Использование неоднородной архитектуры кластера НКС-30Т и техники аппроксимации матрицами малого ранга в прямом методе для решения 3-мерного уравнения Гельмгольца

ИНГГ СО РАН, г.Новосибирск

Соловьев С.А.

Statement of problem

Solve the Helmholtz problem

$$u = ?$$

$$\Delta u + \frac{(2\pi\nu)^2}{V^2} u = f$$

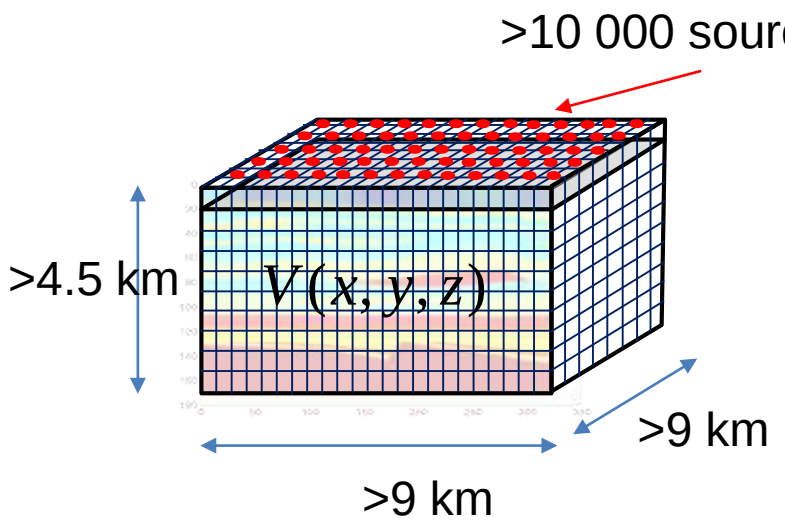

- ✓ Velocity model *V*= 2300м/с ... 6000м/с
- ✓ Frequency 1,…, 16Гц

- ✓ Parallelepipedal grid, step is 30m
- ✓ Perfect Matching Layer (PML)
- ✓ Finite difference approximation

