Parallel Space-Time Likelihood Optimization for Air Pollution Prediction on Large-Scale Systems

Presenter: Mary Lai O. Salvaña, Ph.D.

Joint work with Sameh Abdulah, Hatem Ltaief, Ying Sun, Marc Genton and David Keyes

Extreme Computing Research Center
King Abdullah University of Science and Technology

Platform for Advanced Scientific Computing (PASC) Conference 2022
Congress Center Basel, Switzerland

June 27, 2022

I. Motivation

Modern-Era Retrospective Analysis for Research and Applications, version 2 (MERRA-2) reanalysis log particulate matter (PM) data on January 1, 2016

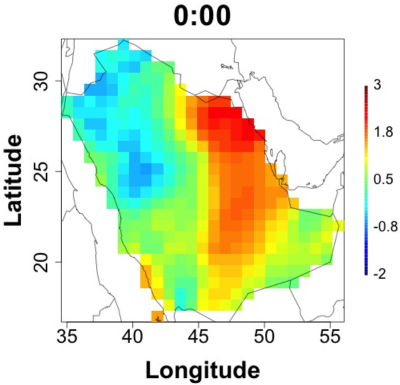

Fig. 1: Middle East

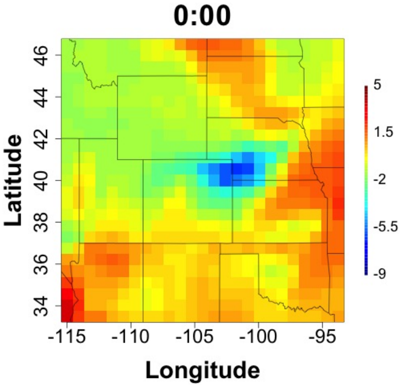

Fig. 2: United States

Overview

II. Preliminaries

II. Preliminaries: Space-Time Geostatistics

- Suppose $\mathbf{Z}_1 \in \mathbb{R}^{n_1}$ is a vector of measurements from n_1 sampled space-time locations, i.e., $\mathbf{Z}_1 = \{Z(\mathbf{s}_1, t_1), \ldots, Z(\mathbf{s}_{n_1}, t_{n_1})\}^\top$, where $Z(\mathbf{s}, t) \in \mathbb{R}$ is the measurement at sampled space-time location $(\mathbf{s}, t) \in \mathbb{R}^d \times \mathbb{R}$.
- Suppose $\mathbf{Z}_2 \in \mathbb{R}^{n_2}$ is a vector of missing measurements at n_2 unsampled space-time locations, i.e., $\mathbf{Z}_2 = \{Z(\tilde{\mathbf{s}}_1, \tilde{t}_1), \ldots, Z(\tilde{\mathbf{s}}_{n_2}, \tilde{t}_{n_2})\}^\top$, where $Z(\tilde{\mathbf{s}}, \tilde{t}) \in \mathbb{R}$ is the missing measurement at unsampled space-time location $(\tilde{\mathbf{s}}, \tilde{t}) \in \mathbb{R}^d \times \mathbb{R}$.
- Suppose $\mathbf{Z}_1$ and $\mathbf{Z}_2$ are jointly Gaussian, i.e.,

$$\begin{bmatrix} \mathbf{Z}_1 \\ \mathbf{Z}_2 \end{bmatrix} \sim \mathcal{N}_{n_1+n_2}\left(\begin{bmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{bmatrix} \right), \tag{1}$$

where $\boldsymbol{\mu}_1 \in \mathbb{R}^{n_1}$ and $\boldsymbol{\mu}_2 \in \mathbb{R}^{n_2}$ are the mean vectors of $\mathbf{Z}_1$ and $\mathbf{Z}_2$, respectively, $\boldsymbol{\Sigma}_{11}$ and $\boldsymbol{\Sigma}_{22}$ are the covariance matrices of $\mathbf{Z}_1$ and $\mathbf{Z}_2$, respectively, and $\boldsymbol{\Sigma}_{12} = \boldsymbol{\Sigma}_{21}^\top$ is the cross-covariance matrix of $\mathbf{Z}_1$ and $\mathbf{Z}_2$.

