

# Solving Multiscale Linear Programs Using the Simplex Method in Quadruple Precision

Ding Ma and Michael A. Saunders

**Abstract** Systems biologists are developing increasingly large models of metabolism and integrated models of metabolism and macromolecular expression. These Metabolic Expression (ME) models lead to sequences of multiscale linear programs for which small solution values of order  $10^{-6}$  to  $10^{-10}$  are meaningful. Standard LP solvers do not give sufficiently accurate solutions, and exact simplex solvers are extremely slow. We investigate whether double-precision and quadruple-precision simplex solvers can together achieve reliability at acceptable cost.

A double-precision LP solver often provides a reasonably good starting point for a Quad simplex solver. On a range of multiscale examples we find that 34-digit Quad floating-point achieves exceptionally small primal and dual infeasibilities (of order  $10^{-30}$ ) when no more than  $10^{-15}$  is requested. On a significant ME model we also observe robustness in almost all (even small) solution values following relative perturbations of order  $10^{-6}$  to non-integer data values.

Double and Quad Fortran 77 implementations of the linear and nonlinear optimization solver MINOS are available upon request.

**Keywords** Flux balance analysis • Metabolic expression model • Multiscale linear program • Simplex method • Quadruple precision • Gfortran libquadmath • MINOS

## 1 Introduction

We consider the solution of large, multiscale linear programs (LPs) of the form

$$\min_x c^T x \text{ s.t. } \ell \leq \begin{pmatrix} x \\ Ax \end{pmatrix} \leq u, \quad (1)$$

---

D. Ma • M.A. Saunders (✉)

Department of Management Science and Engineering, Stanford University, Stanford, CA, USA  
e-mail: [dingma@stanford.edu](mailto:dingma@stanford.edu); [saunders@stanford.edu](mailto:saunders@stanford.edu); [metsaunders@gmail.com](mailto:metsaunders@gmail.com)

© Springer International Publishing Switzerland 2015  
M. Al-Baali et al. (eds.), *Numerical Analysis and Optimization*, Springer  
Proceedings in Mathematics & Statistics 134, DOI 10.1007/978-3-319-17689-5\_9

223

where  $A$  is a sparse matrix whose entries, like the variables in  $x$ , may be of widely varying magnitude. Such problems arise in systems biology in the modeling of biochemical reaction networks, notably Metabolic Expression (ME) models [20, 36]. Reliable solution methods are of such importance to systems biologists that *exact simplex solvers* have been employed [20], even though the typical solution time for an exact solver is measured in weeks for genome-scale models (compared to minutes for a conventional solver using double-precision floating-point arithmetic).

Exact solvers are based on rational arithmetic. There has been considerable work on their implementation and application to important problems [1, 2, 18, 32]. The use of quadruple-precision floating-point has also been mentioned in passing [1, 18].

Let Single, Double, and Quad denote the main floating-point options, with about 7, 15, 16, and 34 digits of precision, respectively. Single is not useful in the present context, and Double may not ensure adequate accuracy. This is the reason for our work. On today's machines, Double is implemented in hardware, while Quad (if available) is typically implemented in a software library such as libquadmath [8]. Fortunately, the GCC Fortran compiler now makes Quad available via the `real(16)` data type. We have therefore been able to make a Quad version of the Fortran 77 linear and nonlinear optimization solver MINOS [26, 27] using the gfortran compiler. Our aim is to explore combined use of the Double and Quad MINOS simplex solvers for the solution of large multiscale linear programs. We seek greater efficiency than is normally possible with exact simplex solvers.

Kahan [16] notes that “*carrying somewhat more precision in the arithmetic than twice the precision carried in the data and available for the result will vastly reduce embarrassment due to roundoff-induced anomalies.*” He further notes that Quad precision is unlikely to be adopted widely in the foreseeable future because of the cost in CPU time and memory (especially cache) relative to Double, but in terms of finding ways to avoid unexpected total loss of accuracy, “*default evaluation in Quad is the humane option.*”

We apply the “humane” approach to difficult LP problems by using the Double simplex solver first, saving the final solution, and warm-starting Quad simplex from that point. For a sequence of related problems, warm-starting each problem in Quad is simplest, but warm-starting in Double and then in Quad may be more efficient.

## 2 Motivating Applications

The Constraint-Based Reconstruction and Analysis (COBRA) approach [31, 33] has been successfully applied to biological processes such as metabolism and macromolecular synthesis, which when integrated result in inherently multiscale models. In the COBRA approach, a biochemical network is represented by a stoichiometric matrix  $S$  with  $m$  rows corresponding to metabolites and  $n$  columns representing reactions. Mathematically,  $S$  is part of the ordinary differential equation that governs the time-evolution of concentrations in the network:

$$\frac{d}{dt}x(t) = Sv(t), \quad (2)$$

where  $x(t) \in \mathbf{R}^m$  is a vector of time-dependent concentrations and  $v(t) \in \mathbf{R}^n$  is a vector of reaction fluxes. With  $c^T v$  being a biologically motivated objective function (such as maximizing the growth rate at steady state), the constraint-based approach constructs the following LP:

$$\max_v c^T v \quad (3a)$$

$$\text{s.t. } Sv = 0, \quad (3b)$$

$$l \leq v \leq u, \quad (3c)$$

where growth is defined as the biosynthetic requirements of experimentally determined biomass composition, and biomass generation is a set of reaction fluxes linked in appropriate ratios [31].

