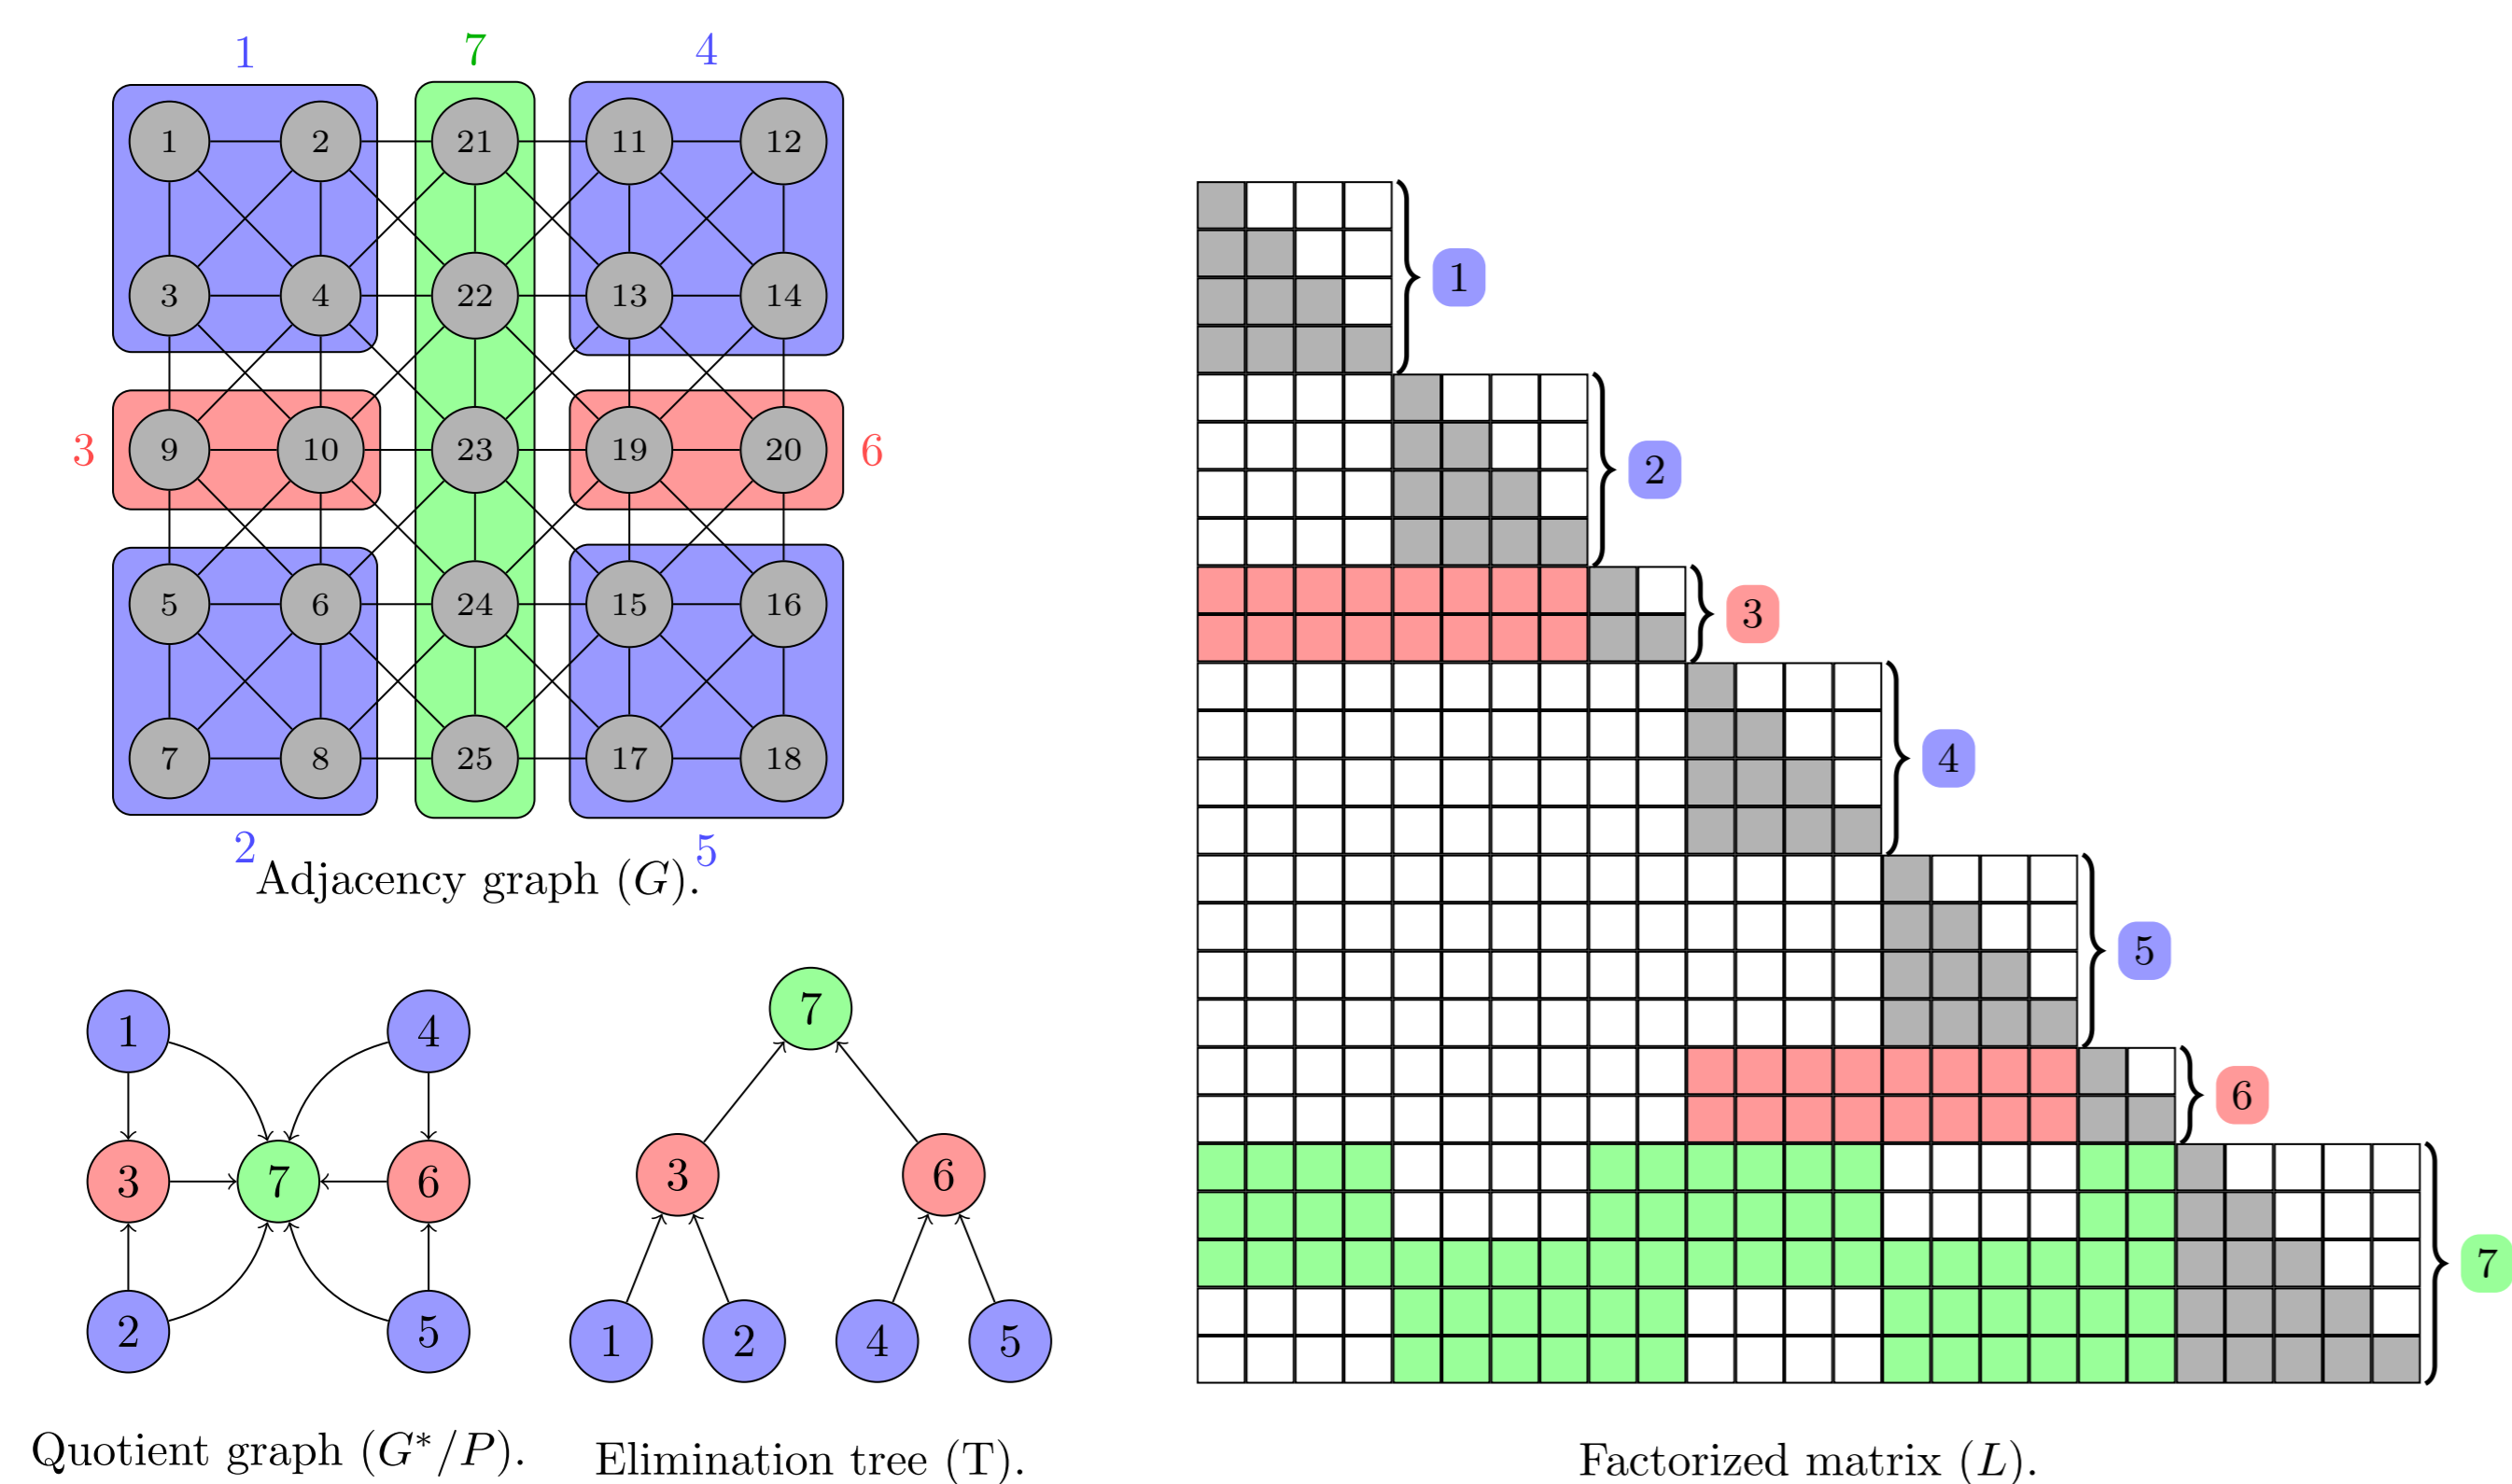
### Direct methods for sparse matrices



#### Analysis (on the graph of the matrix):

- Ordering: minimize fill-in while maximizing parallelism
- Symbolic factorization: predict the structure of the factorized matrix
- Data distribution: optimize the data distribution to accelerate the factorization
- Numerical factorization: decomposition of  $A$  into  $LU$ ,  $LL^T$ , or  $LDL^T$
- Triangular solve

### Symbolic factorization



### Numerical factorization

**Algorithm 1:** Right looking blocked sequential factorization:  $A = LL^T$ .

```

for k = 1 to N do
  /* Factorize the column block */
  Factorize  $A_{k,k}$  in  $L_{k,k} \cdot L_{k,k}^T$ ;
  Solve  $L_{(1-b_k),k} \cdot L_{k,k}^T = A_{(1-b_k),k}$ ;
  /* Trailing supernodes updates */
  for j = 1 to  $b_k$  do
    for i = 1 to  $b_k$  do
       $A_{(i,j)} = A_{(i,j)} - L_{(i,k)} \cdot L_{k,(j)}^T$ ;
    end
  end
end
    
```