Solve the symmetric complex sparse SLAE

$$AX = B, \quad X = \{x_1, \ldots, x_{nrhs}\}, \quad \dim(A) = n \times n$$

n> **20***10^6, *nrhs*>**10^4**

Proposed algorithm, key points

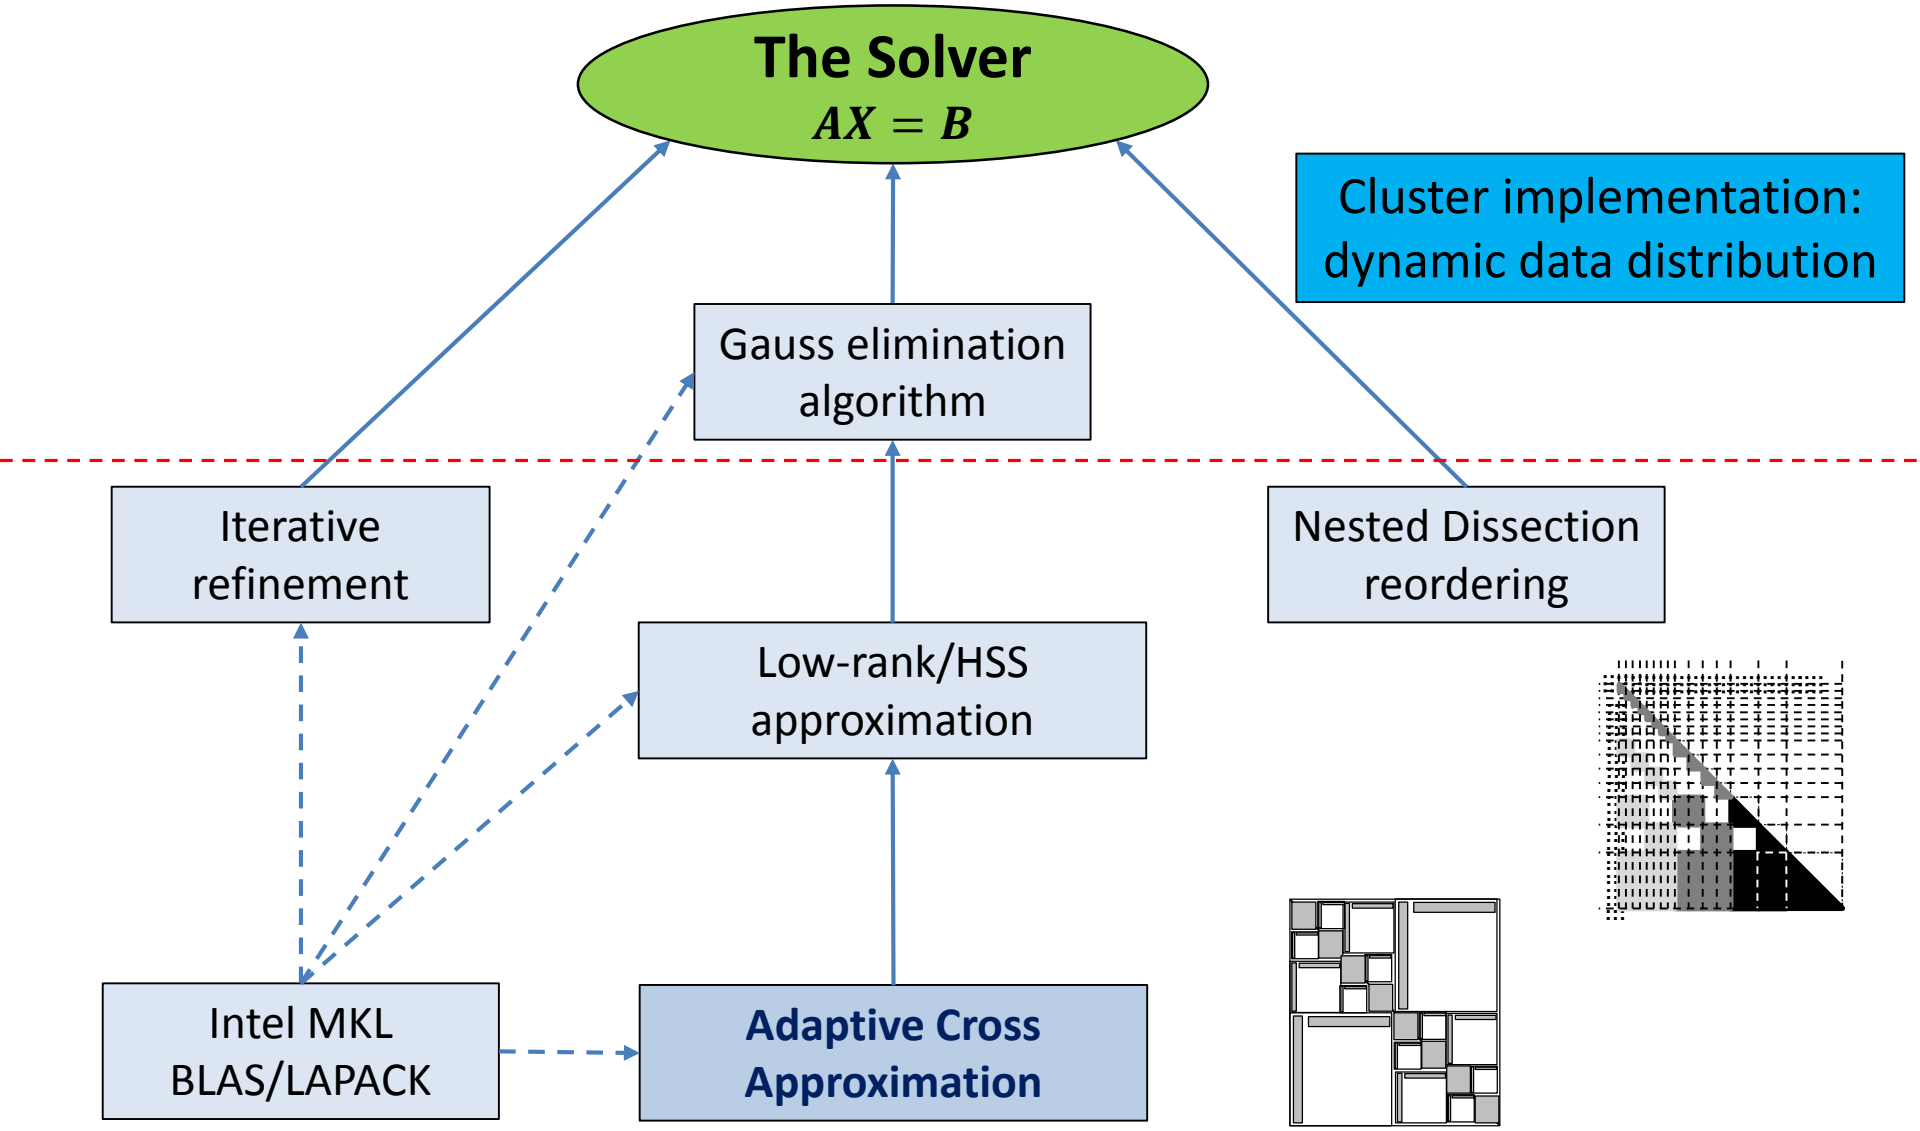

Nested Dissection reordering

$\hat{A} = PAP^t$

$\hat{A} = \hat{L}\hat{D}\hat{L}^t$

A

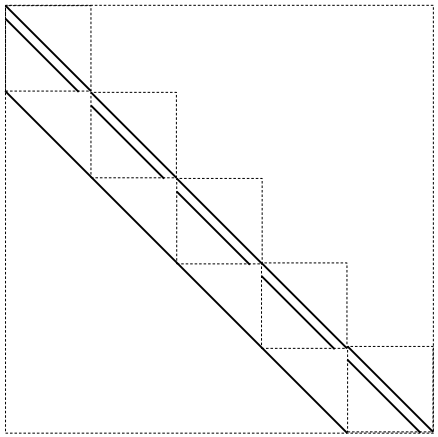

$\hat{A}$

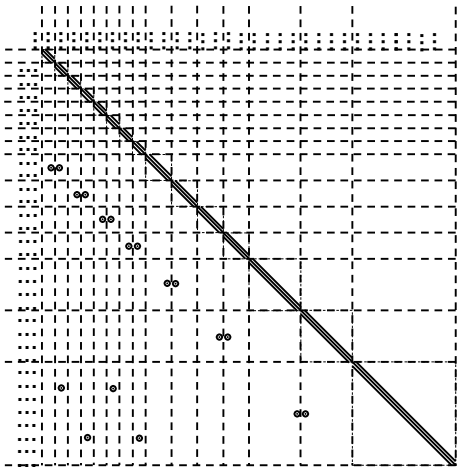

