A Parallel Implementation of the Revised Simplex Algorithm Using OpenMP: Some Preliminary Results

Ploskas Nikolaos[1], Samaras Nikolaos[1], and Margaritis Konstantinos[1]

Department of Applied Informatics, University of Macedonia, 156 Egnatia Str., 54006 Thessaloniki, Greece { ploskas@uom.gr, samaras@uom.gr, kmarg@uom.gr}

Summary. Linear Programming (LP) is a significant research area in the field of operations research. The simplex algorithm is the most widely used method for solving Linear Programming problems (LPs). The aim of this paper is to present a parallel implementation of the revised simplex algorithm. Our parallel implementation focuses on the reduction of the time taken to perform the basis inverse, due to the fact that the total computational effort of an iteration of simplex type algorithms is dominated by this computation. This inverse does not have to be computed from scratch at any iteration. In this paper, we compute the basis inverse with two well-known updating schemes: (i) The Product Form of the Inverse (PFI) and (ii) A Modification of the Product Form of the Inverse (MPFI); and incorporate them with revised simplex algorithm. Apart from the parallel implementation, this paper presents a computational study that shows the speedup among the serial and the parallel implementations in large-scale LPs. Computational results with a set of benchmark problems from Netlib, including some infeasible ones, are also presented. The parallelism is achieved using OpenMP in a shared memory multiprocessor architecture.

Key words: Linear Programming, Revised Simplex Method, Basis Inverse, Parallel Computing, OpenMP.

1 Introduction

Linear Programming (LP) is the process of minimizing or maximizing a linear objective function $z = \sum_{i=1}^{n} c_i \cdot x_i$ subject to a number of linear equality and inequality constraints. Several methods are available for solving LPs, among which the simplex algorithm is the most widely used. We assume that the problem is in its general form. Formulating the linear problem, we can describe it as shown in (LP1).

$$\begin{aligned} &min && c^T x \\ &subject\ to && Ax = b \\ & && x \geq 0 \end{aligned} \qquad (LP1)$$

where $A \in R^{m \times n}$, $(c, x) \in R^n$, $b \in R^m$, and T denotes transposition. Without loss of generality we assume that A has full rank, $rank(A) = m$, where $(m < n)$. The simplex method searches for an optimal solution by moving from one feasible solution to another, along the edges of the feasible region. The dual problem associated with the linear problem in (LP1) is shown in (DP).

$$\begin{aligned} &min && b^T w \\ &subject\ to && A^T w + s = c \\ & && s \geq 0 \end{aligned} \qquad (DP)$$

where $w \in R^m$ and $s \in R^n$. As in the solution of any large scale mathematical system, the computational time for large LPs is a major concern. Parallel programming is a good practice for solving computationally intensive problems in operations research. The application of parallel processing for LP has been introduced in the early 1970s. However, only since the beginning of the 1980s attempts have been made to develop parallel implementations. A lot of architectural features have been used in practice. Preliminary parallel approaches were developed for network optimization, direct search methods and global optimization. A growing number of optimization problems demand parallel computing capabilities. Any performance improvement in the parallelization of the revised simplex would be of great interest.

One of the earliest parallel tableau simplex methods on a small-scale distributed memory Multiple-Instruction Multiple-Data (MIMD) machines is the one introduced by [Fin87]. Stunkel [Stu88] implemented both the tableau and the revised simplex method on a 16-processor Intel hypercube computer, achieving a speedup of between 8 and 12 for small problems from the Netlib set [Gay85]. Helgason, Kennington and Zaki [HKZ88] proposed an algorithm to implement the revised simplex using sparse matrices methods on shared memory MIMD computer. Furthermore, Shu and Wu [SW93] and Shu [Shu95] parallelized the explicit inverse and the LU decomposition of the basis simplex algorithms. Hall and McKinnon [HM96] [HM98] have implemented two parallel schemes for the revised simplex method. The first of Hall and McKinnon's parallel revised simplex implementations was ASYNPLEX [HM96]. In this implementation one processor is devoted to the basis inversion and the remaining processors perform simplex iterations. ASYNPLEX was implemented on a Cray T3D, achieving a speedup of between 2.5 and 4.8 for four modest Netlib problems. The second of Hall and McKinnon's parallel revised simplex implementations was PARSMI [HM98]. PARSMI was tested on modest problems from the Netlib set, resulting in a speedup of between 1.7 and

1.9. Hall [Hal05] implemented a variant of PARSMI on a 8-processor shared memory Sun Fire E15k, leading in a speedup of between 1.8 and 3.

