ORSA Journal on Computing

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Commentary—Theory and Practice for Interior-Point Methods

Michael J. Todd,

ORSA Journal on Computing
Vol. 6, No. 1, Winter 1994

0899-1499/94/0601-0028 $01.25

COMMENTARY

Theory and Practice for Interior-Point Methods

MICHAEL J. TODD / *School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York 14853; Email: miketodd@orie.cornell.edu*

The last decade has been a fascinating time for researchers interested in linear programming and its extensions. It opened with Narendra Karmarkar's theoretical paper on the computational complexity of linear programming problems and his claims that his method solved problems 50 times faster than the simplex method. After 10 years of development in both interior-point and simplex methods, we are now at a point at which both approaches can handle problems that seemed intractable before, but where interior-point methods seem to be superior to simplex algorithms for many very large-scale sparse linear programming problems. Lustig, Marsten, and Shanno (LMS) provide an excellent overview of these advances, particularly with respect to their computational significance. Of course, they have been major contributors to the development of interior-point algorithms as well as of sophisticated implementations with their OB1 code. Their experience leads to many interesting insights and also raises some questions: why does the primal-dual method allow such long steps (99.95% of the way to the boundary, for instance, whereas primal or dual affine-scaling methods seem limited to, say, 95%)? Do the primal and dual "help each other?" And why does the primal-dual affine-scaling algorithm perform poorly, while even tiny centering components render it highly efficient?

In these remarks I want to try to indicate the interplay between theory and computational practice, highlight theoretical developments with (potentially) important consequences and point out some major remaining tasks.

1. Projective Methods

Karmarkar's original algorithm used a projective transformation at each iteration, although his computational experiments were apparently mostly conducted with the affine-scaling variant. There has been a huge number of theoretical papers concerned with projective methods, extending the applicability of Karmarkar's method, making it easier to apply to problems in standard form, and studying the computational complexity of several of its variants. As LMS correctly point out, there are no efficient and robust implementations of this methodology. Besides the papers cited by LMS, computational results have been given by Fraley and Vial,[6] Todd,[20] and Anstreicher and Watteyne.[2] In addition, experiments by this author using a phase I-phase II projective variant on the NETLIB problems typically required about twice the number of iterations of OB1's predictor-corrector code, with a number of failures to attain full accuracy. Why do these methods not perform well in practice, and have they provided any useful lessons for the design of efficient implementations?

Anstreicher and Watteyne[2] conclude that the failure of projective methods to compete favorably with affine-scaling and primal-dual methods is connected with the necessary procedure to generate and update lower bounds on the optimal value; this procedure performs particularly poorly on instances where the dual feasible region has no interior. My own experience supports this conclusion; even if the instance is perturbed slightly to ensure interior feasible solutions whenever the original problem has an optimal solution, frequently several "centering" iterations, with constant objective function value, are necessary before the initial lower bound is first updated.

Given that projective methods are not competitive practically, have they been important to the development of practically efficient codes? My answer is yes; several key ideas were first proposed or developed in the context of projective methods. Such concepts as scaling to center the current iterate, using a centering component in the search direction, proving complexity bounds using a so-called potential function and developing extensions to handle problems for which no feasible starting solution is known were all first investigated in the context of projective methods.

Some of these ideas have been combined to yield other potential-reduction methods, avoiding the necessity of using projective transformations. Such methods were first devised by Gonzaga.[7] His paper describes these methods and variants designed to allow large steps. Again, no implementations of these methods for general use have been developed. But even if they are not practically efficient, the concepts underlying them have been fruitful in their influence on other more practical methods.

2. Primal-Dual Methods

Primal-dual approaches were investigated quite early on in interior points, although the first versions of all variants—projective, affine-scaling, path-following, and potential-reduction—of interior-point methods were developed either in a primal-only or dual-only setting. Megiddo[12] proposed the framework of a symmetric primal-dual algorithm first in 1986, and Kojima, Mizuno, and Yoshise developed the first primal-dual path-following method in early 1987.[9] Their method had a number of very interesting features. They made a large reduction in the barrier parameter at each iteration, allowed their iterates to lie in a wide neighborhood of the so-called central path, and took only a single Newton step before making another large reduction in the barrier parameter. These are properties that are important to the practical success of an implementation of such an algorithm. Nevertheless, they were able to obtain an $O(nL)$-iteration bound on the complexity of solving a linear programming problem with n nonnegative variables and integer data with length L.

At the same time, there were (primal-only or dual-only) path-following methods that needed only $O(\sqrt{n}\,L)$ iterations, and soon primal-dual methods with the same complexity were developed by Kojima, Mizuno, and Yoshise[10] and by Monteiro and Adler.[18] Although these methods and their analyses were very elegant, the improved complexity came at the expense of making tiny changes in the barrier parameter at each iteration (small-step methods), and requiring the iterates to lie in a very constrained neighborhood of the central path. Clearly neither of these properties was desirable in a practical code. The divergence of the theory and practice of interior-point methods was now unfortunately increasing.