$\hat{L}$

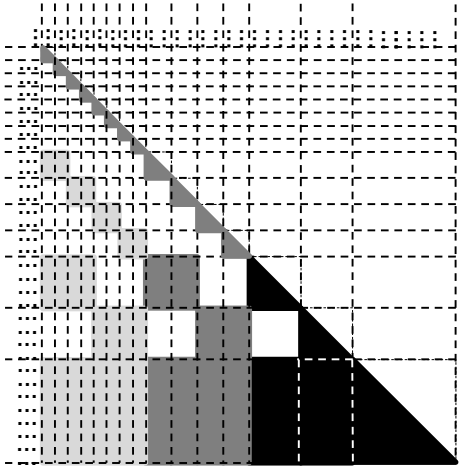

Low-rank approximation / HSS format

Fill-in of L-factor

$\hat{L}$

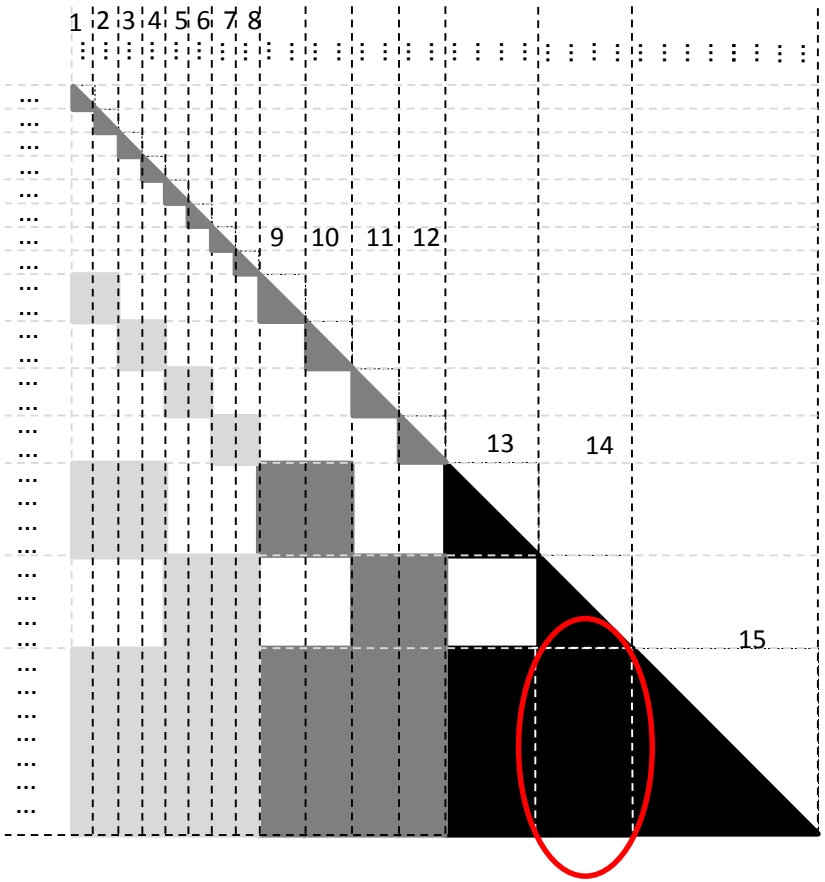

Exact (machine) arithmetic

$\tilde{L}$

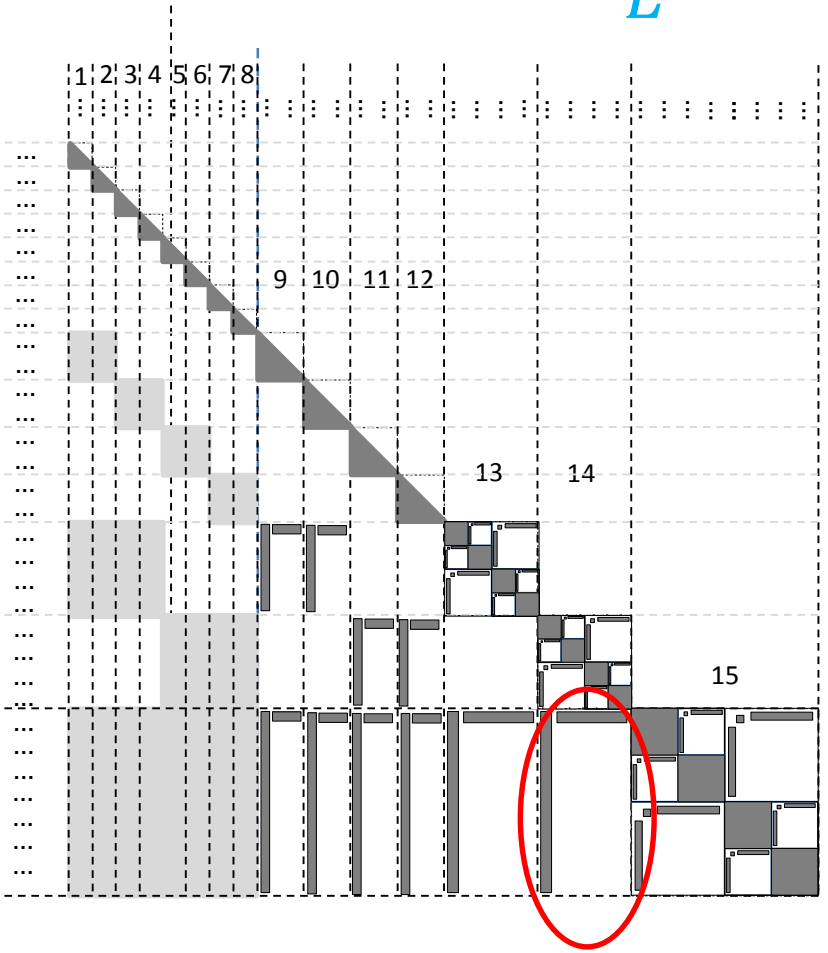

Low-rank approximation

Cluster implementation: Small problem

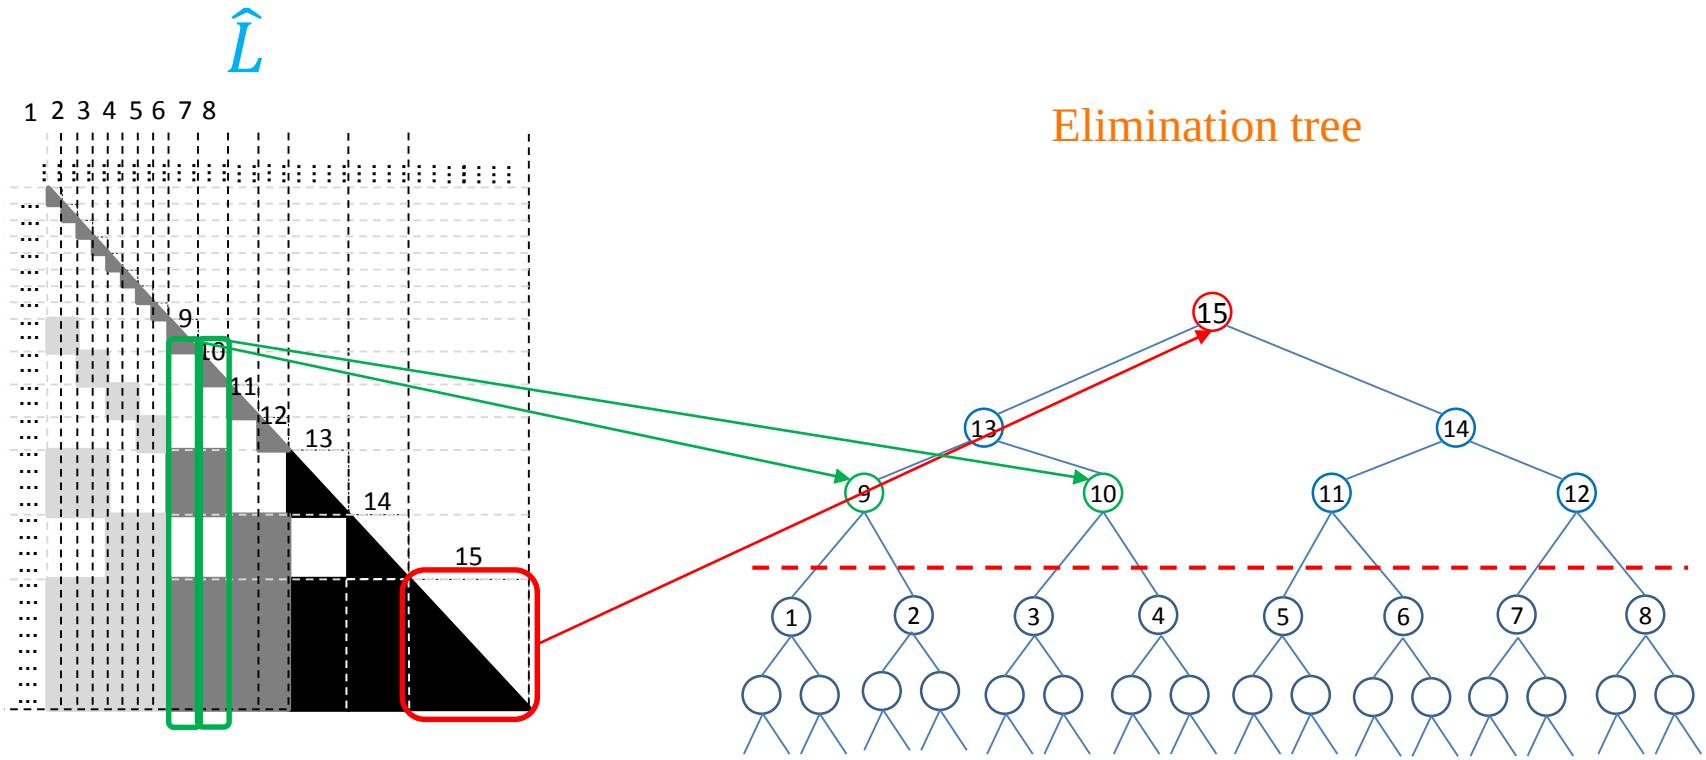

Cluster implementation: Small problem