Simplex algorithms for general LPs on Single Instruction Multiple Data (SIMD) have been reported by [ABK89]. Luo and Reijns [LR92] presented an implementation of the revised simplex method, achieving a speedup of more than 12, when solving modest Netlib problems on 16 transputers. Eckstein et al [EBPG95] implemented a parallelization of standard and revised simplex method in a CM2 machine. Lentini et al [LRTG95] worked on the standard simplex method with the tableau stored as a sparse matrix, resulting in a speedup of between 0.5 and 2.7, when solving medium sized Netlib problems on four transputers. Thomadakis and Liu [TL96] worked on the standard simplex method on MasPar MP-1 and MP-2 machines, achieving a speedup of up to three, when solving large randomly-generated problems. Badr et al [BMPSS06] implemented a dense standard simplex method on eight computers, leading in a speedup of five when solving small random dense LPs. Previous attempts to develop simplex implementations with the aim of exploiting high performance computing architectures are reviewed by [Hal10]. Finally, computational results for parallelizing the network simplex method are reported in [CEFM88] [BH94] [Pet90].

The use of GPUs for general purpose computations is a quite recent topic, which was applied to linear programming. Greeff [Gre04] implemented the revised simplex method on a GPU using OpenGL and Cg and was able to achieve a speedup of up to 11.4 over an identical CPU implementation. Jung and O'Leary [JO08] and Owens et al [OHLGS08] also presented an implementation using Cg and OpenGL. Spampinato and Elster [SE93] proposed a GPU implementation of the revised simplex method, based on the CUDA architecture and achieved a speedup of up to 2.5. Recently, Bieling, Peschlow and Martini [BPM10] also presented an implementation of the revised simplex algorithm and achieved a speedup of up to 10.

This paper presents a parallelization of the revised simplex algorithm on a shared memory multiprocessor architecture. The focus of this parallelization is on the basis inverse. The structure of the paper is as follows. In Section 2, the revised simplex algorithm is described and presented. In Section 3, two methods that have been widely used for basis inversion are analyzed. Section 4 presents the parallel revised simplex algorithm and Sect. 5 gives the computational results. Finally, the conclusions of this paper are outlined in Sect. 6.

2 Revised Simplex Algorithm

The linear problem in (LP1) can be written as shown in (LP2).

$$\begin{aligned} min \quad & c_B^T x_B + c_N^T x_N \\ subject\,to \quad & A_B x_B + A_N x_N = b \\ & x_B, x_N \geq 0 \end{aligned} \qquad (LP2)$$

In (LP2), A_B is a $m \times m$ non-singular sub-matrix of A, called basic matrix or basis. The columns of A which belong to subset B are called basic and those which belong to N are called non basic. The solution $x_B = (A_B)^{-1}b, x_N = 0$ is called a basic solution. A solution $x = (x_B, x_N)$ is feasible iff $x > 0$. Otherwise, (LP2) is infeasible. The solution of (DP) is computed by the relation $s = c - A^T w$, where $w = (c_B)^T (A_B)^{-1}$ are the simplex multipliers and s are the dual slack variables. The basis A_B is dual feasible iff $s \geq 0$.

