Korea Institute of
Science and Technology Information
www.kisti.re.kr

My Internship Experience:
Development and Research Projects

[2024.04.22 – 2024.07.19]

Presenter: Chanyoung Ahn

Overview of projects

1. **Lecture Materials**: developed three lecture materials on MPI
2. **Git Website**: documentation for PaScaL_TDMA library
3. **Benchmark**: performance benchmark of large sparse matrix using three libraries

1 Lecture Materials

https://github.com/cold-young/2024_KISTI_Intern/

Developed three lecture materials on MPI

* MPI 병렬컴퓨팅교육 운영 2회 (강지훈, 권오경) : 상반기 1회 완료

파이썬을 이용한 MPI 병렬프로그래밍

Wrote up 6+ jupyter notebook materials for MPI Workshop

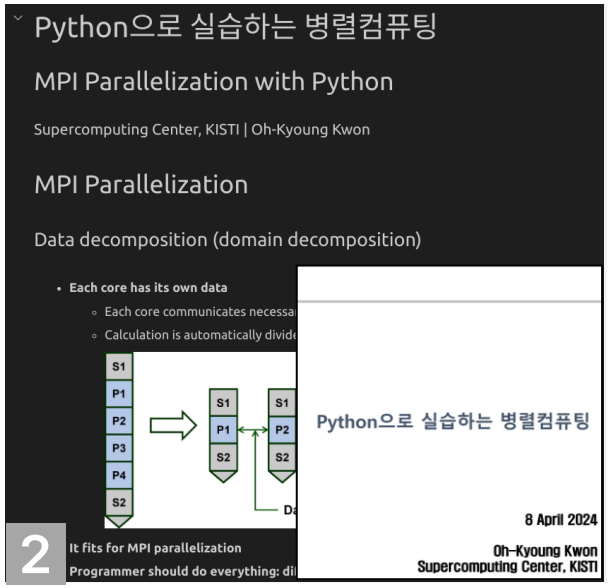

Python으로 실습하는 병렬컴퓨팅

Wrote up a jupyter notebook material for MPI Workshop

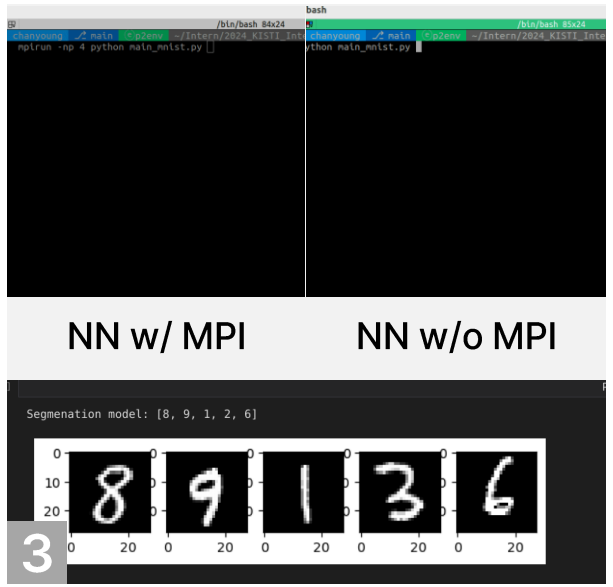

MNIST segmentation NN w/ MPI

Converted NN w/ MPI example to Python from C code

Overview of projects

1 Lecture Materials: developed three lecture materials on MPI

2 Git Website: documentation for PaScaL_TDMA library

3 Benchmark: performance benchmark of large sparse matrix using three libraries

엑사스케일 선형 해석자 수치 라이브러리 개발 / HPC Library for Exascale Linear Solvers

1단계 - 2차 연도(2024)													
추진내용	추진 일정												책임기관 (소속기관)
	1	2	3	4	5	6	7	8	9	10	11	12	
2 수치 라이브러리 통합 저장소 구축													강지훈 (KISTI)
3 국가슈퍼컴퓨터 5호기 대상 최적 병렬화 및 성능 평가													