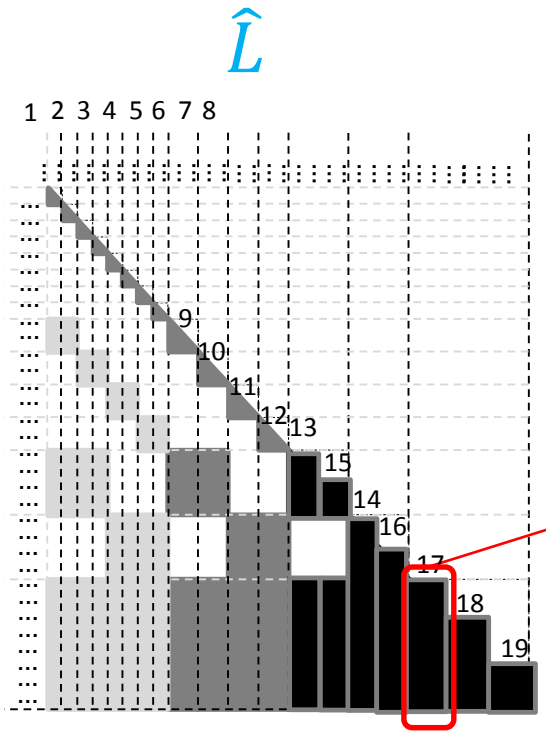

Elimination tree

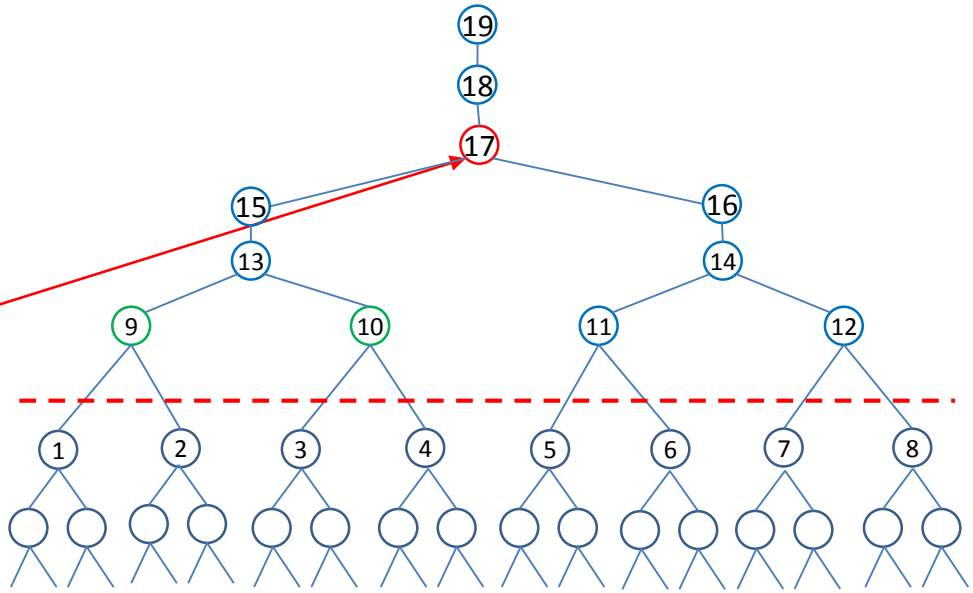

Cluster implementation: Small problem

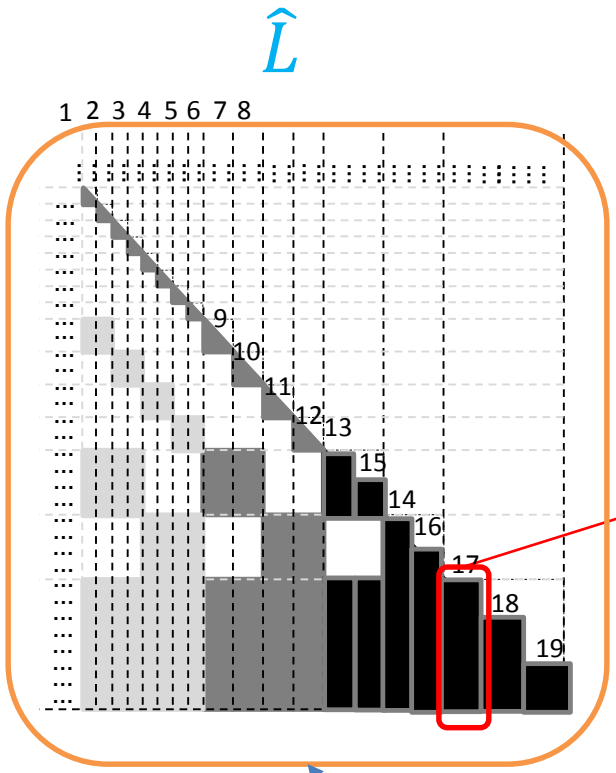

Elimination tree

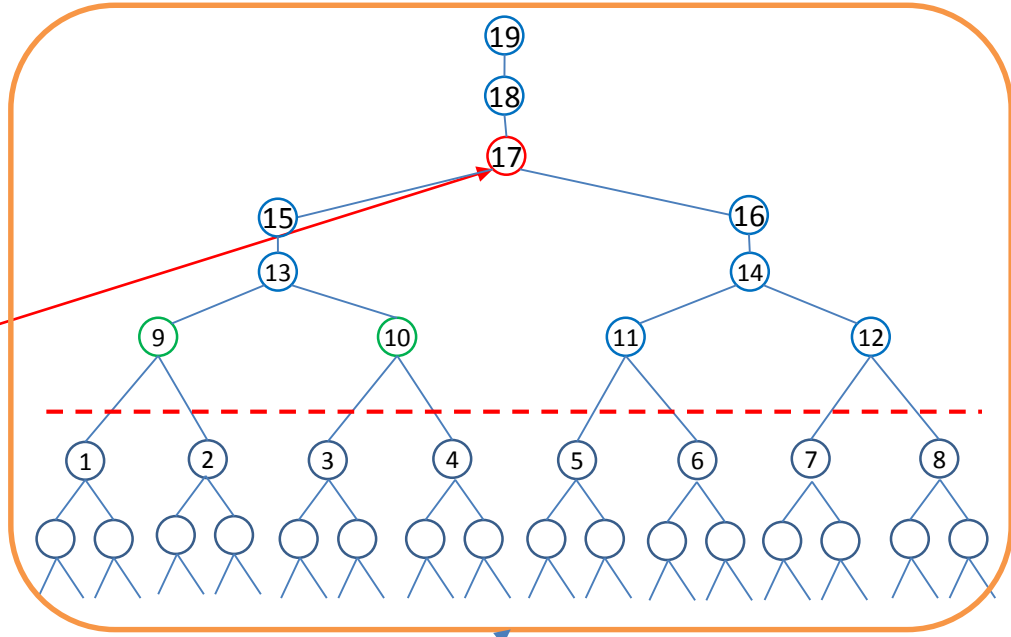

Total free RAM of all cluster nodes

Cluster implementation: Big problem

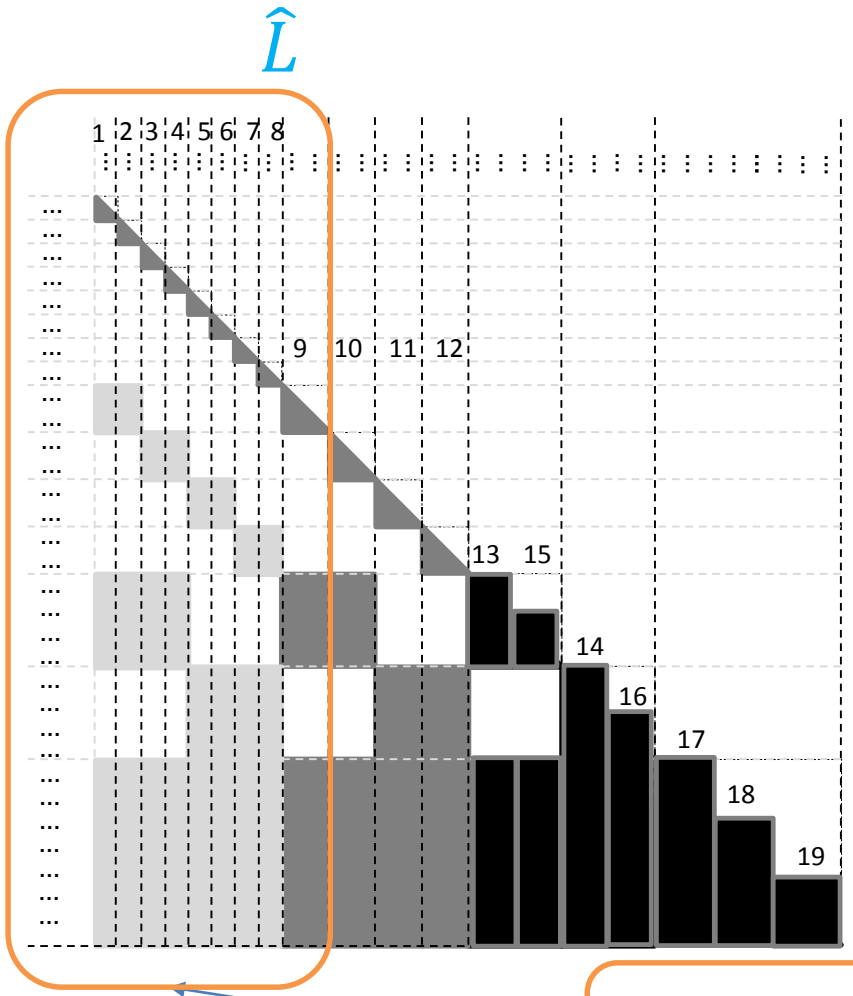

Elimination tree

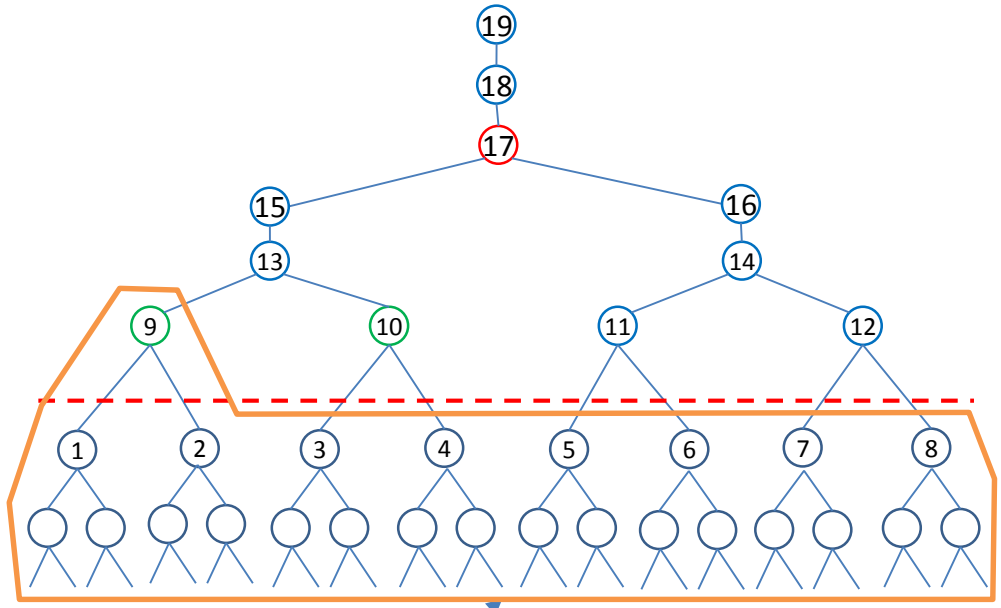