The following applications have motivated our work.

**Flux Balance Analysis (FBA).** FBA is a mathematical and computational approach widely used for studying biochemical reaction networks [30, 31]. The biochemical networks reconstructed in FBA with a linear objective function are essentially LPs as in (3), where the fluxes in vector  $v$  may have widely varying values in the range 0–100 say, with small values such as  $v_j = 10^{-10}$  being meaningful. With the increasingly large, sparse, and multiscale nature of biochemical networks, a Quad solver has become more necessary, practical, and even efficient.

**ME models (FBA with Coupling Constraints).** FBA has been used by Thiele et al. [36] for the first integrated stoichiometric multiscale model of metabolism and macromolecular synthesis for *Escherichia coli* K12 MG1655. The model modifies (3) by adding constraints that couple enzyme synthesis and catalysis reactions to (3b). Coupling constraints of the form

$$c_{\min} \leq \frac{v_i}{v_j} \leq c_{\max} \quad (4)$$

become linear constraints

$$c_{\min} v_j \leq v_i, \quad v_i \leq c_{\max} v_j \quad (5)$$

for various pairs of fluxes  $v_i, v_j$ . They are linear approximations of nonlinear constraints and make  $S$  in (3b) even less well-scaled because of large variations in reaction rates. Quad precision is evidently more appealing in this case.

**ME Models with Nonlinear Constraints.** As coupling constraints are often functions of the organism's growth rate  $\mu$ , O'Brien et al. [29] consider growth-rate optimization nonlinearly with the single  $\mu$  as the objective in (3a) instead of via a linear biomass objective function. Nonlinear constraints of the form

$$\frac{v_i}{v_j} \leq \mu \quad (6)$$

represented as

$$v_i \leq \mu v_j \quad (7)$$

are added to (3b), where  $v_i, v_j, \mu$  are all variables. Constraints (7) are linear if  $\mu$  is fixed at a specific value  $\mu_k$ . O'Brien et al. [29] employ a binary search on a discrete set of values within an interval  $[\mu_{\min}, \mu_{\max}]$  to find the largest  $\mu_k \equiv \mu^*$  that keeps the associated linear program feasible. Thus, the procedure requires reliable solution of a sequence of related LPs.

**Flux Variability Analysis (FVA).** After FBA (3) returns an optimal objective value  $c^T v^* = Z_0$  (3a), FVA examines how far a particular flux  $v_j$  can vary within the feasible region without changing the optimal objective significantly (if  $\gamma \approx 1$ ):

$$\begin{aligned} \max_v \text{ or } \min_v \quad & v_j \\ \text{s.t.} \quad & Sv = 0, \\ & c^T v \geq \gamma Z_0, \\ & l \leq v \leq u, \end{aligned} \quad (8)$$

where  $0 < \gamma < 1$ . Potentially  $2n$  LPs (8) are solved if all reactions are of interest, with warm starts being used when  $j$  increases to  $j + 1$  [12].

**Other Challenging LPs.** A set of difficult LP problems has been collected by Mészáros [22], who names them *problematic* and notes that “*modeling mistakes made these problems “crazy,” but they are excellent examples to test numerical robustness of a solver.*” Our procedure for handling these *problematic* problems seems appropriate for the systems biology models as well.

### 3 Algorithm and Implementation

The primal simplex solver in MINOS includes geometric-mean scaling of the constraint matrix, the EXPAND anti-degeneracy procedure [10, 14], and partial pricing (but no steepest-edge pricing, which would generally reduce total iterations and time). Basis LU factorizations and updates are handled by LUSOL [9, 21]. Cold starts use a Crash procedure to find a triangular initial basis. Basis files are used to preserve solutions between runs.

For Double MINOS, floating-point variables are declared double precision ( $\approx 15$  digits). For Quad MINOS, they are real(16) ( $\approx 34$  digits). The LP data  $A, b, c, S, \ell, u$  are stored in Quad even though they are not known to that precision. This allows operations such as  $Ax$  and  $A^T y$  to be carried out directly on the elements of  $A$  and the Quad vectors  $x, y$ . If  $A$  were stored in Double, such products would require each entry  $A_{ij}$  to be converted from Double to Quad at runtime.