2 Git Website

https://xccels.github.io/PaScaL_TDMA/

Documentation for PaScaL_TDMA library

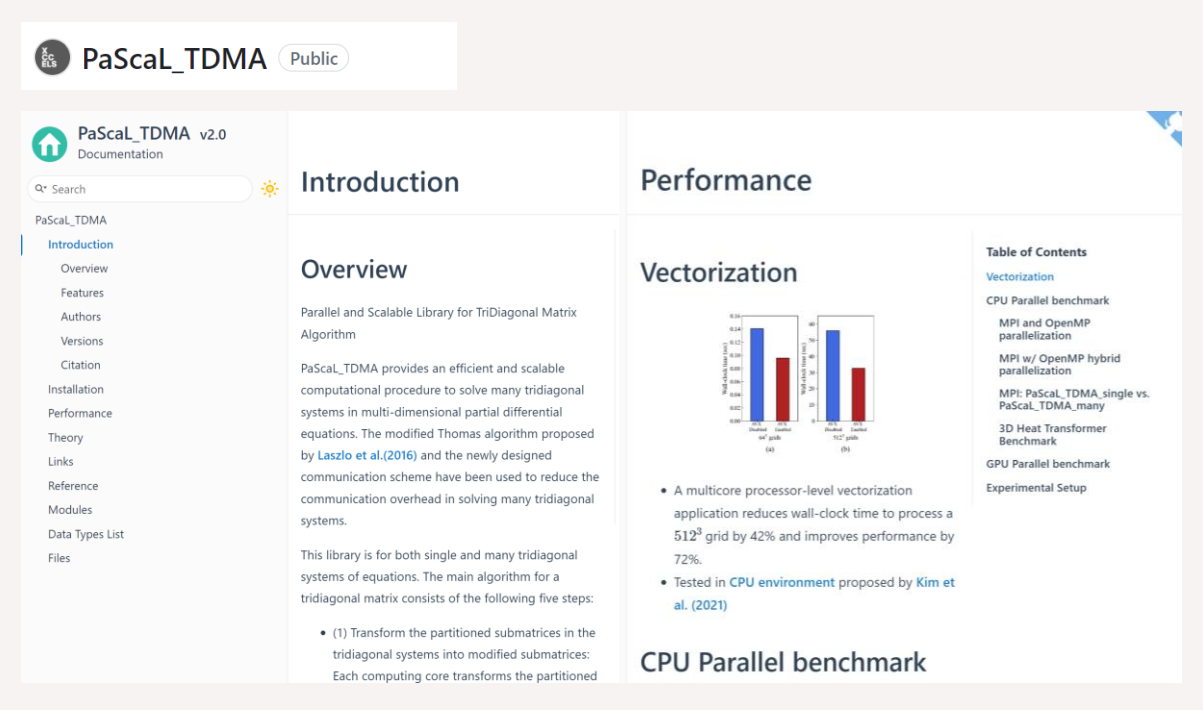

Why

Second year targets of
엑사스케일 선형 해석자 수치 라이브러리 개발

2 수치 라이브러리 통합 저장소 구축	강지훈 (KISTI)

Role

- Develop GitHub repository & PaScaL_TDMA Library documentation site

Result

- Created xccels/PaScaL_TDMA repository
- Wrote up 5+ documentation web pages

3 Benchmark

Performance benchmark of large sparse matrix using three libraries

Parallel performance benchmark of large sparse matrix
using SuperLU-DIST, Strumpack and PaScaL_TDMA

*1안찬영, 1권오경 #1강지훈

CDE

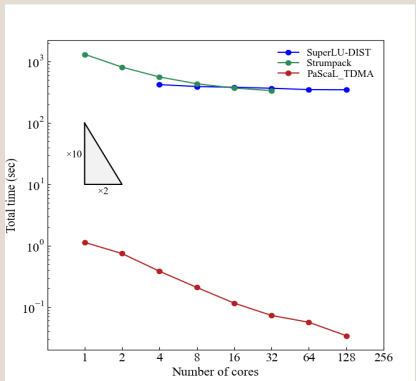

	Total time (sec.)
SuperLU-DIST	348.092226
Strumpack	337.743508
PaScaL_TDMA	**0.034142**

Why

Second year targets of
엑사스케일 선형 해석자 수치 라이브러리 개발

3	국가슈퍼컴퓨터 5호기 대상 최적 병렬화 및 성능 평가	강지훈 (KISTI)

Role

- Parallel benchmark of large sparse matrix using three libraries

Result

- PaScaL_TDMA is best suited to compute large tridiagonal matrices.
- (ongoing) Conference poster for targeting CDE in 2024

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

*1안찬영, 1권오경 #1강지훈

Problem We need to find suitable sparse matrix solvers that enhance efficiency.

Challenge Sparse matrix solver encounter challenges in computational efficiency, becoming a computational bottleneck.

Solution We provide a comparative performance benchmark of three parallel solvers.

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Experimental setup

Benchmark total and factorization times with nine large sparse matrices

Comparative performance benchmark of three parallel solvers; SuperLU-DIST, Strumpack, and PaScaL_TDMA

* All the computations were executed on the Nurion manycore cluster at KISTI.

	r/c	nnz	ratio
g7jac200	59,310	717,620	2.38E-04
torso1	116,158	8,516,500	6.31E-04
twotone	120,750	1,206,265	8.40E-05
stomach	213,360	3,021,648	6.64E-05
ML_Laplace	377,002	27,582,698	1.95E-04
RM07R	381,689	37,464,962	2.57E-04
mc2depi	525,825	2,100,225	7.60E-06
ecology1	1,000,000	4,996,000	5.00E-06
256^3 TDM	16,777,216	50,331,646	1.79E-07