II. Preliminaries: Space-Time Geostatistics

- Prediction for $\mathbf{Z}_2$ can be performed using the concept of conditional distribution of Gaussian processes, i.e.,

$$\mathbf{Z}_2|\mathbf{Z}_1 \sim \mathcal{N}_{n_2}\{\boldsymbol{\mu}_2 + \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}(\mathbf{Z}_1 - \boldsymbol{\mu}_1), \boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}\}. \quad (2)$$

  This means that the best prediction for the missing measurements vector $\mathbf{Z}_2$, denoted $\hat{\mathbf{Z}}_2$, is

$$\hat{\mathbf{Z}}_2 = \boldsymbol{\mu}_2 + \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}(\mathbf{Z}_1 - \boldsymbol{\mu}_1). \quad (3)$$

- Prediction relies on the mean vectors and covariance matrices.
- In practice, mean vectors and covariance matrices are estimated from the given measurements, $\mathbf{Z}_1$.
- This is done by choosing and fitting a parametric mean function and a parametric covariance function to the data.

II. Preliminaries: Space-Time Geostatistics

- Suppose $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \mathbf{0}$.
- To model the covariance matrices $\boldsymbol{\Sigma}_{11}$, $\boldsymbol{\Sigma}_{22}$, and $\boldsymbol{\Sigma}_{12}$, we need a parametric space-time covariance function, denoted $C(\mathbf{s}_1, \mathbf{s}_2; t_1, t_2|\boldsymbol{\theta})$, such that

$$\begin{aligned}
\boldsymbol{\Sigma}_{11}(\boldsymbol{\theta}) &= \{C(\mathbf{s}_i, \mathbf{s}_j; t_i, t_j|\boldsymbol{\theta})\}_{i,j=1}^{n_1}, \\
\boldsymbol{\Sigma}_{22}(\boldsymbol{\theta}) &= \{C(\tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_j; \tilde{t}_i, \tilde{t}_j|\boldsymbol{\theta})\}_{i,j=1}^{n_2}, \text{ and} \\
\boldsymbol{\Sigma}_{12}(\boldsymbol{\theta}) &= [\{C(\mathbf{s}_i, \tilde{\mathbf{s}}_j; t_i, \tilde{t}_j|\boldsymbol{\theta})\}_{i=1}^{n_1}]_{j=1}^{n_2}.
\end{aligned}$$

- $C(\mathbf{s}_1, \mathbf{s}_2; t_1, t_2|\boldsymbol{\theta})$ describes the strength of dependence of the measurements at any two space-time locations $(\mathbf{s}_1, t_1)$ and $(\mathbf{s}_2, t_2)$.

II. Preliminaries: Space-Time Geostatistics

- A popular space-time covariance function proposed in Gneiting (2002) has the form:

$$C(\mathbf{h}, u|\boldsymbol{\theta}) = \frac{\sigma^2}{|u|^{2\alpha}/a_t + 1} \mathcal{M}_\nu \left\{ \frac{\|\mathbf{h}\|/a_s}{(|u|^{2\alpha}/a_t + 1)^{\beta/2}} \right\},$$

  where $\mathbf{h} = \mathbf{s}_1 - \mathbf{s}_2$ and $u = t_1 - t_2$, $\mathcal{M}_\nu$ is the univariate Matérn correlation function with parameter vector $\boldsymbol{\theta} = (\sigma^2, a_s, \nu, a_t, \alpha, \beta)^\top \in \mathbb{R}^6$, such that $\sigma^2 > 0$ is the variance parameter, $\nu > 0$ and $\alpha \in (0,1]$ are the smoothness parameters in space and time, respectively, $a_s, a_t > 0$ are the range parameters in space and time, respectively, and $\beta \in [0,1]$ is the space-time interaction parameter.
- When $\beta = 0$, separable model
- When $\beta > 0$, nonseparable model