To achieve reliability on the Mészáros problems, we developed the following three-step procedure for solving challenging examples of problems (1)–(8):

**Step 1** (Cold start in Double with scaling) Apply Double MINOS with somewhat strict options. Save a final basis file.

**Table 1** MINOS runtime options (defaults and those selected for Steps 1–3)

|                  | Default<br>Double | Step1<br>Double | Step2<br>Quad | Step3<br>Quad |
|------------------|-------------------|-----------------|---------------|---------------|
| Scale option     | 2                 | 2               | 2             | 0             |
| Feasibility tol  | 1e-6              | 1e-7            | 1e-15         | 1e-15         |
| Optimality tol   | 1e-6              | 1e-7            | 1e-15         | 1e-15         |
| LU factor tol    | 100.0             | 10.0            | 10.0          | 5.0           |
| LU update tol    | 10.0              | 10.0            | 10.0          | 5.0           |
| Expand frequency | 10,000            | 100,000         | 100,000       | 100,000       |

**Step 2** (Warm start in Quad with scaling) Start Quad MINOS from the saved file with stricter Feasibility and Optimality tolerances. Save a final basis file.

**Step 3** (Warm start in Quad without scaling) Start Quad MINOS from the second saved file with no scaling but stricter LU tolerances.

Steps 1 and 2 are “obvious” and should usually be sufficient. In case Step 2 is interrupted, Step 3 provides some insurance and ensures that the Feasibility and Optimality tolerances are imposed upon the original problem (not the scaled problem).

Table 1 shows the default runtime options for Double MINOS and the options chosen for Steps 1–3 above. The Feasibility tolerance  $\delta_1$  is applied in absolute form. Thus, a (possibly scaled) solution  $v$  is considered feasible for problem (3) if  $\ell - \delta_1 \leq v \leq u + \delta_1$ . The Optimality tolerance  $\delta_2$  is applied in a relative way. If the current basic solution is of the form  $Sv \equiv Bv_B + Nv_N = b$  and if  $B^T y = c_B$  for the nonsingular basis matrix  $B$ , the current  $v$  is considered optimal if  $z \equiv c - A^T y$  has the correct sign to within the tolerance  $(1 + \|y\|_\infty)\delta_2$ .

For conventional Double solvers it is reasonable to set  $\delta_1$  and  $\delta_2$  in the range  $10^{-6}$  to  $10^{-8}$ . For Quad MINOS we set  $\delta_1 = \delta_2 = 10^{-15}$  to be sure of capturing accurately any fluxes  $v_j$  as small as  $O(10^{-10})$ .

## 4 Numerical Results

We report results from Double and Quad versions of the primal simplex solver in MINOS. All runs were on a 2.93 GHz Apple iMac with quad-core Intel i7, using the gfortran compiler with -O flag. Double MINOS uses 64-bit hardware floating-point throughout. Quad MINOS uses 128-bit software floating-point throughout via gfortran’s libquadmath library.

We applied our three-step procedure to three sets of LP problems. Table 2 lists the problem dimensions and the norms of the optimal primal and dual solution vectors  $x^*$ ,  $y^*$ . Table 3 summarizes the results of Steps 1–3 for each problem.

All problems were input from files in the classical MPS format of commercial mathematical programming systems [24] with 12-character fields for all data values. This was a fortuitous limitation for the ME models, as we mention below. The MPS files for these 14 LP models are downloadable from [25].

**Table 2** Three pilot models from Netlib [28], eight Mészáros *problematic* LPs [22], and three ME biochemical network models [19, 35, 36]

| Model    | $m$    | $n$    | $\text{nnz}(A)$ | $\max  A_{ij} $ | $\ x^*\ _\infty$ | $\ y^*\ _\infty$ |
|----------|--------|--------|-----------------|-----------------|------------------|------------------|
| pilot4   | 411    | 1,000  | 5,145           | 2.8e+04         | 9.6e+04          | 2.7e+02          |
| pilot    | 1,442  | 3,652  | 43,220          | 1.5e+02         | 4.1e+03          | 2.0e+02          |
| pilot87  | 2,031  | 4,883  | 73,804          | 1.0e+03         | 2.4e+04          | 1.1e+01          |
| de063155 | 853    | 1,488  | 5,405           | 8.3e+11         | 3.1e+13          | 6.2e+04          |
| de063157 | 937    | 1,488  | 5,551           | 2.3e+18         | 2.3e+17          | 6.2e+04          |
| de080285 | 937    | 1,488  | 5,471           | 9.7e+02         | 1.1e+02          | 2.6e+01          |
| gen1     | 770    | 2,560  | 64,621          | 1.0e+00         | 3.0e+00          | 1.0e+00          |
| gen2     | 1,122  | 3,264  | 84,095          | 1.0e+00         | 3.3e+00          | 1.0e+00          |
| gen4     | 1,538  | 4,297  | 110,174         | 1.0e+00         | 3.0e+00          | 1.0e+00          |
| l30      | 2,702  | 15,380 | 64,790          | 1.8e+00         | 1.0e+09          | 4.2e+00          |
| iprob    | 3,002  | 3,001  | 12,000          | 9.9e+03         | 3.1e+02          | 1.1e+00          |
| TMA_ME   | 18,210 | 17,535 | 336,302         | 2.1e+04         | 5.9e+00          | 1.1e+00          |
| GlcAerWT | 68,300 | 76,664 | 926,357         | 8.0e+05         | 6.3e+07          | 2.4e+07          |
| GlcAlift | 69,529 | 77,893 | 928,815         | 2.6e+05         | 6.3e+07          | 2.4e+07          |