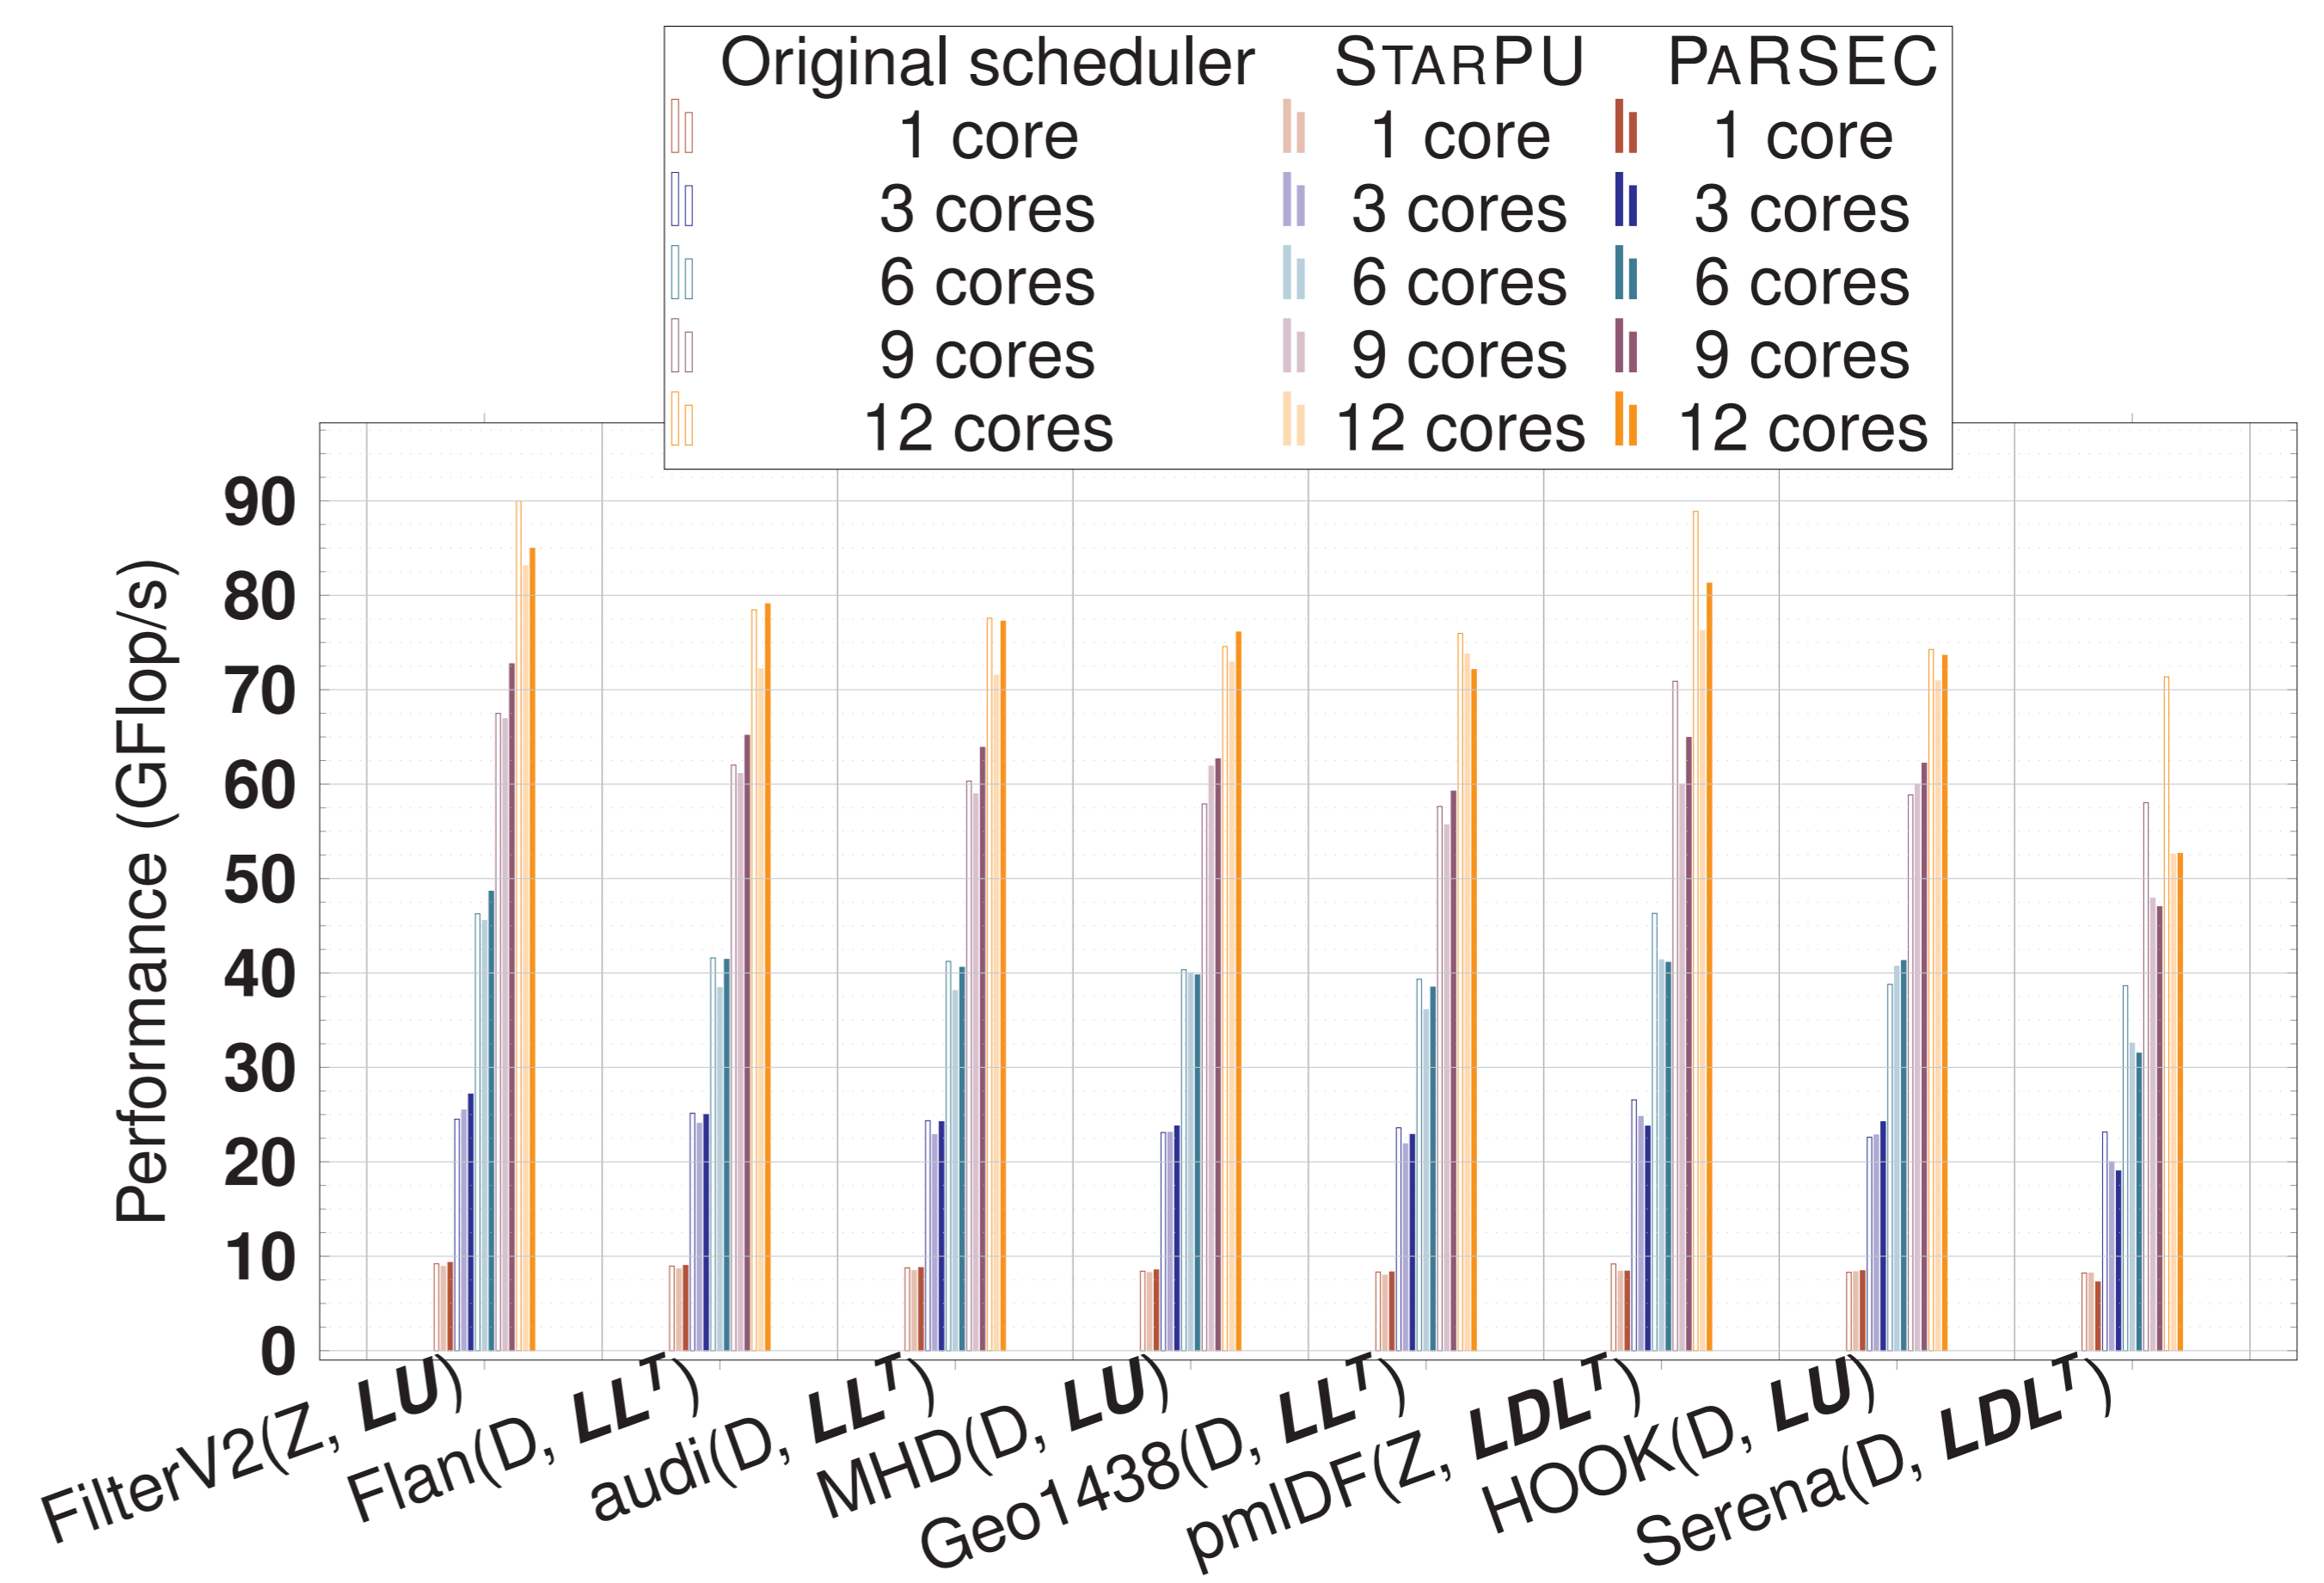
### DAG schedulers considered

- STARPU: <http://starpu.gforge.inria.fr>
  - Dynamic Task Discovery
  - Computes cost models on the fly
  - Multiple kernels on the accelerators
  - Multiple scheduling strategies: Minimum Completion Time, Local Work Stealing, ...
- PARSEC (formerly DAGUE): <http://icl.utk.edu/parsec>
  - Parameterized Task Graph
  - Only the most compute intensive kernel on accelerators
  - Scheduling strategy based on static performance model
  - GPU multi-stream enabled