The introduction of potential-reduction methods into primal-dual settings, by several authors including Kojima, Mizuno, Tanabe, Todd, Ye, and Yoshise, mitigated some of the disadvantages of these theoretical developments, but the parameters chosen in the potential function still tended to lead to small steps and hence inefficient practical methods. Moreover, longer steps required an approximate line search of a potential function, which is computationally expensive for large-scale problems.

Meanwhile, highly efficient implementations, based on the path-following ideas, but taking much larger steps, were being developed and showed excellent numerical behavior, although not even global convergence was proved. One particularly important refinement, developed by Lustig and then explained nicely by LMS, was driven by practical needs: a very elegant way to deal with problems where no interior feasible solutions were available to initialize the algorithms. Thus were born so-called infeasible-interior-point methods, with outstanding practical performance but no theoretical foundation at the time.

2.1. Superlinear Convergence

An important theoretical advance in the last few years has been the study of superlinear or quadratic convergence of primal-dual algorithms. Zhang, Tapia, and Dennis[28] gave sufficient conditions for such rates of convergence in terms of the centering parameter and step size chosen. Further papers described methods that attained such convergence together with polynomiality, culminating in the recent works of Ye et al.[25] and Mehrotra.[13] It may be that quadratic convergence is only manifested at an iteration when the algorithm is about to be halted anyway because sufficient accuracy has already been attained. Nevertheless, as El-Bakry, Tapia, and Zhang[4] point out, the presence of asymptotic superlinear or quadratic convergence usually implies that fast (and indeed increasingly fast) linear convergence is observed in practical computation. Thus the conditions of Zhang, Tapia, and Dennis are important to bear in mind in designing efficient implementations.

2.2. Predictor-Corrector Variants

The performance of practical primal-dual algorithms seems to be considerably improved by the incorporation of higher-order information, as in Mehrotra's predictor-corrector method[14] and its modification in OB1, described by LMS. This was motivated by theoretical work on higher derivatives of trajectories in the feasible region, by Megiddo,[12] Bayer and Lagarias[3] and (with computational results) Adler et al.[1] One of the nicest features of Mehrotra's approach is that solving for the affine-scaling directions gives a good indication of how well the algorithm is progressing, and hence of how fast the barrier parameter μ (and hence the centering component in the search direction) can be decreased.

2.3. Choice of Step Size

Because asymptotically one wants to employ full Newton steps as well as no centering in order to achieve superlinear convergence, it might also be worthwhile to use the affine-scaling step computed in the predictor-corrector method to adaptively change the fraction of the way to the boundary taken.

An alternative way to change the step size adaptively is to use a potential function. One might monitor the reduction obtained in a suitable potential function when choosing step sizes a fixed fraction of the way to the boundary. If the reduction is small, the fraction could be decreased at the next iteration, whereas if it is large, the fraction could be increased. Again, this might allow the conditions for superlinear convergence to hold asymptotically.

Note that Mehrotra[14] advocates the use of a potential function to assure global convergence of his predictor-corrector method, and also proposes an adaptive step-size rule.

2.4. Linear Algebra

As LMS point out, the main computational burden at each iteration is to perform the factorization required to solve a suitable linear system and thus compute the search directions. In OB1, a Cholesky factorization is used. Another approach, recommended by Fourer and Mehrotra[5] and Vanderbei and Carpenter,[23] is to solve the so-called augmented or Karush-Kuhn-Tucker system directly. This allows a less *ad hoc* treatment of free variables (in the augmented system, they correspond to zero diagonal elements) and gives more freedom in pivot selection for sparsity and stability. There has been considerable recent interest in the numerical linear algebra community in the efficient solution of such systems.

The importance of pre-processing to interior-point codes is rightly stressed by LMS. Because the linear systems in such methods involve all the data of the problem (in contrast to simplex methods), problem reductions are even more important.

3. Recent Convergence Theory

The divergence of practical methods from those being investigated theoretically has happily been decreasing recently. LMS quote the important global convergence results of Kojima, Megiddo, and Mizuno[8] and the related polynomial complexity result of Mizuno.[17] Before Mizuno,[17] Zhang[27] had given the first polynomial-time convergence result for a practical infeasible-interior-point method. Since then, there has been considerable further theoretical analysis of such methods, and superlinear convergence of such an algorithm has been established by Wright.[24]

In order to fully analyze infeasible-interior-point methods, it is necessary to describe their behavior on primal or dual infeasible problems. Some of the papers just cited treat this aspect to some extent, but the results are not completely satisfactory. Basically, the methods search very hard for optimality, and only detect infeasibility by diverging. Guaranteed detection is only available theoretically by starting with a very inefficient initial point. A new homogeneous and self-dual formulation due to Ye, Todd, and Mizuno[26] treats optimality and infeasibility more symmetrically and may well prove to be significant computationally.

