

Peta-scale General Solver for Semidefinite Programming Problems with over Two Million Constraints

High-Performance General Solver for Extremely Large-scale Semidefinite Programming Problems (SDPs)

1. **Mathematical Programming**: one of the most central problems in mathematical optimization.
 2. **Many Applications**: combinatorial optimization, control theory, structural optimization, quantum chemistry, sensor network location, data mining, etc.

- **SDPARA** is a parallel implementation of the interior-point method for Semidefinite Programming (SDP)
- Parallel computation for **two major bottleneck parts**
 - **ELEMENTS** ⇒ Computation of Schur complement matrix (SCM)
 - **CHOLESKY** ⇒ Cholesky factorization of Schur complement matrix (SCM)
- **SDPARA** could attain high scalability using **16,320 CPU cores** on the TSUBAME 2.0 supercomputer and some techniques of processor affinity and memory interleaving when the generation of SCM constituted a bottleneck.
- With **4,080 NVIDIA M2050 GPUs** on the TSUBAME 2.0 supercomputer, our implementation achieved **1.018 PFlops** in double precision for a large-scale problem with over two million constraints.

SemiDefinite Programming (SDP)

Primal
Dual

The number of constraints of the primal problem

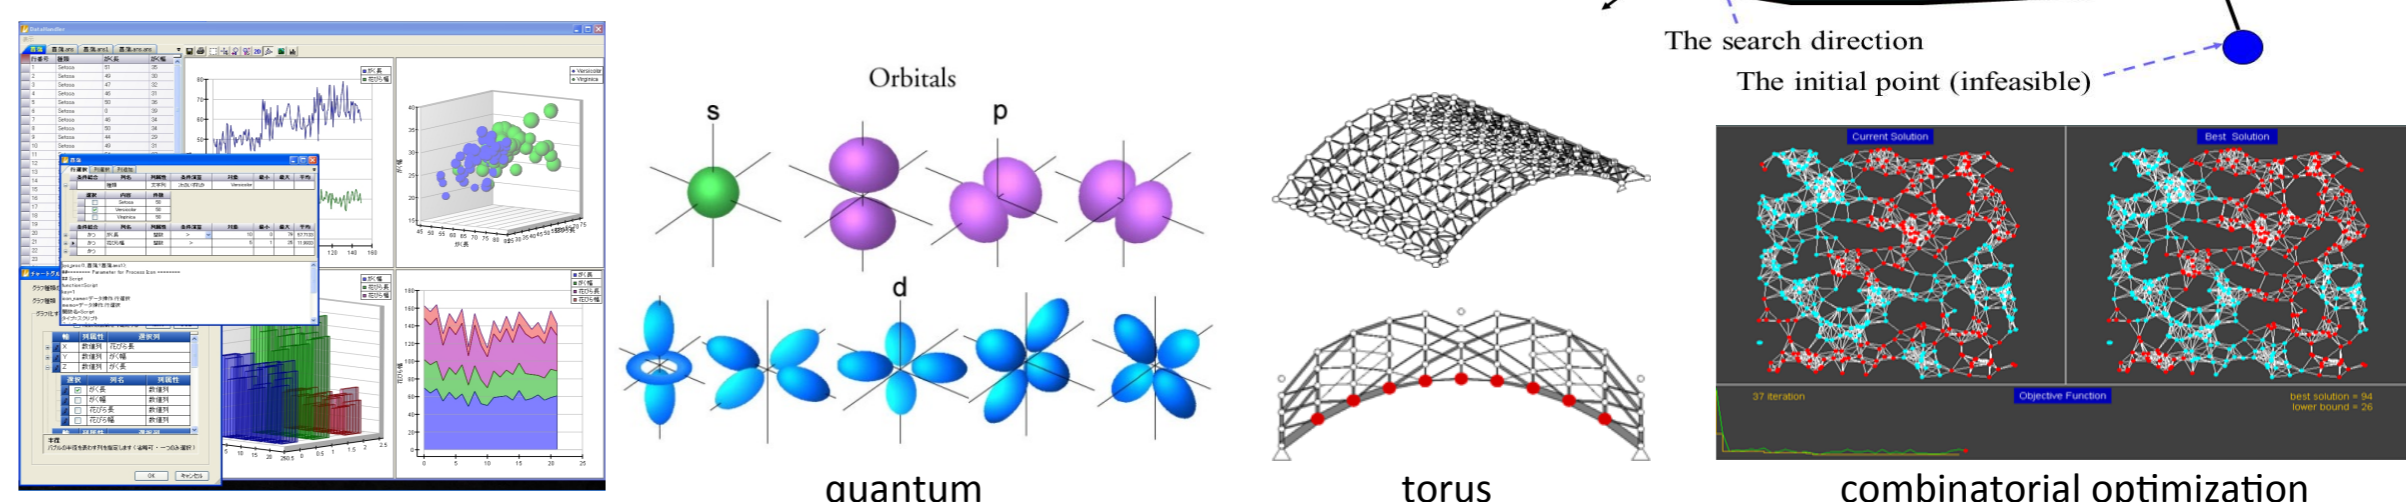
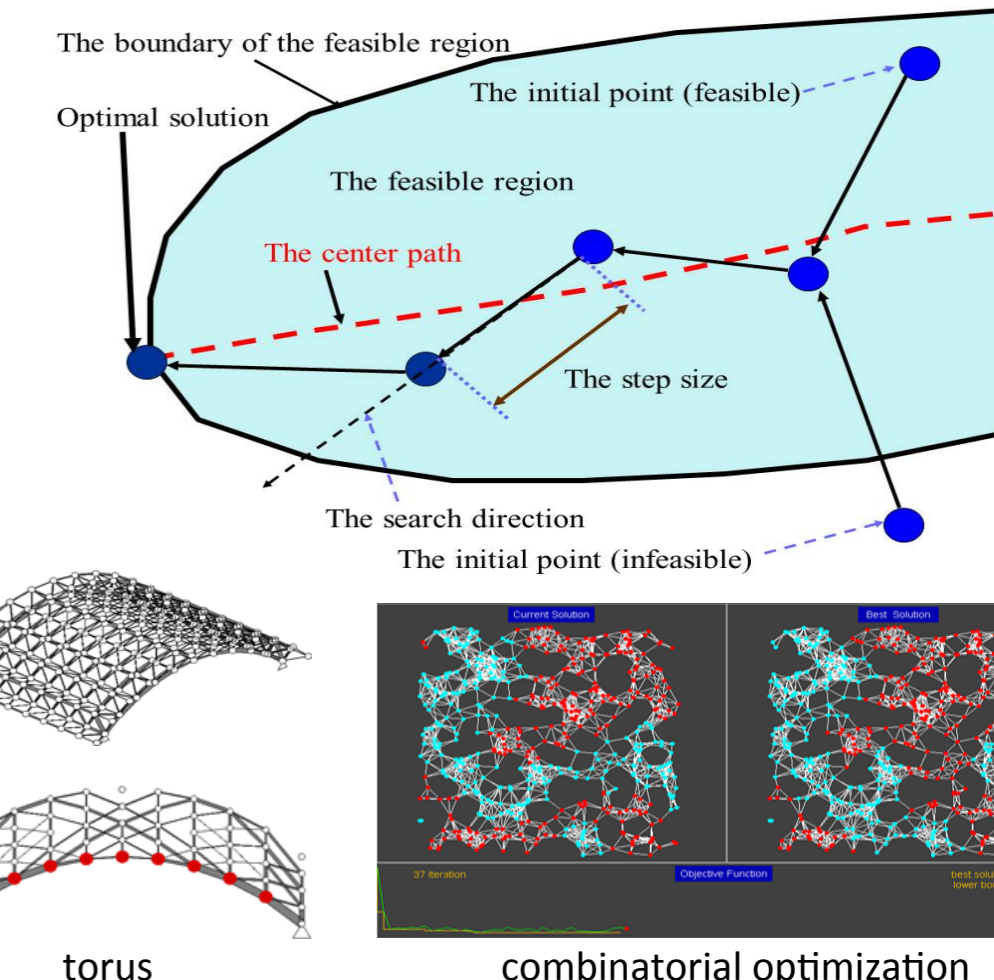
$$\begin{aligned} & \text{minimize } A_0 \bullet X \\ & \text{subject to } A_p \bullet X = b_p \quad (p = 1, 2, \dots, m), \quad X \in S_+^n \end{aligned} \quad (1)$$

