10th World Congress on Structural and Multidisciplinary Optimization
May 19-24, 2013, Orlando, FL, USA

A Matrix-Free Approach to Large-Scale Structural Optimization

Andrew B. Lambe[1], Joaquim R. R. A. Martins[2]

[1] University of Toronto, Toronto, Canada, lambe@utias.utoronto.ca
[2] University of Michigan, Ann Arbor, Michigan, USA, jrram@umich.edu

Abstract

In many problems within structural and multidisciplinary optimization, the computational cost is dominated by computing gradient information for the objective and all constraints. Analytic or semi-analytic methods are often employed to reduce this cost. Due to the structure of the equations for evaluating the gradients analytically, the direct method is most efficient for problems with a small number of design variables, while the adjoint method is most efficient for problems with a small number of constraints. However, if the problem contains both a large number of design variables and a large number of constraints, both methods are inefficient because both require many systems of linear equations to be solved. A common example of this kind of problem is minimum-mass structural design subject to failure constraints.

A common solution to this computational bottleneck is constraint aggregation, i.e., combining a large number of constraints into a composite function that models the same feasible domain. The most ubiquitous example of this is the Kreisselmeier-Steinhauser (KS) function. While reducing the number of constraints allows us to apply the adjoint method efficiently, the KS function provides a conservative estimate of the feasible domain so the final design will not necessarily be locally optimal with respect to the original constraints. Furthermore, using the KS function reduces the conditioning of the optimization problem and can actually degrade the performance of gradient-based optimization algorithms.

In this work, we attempt to get around the gradient computation cost by developing an optimizer that does not require frequent computation of the full constraint Jacobian. Instead, the optimizer may only access products of the Jacobian matrix or its transpose with appropriate vectors. Because the cost of each matrix-vector product is equivalent to either the cost of obtaining either a single constraint gradient or the derivatives of all functions with respect to a single variable, we can reduce the total optimization cost by keeping the number of matrix-vector products low. For very large structural models, these matrix-vector products can be computed efficiently using parallel methods. Furthermore, using only matrix-vector products gives us the flexibility to reduce or avoid constraint aggregation and its associated issues.

The key ingredient to keeping the number of matrix-vector products small is to maintain an estimate of the Jacobian by using a quasi-Newton method. Quasi-Newton methods have been frequently used to estimate problem Hessian information in optimization software, but their use in estimating constraint Jacobian information is rare in the optimization literature. These methods seem well-suited to problems with many constraints and where the Jacobian is not sparse. We present results for two structural optimization problems solved by an augmented Lagrangian algorithm to verify the effectiveness of the Jacobian approximations. We also show trends for how the computational cost scales with problem size and how the matrix-free approach can substantially reduce the cost of the optimization.