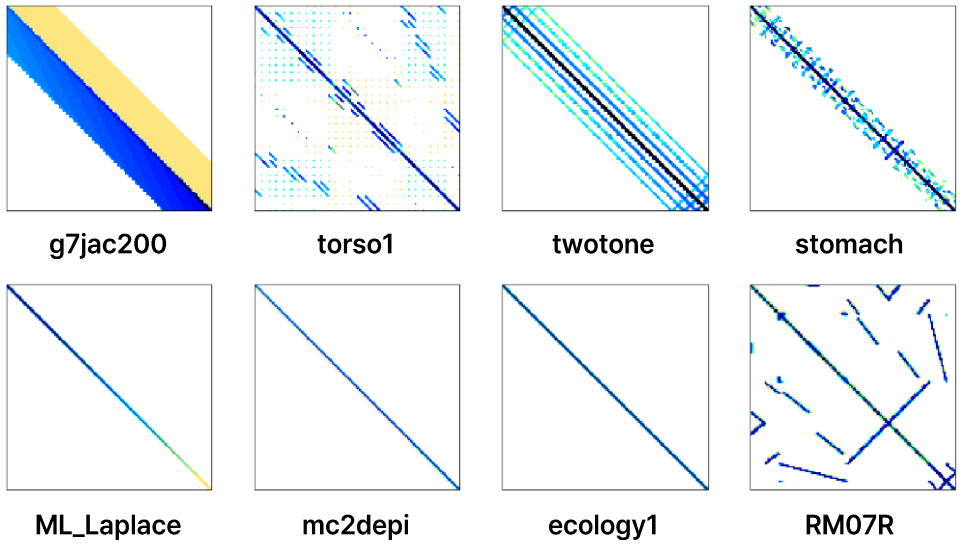

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Overall performance

Benchmark of Three Parallel Solvers on a 256^3 Tri-diagonal Matrix

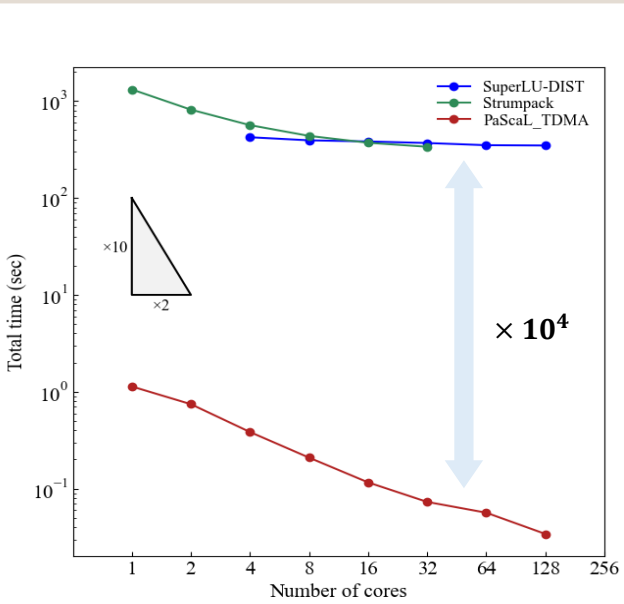

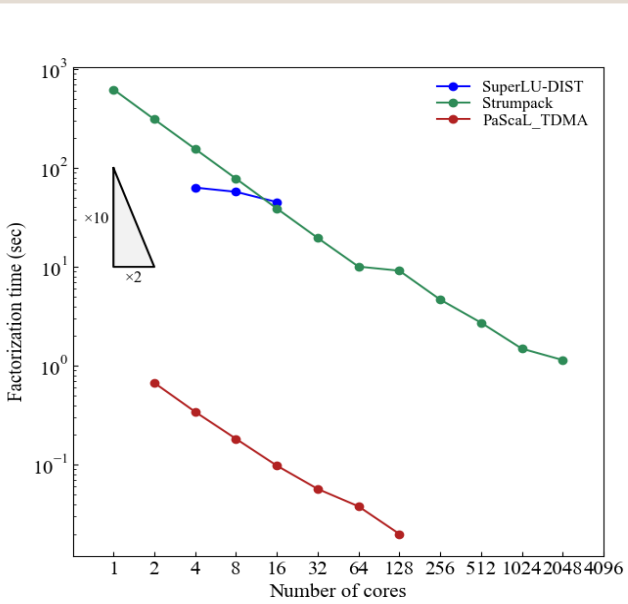

	Total time (sec)	speed up (2^5)
SuperLU-DIST	348.092226	1.2
STRUMPACK	337.743508	3.9
PaScaL_TDMA	**0.034142**	**15.5**

* A 256^3 tri-diagonal matrix was tested in Nurion normal nodes, with a range of 1–4096 cores.

- PaScaL_TDMA computes the 256^3 tri-diagonal matrix in the shortest total time of 0.034142 seconds.
- The solver also shows the best improvement in computation speed due to parallelization.

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Baseline 1: SuperLU-DIST (1/3)