Dimensions of  $m \times n$  constraint matrices  $A$  ( $= S$  for the ME models), and size of the largest optimal primal and dual variables  $x^*$ ,  $y^*$

### The Pilot Problems

These are economic models developed by Prof George Dantzig's group in the Systems Optimization Laboratory at Stanford University during the 1980s. They are available from Netlib [28]. For the middle example (pilot), MINOS required about 24 h on a DEC MicroVAX II during 1987, and did not perform reliably until the EXPAND anti-degeneracy procedure was developed.

Line 1 for pilot in Table 3 shows that Double MINOS with cold start and scaling required 16,060 primal simplex iterations and 5.7 CPU seconds. The final unscaled primal solution  $x$  satisfied the bounds  $\ell$  and  $u$  in (1) to within  $O(10^{-6})$ , and the dual solution  $y$  satisfied the optimality conditions to within  $O(10^{-3})$ .

Line 2 for pilot shows that Quad MINOS starting from that point with scaling needed only 29 iterations and 0.7 s to obtain a very accurate solution (where  $\text{Pinf} = 10^{-99}$  means that the maximum primal infeasibility was 0.0).

Line 3 for pilot shows that in the "insurance" step, Quad MINOS warm-starting again but with no scaling gave a full quad-precision solution at almost no cost: maximum infeasibilities 0.0 and  $O(10^{-32})$ . The final Double and Quad objective values differ in the fourth significant digit, as suggested by removal of Step 1's  $O(10^{-3})$  dual infeasibility.

Results for the bigger problem pilot87 are analogous.

**Table 3** Iterations and runtimes in seconds for Step 1 (Double MINOS) and Steps 2 and 3 (Quad MINOS)

| Model    | Iterations | Times    | Final objective   | Pinf | Dinf |
|----------|------------|----------|-------------------|------|------|
| pilot4   | 1,571      | 0.1      | -2.5811392602e+03 | -05  | -13  |
|          | 6          | 0.0      | -2.5811392589e+03 | -39  | -31  |
|          | 0          | 0.0      | -2.5811392589e+03 | -99  | -30  |
| pilot    | 16,060     | 5.7      | -5.5739887685e+02 | -06  | -03  |
|          | 29         | 0.7      | -5.5748972928e+02 | -99  | -27  |
|          | 0          | 0.2      | -5.5748972928e+02 | -99  | -32  |
| pilot87  | 19,340     | 15.1     | 3.0171038489e+02  | -09  | -06  |
|          | 32         | 2.2      | 3.0171034733e+02  | -99  | -33  |
|          | 0          | 1.2      | 3.0171034733e+02  | -99  | -33  |
| de063155 | 921        | 0.0      | 1.8968704286e+10  | -13  | +03  |
|          | 78         | 0.1      | 9.8830944565e+09  | -99  | -17  |
|          | 0          | 0.0      | 9.8830944565e+09  | -99  | -24  |
| de063157 | 488        | 0.0      | 1.4561118445e+11  | +20  | +18  |
|          | 476        | 0.5      | 2.1528501109e+07  | -27  | -12  |
|          | 0          | 0.0      | 2.1528501109e+07  | -99  | -12  |
| de080285 | 418        | 0.0      | 1.4495817688e+01  | -09  | -02  |
|          | 132        | 0.1      | 1.3924732864e+01  | -35  | -32  |
|          | 0          | 0.0      | 1.3924732864e+01  | -99  | -32  |
| gen1     | 369,502    | 205.3    | -1.6903658594e-08 | -06  | -12  |
|          | 246,428    | 9,331.3  | 1.2935699163e-06  | -12  | -31  |
|          | 2,394      | 81.6     | 1.2953925804e-06  | -45  | -30  |
| gen2     | 44,073     | 60.0     | 3.2927907828e+00  | -04  | -11  |
|          | 1,599      | 359.9    | 3.2927907840e+00  | -99  | -29  |
|          | 0          | 10.4     | 3.2927907840e+00  | -99  | -32  |
| gen4     | 45,369     | 212.4    | 1.5793970394e-07  | -06  | -10  |
|          | 53,849     | 14,812.5 | 2.8932268196e-06  | -12  | -30  |
|          | 37         | 10.4     | 2.8933064888e-06  | -54  | -30  |
| 130      | 1,229,326  | 876.7    | 9.5266141574e-01  | -10  | -09  |
|          | 275,287    | 7,507.1  | -7.5190273434e-26 | -25  | -32  |
|          | 0          | 0.2      | -4.2586876849e-24 | -24  | -33  |
| iprob    | 1,087      | 0.2      | 2.6891551285e+03  | +02  | -11  |
|          | 0          | 0.0      | 2.6891551285e+03  | +02  | -31  |
|          | 0          | 0.0      | 2.6891551285e+03  | +02  | -28  |