Total free RAM of all cluster nodes

Cluster implementation: Big problem

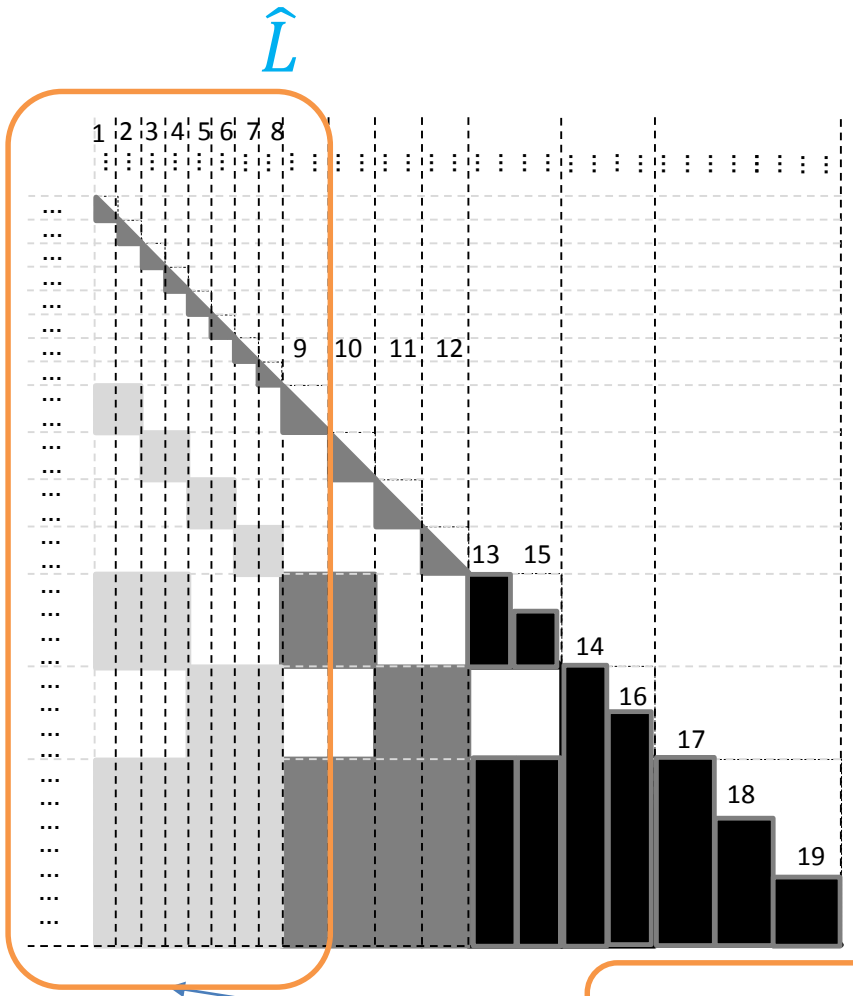

Elimination tree

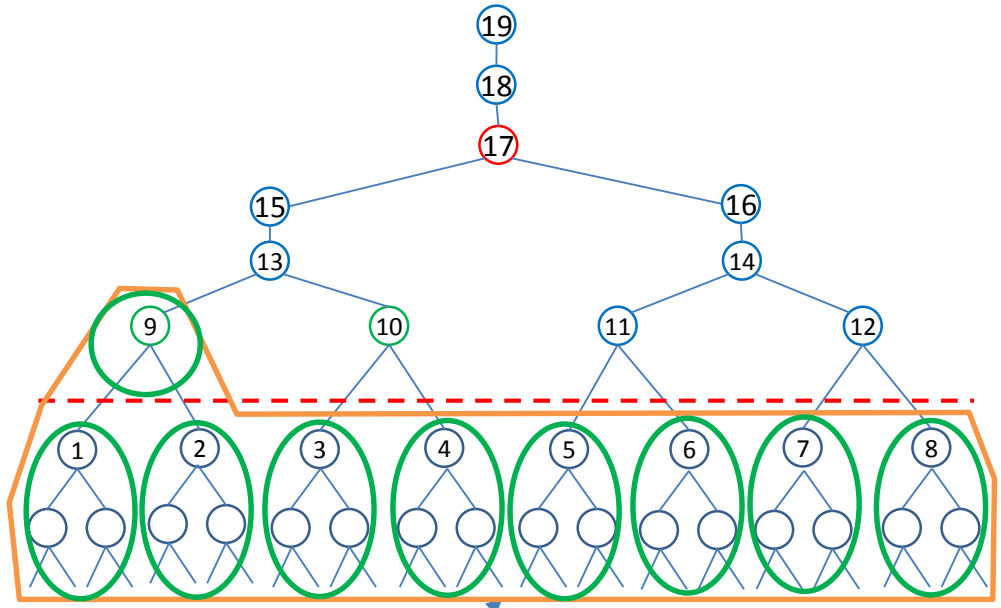

Total free RAM of all cluster nodes

We have 9 cluster nodes

Overcome the memory limitation (dynamic distribution):

"True" L-factors don't fit all RAM

1-st. step. Each cluster node factorizes one elimination node or one sub-tree.