II. Preliminaries: Space-Time Geostatistics

Simulated space-time realizations
from the space-time covariance function model

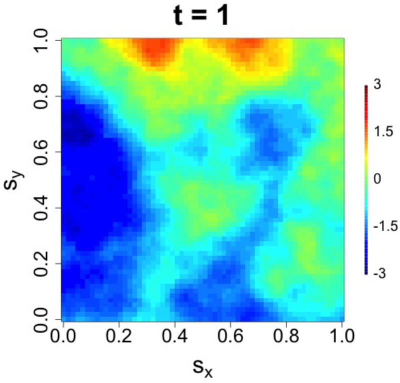

Fig. 3: Nonseparable

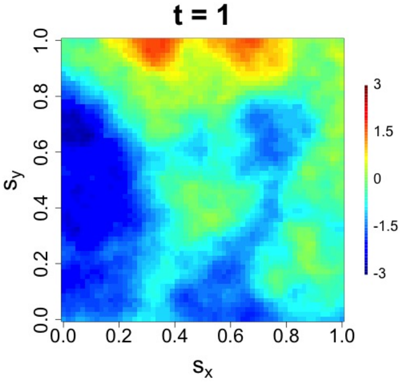

Fig. 4: Separable

II. Preliminaries: Spatio-Temporal Geostatistics

Gaussian log-likelihood function

$$l(\boldsymbol{\theta}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log|\boldsymbol{\Sigma}(\boldsymbol{\theta})| - \frac{1}{2}\mathbf{Z}^\top\boldsymbol{\Sigma}(\boldsymbol{\theta})^{-1}\mathbf{Z},$$

Here $|\boldsymbol{\Sigma}(\boldsymbol{\theta})|$ is the determinant of $\boldsymbol{\Sigma}(\boldsymbol{\theta})$.

The grand challenge in large-scale Gaussian geostatistical modeling lies in the inversion of $\boldsymbol{\Sigma}(\boldsymbol{\theta})$.

STATISTICS @KAUST

II. Preliminaries: The ExaGeoStat Software

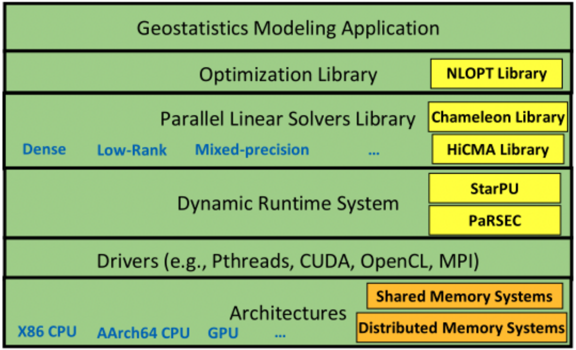

Fig. 5: The ExaGeoStat software layers for geostatistics applications. Source: Huang et al. (2021)

II. Preliminaries: State-of-the-Art Dense Linear Algebra Libraries

Data layout formats

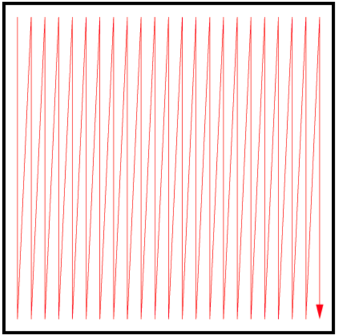

Fig. 6: LAPACK: Column-major data layout format.

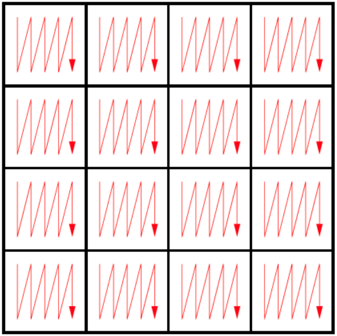

Fig. 7: Chameleon: Tile data layout format.

II. Preliminaries: Directed Acyclic Graph (DAG)

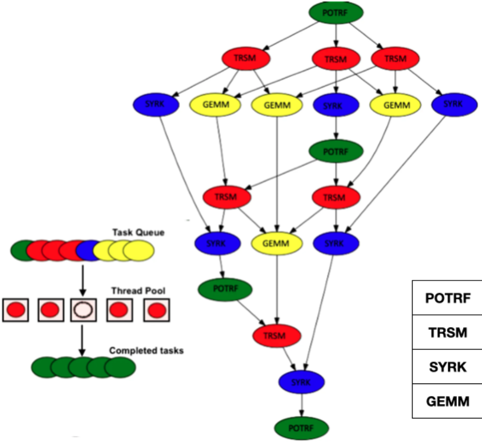

POTRF	**PO**sitive-definite **TR**iangular Cholesky **F**actorization
TRSM	**TR**iangular **S**olve **M**atrix operation
SYRK	**SY**mmetric **R**ank-**K** operation
GEMM	**GE**neral **M**atrix-**M**atrix operation

Fig. 8: An example of task-based DAG to perform Cholesky factorization of 4-tiles by 4-tiles matrix.