(continued)

**Table 3** (continued)

| Model    | Iterations | Times    | Final objective   | Pinf       | Dinf       |
|----------|------------|----------|-------------------|------------|------------|
| TMA_ME   | 12,225     | 37.1     | 8.0051076669e-07  | -06        | -05        |
|          | 685        | 61.5     | 8.7036315385e-07  | -24        | -30        |
|          | 0          | 6.7      | 8.7036315385e-07  | <b>-99</b> | <b>-31</b> |
| GlcAerWT | 62,856     | 9,707.3  | -2.4489880182e+04 | +04        | -05        |
|          | 5,580      | 3,995.6  | -7.0382449681e+05 | -07        | -26        |
|          | 4          | 60.1     | -7.0382449681e+05 | <b>-19</b> | <b>-21</b> |
| GlcAlift | 134,693    | 14,552.8 | -5.1613878666e+05 | -03        | -01        |
|          | 3,258      | 1,067.1  | -7.0434008750e+05 | -09        | -26        |
|          | 2          | 48.1     | -7.0434008750e+05 | <b>-20</b> | <b>-22</b> |

Pinf and Dinf = final maximum primal and dual infeasibilities ( $\log_{10}$  values tabulated). Problem iprob is infeasible. Bold figures show Pinf and Dinf at the end of Step 3. Pinf =  $10^{-99}$  means Pinf = 0. Note that  $\text{Pinf}/\|x^*\|_\infty$  and  $\text{Dinf}/\|y^*\|_\infty$  are all  $O(10^{-30})$  or smaller, even though only  $O(10^{-15})$  was requested. This is an unexpectedly favorable empirical finding

### *The Mészáros Problems*

The *problematic* LPs were provided as MPS files by Ed Klotz [17]. The first two problems have unusually large entries in the constraint matrix  $A$ . The Step 1 Double MINOS solution has at best one digit of precision in the objective value for de063155, and is quite erroneous for de063157. Nevertheless, the Steps 2 and 3 Quad solutions are seen to be highly accurate when the solution norms are taken into account.

The gen\* problems come from image reconstruction, with no large entries in  $A$ ,  $x$ ,  $y$ , but highly degenerate primal solutions  $x$ . (In both Steps 1 and 2 for gen1, 60% of the iterations made no improvement to the objective, and the final solution has 30% of the basic variables on their lower bound.) For gen1, warm-starting Quad MINOS from the Step 1 basis gave an almost feasible initial solution (266 basic variables outside their bounds by more than  $10^{-15}$  with a sum of infeasibilities of only  $O(10^{-8})$ ), yet nearly 250,000 iterations were needed in Step 2 to reach optimality. These examples show that Quad precision does not remove the need for a more rigorous anti-degeneracy procedure (such as Wolfe's method as advocated by Fletcher [6]), and/or steepest-edge pricing [7], to reduce significantly the total number of iterations.

Problem I30 behaved similarly (80% degenerate iterations in Steps 1 and 2). The tiny objective value is essentially zero, so we can't expect the Steps 2 and 3 objectives to agree in their leading digits.

Problem iprob is an artificial one that was intended to be feasible with a very ill-conditioned optimal basis, but the MPS file provided to us contained low-precision data (many entries like 0.604 or 0.0422). Our Double and Quad runs agree that the problem is infeasible. This is an example of Quad removing some doubt that was inevitable with just Double.



### *The Systems Biology ME Problems*

Like the gen\* problems, the ME models showed 40–60% degenerate iterations in Step 1, but fortunately not so many total iterations in Step 2. This is important for FVA and for ME with nonlinear constraints, where there are many warm starts.

**Problem TMA\_ME** developed by Lerman [19] has some large matrix entries  $|S_{ij}|$  and many small solution values  $v_j$  that are meaningful to systems biologists. The ME part of the model also contains small matrix entries. In Step 1, almost all iterations went on finding a feasible solution, and the objective then had the correct order of magnitude.

This was the first ME model that we used for Quad experiments (in April 2012). The data  $S, c, \ell, u$  in (3) came as a Matlab structure with  $c_j = 0, l_j = 0, u_j = 1,000$  for most  $j$ , except  $c_{17533} = 1$  (meaning maximize flux  $v_{17533}$ ), four variables had smaller positive upper bounds, the last variable had moderate positive bounds, and 64 variables were fixed at zero. We output the data to a plain text file. Most entries of  $S$  are integers (represented exactly), but about 5,000  $S_{ij}$  values are of the form  $8.037943687315e-01$  or  $3.488862338191e-06$  with 13 significant digits. The text data was read into Double and Quad versions of a prototype Fortran 90 implementation of SQOPT [11].