In each iteration, simplex algorithm interchanges a column of matrix A_B with a column of matrix A_N and constructs a new basis $A_{\overline{B}}$. Any iteration of simplex type algorithms is relatively expensive. The total work of an iteration of simplex type algorithms is dominated by the determination of the basis inverse. This inverse however, does not have to be computed from scratch during each iteration. Simplex type algorithms maintain a factorization of basis and update this factorization in each iteration. There are several schemes for updating basis inverse. Two well-known schemes are (i) the Product Form of the Inverse (PFI) and (ii) a Modification of the Product Form of the Inverse, developed by Benhamadou [Ben02]. These methods, in order to compute the new basis, use only information about the entering and leaving variables along with the current basis. A formal description of the revised simplex algorithm is given in Table 1.

3 Methods Used for Basis Inversion

The revised simplex algorithm differs from the original method. The former uses the same recursion relations to transform only the inverse of the basis in each iteration. It has been implemented to reduce the computation time of the basis inversion and is particularly effective for sparse linear problems. In this section, we will review two methods that have been widely used for basis inversion: (i) the Product Form of the Inverse and (ii) a Modification of the Product Form of the Inverse.

Table 1. Revised Simplex Algorithm

Step 0. *(Initialization).*
Start with a feasible partition (A_B, A_N). Compute $(A_B)^{-1}$ and vectors x_B, w and s_N.
Step 1. *(Test of optimality).*
if $s_N \geq 0$ then STOP. The linear problem is optimal.
else
Choose the index l of the entering variable using a pivoting rule.
Variable x_l enters the basis.
Step 2. *(Minimum ratio test).*
Compute the pivot column $h_l = (A_B)^{-1} A_l$.
if $h_l \leq 0$ then STOP. The linear problem is unbounded.
else
Choose the leaving variable $x_{B[r]} = x_k$ using the following relation:
$x_{B[r]} = \frac{x_{B[r]}}{h_{il}} = min\left\{\frac{x_{B[i]}}{h_{il}} : h_{il} < 0\right\}$
Step 3. *(Pivoting).*
Swap indices k and l. Update the new basis inverse $(A_{\overline{B}})^{-1}$, using PFI or MPFI.
Go to Step 1.

3.1 Product Form of the Inverse

The PFI scheme, in order to compute the new basis, uses information only about the entering and leaving variables along with the current basis. The new basis inverse can be updated at any iteration using the (1).

$$(A_{\overline{B}})^{-1} = (A_B E)^{-1} = E^{-1}(A_B)^{-1} \tag{1}$$

where E^{-1} is the inverse of the eta-matrix and can be computed by (2).

$$E^{-1} = I - \frac{1}{h_{rl}}(h_l - e_l)\, e_l^T = \begin{bmatrix} 1 & & -h_{1l} & & \\ & \ddots & \vdots & & \\ & & 1/h_{rl} & & \\ & & \vdots & \ddots & \\ & & -h_{ml}/h_{rl} & & 1 \end{bmatrix} \tag{2}$$

If the current basis inverse is computed using regular multiplication, then the complexity of the PFI is $\Theta\left(m^3\right)$.

3.2 A Modification of Product Form of the Inverse

MPFI updating scheme has been presented by Benhamadou [Ben02]. The key idea is that the current basis inverse $(A_{\overline{B}})^{-1}$ can be computed from the previous inverse $(A_B)^{-1}$ using a simple outer product of two vectors and one matrix addition, as shown in (3).

$$(A_{\overline{B}})^{-1} = \left(A_{\overline{B}_{r.}}\right)^{-1} + v \otimes (A_{B_{r.}})^{-1} \tag{3}$$

The updating scheme of the inverse is shown in (4).

$$(A_{\overline{B}})^{-1} = \begin{vmatrix} b_{11} & \cdots & b_{1m} \\ \vdots & \ddots & \vdots \\ 0 & 0 & 0 \\ \vdots & \ddots & \vdots \\ b_{m1} & \cdots & b_{mm} \end{vmatrix} + \begin{vmatrix} -\frac{h_{1l}}{h_{rl}} \\ \vdots \\ -\frac{1}{h_{rl}} \\ \vdots \\ -\frac{h_{ml}}{h_{rl}} \end{vmatrix} (A_B)^{-1} : \begin{vmatrix} b_{r1} & \cdots & b_{rr} & \cdots & b_{rm} \end{vmatrix} \tag{4}$$

The outer product requires m^2 multiplications and the addition of two matrices requires m^2 additions. The total cost of the above method is $2m^2$ operations (multiplications and additions). Hence, the complexity is $\Theta(m^2)$.