2-nd. step. Two cluster nodes stores factors, seven – factorizes next elimination nodes.

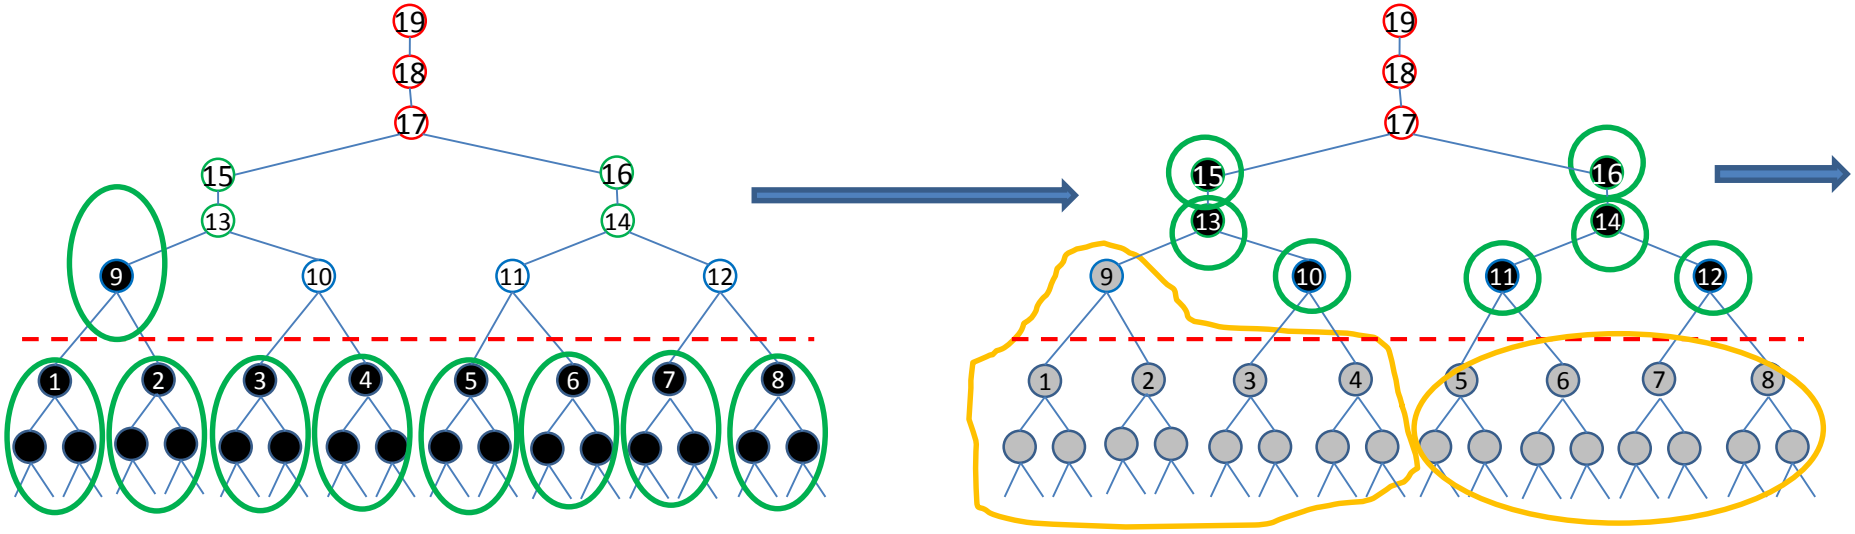

We have 9 cluster nodes

- cluster node perform factorization

- cluster node stores factors

Overcome the memory limitation (dynamic distribution):

"True" L-factors don't fit all RAM

i-st. step. some cluster nodes factorize elimination nodes, another – store factors.

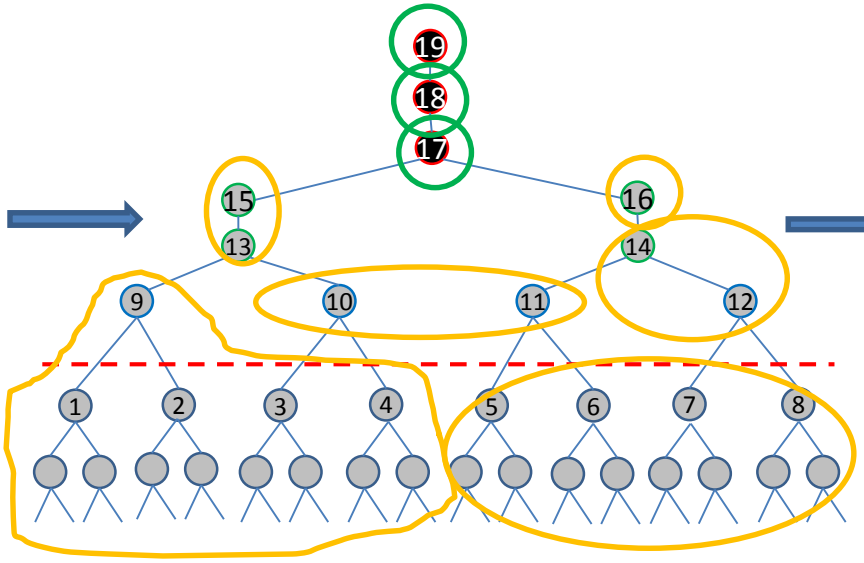

Last step. All cluster nodes stores factors

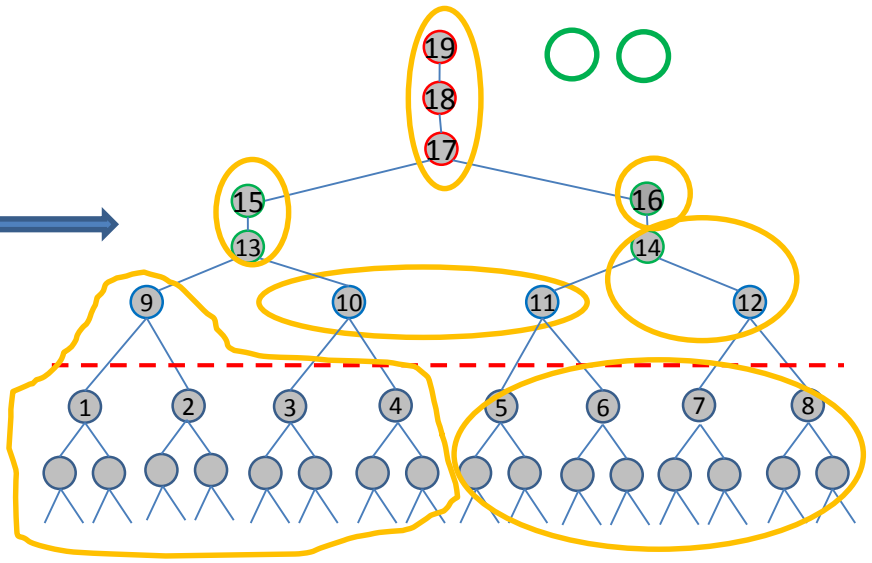

We have 9 cluster nodes

- cluster node perform factorization

- cluster node stores factors

Numerical experiments, test description

Cluster name	#nodes on cluster	RAM size on node (GB)	Proc. type (Intel) processor number, code name	#cores on node
NSK-30T	1 4	1000 96	Intel(R) Xeon(R) CPU E7-4870 @ 2.40GHz,SSE4 Intel(R) Xeon(R) CPU X5675 @ 3.07GHz, SSE4	8x10 2x6