Optimization of –r –c Parameters for RM07R

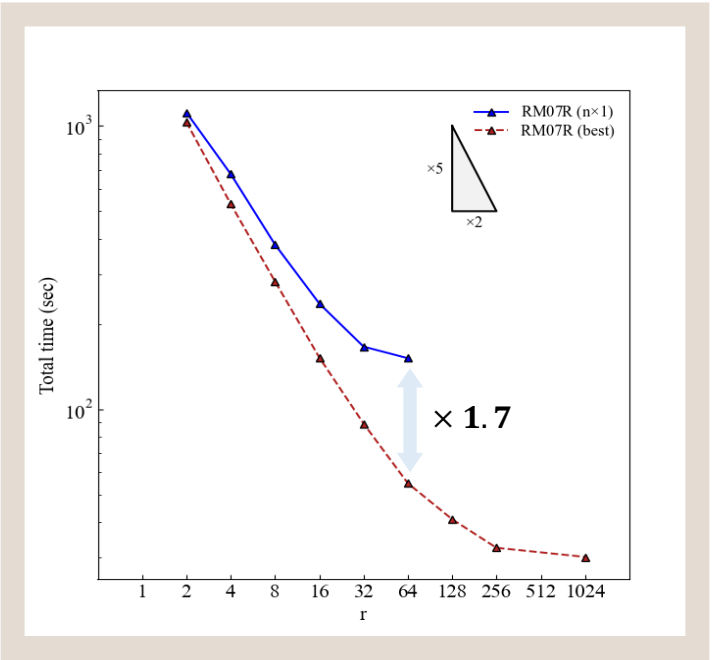

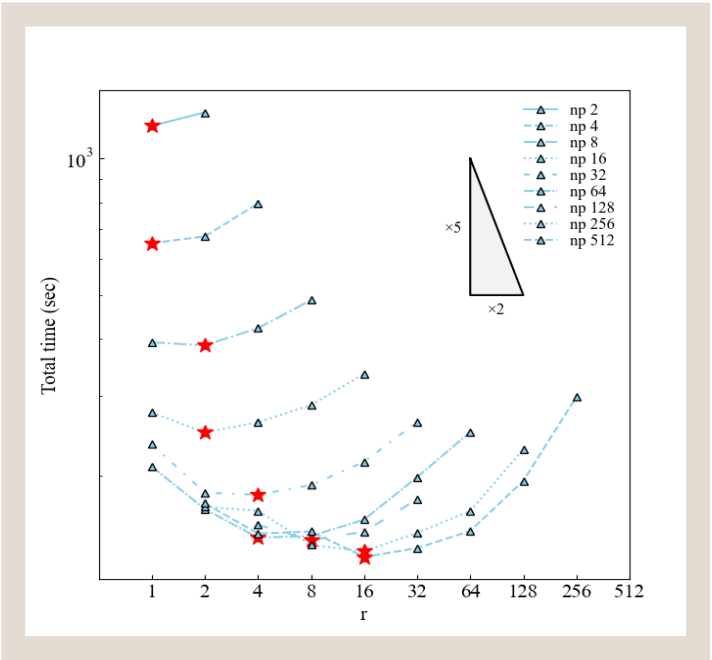

np	r	c
2	1	2
4	1	4
8	2	4
16	2	8
32	4	8
64	4	16
128	8	16
256	16	16
512	16	32

- Two parameters, -r and -c, affect the performance of total and factorization times.
- As the number of cores increases, the optimized value of parameter -c tends to rise from 1 to 2, 4 ... 16.

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Baseline 1: SuperLU-DIST (2/3)

Benchmark of SuperLU-DIST on Six Large Sparse Matrices

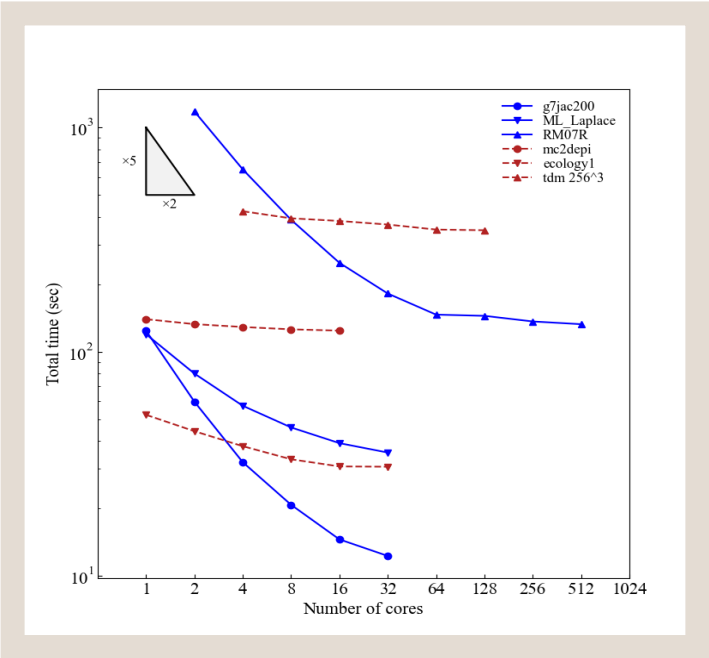

(a) Total Time

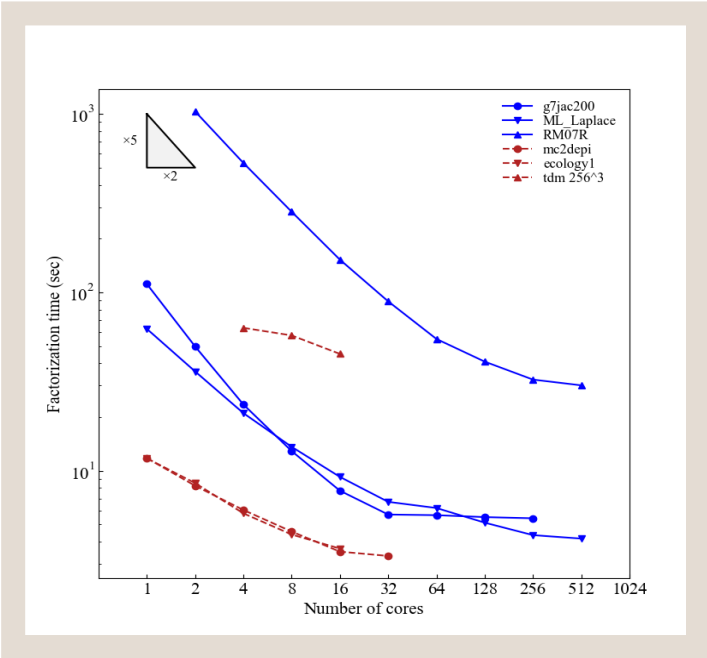

(b) Factorization Time

	r/c	ratio	speed up (2^5)
g7jac200	59,310	2.38E-04	10.1
RM07R	381,689	2.57E-04	8.1
ML_Laplace	377,002	1.95E-04	3.4
mc2depi	525,825	7.60E-06	1.1
ecology1	1,000,000	5.00E-06	1.7
256^3 TDM	16,777,216	1.79E-07	1.2

* These six matrices have been tested in Nurion normal nodes with a range of 1–1024 cores.

$$\text{nnz\_ratio} = \frac{\text{The number of nonezero}}{\text{The number of all elements}}$$

- The efficiency of parallelization tends to decrease as the nonzero ratio decreases in SuperLU-DIST.