II. Preliminaries: Particle Swarm Optimization

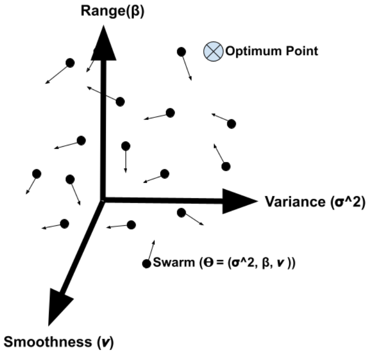

Fig. 9: Iteration 0

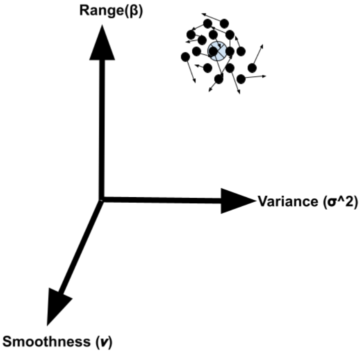

Fig. 10: Iteration N

PSO-based MLE optimization of a univariate purely spatial covariance function.

III. Contributions

III. Contributions

1. Implementation of high-performance space-time model on large-scale systems
2. Visualization of space-time random fields generated by the high-performance implementation
3. Incorporation of the PSO algorithm to the MLE operation to utilize the execution performance on distributed environments
4. Illustration of benefits of flexible vs. simple space-time model via large scale space-time experiments
5. Application to air pollution datasets from the Middle East and US

IV. The Proposed Two-Level Parallelization Framework

IV. Proposed: Parallel Optimization Strategy

- *MPI_COMM_WORLD*: MPI default communicator
- *MPI_Comm_split*: partition the default communicator into disjoint subgroups associated with different sub-communicators
- PPSwarm algorithm

We split the default communicator into a set of sub-communicators where each can be used to evaluate a single log-likelihood function solution.

STATISTICS @KAUST

IV. Proposed: Framework

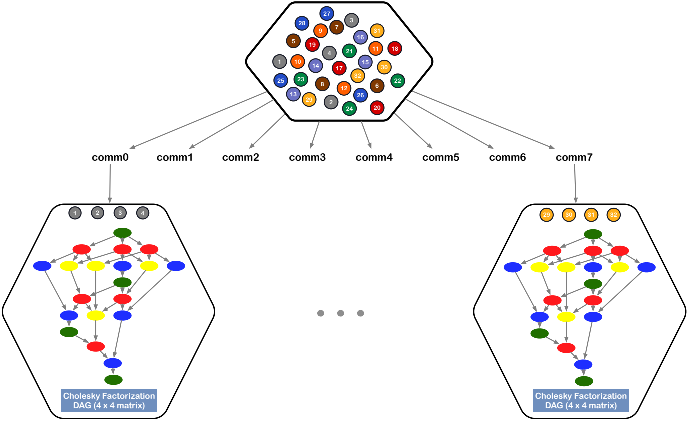

Fig. 11: Testcase using 32 nodes and 8 MPI sub-communicators. Each sub-communicator includes 4 nodes that estimate the log-likelihood function with a certain set of parameters in parallel using the StarPU runtime system.

V. Performance Results & Analysis

V. Performance Results & Analysis

- Synthetic data from the space-time model with $\beta \in \{0.1, 0.5, 1\}$
  - 400 spatial locations
  - 100 temporal locations
  - 40,000 × 40,000 covariance matrix size
- Real data of particulate matter
  - 550 spatial locations
  - 730 temporal locations
  - 401,500 × 401,500 covariance matrix size
- The performance is tested on an Intel-based Cray XC40 system with 6,174 compute nodes, each of which has two 16-core Intel Haswell CPUs at 2.30 GHz and 128 GB of memory.
- All the experiments were conducted on the whole number of cores with different nodes.

