Parallelized sparse inverse covariance matrix estimation

Author and Presenter:
Aryan Eftekhari (Università della Svizzera italiana, Switzerland)

Co-authors:
Matthias Bollhöfer (TU Braunschweig, Germany)
Olaf Schenk (Università della Svizzera italiana, Switzerland)

Abstract: We consider the problem of estimating sparse inverse covariance matrices for high-dimensional datasets using the ℓ_1-regularized Gaussian maximum likelihood method. This problem is particularly challenging as the required computational resources increase superlinearly with the number of random variables. We introduce a performant and scalable algorithm, which builds on recent advancements of second-order methods for sparse problems [1, 2]. The routine leverages the intrinsic parallelism in the linear algebra operations and exploits the underlying sparsity of the problem. The computational bottlenecks are identied, and the respective subroutines are parallelized using a hybrid MPI-OpenMP approach. Numerical examples conducted at the Swiss National Supercomputing Center (Cray XC40) show that, in comparison to the state-of-the-art algorithms, the proposed routine provides signicant speed-up with scalability up to 128 nodes. The developed framework is used to approximate the sparse inverse covariance matrix for datasets with up to 10 million random variables.

References

[1] Cho-Jui Hsieh, Mátyás A. Sustik, Inderjit S. Dhillon and Pradeep K. Ravikumar. *Sparse inverse covariance matrix estimation using quadratic approximation*, In: J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, vol. 24, pp. 2330-2338; Neural Information Processing Systems Foundation, 2011.

[2] Cho-Jui Hsieh, Mátyás A. Sustik, Inderjit S. Dhillon, Pradeep K. Ravikumar and Russell A. Poldrack. *BIG & QUIC: Sparse inverse covariance estimation for a million variables*, In: C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, vol. 26, pp. 3165-3173; Neural Information Processing Systems Foundation, 2013.