For the present paper, we used the same Matlab data to generate an MPS file for input into MINOS. Since this is limited to six significant digits, the values in the preceding paragraph were rounded to  $8.03794e-01$  and  $3.48886e-06$  and in total about 5,000  $S_{ij}$  values had  $O(10^{-6})$  relative perturbations of this kind. We have been concerned that such data perturbations could alter the FBA solution greatly because the final basis matrices could have condition number as large as  $10^6$  or even  $10^{12}$  (as estimated by LUSOL). In comparing Quad SQOPT on Matlab data with Quad MINOS on MPS data, we fortunately observe that the final objective values for TMA\_ME agree to five digits and match the results from SoPlex [34] and the exact simplex solver QSopt\_ex [32], as reported by Lerman [19]:

|              | Optimal objective |             |
|--------------|-------------------|-------------|
| SoPlex 80bit | 8.703671403e-07   | Matlab data |
| QSopt_ex     | 8.703646169e-07   | Matlab data |
| Quad SQOPT   | 8.703646169e-07   | Matlab data |
| Quad MINOS   | 8.703631539e-07   | MPS data    |

More importantly, for the most part *even small solution values* are perturbed in only the 5th or 6th significant digit. Let  $v$  and  $w$  be the solutions obtained by the two Quad solvers on slightly different data. Some example solution values follow:

| $j$              | 107          | 201          | 302          |             |
|------------------|--------------|--------------|--------------|-------------|
| Quad SQOPT $v_j$ | 2.336815e-06 | 8.703646e-07 | 1.454536e-11 | Matlab data |
| Quad MINOS $w_j$ | 2.336823e-06 | 8.703632e-07 | 1.454540e-11 | MPS data    |

Among all  $j$  for which  $\max(v_j, w_j) > 10^{-15}$  (the feasibility tolerance), the largest relative difference  $|v_j - w_j|/\max(v_j, w_j)$  was less than  $10^{-5}$  for all but 31 variables. For 22 of these pairs, either  $v_j$  or  $w_j$  was primal or dual degenerate (meaning one of them was zero and there are alternative solutions with the same objective value). The remaining nine variables had these values:

| $j$   | $v_j$      | $w_j$      | Relative difference |
|-------|------------|------------|---------------------|
| 16383 | 6.0731e-07 | 2.0374e-06 | 0.70                |
| 16459 | 1.7090e-06 | 2.1778e-06 | 0.22                |
| 16483 | 2.4675e-06 | 5.9936e-07 | 0.76                |
| 16730 | 1.4432e-06 | 7.8685e-07 | 0.46                |
| 17461 | 1.7090e-06 | 2.1778e-06 | 0.22                |
| 17462 | 2.4675e-06 | 5.9936e-07 | 0.76                |
| 17478 | 6.0731e-07 | 2.0374e-06 | 0.70                |
| 17507 | 1.4432e-06 | 7.8685e-07 | 0.46                |
| 17517 | 8.7036e-07 | 2.9740e-06 | 0.71                |

We see that the  $v_j, w_j$  values are quite small (the same magnitude as the data perturbation), and for each of the nine pairs there is about one digit of agreement. In general we could expect thousands of small solution pairs to differ this much, yet for almost *all* 17,535 pairs, there are at least five digits of agreement.

These observations about two forms of problem TMA\_ME are welcome empirical evidence of the robustness of this particular multiscale model. Quad solvers can help evaluate the robustness of future (increasingly large) models of metabolic networks by enabling similar comparison of high-accuracy solutions for slightly different problems.

**Problem GlcAerWT** is an ME model from the detailed study by Thiele et al. [36]. Difficulties with solving TMA\_ME and GlcAerWT led to the lifting technique of Sun et al. [35] (and to problem GlcAlift).

After 33,000 iterations on GlcAerWT, Double MINOS began to report singularities every 50–100 iterations following updates to the basis LU factors. After another 30,000 iterations, MINOS terminated Step 1 with maximum infeasibility  $O(10^4)$ . Step 2 required significant work to achieve a reasonably accurate solution. Step 3 quickly confirmed the final objective value with high accuracy considering the  $O(10^7)$  primal and dual solutions norms.

**Problem GlcAlift** is a reformulated version of GlcAerWT in which some large matrix entries  $c_{\max}$  in (5) have been reduced via the lifting technique [35]. In Step 1, Double MINOS again reported frequent singularities and required twice as many iterations as GlcAerWT, but a near-optimal solution was found (allowing for the primal and dual solution norms of  $O(10^7)$ ), and Steps 2–3 were more efficient and accurate. The objective function for both GlcA models is to maximize variable  $v_{60069}$ . The fact that the Step 1 objective values have no correct digits illustrates

the challenge these models present and emphasizes the benefits that Quad precision offers. Theoretically the optimal objectives for GlcAerWt and GlcAlift should agree. We assume that the limited data precision in the MPS files is responsible for only three-digit agreement. Fortunately the Tomlab interface used by Thiele et al. [36] allows full double-precision data [37]. We can do the same for MINOS, as we did for SQOPT.