$$\begin{aligned} & \text{maximize } \sum_{p=1}^m b_p z_p \\ & \text{subject to } \sum_{p=1}^m A_p z_p + Y = A_0, \quad Y \in S_+^n \end{aligned} \quad (2)$$

The matrix size

S^n : $n \times n$ symmetric matrices
 $X \bullet Y = \sum_{i=1}^n \sum_{j=1}^n X_{ij} Y_{ij}$
 $X \in S_+^n$ (S_+^{++}): $X \in S^n$ is positive semidefinite (or positive definite)
 $A_p \in S^n$ ($p = 0, 1, \dots, m$) and $b \in \mathbb{R}^m$

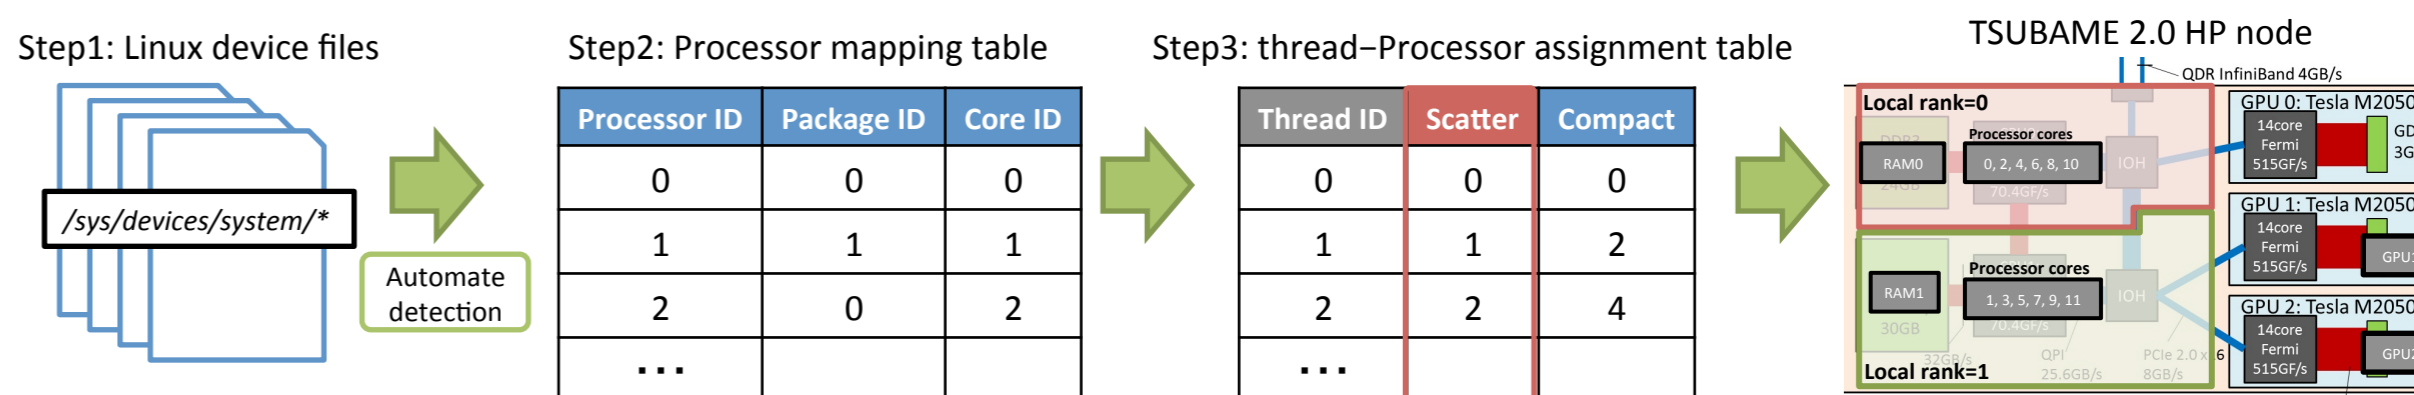
1. Choose an initial point (X, Y, z) with $X \succ 0, Y \succ 0$.
2. Compute Search Direction (dX, dY, dz) .
3. Compute Step Length α_p, α_d .
 $X + \alpha_p dX \succ 0, Y + \alpha_d dY \succ 0$
4. Update (X, Y, z)
 $\leftarrow (X + \alpha_p dX, Y + \alpha_d dY, z + \alpha_d dz)$.
5. Goto 2 if (X, Y, z) is not close to optimal.



