SOLVING LARGE-SCALE SEMIDEFINITE PROGRAMS IN PARALLEL

M. V. NAYAKKANKUPPAM

AUTHOR'S RESPONSE TO REFEREE #1

We thank the referee for a thorough review and constructive criticisms, which we address below.

Error measures: Measures of primal/dual infeasibility and complementarity are appropriate for primal-dual interior point methods, but not for the bundle method. The reformulation is an unconstrained one; by construction, iterates are always dual feasible for the original SDP. The method is a purely dual one; a primal iterate is never explicitly computed or updated. However, the sequence of optimal subgradients W^k of $\lambda_{\max}(\cdot)$ generated by the method converges to a primal solution X in the limit (proved by Helmberg). Thus, the norm of the subgradient $AW^k - b$ (of f) is simultaneously an estimate of primal infeasibility and optimality (complementarity). This norm, upon termination, is of the same order as the δ used in the termination criterion, but this value itself doesn't yield a usable error estimate on the optimal value. For the quantum chemistry problems, the relative errors in the optimal value are given in Table 2; for the other problems, it's less than 0.1%.

Usefulness of low accuracy solutions: Although the ground state energy calculations require 6-7 digits, there are similar SDP's arising from dipole moment calculations for which the accuracy produced by the method is sufficient (although these test problems are not readily available). In combinatorial optimization, the accuracy provided by the bundle method is already quite useful in branch-and-bound algorithms.

Choice of test set: Most problems in the DIMACS test suite are not readily reformulated as eigenvalue optimization problems; this requires a compact primal feasible set, with a known upper bound on an optimal solution. The only significant problem class exlcuded are the max-cut instances, which are not well handled by our data distribution scheme, as explained in Section 10.

Besides, a key contribution of the paper is the handling of block structure, so block structured test problems are required to illustrate the performance of the method. Outside of the chosen problems, the standard test sets contain few instances (perhaps none) that are simultaneously block structured and amenable to reformulation as eigenvalue optimization.

Finally, we disagree with the referee's claim that most of the chosen problems can be solved by primal-dual interior-point codes on a desktop PC. The last paragraph of Section 9.6 should convince the referee otherwise.

Scalability: We didn't tabulate parallel efficiencies because the table already overflows the page and has an overwhelming amount of numbers in it. We included speed-up plots to give a quick pictorial idea, and the efficiencies

are readily computed from the timing data, if necessary. Except in the rare instance of a superscalable algorithm, it is generally to be expected that efficiency decreases (to the point of slow-down, rather than speed-up) with increasing processor count, if the problem size is held fixed. [As an aside, it then makes sense to consider scaled speed-up or scaled efficiency (increasing problem size and processor count simultaneously). This is easy to do in PDE applications (one simply refines the mesh size), but is obviously not applicable in optimization where each test problem stands on its own.] The paper makes guarded and carefully worded claims regarding scalability: *acceptable parallel scalability on sufficiently large problems*. The limits of scalability on the chosen problems are clearly mentioned in the text. We still maintain that these are toy instances in the realm of quantum chemistry. Larger atomic systems, or discretization of the chosen systems in higher dimensional bases, will result in much bigger problems on which efficiencies would be quite good even on 64 processors. Such problems will easily exceed the capabilities of standard primal-dual interior-point methods — even parallel versions.

A paragraph has been added in Section 9.5 to explain the anomalies in Table 7.

Comparison with other methods: Earlier works in the literature (citations have been added) show the relative strengths and weakness of interior-point methods versus the spectral bundle method; therefore a further comparison of these two types of algorithms is redundant. Regarding a comparison of software codes, this is the only parallel implementation of the spectral bundle method as far as we know. But in any case, this paper is not about benchmarking software codes. The goal of the paper is to enhance the applicability and performance of the bundle method on problems that are particularly amenable to solution by subgradient methods.

The claim that we obtain "faster solution times than previously possible" appears only in the abstract, and is not altogether unreasonable. The proposed algorithmic enhancements (particularly relating to block structure) combined with the parallel implementation allows solution times in tens of minutes in contrast with the tens of hours achieved in recent cited work (Zhao et al. 2004, Fukuda et al. 2005); see Section 9.6.

Appropriateness for the journal: The paper was submitted to a special issue of Mathematical Programming focusing on large-scale nonlinear and semidefinite programming, including parallel computing aspects. This revised version is being submitted in deference to the editor's decision.