Strong Linear Relaxations for Global Optimization Problems with Multilinear Terms

Next-Generation Solvers for Mixed-Integer Nonlinear Programs: Structure, Search, and Implementation

Jeff Linderoth, James Luedtke, Mahdi Namazifar
Department of Industrial and Systems Engineering
University of Wisconsin-Madison
Madison, WI 53706, USA

Andrew Miller
IMB, Université Bordeaux 1
RealOpt, INRIA Bordeaux
3405 Talence, France

Abstract

Multilinear functions appear in many optimization problems, including blending and electricity transmission. Solving these problems to provable global optimality remains a significant challenge for modern software. Algorithms are based on relaxing the nonconvexities created by the multilinear terms and branching in a manner to ensure convergence of the relaxation to the optimal solution.

By an appropriate reformulation, any optimization problem containing multilinear functions can be relaxed by obtaining a convex outer approximation of the nonconvex set:

$$X \stackrel{\text{def}}{=} \left\{(x,z) \in H \times \mathbb{R}^m \mid z_t = \prod_{j \in J_t} x_j \ \forall i = 1, \ldots m\right\},$$

where $H = \{x \in \mathbb{R}^n \mid \ell \leq x \leq u\}$ is a box. A common technique for creating a relaxation of X is to decompose the products into bilinear terms and then use a relaxation known as the McCormick (1976) envelope for each term separately. Alternatively, using the results of Rikun (1997) and Sherali (1997), the convex hull of X can be explicitly described as a polyhedron in a higher-dimensional space, but this formulation includes 2^n decision variables, limiting its direct use to low-dimensional problems. This motivates two research questions. First, how much improvement is possible when using the full convex hull relaxation? Second, can we obtain a linear relaxation of manageable size that obtains much of the benefit of the full convex hull?

We first consider the possible improvement the convex hull formulation can provide for bilinear functions. We obtain the surprising result that for any bilinear function having positive coefficients on the terms, the McCormick relaxation is guaranteed to provide a relaxation error that is within factor of $2 - 2/\chi$ of the relaxation error of the convex hull, where χ is the coloring number of graph associated with the bilinear function. To our knowledge, this is the first approximation ratio-type result for comparing the strength of relaxations of global optimization problems. This result suggests that the potential for improvement from the convex hull formulation is greater for denser graphs. For more general functions, including bilinear functions having positive and negative coefficients, and higher-order multilinear functions we find that the potential for improvement by using the convex hull formulation can be even greater.

These theoretical results provide insight into how to construct relaxations that achieve a balance between size and strength. Our idea is to choose subsets of variables that form a *covering* of the multilinear terms and then construct the convex hull formulation independently for each subset of variables. By limiting the size of the variable subsets, we are able to control the size of the resulting relaxation. Choosing the "best" covering is in general a difficult problem, but we devise a heuristic strategy that consisently provides better relaxations than simpler schemes. We find that it is indeed possible to obtain relaxations that are closer in size to the McCormick relaxation while yielding bounds that are closer to those of the full convex hull. Our proposed scheme significantly outperforms the standard McCormick relaxation approach that is used in most global optimization solvers, and on some classes of instances also improves upon relaxations obtained by semidefinite programming.