The major bottleneck parts of PDIPM for the various types of SDP problems

1. **CHOLESKY**: Sparse SCM
 e.g.) the sensor network location problem and the polynomial optimization problem
 Parallel Sparse Cholesky factorization of sparse SCM
2. **ELEMENTS**: $m < n$ (not $m \gg n$), and Fully Dense SCM
 e.g.) the quantum chemistry problem and the truss topology problem
 Efficient Hybrid (MPI-OpenMP) parallel computation
 Automatic configuration for CPU Affinity and memory interleaving
3. **CHOLESKY**: $m \gg n$, and Fully Dense SCM
 Massively parallel dense Cholesky factorization using GPUs
 e.g.) the combinatorial optimization problem the quadratic assignment problem (QAP)
 Overlapping techniques for computation, PCI-Express comm., and MPI comm.
 Our implementation achieved **1.018 PFlops** in double precision for large-scale Cholesky factorization using 2,720 CPUs and 4,080 GPUs.

Automate configuration of CPU Affinity and memory allocation policy



Parallel Computation for ELEMENTS

```

set_mempolicy(interleaving)
for thread_id = 0, 1, ..., p-1 do
  packages ← packages ∪ { get_package_id(thread_id, scatter) }
end
set_mempolicy(packages, MPOL_INTERLEAVE) // set

B = 0
for l = 1, 2, ..., h do
  for j ∈ {j | F_j ≠ 0} parallel(MPI+thread) do
    set_affinity(scatter)
    for i ∈ {i | F_i ≠ 0} do
      Selects F_1, F_2 or F_3 and compute B_ij^l
      B_ij = B_ij + B_ij^l
    end
  end
  unset_affinity()
end
  
```

#omp parallel {
 thread_id ← omp_get_thread_num()
 pinned(thread_id, scatter)
 ...
 unpinned()
}