### Matrix collection

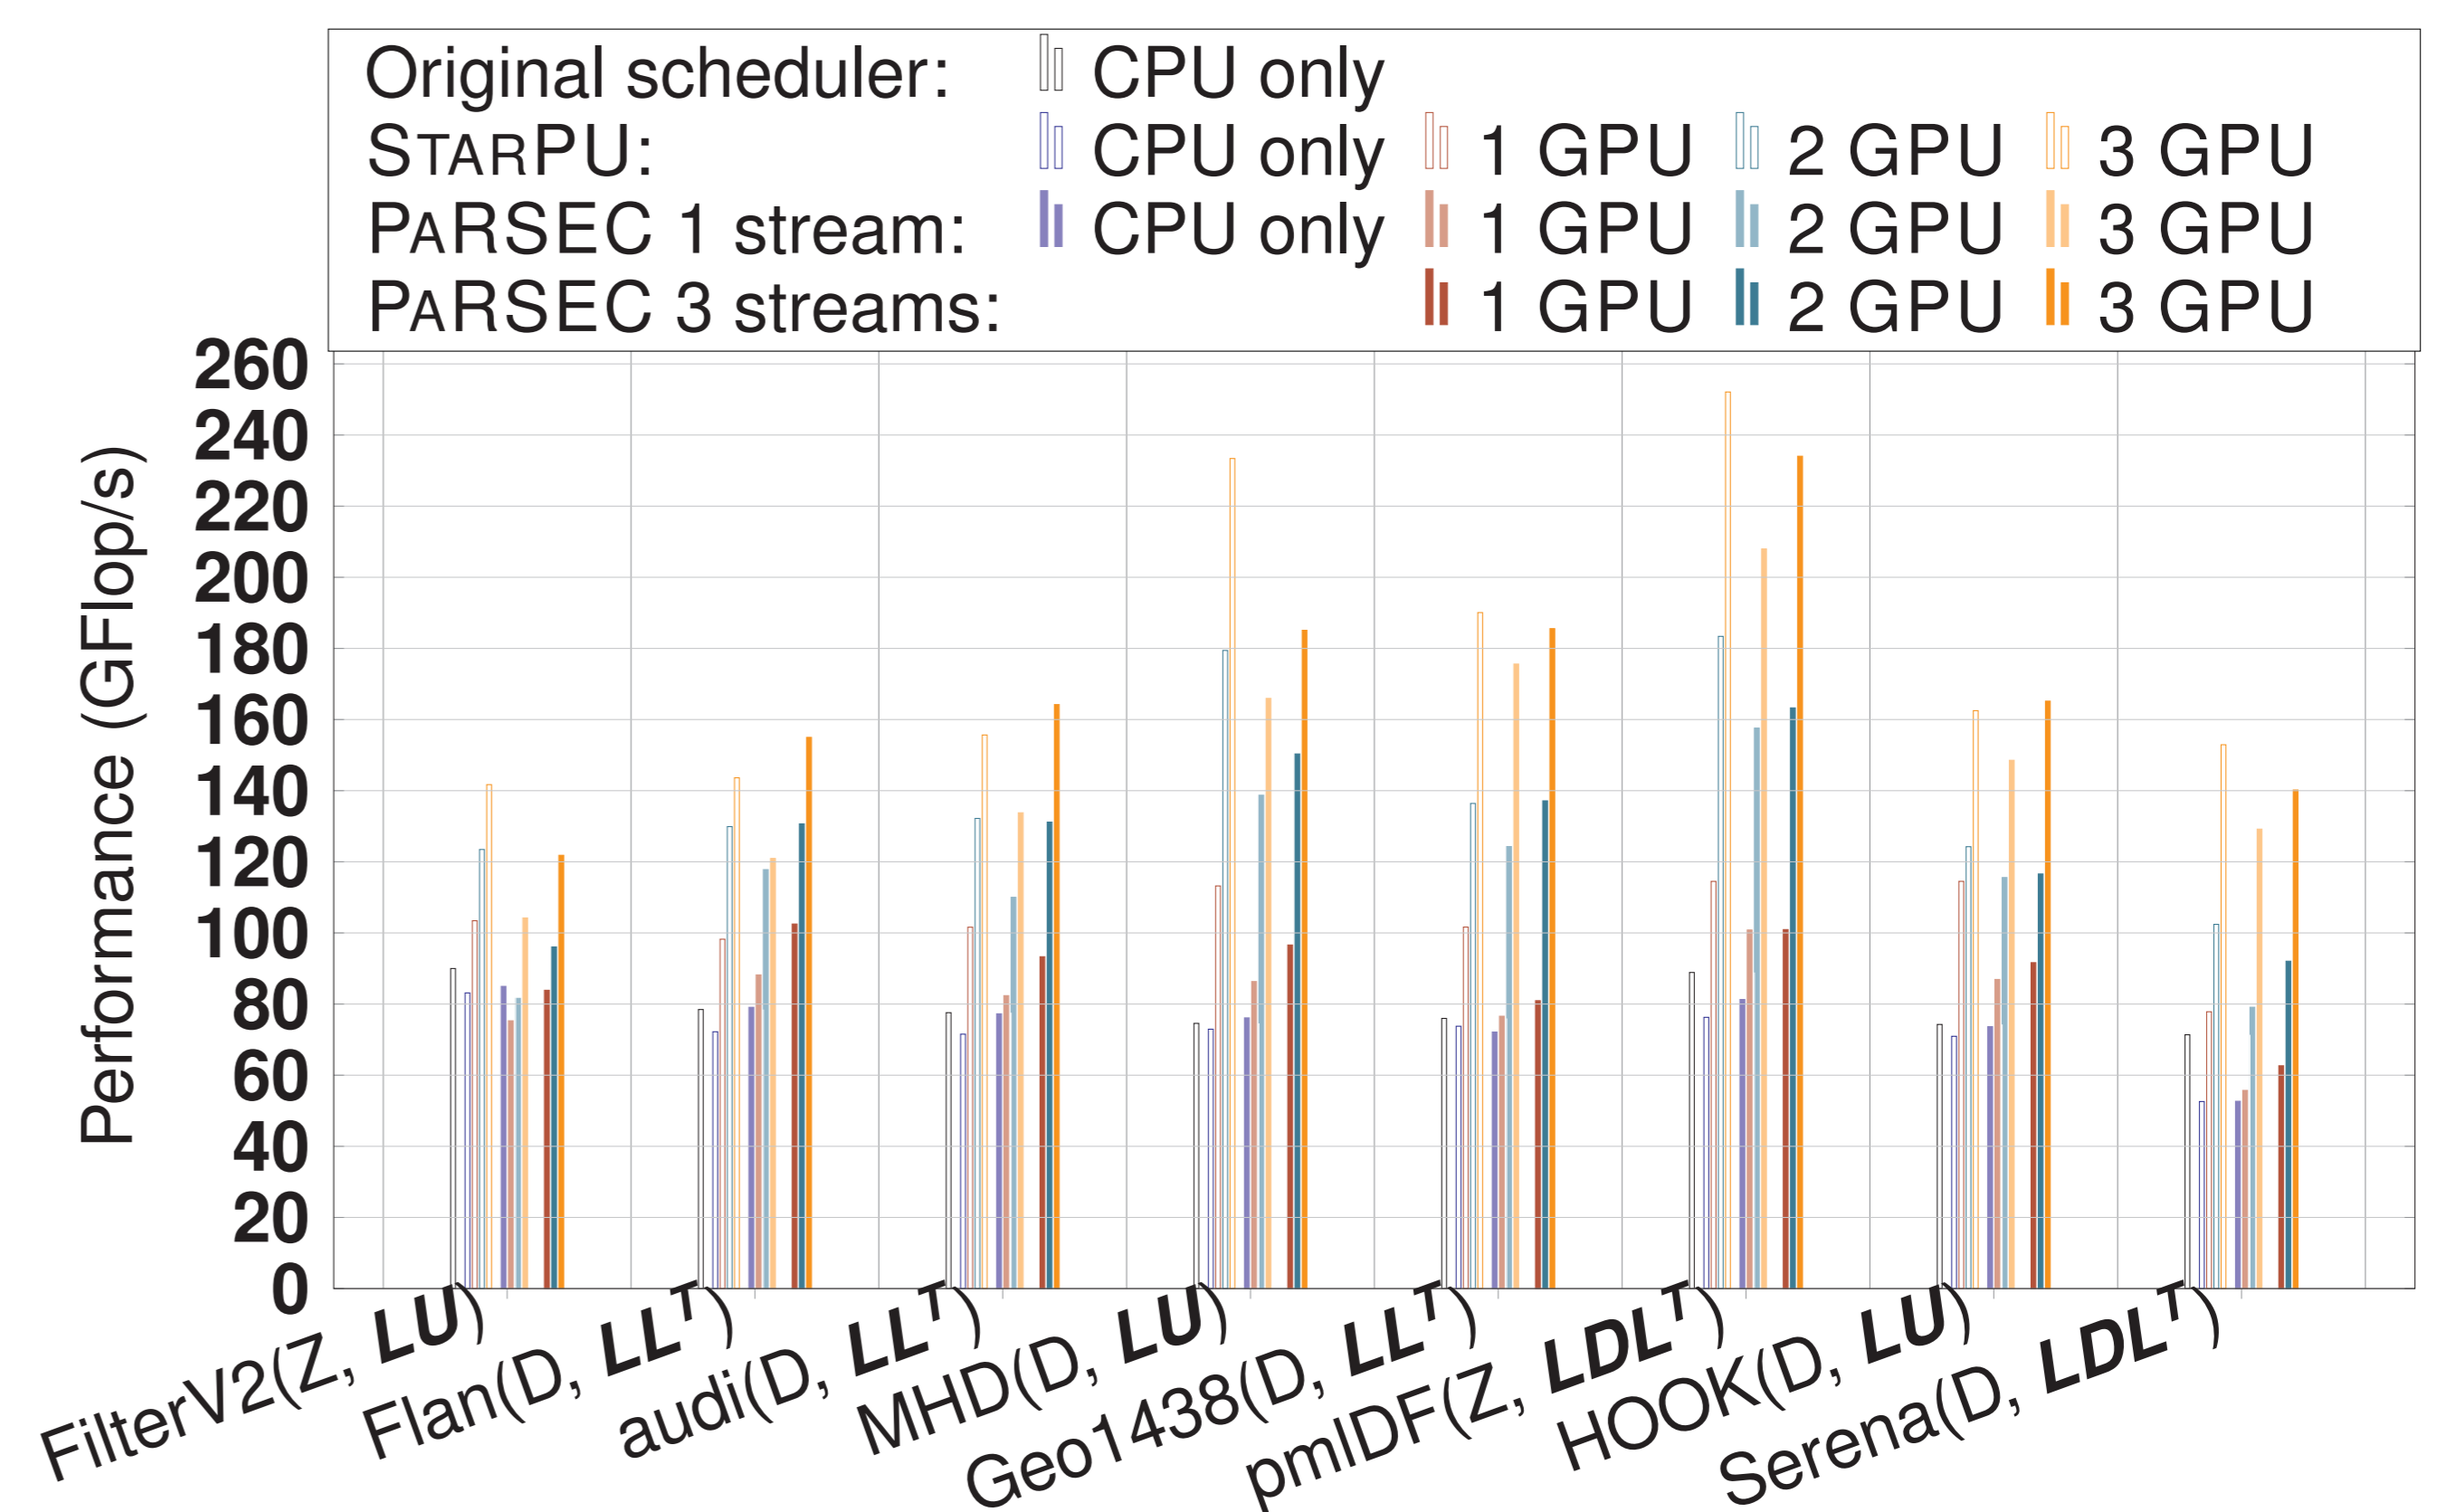
Matrix	Prec	Method	Size	nnz <sub>A</sub>	nnz <sub>L</sub>	TFlop
FilterV2	Z	LU	0.6e+6	12e+6	536e+6	3.6
Flan	D	LL <sup>T</sup>	1.6e+6	59e+6	1712e+6	5.3
Audi	D	LL <sup>T</sup>	0.9e+6	39e+6	1325e+6	6.5
MHD	D	LU	0.5e+6	24e+6	1133e+6	6.6
Geo1438	D	LL <sup>T</sup>	1.4e+6	32e+6	2768e+6	23
Pmldf	Z	LDL <sup>T</sup>	1.0e+6	8e+6	1105e+6	28
Hook	D	LU	1.5e+6	31e+6	4168e+6	35
Serena	D	LDL <sup>T</sup>	1.4e+6	32e+6	3365e+6	47

### CPU scaling (2 hexa-core Intel Xeon X5650)



- Same performances as PASTIX original finely tuned scheduler
- 80 GFlop/s out of 12\*4\*2.67GHz = 128 GFlop/s practical peak
- Can handle heterogeneous system more easily


### GPU scaling (2 hexa-core Intel Xeon X5650 + 3 NVIDIA M2070)



- Up to 160 GFlop/s (x3.1) using 3 GPUs with STARPU
- Up to 120 GFlop/s (x2.4) using 3 GPUs with PARSEC (1 stream)
- Up to 155 GFlop/s (x2.7) using 3 GPUs with PARSEC (3 streams)

### ANR Anemos: MHD simulation in Tokamaks

- Simulating the instabilities and their control in fusion reactors
- Designing elements of the ITER project



### AlgoTech: a successful switch to HPC

- HPC-PME Initiative (HPC for SMEs) by Bpifrance, GENCI and Inria
- AlgoTech is an innovative SME specialized in edition of electrical diagrams and cabling software for civil engineering, industrial machines and on-board systems
- AlgoTech has assessed and realized the technological leap needed for developing accelerated software using the PaStiX solver