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Baseline 1: SuperLU-DIST (3/3)

Benchmark of SuperLU-DIST on Two Large Sparse Matrices

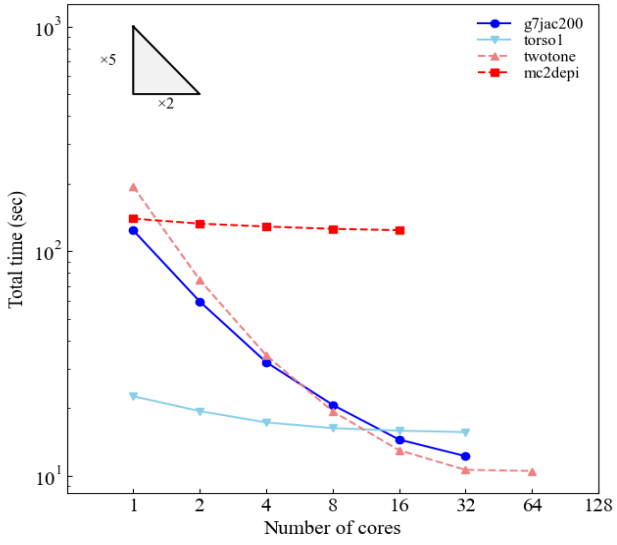

(a) Total Time in two outliers

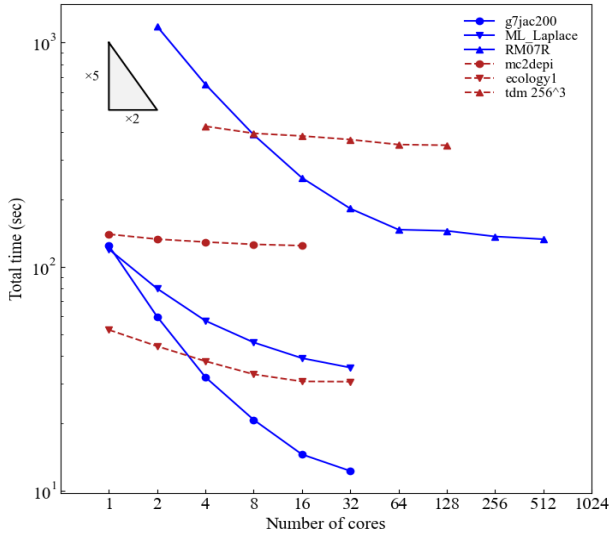

(b) Total Time in six matrices

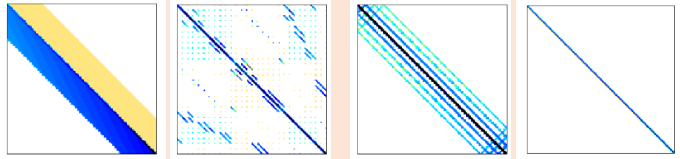

	r/c	ratio	speed up (2^5)
torso1	116,158	6.31E-04	1.44
g7jac200	59,310	2.38E-04	10.1
RM07R	381,689	2.57E-04	8.1
ML_Laplace	377,002	1.95E-04	3.4
twotone	120,750	8.40E-05	18.2
stomach	213,360	6.64E-05	3.0
mc2depi	525,825	7.60E-06	1.1
ecology1	1,000,000	5.00E-06	1.7
256^3 TDM	16,777,216	1.79E-07	1.2

- The parallelization of two outlier matrices tends to increase as the nonzero ratio decreases in SuperLU-DIST.

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Baseline 1 & 2: SuperLU-DIST vs. STRUMPACK