## 5 Discussion

While today’s advanced LP solvers such as CPLEX, Gurobi, Mosek, and Xpress [3, 5, 13, 23] are effective on a wide range of large and challenging linear optimization models, the study by Thiele et al. [36] emphasizes the need for improved reliability in solving FBA and ME models in systems biology. Fortunately, reformulation [35] and careful use of the commercial solvers CPLEX and Gurobi permitted successful analysis in Thiele et al. [36] of the GlcAerWT and GlcAlift models discussed here, and we encourage this approach for Step 1 of our proposed three-step procedure (Section 3). The bulk of the work in solving multiscale LP problems can still be performed there by conventional Double solvers (possibly including Barrier solvers with simplex “cross-over” to provide a basis).

Our aim has been to demonstrate that the Step 2 and 3 warm starts with Quad solvers will be acceptably efficient, and that the accuracy achieved exceeds requirements by a very safe margin. The “humane” approach of Kahan [16]—use of Quad LP solvers—is certainly more efficient than applying exact simplex solvers, even though the latter have proved their value in several applications [1, 2, 18, 20].

An intriguing question remains concerning the bold figures in Table 3. The primal and dual solutions obtained with Quad precision are *substantially more accurate than the  $10^{-15}$  requested*. The same has been true for all of the classic set of Netlib problems [28] that we have run. Kahan [16] explains that “perturbations get amplified by singularities near the data.” He describes a “pejorative surface” of data points where singularity exists, and expects loss of accuracy as data approaches the surface. The volume surrounding the pejorative surface is the danger zone, but: *“Arithmetic precision is usually extravagant enough if it is somewhat more than twice as [great] as the data’s and the desired result’s. Often that shrunken volume contains no data.”* We can surmise that Kahan has anticipated our observed situation, wherein LP problems defined with double-precision data appear unlikely to be too ill-conditioned for a Quad solver.

We believe that quadruple-precision solutions are now practical for multiscale LP applications such as FBA and FVA models in systems biology [12, 30, 31, 36], and that they justify increased confidence as systems biologists build ever-larger models to explore new hypotheses about metabolism and macromolecular synthesis. Our three-step procedure of Section 3 allows combined use of Double and Quad solvers and should lead to solutions of exceptional accuracy in other areas of computational science involving multiscale optimization problems. For example, Dattorro [4] has

derived an approach to analog filter design that requires a Quad LP or nonlinear solver to deal with a wide range of frequencies (which must be raised to the fourth power). We look forward to implementing this approach, as well as treating the nonlinear constraints (7) directly to take advantage of the nonlinear algorithms in Quad MINOS.

## Funding

This work was supported by the National Institute of General Medical Sciences of the National Institutes of Health [award U01GM102098]. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

**Acknowledgements** We thank Ronan Fleming, Philip Gill, Ed Klotz, Joshua Lerman, Edward O’Brien, Yuekai Sun, Ines Thiele, and Elizabeth Wong for much help during the course of this work and for valuable comments on the manuscript. Joshua Lerman at UC San Diego provided the model named TMA\_ME here (originally model\_final\_build\_unscaled.mat) and advised us of the final objective values obtained by SoPlex and QSopt\_exact. Ines Thiele at the University of Luxembourg provided the GlcAerWT model and insight into the ME coupling constraints. Yuekai Sun at Stanford University created the reformulated version named GlcAlift here. Ed Klotz of IBM in Carson City NV provided MPS files for the Mészáros *problematic* LPs and lengthy discussions of their properties. A referee also provided valuable feedback.

## References

1. Applegate, D.L., Cook, W., Dash, S., Espinoza, D.G.: Exact solutions to linear programming problems. *Oper. Res. Lett.* **35**, 693–699 (2007)
2. Cook, W., Dash, S., Espinoza, D.G.: Exact solutions to linear programming problems. *Oper. Res. Lett.* **35**, 693–699 (2007)
3. CPLEX: IBM ILOG CPLEX optimizer. <http://www.ibm.com/software/commerce/optimization/cplex-optimizer/> (2014)
4. Dattorro, J.: Private Communication. Stanford University, Stanford (2014)
5. FICO Xpress Optimization Suite: <http://www.fico.com/en/products/fico-xpress-optimization-suite/> (2015)
6. Fletcher, R.: On Wolfe’s method for resolving degeneracy in linearly constrained optimization. *SIAM J. Optim.* **24**(3), 1122–1137 (2014)
7. Forrest, J.J., Goldfarb, D.: Steepest-edge simplex algorithms for linear programming. *Math. Program.* **57**, 341–374 (1992)
8. GCC libquadmath: The GCC Quad-precision math library application programming interface (API). <http://gcc.gnu.org/onlinedocs/libquadmath/> (2014)
9. Gill, P.E., Murray, W., Saunders, M.A., Wright, M.H.: Maintaining LU factors of a general sparse matrix. *Linear Algebra Appl.* **88/89**, 239–270 (1987)
10. Gill, P.E., Murray, W., Saunders, M.A., Wright, M.H.: A practical anti-cycling procedure for linear and nonlinear programming. *Math. Program.* **45**, 437–474 (1989)
11. Gill, P.E., Murray, W., Saunders, M.A.: SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Rev.* **47**(1), 99–131 (2005). SIGEST article