Fig. CPU time for Electronic4

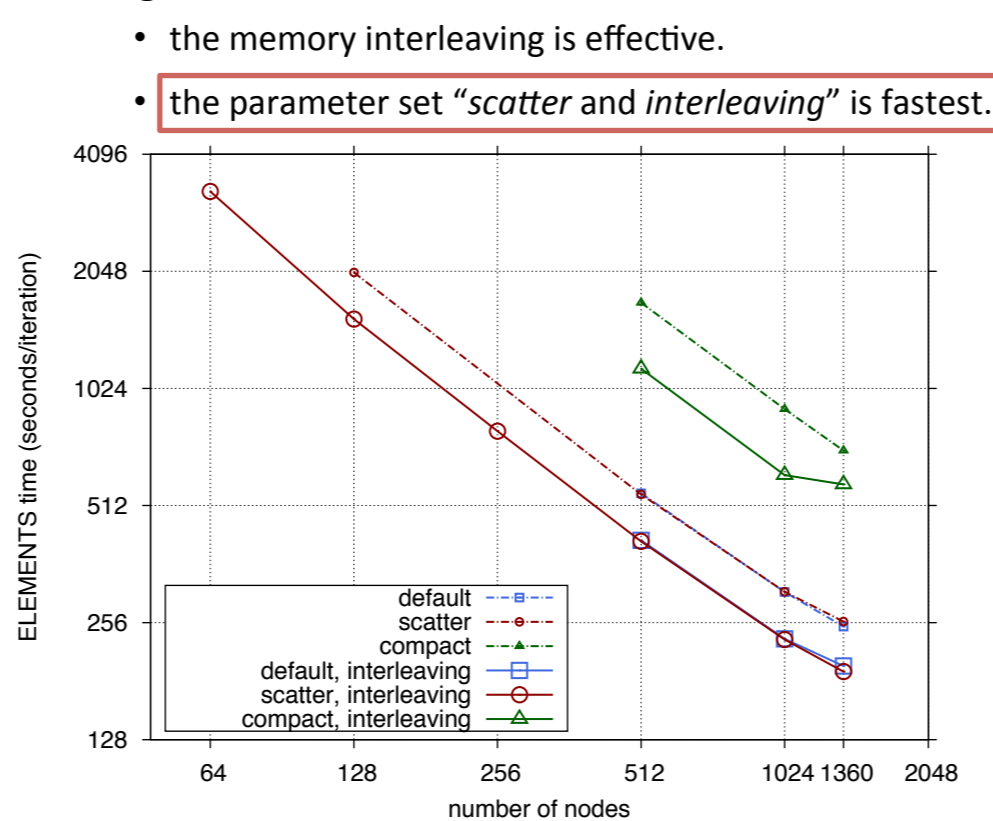
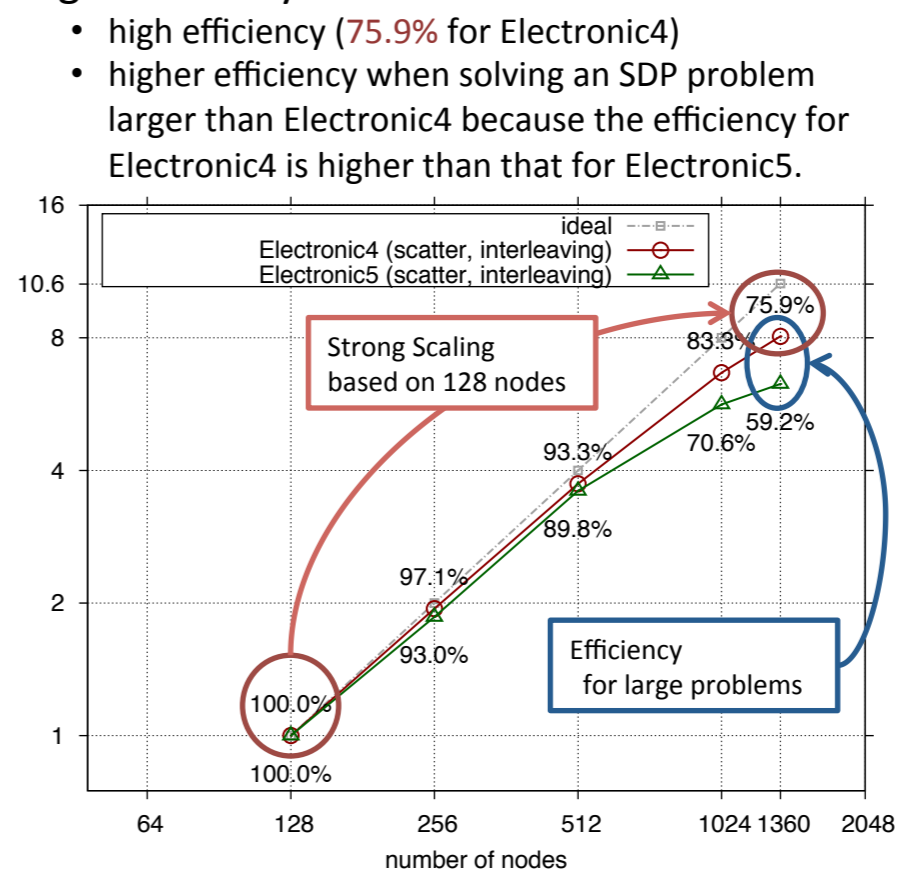
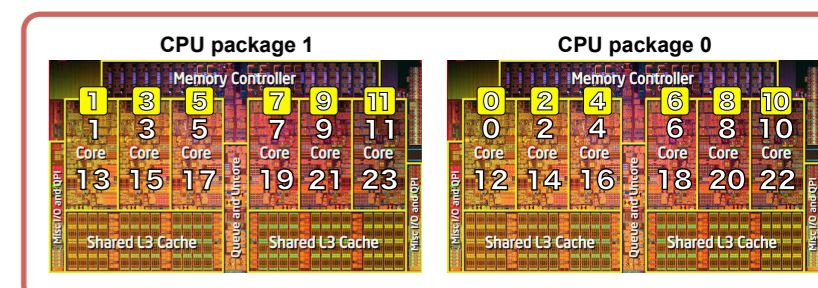