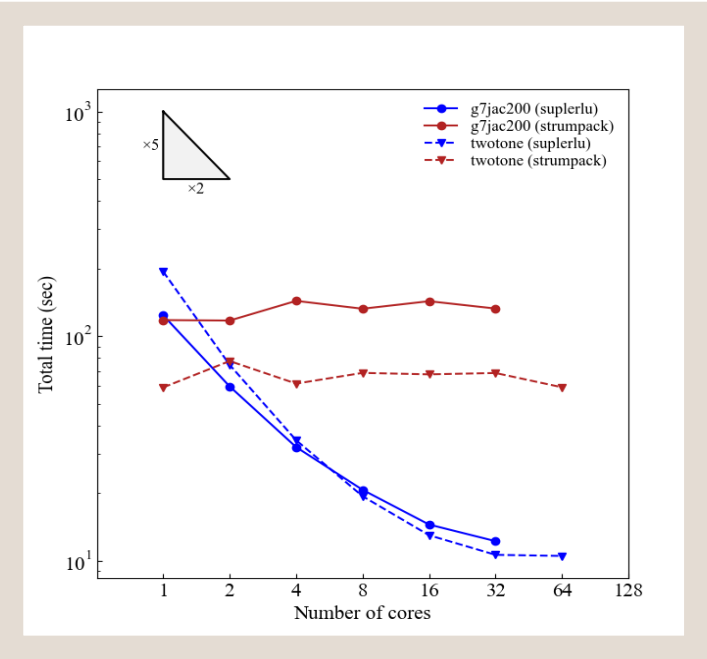

(a) SuperLU > STRUMPACK

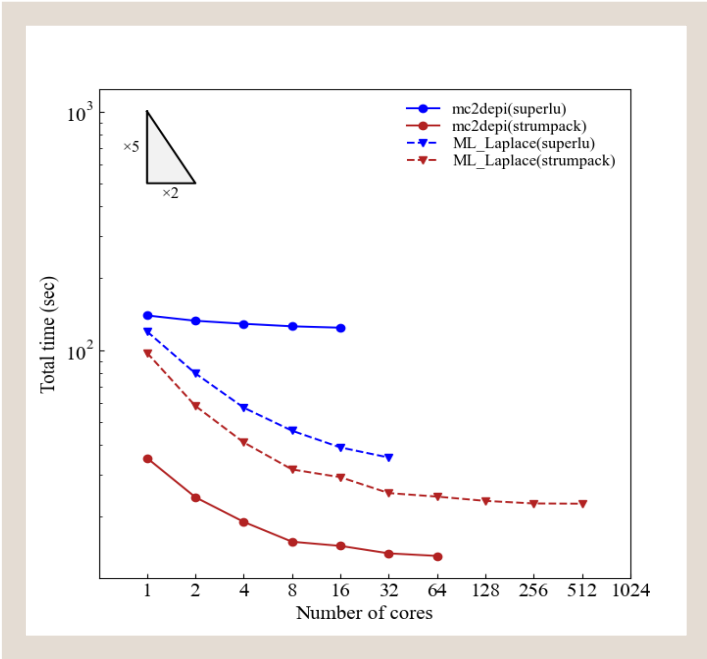

(b) SuperLU < STRUMPACK

g7jac200

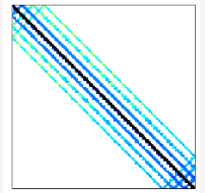

twotone

SuperLU > STRUMPACK

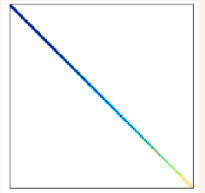

ML_Laplace

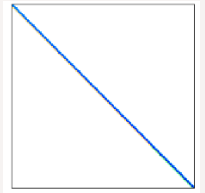

mc2depi

SuperLU < STRUMPACK

- As the diagonal components become denser, SuperLU may tend to outperform STRUMPACK.

3 Benchmark

Parallel Performance Benchmark of Large Sparse Matrix
using SuperLU-DIST, Strumpack, and PaScaL_TDMA

Conclusion & Discussion

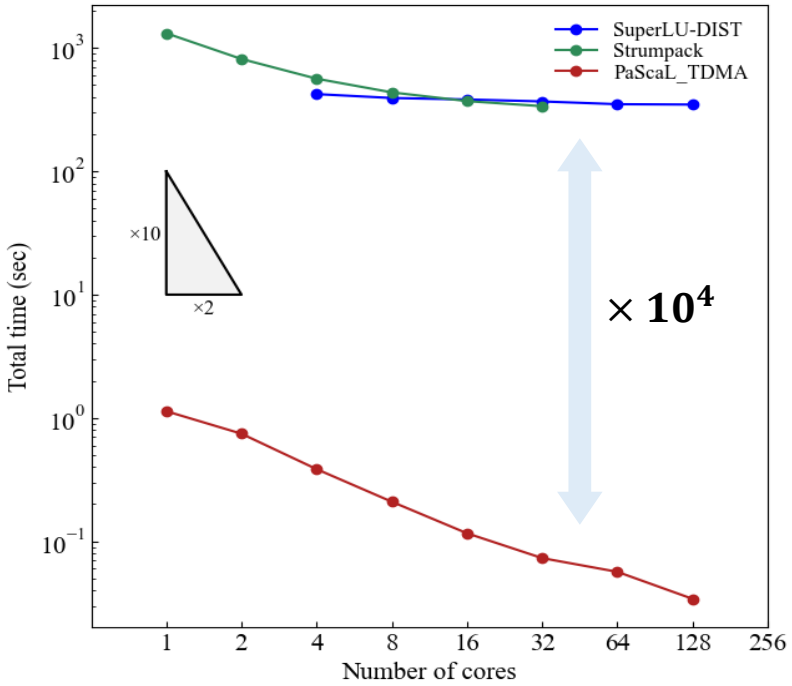

- PaScaL_TDMA performs tri-diagonal matrix computations 10,000 times faster than baseline solvers.
- PaScaL_TDMA remains stable with increased cores and maintains high computation speed.
  (SuperLU-DIST: over 16 cores, STRUMPACK: 64 cores)
- The results of these performance benchmarks highlight the necessity for optimized solvers based on matrix structure.

3 Benchmark

Future Work

- **It is essential to conduct a comprehensive performance comparison of the three solvers, including a time analysis with MPI functionality.**
- **Matrix structure impacts parallelization performance; analyzing this can provide valuable insights for developing more suitable libraries.**
  - Diagonal matrix form reduces parallelization impact in SuperLU-DIST.
  - Increased density of diagonal components decrease parallelization impact in STRUMPACK.

Conclusion

1. Lecture Materials: developed three lecture materials on MPI

   → * MPI 병렬컴퓨팅교육 운영 2회 (강지훈, 권오경) : 상반기 1회 완료

2. Git Website: documentation for PaScaL_TDMA library

3. Benchmark: performance benchmark of large sparse matrix using three libraries

엑사스케일 선형 해석자 수치 라이브러리 개발 / HPC Library for Exascale Linear Solvers

1단계 - 2차 연도(2024)

추진내용	추진 일정												책임기관 (소속기관)
	1	2	3	4	5	6	7	8	9	10	11	12	
수치 라이브러리 통합 저장소 구축													강지훈 (KISTI)
국가슈퍼컴퓨터 5호기 대상 최적 병렬화 및 성능 평가													

평가항목	가중치(%)	연차	연차별 목표(조건/환경)
(정량) 유사 수치 라이브러리 대비 성능	20	(1단계) 1, 2차 연도	삼중대각행렬 해석자에 대해 5호기(NURION KNL)에서 1024노드 (65,536코어, 3PF)까지 경쟁 라이브러리(SuperLU, Strumpack, ScaLAPACK)와의 성능 비교 및 70%의 성능 달성

* 2023 연차보고서 / 1-2. 평가 주안점의 차년도 목표

→ 2 PaScaL_TDMA repo / website

→ 3 Benchmark two baselines on Nurion knl, from 1 node to 32 nodes. (50%)