Our understanding of affine-scaling methods has improved recently, largely as a result of the work of Tsuchiya and Muramatsu,[22] who prove convergence (of the primal or dual version) with long steps making no nondegeneracy assumptions. Their analysis shows that the direction includes some "automatic centering," suggesting why the method has been quite successful in practice.

LMS cite the recent result of Powell, showing that Karmarkar's algorithm may require $\Omega(n)$ iterations to obtain a constant improvement in precision. In a related vein, the recent paper[21] proves that long-step symmetric primal-dual methods, very similar to an earlier version of OB1, may require $\Omega(n^{1/3})$ iterations to achieve a modest constant improvement in the duality gap. This is far better than the linear growth of Powell's example, but certainly asymptotically much worse than the logarithmic growth observed in practical computation, indicating that even interior-point methods may have pathologically bad examples like the simplex method.

4. Open Questions

LMS conclude their article by mentioning three areas under current study. One is the detection of infeasibility discussed in the previous section, whereas another addresses the important topic of a "cross-over" from an interior-point approach to a simplex method. The remaining area is that of warm starts, and is perhaps the weakest aspect of current interior-point methods.

The recent paper of LMS[11] demonstrates their success in solving various perturbations of a given linear programming problem. Interestingly, a variant of OB1 incorporating the global safeguards of Kojima, Megiddo, and Mizuno[8] proved superior to the original code for these problems (and not much worse on all the NETLIB problems). Nevertheless, the algorithm with a warm start still required about a quarter to a third of the iterations needed from a cold start. No comparison with the simplex method was given, but it seems that improved performance will be necessary to compete with simplex re-optimization techniques. Of course, such a capability is crucial in attacking integer programming problems using either a branch-and-bound or cutting-plane approach. (Note that Mitchell and Todd[16] give mildly encouraging results on using an interior-point algorithm in a cutting-plane setting, and further papers of Mitchell and co-workers give more reason for optimism for such a marriage—see for instance [15].)

Finally, although it is not within the scope of LMS's article, I just wish to mention that a great deal of work has been done on extending interior-point methods to various classes of nonlinear programming problems. One of the most general settings in which polynomial time bounds can be established is related to the concept of *self-concordant* functions and has been studied in great depth of Nesterov and Nemirovsky.[19]

Acknowledgments

This research was supported in part by NSF, AFOSR, and ONR through NSF grant DMS-8920550.

References

1. I. Adler, M. G. C. Resende, G. Veiga and N. K. Karmarkar, 1989. An Implementation of Karmarkar's Algorithm for Linear Programming, *Mathematical Programming 44*, 297–335. Errata in *Mathematical Programming 50*, 415, 1991.
2. K. M. Anstreicher and P. Watteyne, 1993. A Family of Search Directions for Karmarkar's Algorithm, *Operations Research 41*, 759–767.
3. D. A. Bayer and J. C. Lagarias, 1989. The Nonlinear Geometry of Linear Programming, Part II: Legendre Transform Coordinates, *Transactions of the American Mathematical Society 314*, 527–581.
4. A. El-Bakry, R. Tapia and Y. Zhang, 1991. Numerical Comparisons of Local Convergence Strategies for Interior-point Methods in Linear Programming, Technical Report TR 91-18, Department of Computational and Applied Mathematics, Rice University, Houston, TX.
5. R. Fourer and S. Mehrotra, 1991. Performance of an Augmented System Approach for Solving Least-Squares Problems in an Interior-Point Method for Linear Programming, *COAL Newsletter 19*, 26–31.
6. C. Fraley and J. P. Vial, 1989. Numerical Study of Projective Methods for Linear Programming, pp. 25–38 in *Optimization: Proceedings of the 5th French-German Conference in Castel-Novel, Varetz, France, October 1988*, S. Dolecki (ed.), Springer Verlag, Berlin, Germany.
7. C. C. Gonzaga, 1991. Large Steps Path-Following Methods for Linear Programming, Part II: Potential Reduction Method, *SIAM Journal on Optimization 1*, 280–292.
8. M. Kojima, N. Megiddo and S. Mizuno, 1993. A Primal-Dual Infeasible Interior Point Algorithm for Linear Programming, *Mathematical Programming* 61, 263–280.
9. M. Kojima, S. Mizuno and A. Yoshise, 1989. A Primal-Dual Interior Point Algorithm for Linear Programming, pp. 29–47 in *Progress in Mathematical Programming: Interior Point and Related Methods*, N. Megiddo (ed.), Springer Verlag, NY.
10. M. Kojima, S. Mizuno and A. Yoshise, 1989. A Polynomial-Time Algorithm for a Class of Linear Complementarity Problems, *Mathematical Programming 44*, 1–26.
11. I. J. Lustig, R. E. Marsten and D. F. Shanno, 1992. Computational Experience with a Globally Convergent Primal-Dual Predictor-Corrector Algorithm for Linear Programming, Technical

Report SOR-92-10, School of Engineering and Applied Science, Dept. of Civil Engineering and Operations Research, Princeton University, Princeton, NJ.