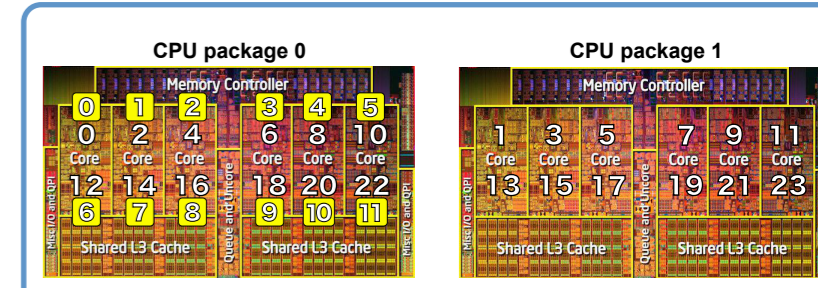
Fig. Scalability for Electronic4 and Electronic5



Scatter-type Affinity distributes OpenMP threads as evenly as possible across the entire system.



Compact-type Affinity binds the (n+1)-th OpenMP thread in a free-thread context as close as possible to the thread context in which the n-th OpenMP thread was bound.

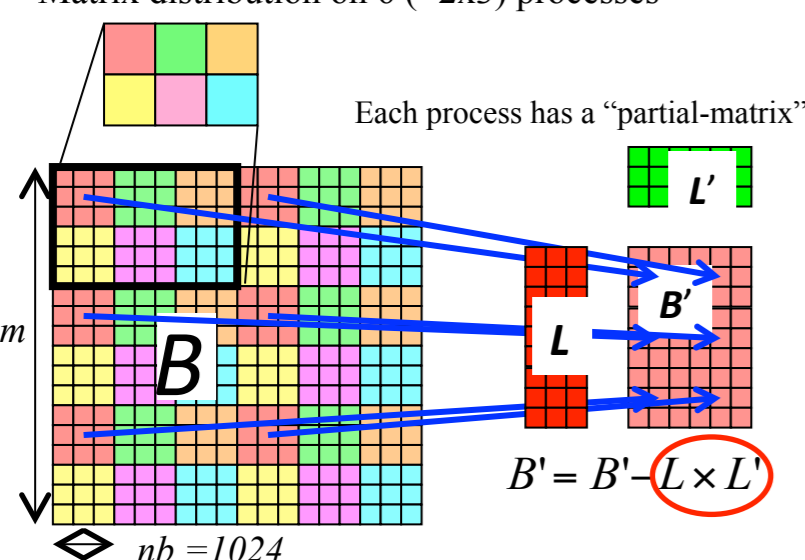


Parallel Computation for CHOLESKY

For problems with $m \gg n$, high performance **CHOLESKY** is implemented for GPU supercomputers. Key for petaflops is **overlapping computation, PCI-Express communication and MPI communication.**

