Comput Optim Appl (2014) 59:399–403
DOI 10.1007/s10589-014-9707-3

COAP 2013 Best Paper Prize

Published online: 21 October 2014

Each year, the editorial board of Computational Optimization and Applications (COAP) selects a paper from the preceding year's publications for the Best Paper Award. In 2013, 91 papers were published by COAP. The recipients of the 2013 Best Paper Award are Miles Lubin, Cosmin Petra and Mihai Anitescu of Argonne National Laboratory, and Julian Hall of the University of Edinburgh for their paper "Parallel distributed-memory simplex for large-scale stochastic LP problems" published in volume 55 pages 571–596. This article highlights the research related to the award winning paper.

The paper describes the development and implementation of techniques for solving very large scale stochastic linear programming (LP) problems using the dual revised simplex method on large multiprocessor architectures. The resulting solver (`PIPS-S`) is notable for being the first parallel implementation of the revised simplex method which is both efficient and scalable. When solving large instances on a multiprocessor cluster, `PIPS-S` was hundreds of times faster than `clp` [7] and the largest problem solved had 10^8 variables.

The development of an efficient parallel implementation of the revised simplex method has attracted the attention of leading researchers over the past 30 years, but with relatively little success [1]. For example, Julian Hall and Ken McKinnon devised and implemented two parallel schemes in the mid 1990s [3,4] which gave limited speedup on general sparse LP problems. A byproduct of this work was the identification of hyper-sparsity in the simplex method and techniques for its exploitation using serial computing [5].

In 2009, the Argonne group led by Mihai Anitescu commenced a project in the optimization of complex energy systems under uncertainty. The project, funded by the US Department of Energy, was aimed at using stochastic programming to systematically mitigate uncertainty associated with renewable sources of energy such as wind and solar in the operation and planning of electric power grids. The group assembled realistic power-grid stochastic optimization models covering large geographic areas. In addition to accounting for network topologies, the models incorporated scenarios for generation levels from wind turbines and solar plants computed by using state-of-the-art numerical weather forecasting. The resulting stochastic programming problems were of unprecedented sizes that could be solved only by means of parallel, distributed-memory computing.

To tackle these problems, Cosmin Petra and Miles Lubin initially developed `PIPS-IPM` [8–11], a solver based on interior-point methods which achieved parallelism by decomposing the linear-algebra operations required at each iteration. The solver was quite successful in solving large-scale stochastic LPs and QPs on a number of high-performance computing platforms, scaling up to hundreds of thousands of cores. However, the well-known difficulties in hot-starting interior-point algorithms led the Argonne group to consider alternative algorithms for situations where a sequence of continuous problems must be solved, for example, within a branch-and-bound algorithm for solving mixed-integer stochastic programming problems. The obvious candidate was the (dual) simplex method, and the hope was that a decomposition based on linear algebra could again be successful.

So it was that on 1 August 2011 Miles sent a message to Julian asking what code might exist from which an efficient parallel dual simplex solver for block angular LP problems could be developed. Julian's immediate response was that there was essentially nothing to work from and that, from his experience of developing a parallel primal simplex solver for row-linked block angular LP problems with Edmund Smith [13], he was acutely aware of the challenge of combining the words "efficient", "parallel", "simplex" and "structured". To be fair, despite his relative inexperience, so was Miles. In a brief flurry of emails, it was concluded that the code was best written from scratch, based on the matrix factorization and linear system solver components of CoinUtils [7]. Thus an exchange of advice developed into a transatlantic collaboration.

Given that Miles had only 3–4 months available for coding, some of the more ambitious ideas for the implementation were shelved. In particular, motivated by Julian's experience with Smith, the product form update was chosen, leaving as a "black box" the solution of systems involving the matrix factorised directly. This approach led to some loss of serial efficiency relative to `clp`, but doubtless paid off in code simplicity and parallel scalability. During the course of the autumn, Julian explained a few computational tricks to Miles who, otherwise, did all the work. Overall efficiency was also enhanced by exploiting the structure of the RHS of linear systems. Without this block hyper-sparsity, scalability might have been improved at the expense of performance relative to an efficient serial solver. This trade-off between absolute efficiency and scalable performance is a recurrent feature of parallel simplex research and a reason why some superficially impressive work is of little practical value.

A feature of the stochastic LP problems which made their solution tractable was that the solution of the problem with half the number of scenarios could be used to obtain a "crash" basis from which only a relatively small number of iterations were required to obtain an optimal solution. Additionally, the efficiency of the implementation benefited greatly from recent advances in high-performance computing architectures (namely, low-latency interconnects).

In the context of power systems, this algorithmic advancement helps build a foundation for developing computational infrastructure to support real-time scheduling and dispatch accounting for the high levels of renewable energy penetration anticipated in the US by 2030 (between 20 and 30 %).

