Scalable Stochastic Programming

Stochastic Optimization of Complex Systems

Mihai Anitescu
Mathematics and Computer Science Division (MCS),
Argonne National Laboratory (ANL)
9700 South Cass Ave, Argonne, IL, 60439, USA

Abstract

Optimal design, planing and management of nationally critical complex energy systems such as the power electric grid require the optimization of large-scale interconnected sub-systems in the presence of multiple sources of uncertainty. For complex energy systems, the source of uncertainty can be incomplete information or reductive modeling of weather conditions, consumer demand, market prices, etc. In this work we investigate scalable approaches for one framework for decision-making under uncertainty: stochastic programming (SP) with recourse.

Our methodology relies on approximating the underlying uncertainty of the SP problems via sampling, and solving the corresponding sample average approximation (SAA) problem using an interior-point method with a specialized linear algebra layer. SAA SP problems have a nested block half-arrow shaped Jacobian and a block separable objective function which allows the Hessian to be permuted in a nested block arrow shape form. We use the direct Schur complement (DSC) method to make use of the particular structure of the Hessian. This method not only succeeds where general purpose linear solvers fail but is also suitable for a scenario-based tree-like parallelization. The only bottleneck in the parallel execution flow is the factorization of the dense Schur complement of the first stage Hessian block. We designed the preconditioned Schur complement (PSC) method that uses preconditioned Krylov subspace iterative methods to hide the latency caused by the direct factorization of the Schur complement. An analysis of the spectral properties of the preconditioner reveals that the eigenvalues of the preconditioned matrix cluster around 1 exponentially fast with the size of the subset of the scenarios incorporated in the preconditioner.

The preconditioning technique however suffers from a memory bottleneck as does the classical Schur-complement. We overcome this issue by using a novel approach for distributing and directly solving the dense Schur-complement systems in parallel in a distributed memory environment. To efficiently use all the cores inside a node, our code, PIPS, implements a hybrid parallel programming model which uses multithreaded sparse linear solver WSMP together with OpenMP to achieve shared-memory parallelism for the second-stage computations. The solver has been used to solve a stochastic economic dispatch problem with wind power integration and transmission constraints on 131,072 cores (32,768 nodes) of the "Intrepid" Blue Gene/P system at Argonne National Laboratory.