Data Decomposition

- The dense matrix $B(m \times m)$ is distributed with 2D Block-Cyclic distribution with block size nb
 Matrix distribution on 6 (=2x3) processes



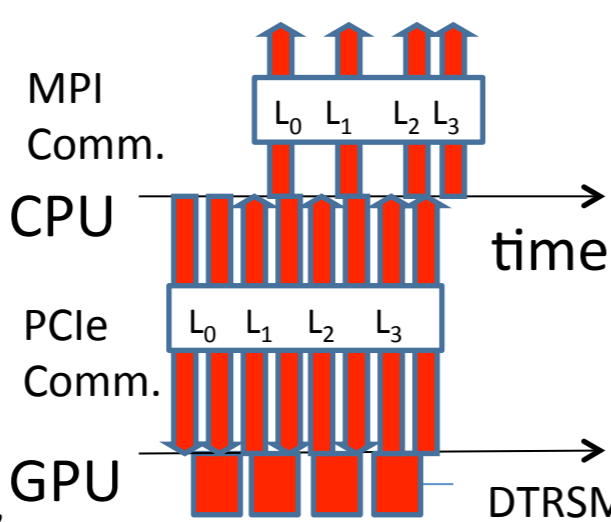
Basic Algorithm

- For $k = 0, 1, 2, \dots, \lceil m/nb \rceil - 1$
1. Diagonal block factorization
 2. Panel factorization
 → Compute L by GPU DTRSM
 3. Broadcast L, transpose L (L') and broadcast L'
 4. Update $B' = B' - L \times L'$ with fast GPU DGEMM

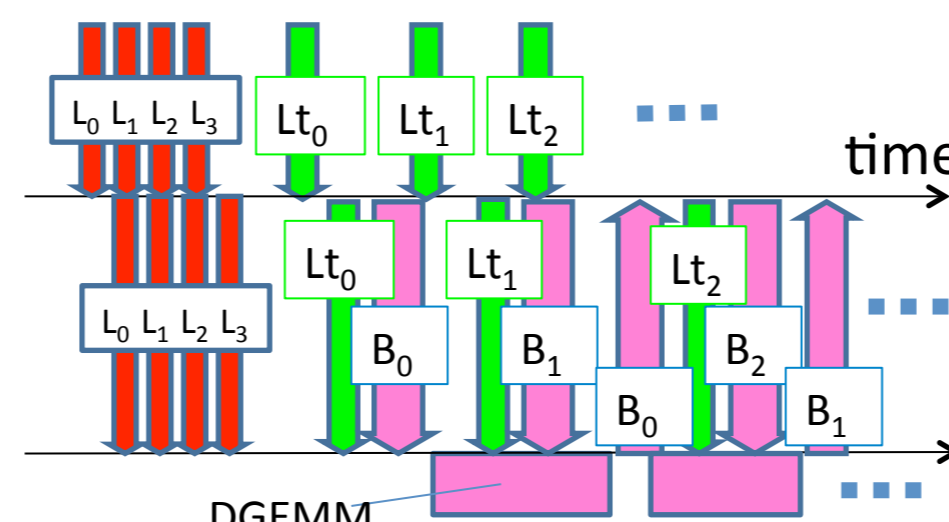
Design Strategy

- Our target is large problems with $m > 2$ million
 → GPU memory is too small. Matrix data are usually placed on host memory
- Blocksize nb should be sufficiently large to mitigate CPU-GPU PCIe communication
 - We still suffer from heavy PCIe communication
 → $nb=1024$
 → Overlap GPU comp., PCIe comm., and MPI comm. in Step 2, 3 and 4