Total free RAM ~ 1245GB = 90%*(1000+4*96)

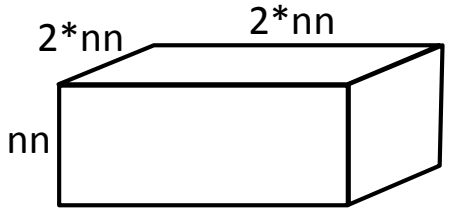

nn=100, 150, 170, 180, 190, 200, 220 :

- ✓ 3D domain, v=2400m/s., nu=4 Hz
- ✓ hx=hy=hz=30m,
- ✓ PML: width=10points, d0=450
- ✓ 27-point stencil;
- ✓ Eps_lowrank=10^(-6)
- ✓ Save L-factor after factorization step, **IBRIX**

Numerical experiments, test description

Cluster name	#nodes on cluster	RAM size on node (GB)	Proc. type (Intel) processor number, code name	#cores on node
NSK-30T	1 4	1000 96	Intel(R) Xeon(R) CPU E7-4870 @ 2.40GHz,SSE4 Intel(R) Xeon(R) CPU X5675 @ 3.07GHz, SSE4	8x10 2x6

Total free RAM ~ 1245GB = 90%*(1000+4*96)

nn=100, 150, 170, 180, 190, 200, 220

OMP_NUM_THREADS=8

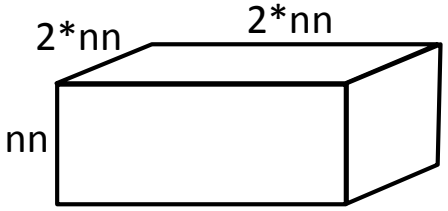

	100	150	170	180	190	200	220
n_eq, *10^6	5.8	17	24	29	33	38	50
Memory, GB	6+77	18+240	24+367	29+435	34+528	40+625	57+853
Time_fct	2 134	11 800	15 300	21 600	27 880	28 900	39 000
Time_slv	7	100	141	291		2661	----

Memory swapping, etc.

Solution step: memory swapping

Reading L-factor while solving LDL^tx=b

	total	used	free	shared	buffers	cached
Mem:	990	218	772	0	0	217
-/+ buffers/cache:		1	989			
Swap:	127	0	127			

	total	used	free	shared	buffers	cached
Mem:	990	990	0	0	0	318
-/+ buffers/cache:		671	319			
Swap:	127	**2**	125			

Memory swapping, etc.

	total	used	free	shared	buffers	cached
Mem:	990	990	0	0	0	217
-/+ buffers/cache:		772	217			
Swap:	127	**6**	121			

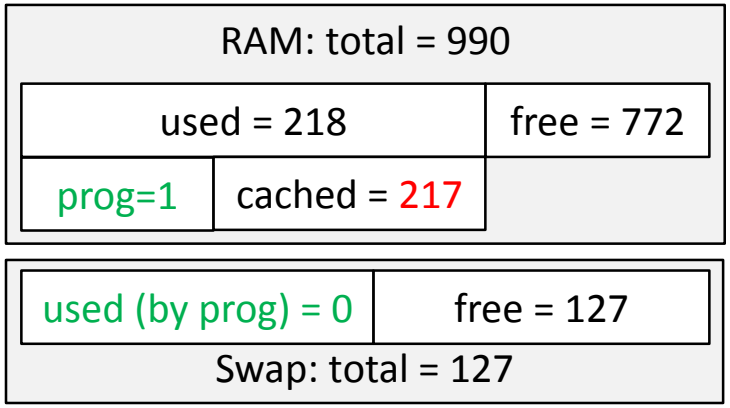

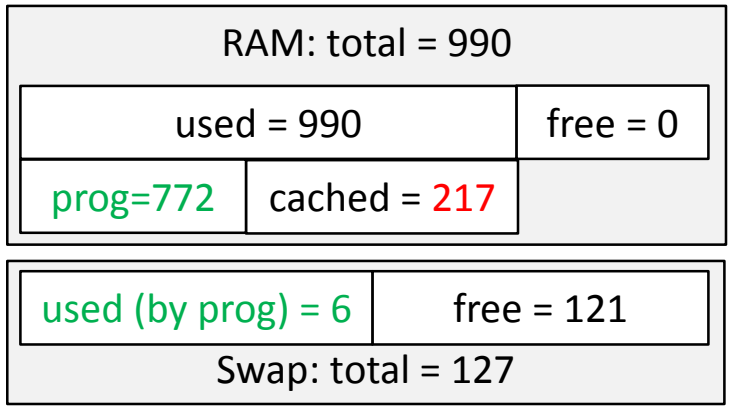

Solution step: memory swapping

Read L-factor before solve step

	total	used	free	shared	buffers	cached
Mem:	990	218	772	0	0	217
-/+ buffers/cache:		1	989			
Swap:	127	0	127			

Reading data…

	total	used	free	shared	buffers	cached
Mem:	990	990	0	0	0	313
-/+ buffers/cache:		676	314			
Swap:	127	0	127			

disk cache decreased

Solve step

	total	used	free	shared	buffers	cached
Mem:	990	990	0	0	0	210
-/+ buffers/cache:		779	211			
Swap:	127	0	127			

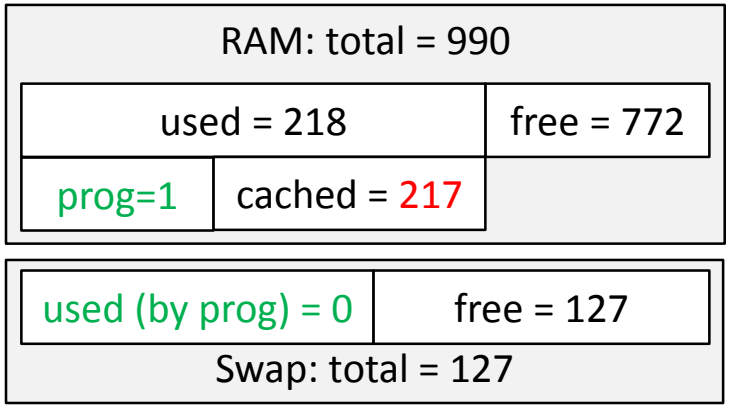

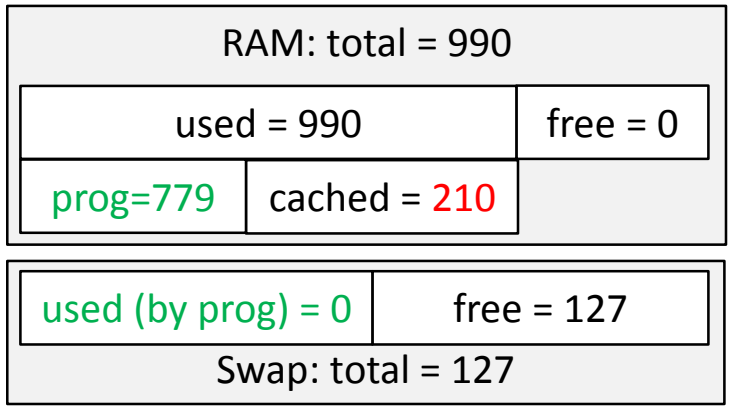

Numerical experiments, test description

Cluster name	#nodes on cluster	RAM size on node (GB)	Proc. type (Intel) processor number, code name	#cores on node
NSK-30T	1 4	1000 96	Intel(R) Xeon(R) CPU E7-4870 @ 2.40GHz,SSE4 Intel(R) Xeon(R) CPU X5675 @ 3.07GHz, SSE4	8x10 2x6

nn=220 :