4 Parallel Revised Simplex Algorithm

The parallelization of all the individual steps of the revised simplex algorithm is limited and very hard to achieve. However, it is also essential for any algorithm to perform basis inverse in parallel with simplex iterations, otherwise basis inverse will become the dominant step and limit the possible speedup. Our parallel implementation focuses on the reduction of the time taken to perform the basis inverse. The basis inversion is done with the Product Form of the Inverse and a Modification of the Product Form of the Inverse, as described in the previous section.

Both methods take as input the previous basis inverse $(A_B)^{-1}$, the pivot column (h_l), the index of the leaving variable (k) and the number of the constraints (m).

The most time-consuming step of PFI scheme is the matrix multiplication of (1). Our parallel algorithm uses the block matrix multiplication algorithm for this step. This algorithm suggests a recursive divide-and-conquer solution, as described in [Hak93] [HZ83]. This method has significant potential for parallel implementations, especially on shared memory implementations.

Let us assume that we have p processors. Table 2 shows the steps that we used to compute the new basis inverse $(A_{\overline{B}})^{-1}$ with the PFI scheme. Table 3 shows the steps that we used to compute the new basis inverse $(A_{\overline{B}})^{-1}$ with the MPFI scheme.

Table 2. Parallel PFI

Step 0. Compute the column vector: $v = \left[-\frac{h_{1l}}{h_{rl}} \cdots \frac{1}{h_{rl}} \cdots -\frac{h_{ml}}{h_{rl}} \right]^T$ Each processor computes in parallel m/p elements of v. **Step 1.** Replace the r^{th} column of an identity matrix with the column vector v. Each processor assigns in parallel m/p elements to the identity matrix. This matrix is the inverse of the Eta-matrix. **Step 2.** Compute the new basis inverse using (1) with block matrix multiplication. Each processor will compute m/p rows of the new basis.

Table 3. Parallel MPFI

Step 0. Compute the column vector: $v = \left[-\frac{h_{1l}}{h_{rl}} \cdots \frac{1}{h_{rl}} \cdots -\frac{h_{ml}}{h_{rl}} \right]^T$ Each processor computes in parallel m/p elements of v. **Step 1.** (The following steps are computed in parallel) **Step 1.1**. Compute the outer product $v \otimes (A_{B_r})^{-1}$ with block matrix multiplication. **Step 1.2**. Copy matrix $(A_B)^{-1}$ to matrix $(A_{\overline{B}})^{-1}$. Set the r^{th} row of $(A_{\overline{B}})^{-1}$ equal to zero. Each processor computes in parallel m/p rows of $(A_{\overline{B}})^{-1}$. **Step 2.** Compute the new basis inverse using relation (3). Each processor computes in parallel m/p rows of the new basis.

5 Computational Experiments

In this section we report the computational results of running our implementations on a set of LPs available through Netlib. The three most usual approaches to analyzing algorithms are i) worst-case analysis, ii) average-case analysis and iii) experimental analysis. Computational studies have been proven useful tools in order to examine the practical efficiency of an algorithm or even compare algorithms by using the same problem sets. The computational comparison has been performed on a quad-processor Intel Xeon 3.2 GHz with 2 Gbyte of main memory running under Ubuntu 10.10 64-bit and performed on GCC 4.5.2. In the following computational results all reported CPU times were measured in seconds. The algorithms have been implemented using C++ and OpenMP. In all LPs from the Netlib collection, the parallel versions of the simplex algorithm converge to the same solution.