Producer of panel L

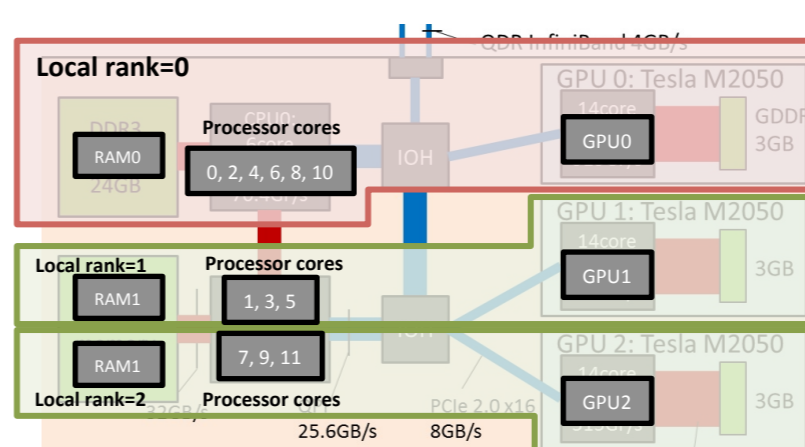


Consumers of panel L

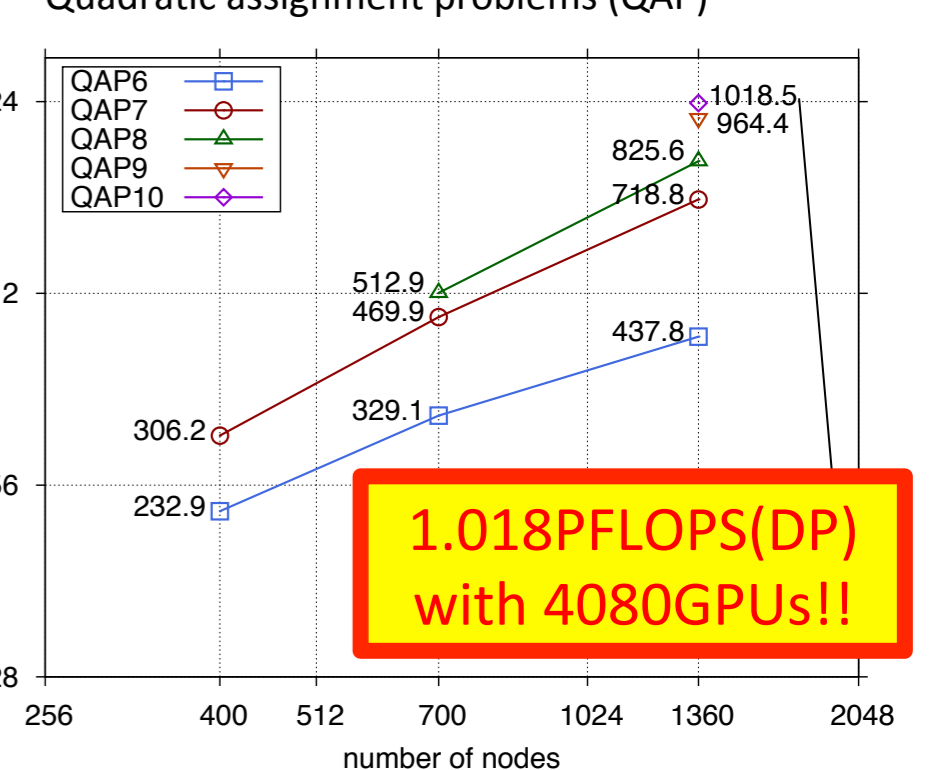


CPU/GPU Affinity

- 1 MPI process per GPU
 → 3 processes per node
- CPU/memory affinity is determined in a GPU centric way



Performance of CHOLESKY on TSUBAME2.0



SDPARA can solve the largest SDP problem

- DNN relaxation problem: QAP10 QAPLIB with **2.33 million constraints**
- Using 1360 nodes, 2720 CPUs, 4080 M2050 GPUs
- **1.018PFLOPS** in CHOLESKY (2.33m x 2.33m)
- The **fastest and largest result** as mathematical optimization problems!!