V. Performance Results & Analysis: Synthetic Data

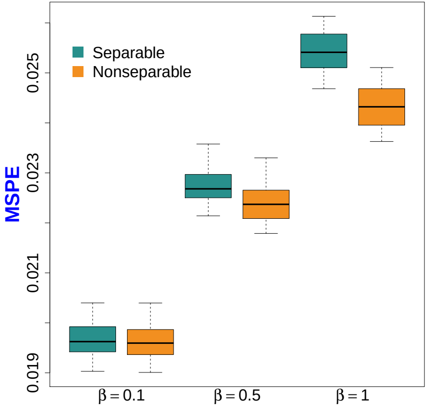

Fig. 12: Boxplots of the prediction errors when fitting a separable and nonseparable model on space-time data with varying degrees of space-time interactions. Weak, moderate, and strong space-time interactions are represented by $\beta = 0.1, 0.5$, and 1, respectively.

V. Performance Results & Analysis: Synthetic Data

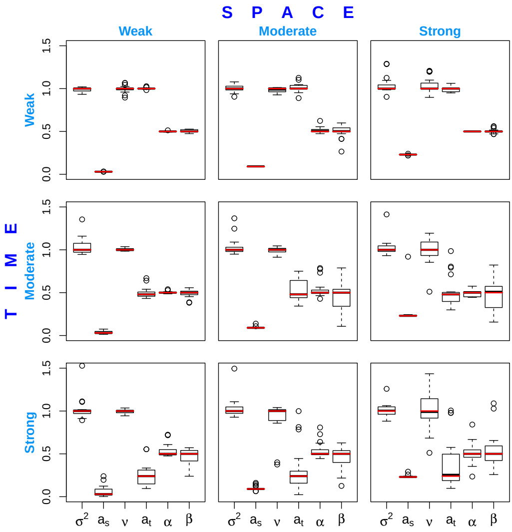

Fig. 13: Boxplots of parameter estimates under varying degrees of space-time dependence. The true parameters are highlighted in red.

V. Performance Results & Analysis: Real Data

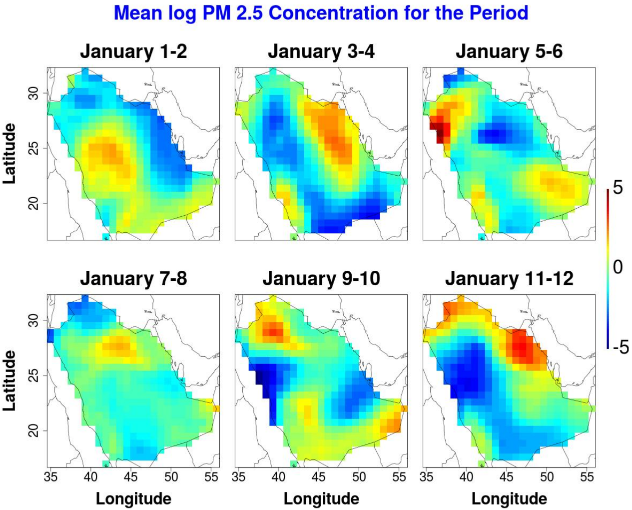

Fig. 14: Visualization of the log PM2.5 dataset after space-time mean removal at the first six time points in 2016 over Saudi Arabia.

V. Performance Results & Analysis: Real Data

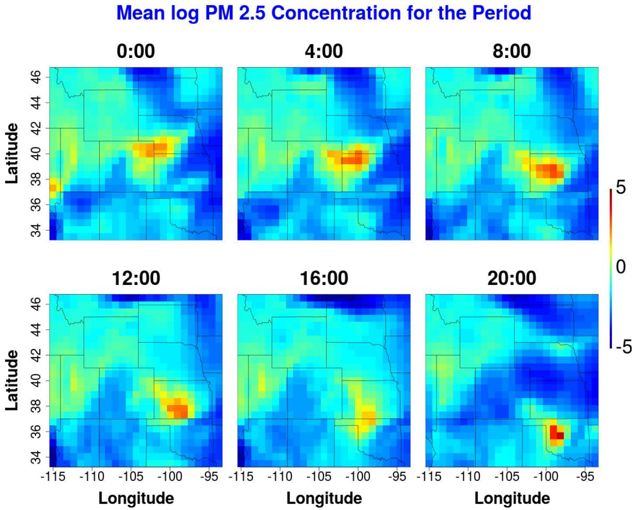

Fig. 15: Visualization of the log PM2.5 dataset after space-time mean removal at four hour intervals on January 1, 2016 over the Midwest US.