5.1 Problem Instances

The test set used in our experiments were the Netlib set of LPs. The Netlib library is a well known suite containing many real world LPs. Ordóñez and Freund [OF03] have shown that 71% of the Netlib LPs are ill-conditioned.
Below there are some useful information about the data set, which was used in the computational study. The first column of Table 4 includes the name of the problem, the second the number of constraints, the third the number of variables, the fourth the non-zero elements of matrix A and the fifth the density of the coefficient matrix A. Let $nnz(A)$ denote the number of non-zeros in the matrix A. The density of matrix A is defined as the ratio of the $nnz(A)$ to the total number of its elements.
All LPs have been presolved. The purpose of the presolve analysis is to improve linear problem's numerical properties and computational characteristics. The last row of each table shows the average value of each column.

5.2 Computational Results

The algorithms described in Sect. 4 have been experimentally implemented. Table 5 presents the results from the execution of the serial and parallel implementations of the above mentioned updating schemes. For each implementation, the table shows the CPU time for the basis inverse and the total CPU time.

Table 4. Statistics of the Benchmarks

Problem	**Constraints**	**Variables**	**Non-Zeros A**	**Sparcity A**
agg	488	163	2410	3.03%
agg2	516	302	4284	2.75%
agg3	516	302	4300	2.76%
bandm	305	472	2494	1.73%
brandy	220	249	2148	3.92%
e226	223	282	2578	4.10%
fffff800	524	854	6227	1.39%
israel	174	142	2269	9.18%
lotfi	153	308	1078	2.29%
sc105	105	103	280	2.59%
sc205	205	203	551	1.32%
scfxm1	330	457	2589	1.72%
scfxm2	660	914	5183	0.86%
scfxm3	990	1371	7777	0.57%
scrs8	490	1169	3182	0.56%
share1b	117	225	1151	4.37%
share2b	96	79	694	9.15%
ship04l	402	2118	6332	0.74%
ship04s	402	1458	4352	0.74%
ship08l	778	4283	12802	0.38%
ship08s	778	2387	7114	0.38%
ship12l	1151	5427	16170	0.26%
ship12s	1151	2763	8178	0.26%
stocfor1	117	111	447	3.44%
klein2	477	54	4585	17.80%
klein3	994	88	12107	13.84%
Total	12362	26284	121282	
Average	474.96	1044.84	4754.88	

Table 5. Basis Inverse and Total Time of the Serial and Parallel Implementations

Problem	Serial implementations PFI Time of basis inverse	Serial PFI Total time	Serial MPFI Time of basis inverse	Serial MPFI Total time	Parallel implementations PFI Time of basis inverse	Parallel PFI Total time	Parallel MPFI Time of basis inverse	Parallel MPFI Total time
agg	2.83	4.58	2.28	4.05	1.52	3.32	1.13	2.95
agg2	3.54	5.65	2.78	4.97	1.81	3.96	1.52	3.69
agg3	3.28	5.61	2.62	5.03	1.85	4.05	1.48	3.78
bandm	1.01	1.62	0.74	1.41	0.84	1.52	0.61	1.29
brandy	1.30	2.76	1.06	2.55	1.04	2.48	0.76	2.22
e226	1.66	3.34	1.38	3.09	1.22	2.82	0.85	2.50
fffff800	6.10	12.77	4.86	11.64	5.18	11.58	4.31	10.89
israel	0.82	1.73	0.63	1.65	0.53	1.48	0.45	1.31
lotfi	0.33	0.85	0.29	0.80	0.25	0.78	0.21	0.68
sc105	0.08	0.09	0.03	0.07	0.01	0.07	0.02	0.06
sc205	0.55	0.91	0.51	0.85	0.40	0.74	0.20	0.56
scfxm1	3.72	7.11	2.96	6.38	3.17	6.31	2.49	5.80
scfxm2	30.76	62.34	24.26	56.40	26.34	58.56	21.22	52.54
scfxm3	109.06	244.48	83.97	219.22	91.67	224.76	72.93	209.23
scrs8	11.17	22.20	8.69	19.81	9.56	20.20	7.34	17.90
share1b	0.15	0.29	0.09	0.26	0.11	0.25	0.08	0.23
share2b	0.05	0.11	0.04	0.10	0 .01	0.10	0.03	0.09
ship04l	5.20	14.53	4.18	13.65	4.43	13.65	3.56	12.90
ship04s	1.76	4.55	1.52	4.31	1.54	4.22	1.21	4.10
ship08l	33.78	94.70	26.30	86.45	30.02	91.90	22.10	82.50
ship08s	7.43	19.07	5.99	17.49	6.50	18.09	5.01	16.95
ship12l	120.21	335.98	89.94	305.95	100.05	317.80	75.64	292.00
ship12s	17.20	43.52	13.45	39.16	12.99	42.84	10.77	36.60
stocfor1	0.04	0.06	0.03	0.05	0.02	0.06	0.02	0.04
klein2	17.36	33.01	13.53	28.80	9.85	25.30	7.55	22.74
klein3	192.50	413.33	148.69	368.50	106.92	326.70	75.30	295.01
Total	571.89	1335.19	440.82	1202.64	417.83	1183.54	316.79	1078.56
Average	22.00	51.35	16.95	46.26	16.07	45.52	12.18	41.48

