Exploiting Low-Rank Structure in Semidefinite Programming by Approximate Operator Splitting

Mario Souto[1], Joaquim D. Garcia [1], Álvaro Veiga[1]

[1] Department of Electrical Engineering,

Laboratory of Applied Mathematical Programming and Statistics (LAMPS),

Pontifical Catholic University of Rio de Janeiro (PUC-Rio)

In contrast to many other convex optimization classes, state-of-the-art semidefinite programming solvers are still unable to efficiently solve large scale instances. This work aims to reduce this scalability gap by proposing a novel proximal algorithm for solving general semidefinite programming problems. The proposed methodology, which is based on the primal-dual hybrid gradient method, allows for the presence of linear inequalities without the need to add extra slack variables and avoids solving a linear system at each iteration. More importantly, it simultaneously computes the dual variables associated with the linear constraints. The main contribution of this work is that it achieves a substantial speedup improvement by effectively adjusting the proposed algorithm in order to exploit the low-rank property inherent to several semidefinite programming problems. This modification is the key element that allows the operator splitting method to efficiently scale to larger instances. Convergence guarantees are presented along with an intuitive interpretation of the algorithm. Additionally, an open-source semidefinite programming solver called `ProxSDP` is made available and its implementation details are discussed. Case studies are presented in order to evaluate the performance of the proposed methodology.