STATISTICS @KAUST

V. Performance Results & Analysis: Real Data

Table 1: A summary of the estimated parameters of the nonseparable (NS) and separable (S) models and their corresponding errors (MSPE) and prediction uncertainty (PU) for the Saudi Arabia and US datasets. MSPE1 and PU1 correspond to Testing Dataset 1, while MSPE2 and PU2 point to Testing Dataset 2. The best model reports the lower MSPE and PU.

	Model	$\hat{\sigma}^2$	$\hat{a}_s$	$\hat{\nu}$	$\hat{a}_t$	$\hat{\alpha}$	$\hat{\beta}$	MSPE1 / PU1	MSPE2 / PU2
Saudi Arabia	NS	1.29	1.34	2.15	1.12	0.14	0.75	0.0017/76	1.08/875
	S	2.61	1.27	2.15	2.04	0.03	0	0.0018/78	1.14/1066
US	NS	0.47	1.33	1.12	6.77	0.72	0.14	0.0028/155	0.05/322
	S	2.12	1.54	1.47	7.99	0.48	0	0.0031/134	0.06/1118

V. Performance Results & Analysis: Real Data

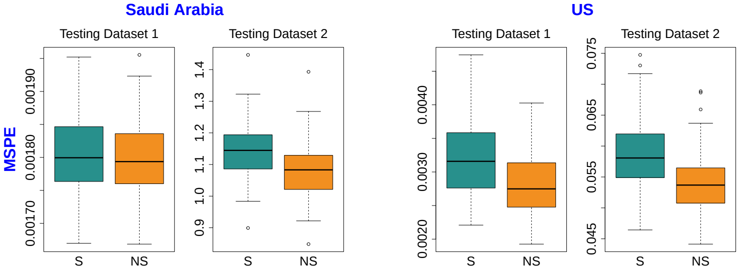

Fig. 16: Boxplots of the prediction errors for the separable (S) and nonseparable (NS) models in the real data pseudo cross-validation study.

V. Performance Results & Analysis: Real Data

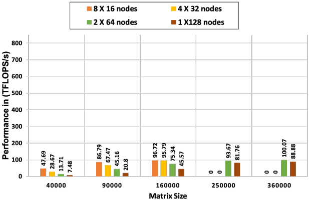

(a) 128 nodes.

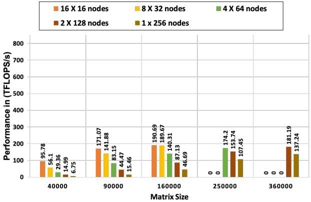

(b) 256 nodes.

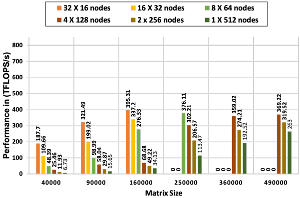

(c) 512 nodes.

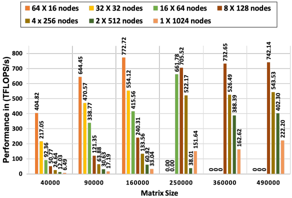

(d) 1024 nodes.

Fig. 17: Performance of a single MLE optimization step using different number of nodes on Shaheen-II Cray XC40 Supercomputer.

STATISTICS @KAUST

V. Performance Results & Analysis: Real Data

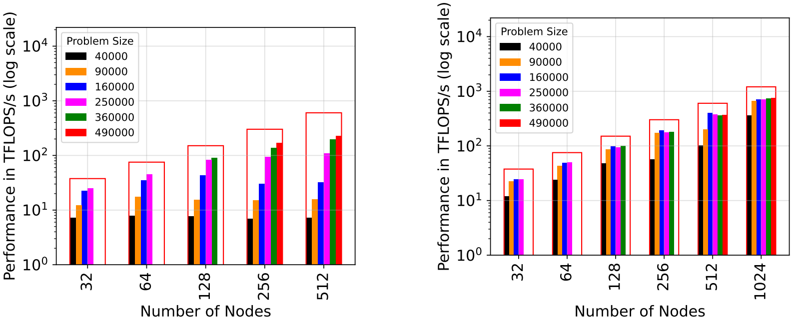

(a) One-level parallelization.

(b) Two-level parallelization.

Fig. 18: Performance of one MLE optimization step using single and n MPI communicators on Shaheen-II Cray XC40 Supercomputers. In (b) x MPI sub-communicators is used where x is tuned for performance.