The paper has already been acknowledged by COIN-OR with the award of the COIN-OR INFORMS 2013 Cup, recognising that COIN-OR libraries, in particular the CoinFactorization class of the CoinUtils package, were a foundation for the implementation of `PIPS-S`.

Another promising idea which offers greater scope for parallelism was considered but not implemented was suboptimization [12] using minor iterations of the standard dual simplex method. This has provided the basis of an efficient parallel implementation of the dual simplex method for general LP problems developed by Huangfu and Hall [2,6]. Huangfu's successful implementation of these techniques within the FICO Xpress dual simplex solver has featured in FICO advertising. Perhaps, after 30 years, the "pipe dream" of parallelising the simplex method is a reality.

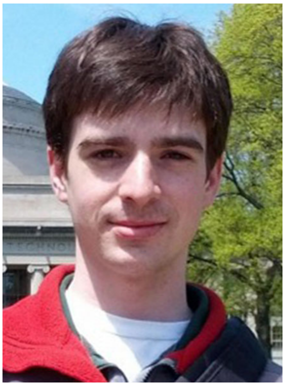

Miles Lubin received his B.A. in Applied Mathematics and M.S. in Statistics from the University of Chicago in 2011, before spending a year at Argonne National Laboratory, where this work was performed. He is currently a Ph.D. candidate in Operations Research at the Massachusetts Institute of Technology.

Julian Hall received his B.A. in Mathematics from the University of Oxford in 1987 and his Ph.D. from the University of Dundee in 1992. Since 1990 he has been employed as a lecturer in the School of Mathematics at the University of Edinburgh.

Cosmin Petra is a computational mathematician in the Mathematics and Computer Science Division at Argonne National Laboratory. He received his B.Sc. degree in Mathematics and Computer Science from "Babes-Bolyai" University, Romania, and his M.S. and Ph.D. degrees in Applied Mathematics from the University of Maryland, Baltimore County, USA. His research interests lie at the intersection of numerical optimization, operations research, statistical analysis, and high-performance scientic computing with a focus on the optimization of large-scale energy systems under uncertainty.

Mihai Anitescu received his Engineer degree in electrical engineering from the "Politehnica" University of Bucharest in 1992 and his Ph.D. in applied and computational mathematics from the University of Iowa in 1997. He is currently jointly appointed as a senior computational mathematician at Argonne National Laboratory and a professor in the Department of Statistics at the University of Chicago.

References

1. Hall, J.A.J.: Towards a practical parallelisation of the simplex method. Comput. Manag. Sci. **7**(2), 139–170 (2010)
2. Hall, J.A.J., Huangfu, Q.: A high performance dual revised simplex solver. In: W, R., et al. (eds.) PPAM 2011, Part I, Volume 7203 of LNCS, pp. 143–151. Springer, Heidelberg (2012)

3. Hall, J.A.J., McKinnon, K.I.M.: PARSMI, a parallel revised simplex algorithm incorporating minor iterations and Devex pricing. In: Waśniewski, J., Dongarra, J., Madsen, K., Olesen, D. (eds.) Applied Parallel Computing, Volume 1184 of Lecture Notes in Computer Science, pp. 67–76. Springer, Berlin (1996)
4. Hall, J.A.J., McKinnon, K.I.M.: ASYNPLEX, an asynchronous parallel revised simplex method algorithm. Ann. Oper. Res. **81**, 27–49 (1998)
5. Hall, J.A.J., McKinnon, K.I.M.: Hyper-sparsity in the revised simplex method and how to exploit it. Comput. Optim. Appl. **32**(3), 259–283 (2005)
6. Huangfu, Q., Hall, J. A. J.: Parallelizing the dual revised simplex method. Technical Report ERGO-14-011, School of Mathematics, University of Edinburgh (2014)
7. Lougee-Heimer, R., et al.: The COIN-OR initiative: open source accelerates operations research progress. ORMS Today **28**(4), 20–22 (2001)
8. Lubin, M., Petra, C. G., Anitescu, M., Zavala, V.: Scalable stochastic optimization of complex energy systems. In Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11, pages 64:1–64:64, ACM, New York (2011)
9. Petra, C.G., Anitescu, M.: A preconditioning technique for Schur complement systems arising in stochastic optimization. Comput. Optim. Appl. **52**, 315–344 (2012)
10. Petra, C.G., Schenk, O., Anitescu, M.: Real-time stochastic optimization of complex energy systems on high performance computers. Comput. Sci. Eng. **99**(PrePrints), 1–9 (2014)
11. Petra, C.G., Schenk, O., Lubin, M., Gärtner, K.: An augmented incomplete factorization approach for computing the Schur complement in stochastic optimization. SIAM J. Sci. Comput. **36**(2), C139–C162 (2014)
12. Rosander, R.R.: Multiple pricing and suboptimization in dual linear programming algorithms. Math. Progr. Study **4**, 108–117 (1975)
13. Smith, E.: Parallel solution of linear programs. PhD thesis, University of Edinburgh (2013)