12. N. MEGIDDO, 1989. Pathways to the Optimal Set in Linear Programming, pp. 131–158 in *Progress in Mathematical Programming: Interior Point and Related Methods*, N. Megiddo (ed.), Springer Verlag, New York. Identical version in: *Proceedings of the 6th Mathematical Programming Symposium of Japan, Nagoya, Japan*, pages 1–35, 1986.
13. S. MEHROTRA, 1993. Quadratic Convergence in a Primal-Dual Method, *Mathematics of Operations Research 18*, 741–751.
14. S. MEHROTRA, 1992. On the Implementation of a Primal-Dual Interior Point Method, *SIAM Journal on Optimization 2*, 575–601.
15. J. E MITCHELL AND B. BORCHERS, 1992. A Primal-Dual Interior Point Cutting Plane Method for the Linear Ordering Problem, *COAL Bulletin, 21*, 13–18.
16. J. E MITCHELL AND M. J. TODD, 1992. Solving Combinatorial Optimization Problems using Karmarkar's Algorithm, *Mathematical Programming 56*, 245–284.
17. S. MIZUNO, 1992. Polynomiality of the Kojima-Megiddo-Mizuno Infeasible Interior Point Algorithm for Linear Programming, Technical Report 1006, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.
18. R. D. C. MONTEIRO AND I. ADLER, 1989. Interior Path Following Primal-Dual Algorithms· Part I· Linear Programming, *Mathematical Programming 44*, 27–41.
19. Y. E. NESTEROV AND A. S. NEMIROVSKY, 1993. *Interior Point Polynomial Methods in Convex Programming: Theory and Algorithms*, SIAM, Philadelphia, PA.
20. M. J. TODD, 1990. The Effects of Degeneracy, Null and Unbounded Variables on Variants of Karmarkar's Linear Programming Algorithm, pp. 81–91 in *Large-Scale Numerical Optimization*, T. F. Coleman and Y. Li (eds.), SIAM, Philadelphia, PA.
21. M. J. TODD, 1993. A Lower Bound on the Number of Iterations of Primal-Dual Interior-Point Methods for Linear Programming, Technical Report 1050, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.
22. T. TSUCHIYA AND M. MURAMATSU, 1992. Global Convergence of the Long-Step Affine Scaling Algorithm for Degenerate Linear Programming Problems, Research Memorandum 423, The Institute of Statistical Mathematics, 4-6-7 Minami-Azabu, Minato-ku, Tokyo 106, Japan.
23. R. J. VANDERBEI AND T. J. CARPENTER, 1993. Symmetric Indefinite Systems for Interior Point Methods, *Mathematical Programming 58*, 1–32.
24. S. J. WRIGHT, 1993. A Superlinear Infeasible-Interior-Point Algorithm for Monotone Nonlinear Complementarity Problems, Technical Report MCS-P344-1292, Mathematical and Computer Science Division, Argonne National Laboratory, Argonne, IL.
25. Y YE, O. GÜLER, R. A. TAPIA AND Y ZHANG, 1993. A Quadratically Convergent $O(\sqrt{n}L)$-Iteration Algorithm for Linear Programming, *Mathematical Programming 59*, 151–162.
26. Y. YE, M. J. TODD AND S. MIZUNO, 1992. An $O(\sqrt{n}L)$-Iteration Homogeneous and Self-Dual Linear Programming Algorithm, Technical Report, Department of Management Science, University of Iowa, Iowa City, IA (to appear in *Mathematics of Operations Research*).
27. Y. ZHANG, 1992. On the Convergence of a Class of Infeasible Interior-Point Methods for the Horizontal Linear Complementarity Problem, Research Report 92-07, Department of Mathematics and Statistics, University of Maryland Baltimore County, Baltimore, MD (to appear in *SIAM Journal on Optimization*).
28. Y. ZHANG, R. A. TAPIA AND J E. DENNIS, 1992. On the Superlinear and Quadratic Convergence of Primal-Dual Interior Point Linear Programming Algorithms, *SIAM Journal on Optimization 2*, 304–324.