V. Performance Results & Analysis: Real Data

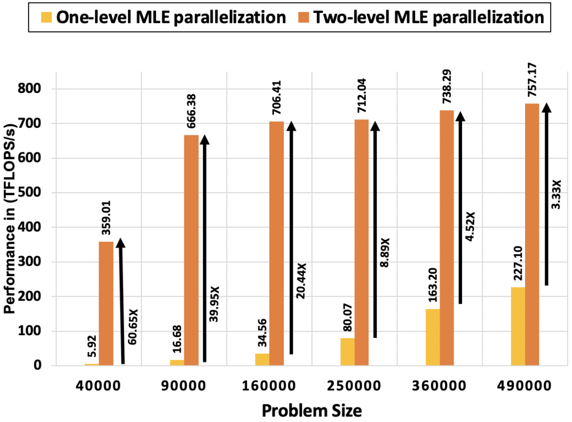

Fig. 19: One-level versus two-level MLE parallelization performance using 1024 nodes on Shaheen-II Cray XC40 system.

V. Performance Results & Analysis: Real Data

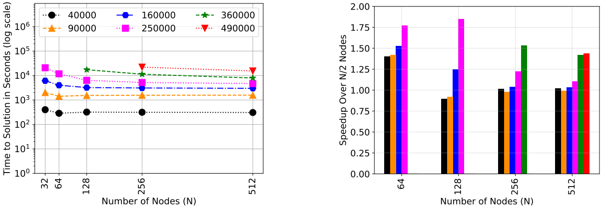

(a) One-level parallelization.

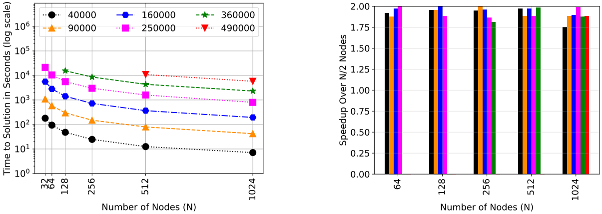

(b) Two-level parallelization.

Fig. 20: Time-to-solution of full MLE operation using 100 optimization iterations on Shaheen-II Cray XC40 system. In (b), we tune x MPI sub-communicators for performance.

STATISTICS @KAUST

VI. Summary

VI. Summary

- Proposed a two-level parallelization framework of geostatistical space-time modeling
  - Upper level: MPI sub-communicators perform independent log-likelihood function evaluation with different sets of parameters via PPSO algorithm
  - Inner level: task-based parallel technique is used to perform linear solver operations on a given set of nodes representing a single MPI sub-communicator
- Demonstrated, through synthetic and real datasets, the merits of the proposed implementation
- Achieved high prediction accuracy with up to 757 TFLOPS/s using 1024 nodes on the KAUST Shaheen-II Cray XC40 system (around 63% of the theoretical peak)
- Other avenues for research: tile low rank and mixed-precision approximations to accelerate further the modeling process

STATISTICS @KAUST

References I

Abdulah, S., Ltaief, H., Sun, Y., Genton, M. G., and Keyes, D. E. (2018). ExaGeoStat: A high performance unified software for geostatistics on manycore systems. *IEEE Transactions on Parallel and Distributed Systems*, 29(12):2771–2784.

Augonnet, C., Thibault, S., Namyst, R., and Wacrenier, P.-A. (2011). StarPU: a unified platform for task scheduling on heterogeneous multicore architectures. *Concurrency and Computation: Practice and Experience*, 23(2):187–198.

CHAMELEON (2021). The Chameleon project: A dense linear algebra software for heterogeneous architectures. Available at `https://project.inria.fr/chameleon/`.

Gneiting, T. (2002). Nonseparable, stationary covariance functions for space–time data. *Journal of the American Statistical Association*, 97(458):590–600.

Huang, H., Abdulah, S., Sun, Y., Ltaief, H., Keyes, D. E., and Genton, M. G. (2021). Competition on spatial statistics for large datasets. *Journal of Agricultural, Biological and Environmental Statistics*, 26(4):580–595.

References II

Vaz, A. I. F. and Vicente, L. N. (2007). A particle swarm pattern search method for bound constrained global optimization. *Journal of Global Optimization*, 39(2):197–219.

Questions?