Korea Institute of
Science and Technology Information
www.kisti.re.kr

Thank you for listening

Presenter: Chanyoung Ahn

Appendix

Diagonal Correlation Factor (Fail)

	r/c	nnz	nonzero ratio	Diagonal correlation factor	superlu speedup	strumpack speedup	PTDMA
g7jac200	59,310	837,936	2.38E-04	0.933829	10.1	0.9	
torso1	116,158	8,516,500	6.31E-04	0.118738	1.4	3.8	
twotone	120,750	1,224,224	8.40E-05	0.985866	18.2	0.9	
stomach	213,360	3,021,648	6.64E-05	0.999528	3.0	2.6	
ML_Laplace	377,002	27,689,972	1.95E-04	0.999949	3.4	3.9	
RM07R	381,689	37,464,962	2.57E-04	0.805946	8.1	error	
mc2depi	525,825	2,100,225	7.60E-06	0.999998	1.1	2.5	
ecology1	1,000,000	4,996,000	5.00E-06	0.999999	1.7	3.0	
256^3 TDM	**16,777,216**	**50331646**	**1.79E-07**	**0.999999**	**1.2**	**3.9**	**15.5**

Pearson Correlation Factor

$$r_{XY} = \frac{\sum_i^n \left(X_i - \overline{X}\right)\left(Y_i - \overline{Y}\right)}{\sqrt{\sum_i^n \left(X_i - \overline{X}\right)^2}\sqrt{\sum_i^n \left(Y_i - \overline{Y}\right)^2}}$$