In order to show more clearly the superiority of parallel implementations over the serial ones, we provide the Table 6. Table 6 presents the speedup obtained by the parallel implementations regarding the CPU time for the basis inverse and the total CPU time, for both PFI and MPFI schemes. We now plot the ratios taken from Table 6 in Fig. 1. The total time is in logarithmic scale.

Table 6. Basis Inverse and Total Time of the Serial and Parallel Implementations

Problem	Speedup			
	PFI		MPFI	
	Basis inverse	Total	Basis inverse	Total
agg	1.86	1.38	2.02	1.37
agg2	1.96	1.43	1.83	1.35
agg3	1.77	1.39	1.77	1.33
bandm	1.20	1.07	1.21	1.09
brandy	1.25	1.11	1.39	1.15
e226	1.36	1.18	1.62	1.24
fffff800	1.18	1.10	1.13	1.07
israel	1.55	1.17	1.40	1.26
lotfi	1.32	1.09	1.38	1.18
sc105	8.00	1.29	1.50	1.17
sc205	1.38	1.23	2.55	1.52
scfxm1	1.17	1.13	1.19	1.10
scfxm2	1.17	1.06	1.14	1.07
scfxm3	1.19	1.09	1.15	1.05
scrs8	1.17	1.10	1.18	1.11
share1b	1.36	1.16	1.13	1.13
share2b	5.00	1.10	1.33	1.11
ship04l	1.17	1.06	1.17	1.06
ship04s	1.14	1.08	1.26	1.05
ship08l	1.13	1.03	1.19	1.05
ship08s	1.14	1.05	1.20	1.03
ship12l	1.20	1.06	1.19	1.05
ship12s	1.32	1.02	1.25	1.07
stocfor1	2.00	1.00	1.50	1.25
klein2	1.76	1.30	1.79	1.27
klein3	1.80	1.27	1.97	1.25
Total	46.55	29.95	37.44	30.38
Average	1.79	1.15	1.44	1.17

From the above results, we observe: (i) the MPFI scheme is in most problems faster than PFI both in serial and in parallel implementation, (ii) using PFI scheme, the speedup gained from the parallelization is of average 1.79 for the