12. Gudmundsson, S., Thiele, I.: Computationally efficient flux variability analysis. *BMC Bioinf.* **11**(489), 3 pp. (2010)
13. Gurobi: Gurobi optimization system for linear and integer programming. <http://www.gurobi.com> (2014)
14. Hall, J.A.J., McKinnon, K.I.M.: The simplest examples where the simplex method cycles and conditions where EXPAND fails to prevent cycling. *Math. Program. Ser. B* **100**, 133–150 (2004)
15. IEEE standard for floating-point arithmetic: IEEE Std 754-2008. IEEE Computer Society (2008)
16. Kahan, W.: Desperately needed remedies for the undebuggability of large floating-point computations in science and engineering. In: *IFIP/SIAM/NIST Working Conference on Uncertainty Quantification in Scientific Computing*, Boulder (2011). <http://www.eecs.berkeley.edu/~wkahan/Boulder.pdf>
17. Klotz, E.: Private Communication. IBM, Carson City (2014)
18. Koch, T.: The final NETLIB-LP results. *Oper. Res. Lett.* **32**, 138–142 (2004)
19. Lerman, J.A.: Private Communication. University of California, San Diego (2012)
20. Lerman, J.A., Hyduke, D.R., Latif, H., Portnoy, V.A., Lewis, N.E., Orth, J.D., Schrimper-Rutledge, A.C., Smith, R.D., Adkins, J.N., Zengler, K., Palsson, B.Ø.: In silico method for modelling metabolism and gene product expression at genome scale. *Nat. Commun.* **3**(929), 10 pp. (2012)
21. LUSOL: Sparse LU factorization package. <http://web.stanford.edu/group/SOL/software/lusol> (2013)
22. Mészáros, C.: A collection of challenging LP problems. [http://www.sztaki.hu/~meszaros/public\\_ftp/lptestset/problematic](http://www.sztaki.hu/~meszaros/public_ftp/lptestset/problematic) (2004)
23. MOSEK: MOSEK Optimization Software. <http://www.mosek.com/> (2014)
24. MPS: Input format for LP data. <http://lpsolve.sourceforge.net/5.5/mps-format.htm> (1960)
25. MPS files: Data files in MPS format for models in Table 2. <http://web.stanford.edu/group/SOL/multiscale/models.html> (2014)
26. Murtagh, B.A., Saunders, M.A.: Large-scale linearly constrained optimization. *Math. Program.* **14**, 41–72 (1978)
27. Murtagh, B.A., Saunders, M.A.: A projected Lagrangian algorithm and its implementation for sparse nonlinear constraints. *Math. Program. Stud.* **16**, 84–117 (1982)
28. Netlib: Netlib collection of LP problems in MPS format. <http://www.netlib.org/lp/data> (1988)
29. O'Brien, E.J., Lerman, J.A., Chang, R.L., Hyduke, D.R., Palsson, B.Ø.: Genome-scale models of metabolism and gene expression extend and refine growth phenotype prediction. *Mol. Syst. Biol.* **9**(693), 13 pp. (2013)
30. Orth, J.D., Thiele, I., Palsson, B.Ø.: What is flux balance analysis? *Nat. Biotechnol.* **28**(3), 245–248 (2010)
31. Palsson, B.Ø.: *Systems Biology: Properties of Reconstructed Networks*. Cambridge University Press, New York (2006)
32. QSopt\_ex: QSopt\_ex: a simplex solver for computing exact rational solutions to LP problems. <http://www.math.uwaterloo.ca/~bico/qsopt/index.html> (2008)
33. Schellenberger, J., Que, R., Fleming, R.M.T., Thiele, I., Orth, J.D., Feist, A.M., Zielinski, D.C., Bordbar, A., Lewis, N.E., Rahmanian, S., et al.: Quantitative prediction of cellular metabolism with constraint-based models: the COBRA Toolbox v2.0. *Nat. Protoc.* **6**(9), 1290–1307 (2011)
34. SoPlex: Soplex: The sequential object-oriented simplex solver. <http://soplex.zib.de> (1996)
35. Sun, Y., Fleming, R.M.T., Thiele, I., Saunders, M.A.: Robust flux balance analysis of multiscale biochemical reaction networks. *BMC Bioinf.* **14**, 240 (2013)
36. Thiele, I., Fleming, R.M.T., Que, R., Bordbar, A., Diep, D., Palsson, B.Ø.: Multiscale modeling of metabolism and macromolecular synthesis in *E. coli* and its application to the evolution of codon usage. *PLoS One* **7**(9), 18 pp. (2012)
37. TOMLAB: Optimization environment for MATLAB. <http://tomopt.com> (2014)