- ✓ 3D domain (13200m x 13200m x 6600m), v=2400m/s., nu=4 Hz
- ✓ hx=hy=hz=30m,
- ✓ PML: width=10points, d0=450
- ✓ 27-point stencil;
- ✓ Grid = 460x460x240 (~50*10^6)
- ✓ Eps_lowrank=**10^(-6.5);** OMP_NUM_THREADS=1, 12

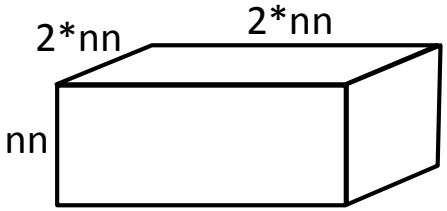

Numerical experiments

Cluster name	#nodes on cluster	RAM size on node (GB)	Proc. type (Intel) processor number, code name	#cores on node
NSK-30T	1 4	1000 96	Intel(R) Xeon(R) CPU E7-4870 @ 2.40GHz,SSE4 Intel(R) Xeon(R) CPU X5675 @ 3.07GHz, SSE4	8x10 2x6

Results:

- ✓ Reordering time = 2312s.=~38m.
- ✓ Factorization time:

  44034s.=~12h. (OMP_NUM_THREADS=12)

  ~48h. (OMP_NUM_THREADS=1)
- ✓ Solution time =196s. = 3-4m. (per 1rhs, OMP_NUM_THREADS=8)
- ✓ Size(L_low)=894G (~849G+~53G)
- ✓ Write time 9803s|612s (~90 MB/s)
- ✓ Read time 6204s…9083s|235s…581s (147…100 MB/s)

Visualization

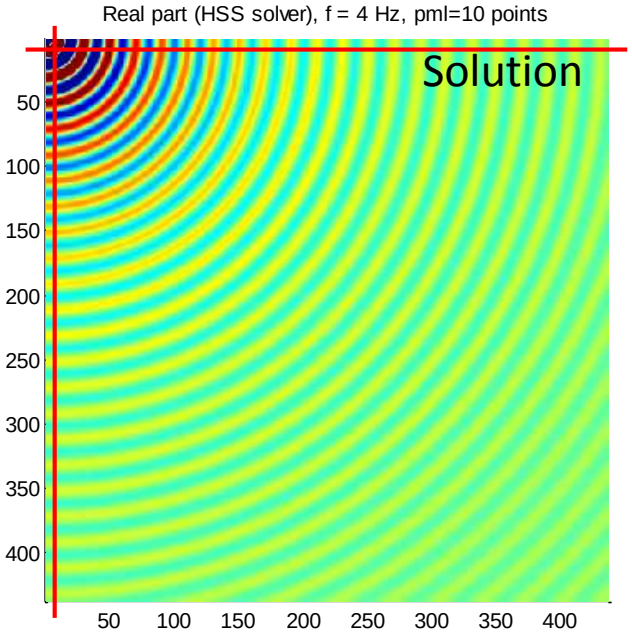

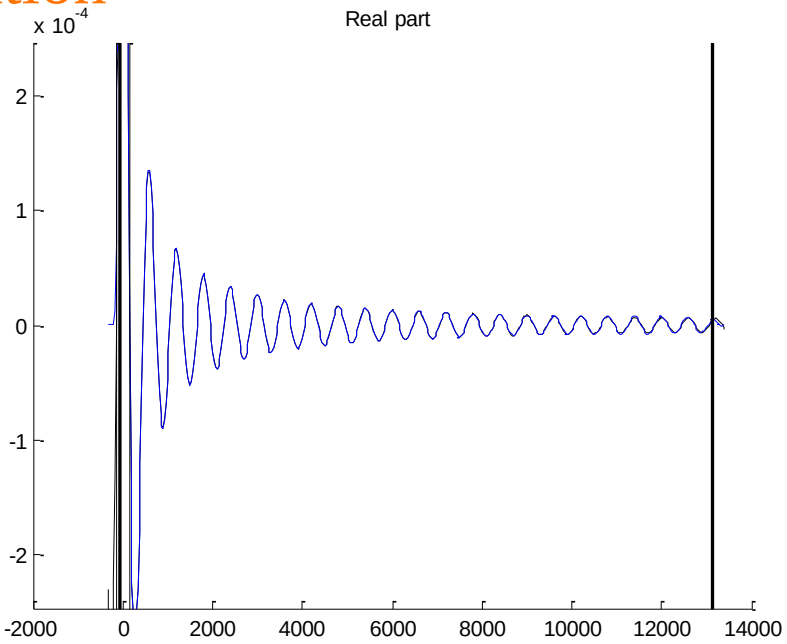

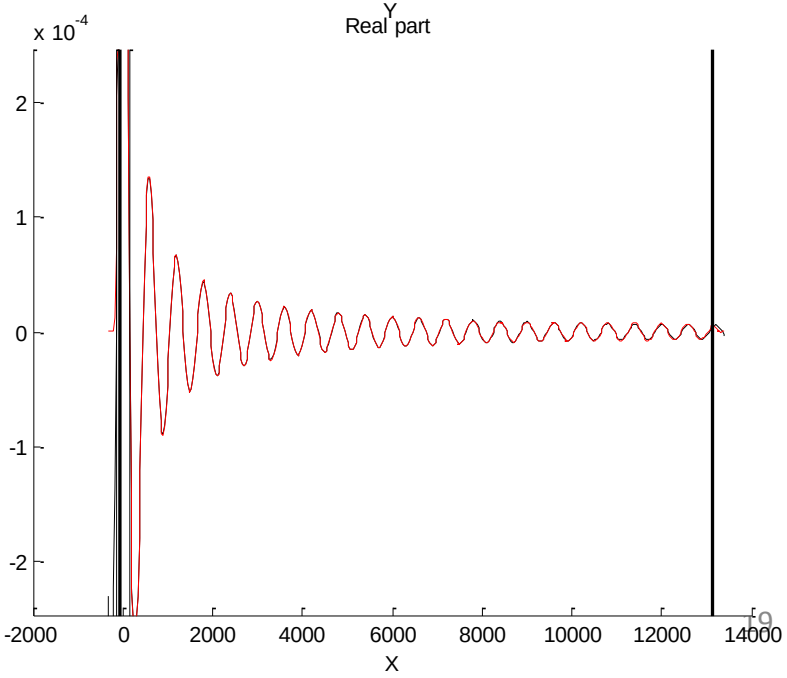

Future work

✓ Solve the real geophysical model on the cluster NSK-30T

Cluster name	#nodes on cluster	RAM size on node (GB)	Proc. type (Intel) processor number, code name	#cores on node
NSK-30T	1 **>20**	1000 96	Intel(R) Xeon(R) CPU E7-4870 @ 2.40GHz,SSE4 Intel(R) Xeon(R) CPU X5675 @ 3.07GHz, SSE4	8x10 2x6

Road map

- ✓ Adapt our solver to handle with **3D Elasticity problem**
- ✓ Use 3D finite-element approximation => solve **sparse matrices with general pattern**
- ✓ Solve **1000*1000*1000** acoustic problem (scalar Helmholtz)

Q&A

BACKUP