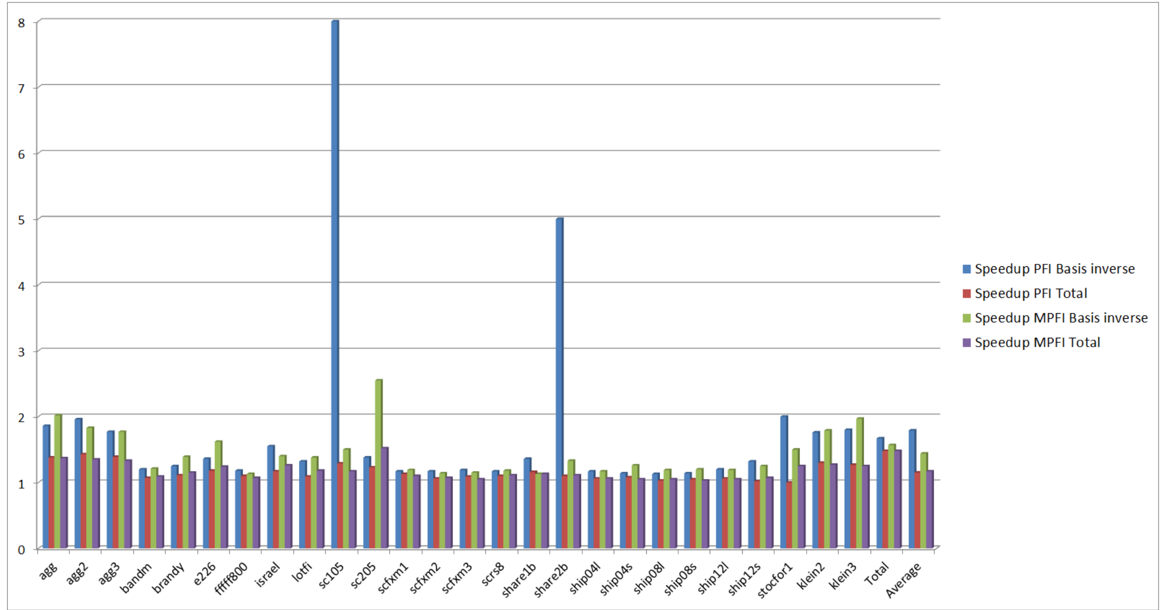

Fig. 1. Basis Inverse and Total Time of the Serial and Parallel Implementations

time of basis inverse and 1.15 for total time, and (iii) using MPFI scheme, the speedup is of average 1.44 for the time of basis inverse and 1.17 for total time.

6 Conclusions

A parallel implementation for the revised simplex algorithm has been described in this paper. Some preliminary computational results on Netlib problems have reported a speedup of average 1.79 and 1.44 regarding the basis inverse procedure, using PFI and MPFI updating schemes respectively. These results could be further improved by performance optimization. In future work, we plan to implement our parallel algorithm combining the Message Passing Interface (MPI) and OpenMP programming models to exploit parallelism beyond a single level. Furthermore, we intend to port our algorithm to a GPU implementation based on the CUDA architecture.

References

[ABK89] Agrawal, A., Blelloch, G.E., Krawitz, R.L., Phillips, C.A.: Four vector-matrix primitives. Proceedings ACM Symposium on Parallel Algorithms and Architectures, 292–302 (1989)

[BMPSS06] Badr, E.S., Moussa, M., Papparrizos, K., Samaras, N., Sifaleras, A.: Some computational results on MPI parallel implementations of dense simplex method. Transactions on Engineering, Computing and Technology, **17**, 228–231 (2006)

[BH94] Barr, R.S., Hickman, B.L.: Parallel Simplex for Large Pure Network Problems: Computational Testing and Sources of Speedup. Operations Research, **42(1)**, 65–80 (1994)

[Ben02] Benhamadou, M.: On the simplex algorithm 'revised form'. Advances in Engineering Software, **33**, 769–777 (2002)

[BPM10] Bieling, J., Peschlow, P., Martini, P.: An efficient GPU implementation of the revised simplex method. Proceedings of IPDPS Workshops, 1–8 (2010)

[CEFM88] Chang, M.D., Engquist, M., Finkel, R., Meyer, R.R.: A Parallel Algorithm for Generalized Networks. Annals of Operations Research, **14(1-4)**, 125–145 (1988)

[EBPG95] Eckstein, J., Boduroglu, I., Polymenakos, L., Goldfarb, D.: Data-Parallel Implementations of Dense Simplex Methods on the Connection Machine CM-2. ORSA Journal on Computing, **7(4)**, 402–416 (1995)

[Fin87] Finkel, R.A.: Large-Grain Parallelism: Three Case Studies. In: Jamieson, H. (ed) Proceedings of Characteristics of Parallel Algorithms. The MIT Press (1987)