Appendix

All Matrices Benchmark in SuperLU-DIST

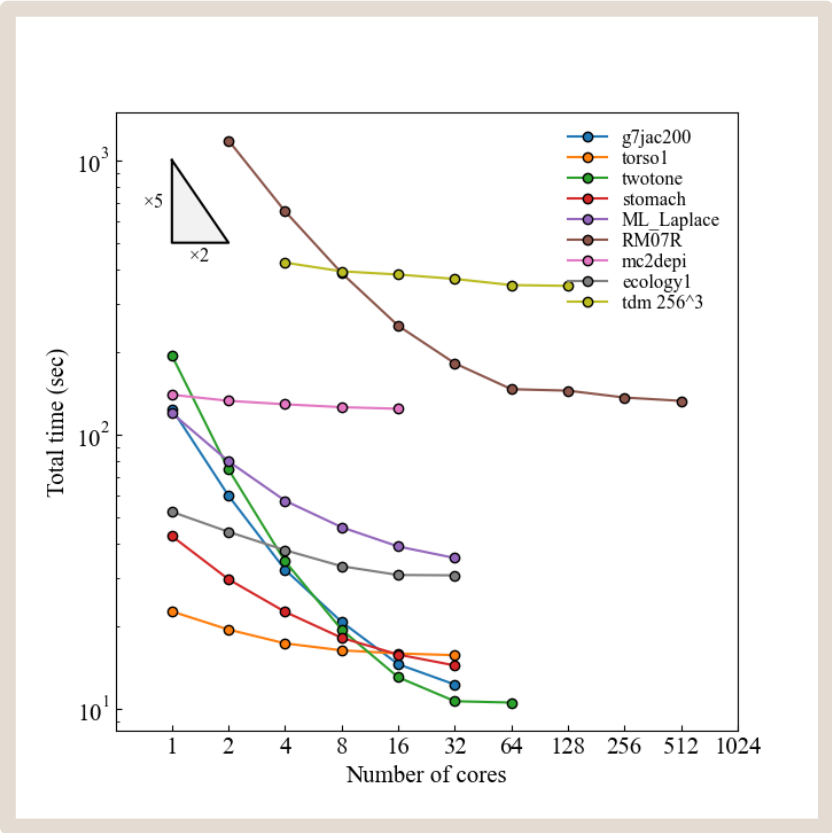

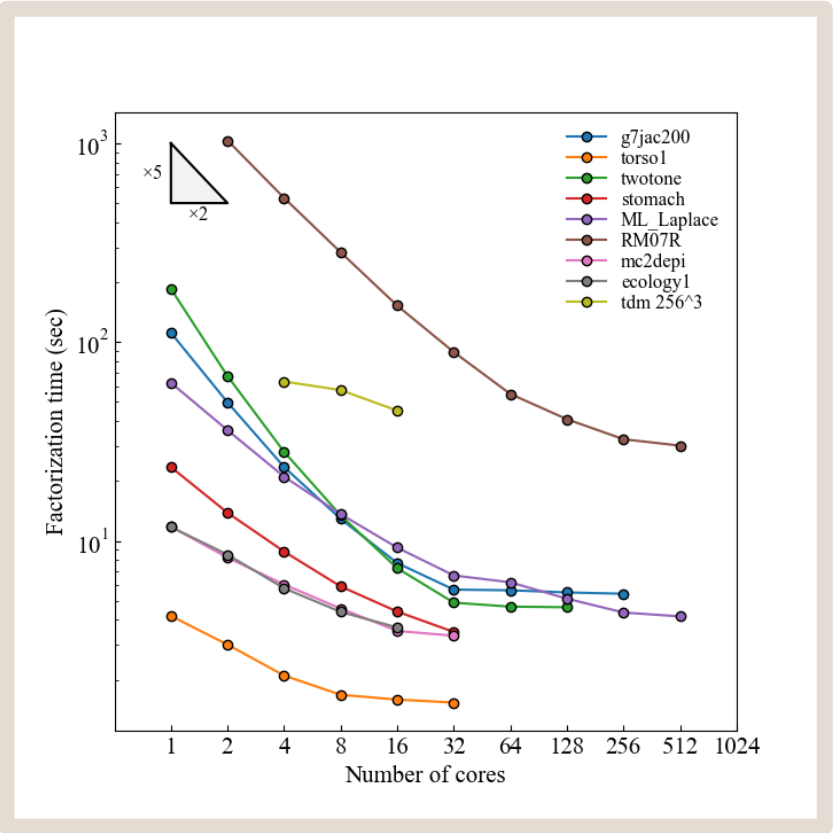

Appendix

All Matrices Benchmark in STRUMPACK

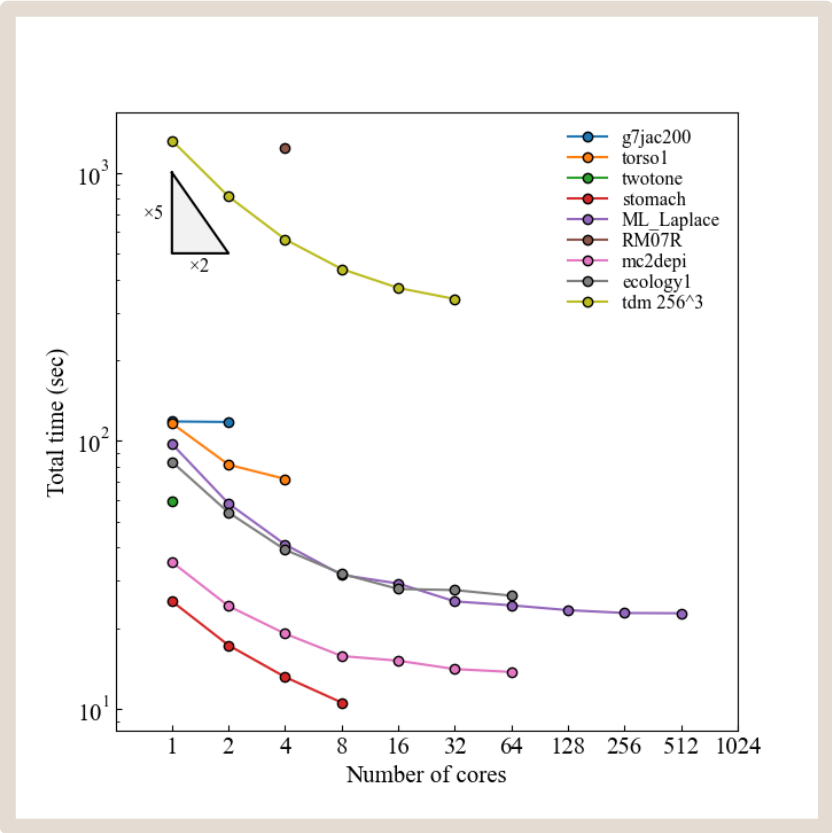

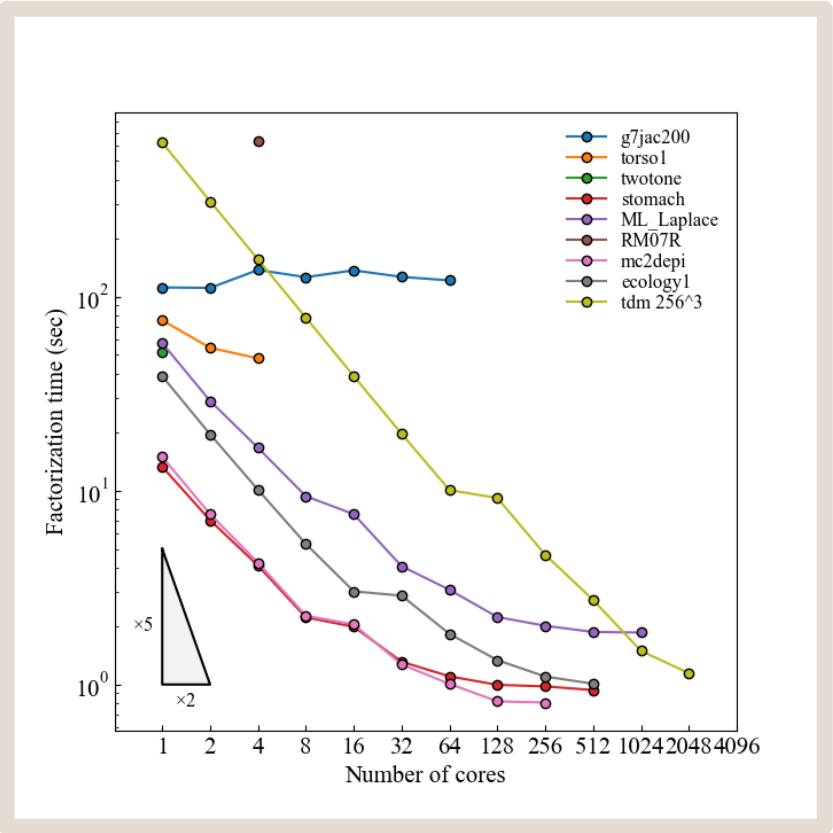

Appendix

Useful links

https://xccels.github.io/PaScaL_TDMA/	Public repository / PaScaL_TDMA Library
https://github.com/cold-young/2024_KISTI_Intern/	Public repository / MPI with segmentation materials, documentation for Supercomputing .....
https://github.com/cold-young/2024-CDE-KISTI/	Private repository / All benchmark files: three benchmark file, raw data, and visualization utils
Google Drive	Google Drive / All benchmark data; Total and factorization times of three solvers