[Gay85] Gay, D.M.: Electronic mail distribution of linear programming test problems. Mathematical Programming Society COAL Newsletter, **13**, 10–12 (1985)

[Gre04] Greeff, G.: The revised simplex method on a GPU. Stellenbosch University, South Africa, Honours Year Project (2004)

[HM96] Hall, J.A.J., McKinnon, K.I.M.: PARSMI, a parallel revised simplex algorithm incorporating minor iterations and Devex pricing. In: Wasniewski, J., Dongarra, J., Madsen, K., Olesen, D. (ed) Applied Parallel Computing. volume 1184 of Lecture Notes in Computer Science, Springer (1996)

[HM98] Hall, J.A.J., McKinnon, K.I.M.: ASYNPLEX an asynchronous parallel revised simplex algorithm. Annals of Operations Research, **81(0)**, 27–50 (1998)

[Hal05] Hall, J.A.J.: SYNPLEX: a task-parallel scheme for the revised simplex method. In Contributed talk at the second international workshop on combinatorial scientific computing (2005)

[Hal10] Hall, J.A.J.: Towards a practical parallelisation of the simplex method. Computational Management Science, **7**, 139–170 (2010)

[Hak93] Hake, J.F.: Parallel Algorithms for Matrix Operations and Their Performance in Multiprocessor Systems. In: Kronsjo, L., Shumsheruddin, D. (ed) Advances in Parallel Algorithms. Halsted Press, New York (1993)

[HKZ88] Helgason, R.V., Kennington, J.L., Zaki, H.A.: A parallelization of the simplex method. Annals of Operations Research, **14(1-4)**, 17–40 (1988)

[HZ83] Horowitz, E., Zorat, A.: Divide-and-Conquer for Parallel Processing. IEEE Trans. Comput., **C-32(6)**, 582–585 (1983)

[JO08] Jung, J.H., O'Leary, D.P.: Implementing an interior point method for linear programs on a CPU-GPU system. Electronic Transaction on Numerical Analysis, **28**, 174–189 (2008)

[LRTG95] Lentini, M., Reinoza, A., Teruel, A., Guillen, A.: SIMPAR: a parallel sparse simplex. Computational and Applied Mathematics, **14(1)**, 49–58 (1995)

[LR92] Luo, J., Reijns, G.L.: Linear programming on transputers. In: van Leeuwen, J. (ed) Algorithms, software, architecture. IFIP transactions A (computer science and technology). Elsevier, Amsterdam, (1992)

[OHLGS08] Owens, J.D., Houston, M., Luebke, D., Green, S., Stone, J.E.: GPU Computing. Proceedings of the IEEE, **96(5)**, 879–899 (2008)

[OF03] Ordóñez, F., Freund, R.: Computational experience and the explanatory value of condition measures for linear optimization. SIAM J. on Optimization, **14(2)**, 307–333 (2003)

[Pet90] Peters, J.: The Network Simplex Method on a Multiprocessor. Networks, **20(7)**, 845–859 (1990)

[Shu95] Shu, W.: Parallel implementation of a sparse simplex algorithm on MIMD distributed memory computers. Journal of Parallel and Distributed Computing, **31(1)**, 25–40 (1995)

[SW93] Shu, W., Wu, M.: Sparse Implementation of Revised Simplex Algorithms on Parallel Computers. Proceedings of Sixth SIAM Conference on Parallel Processing for Scientific Computing, Norfolk (1993)

[SE93] Spampinato, D.G., Elster, A.C.: Linear optimization on modern GPUs. Proceedings of the 2009 IEEE International Symposium on Parallel and Distributed Processing (2009)

[Stu88] Stunkel, C.B.: Linear optimization via message-based parallel processing. Proceedings of International Conference on Parallel Processing, **3**, 264–271 (1988)

[TL96] Thomadakis, M.E., Liu, J.C.: An Efficient Steepest-Edge Simplex Algorithm for SIMD Computers. Proceedings of the International conference on Super-Computing (1996)