

# *A Computational Comparison of Basis Updating Schemes for the Simplex Algorithm on a CPU-GPU System*

Nikolaos Ploskas\* and Nikolaos Samaras†

*Department of Applied Informatics, School of Information Sciences, University of Macedonia, Greece*

August 18, 2013

## **Abstract**

Linear Programming is a significant area in the field of operations research. The computation of the basis inverse is the most time-consuming step in simplex type algorithms and must be carefully designed and implemented. This inverse does not have to be computed from scratch at any iteration, but updating schemes can be applied to accelerate this calculation. In this paper, we perform a computational comparison in which the basis inverse is computed with five different updating schemes. Then, we propose a parallel implementation of two updating schemes on a CPU-GPU System. All these schemes have been implemented under the MATLAB and CUDA environment. Finally, a computational study on randomly generated full dense linear programs is presented to establish the practical value of GPU-based implementation.

Keywords: Simplex Algorithm; Basis Inverse; Graphics Processing Unit; MATLAB; Compute Unified Device Architecture

## **1 Introduction**

Linear Programming (LP) is the process of minimizing or maximizing a linear objective function  $z = \sum_{j=1}^n c_j x_j$  to a number of linear equality and inequality constraints. Simplex algorithm is the most widely used method for solving Linear Programming problems (LPs). Consider the following linear programming problem in the standard form:

\*Email: ploskas@uom.gr; Corresponding author

†Email: samaras@uom.gr

$$\begin{aligned}
& \min && c^T x \\
& \text{subject to} && Ax = b \\
& && x \geq 0
\end{aligned} \tag{LP.1}$$

where  $A \in \mathbb{R}^{m \times n}$ ,  $(c, x) \in \mathbb{R}^n$ ,  $b \in \mathbb{R}^m$ , and  $T$  denotes transposition. We assume that  $A$  has full rank,  $\text{rank}(A) = m, m < n$ . Consequently, the linear system  $Ax = b$  is consistent. The simplex algorithm searches for an optimal solution by moving from one feasible solution to another, along the edges of the feasible set. The dual problem associated with the (LP.1) is:

$$\begin{aligned}
& \min && b^T w \\
& \text{subject to} && A^T w + s = c \\
& && s \geq 0
\end{aligned} \tag{DP.1}$$

where  $w \in \mathbb{R}^m$  and  $s \in \mathbb{R}^n$ .

The past twenty years have been a time of remarkable developments in optimization solvers. Real life LPs tend to be large in size. A growing number of problems demand parallel computing capabilities. The explosion in computational power (CPUs and GPUs) has made it possible to solve large and difficult LPs in a short amount of time. As in the solution of any large scale mathematical system, the computational time for large LP problems is a major concern. The basis inverse dictates the total computational effort of an iteration of simplex type algorithms. This inverse does not have to be computed from scratch at any iteration, but can be updated through a number of updating schemes. All efficient versions of the simplex algorithm work with some factorization of the basis matrix  $B$  or its inverse  $B^{-1}$ .

Dantzig and Orchard-Hays [1] have proposed the Product Form of the Inverse (PFI), which maintains the basis inverse using a set of eta vectors. Benhamadou [2] proposed a Modification of the Product Form of the Inverse (MPFI). The key idea is that the current basis inverse  $(A_{\bar{B}})^{-1}$  can be computed from the previous inverse  $(A_B)^{-1}$  using a simple outer product of two vectors and one matrix addition.

LU decomposition produces generally sparser factorizations than PFI [3]. The LU factorization for the basis inverse has been proposed by Markowitz [4]. Markowitz used LU decomposition to fully invert a matrix, but used the PFI scheme to update the basis inverse during simplex iterations. Bartels and Golub [5] have later proposed a scheme to update a sparse factorization, which was more stable than using PFI. Their computational experiments, however, proved that it was more computationally expensive. Forrest and Tomlin [6] created a variant of the Bartels-Golub method by sacrificing some stability characteristics causing the algorithm to have a smaller growth rate in the number of non-zero elements relative to the PFI scheme. Reid [7] proposed two variants of the Bartels-Golub updating scheme that aim to balance the sparsity and the numerical stability of the factorization. A variant of the Forrest-Tomlin update was proposed by Suhl and Suhl [8]. Other important updating techniques can be found in Saunders [9] and Goldfarb [10]. A full description of most of these updating methods can be found in Nazareth [11] and Chvatal [12].

There have been many reviews and variants of these methods individually, but only a few comparisons between them that are either obsolete or don't compare all these updating schemes. McCoy and Tomlin [13] report the results of some experiments on measuring the accuracy of the PFI scheme, Bartels-Golub method and Forrest-Tomlin scheme. Lim et al. [14] provide a comparative study between Bartels-Golub method, Forrest-Tomlin method and Reid method. Badr et al. [15] perform a computational evaluation of PFI and MPFI updating schemes. Ploskas et al. [16] compare PFI and MPFI updating schemes both on their serial and their parallel implementation.

Originally, GPUs used to accelerate graphics rendering. GPUs have gained recently a lot of popularity and High Performance Computing applications have already started to use them. The computational capabilities of GPUs exceeds the one of CPUs. GPU is utilized for data parallel and computationally intensive portions of an algorithm. NVIDIA introduced Compute Unified Device Architecture (CUDA) in late 2006. CUDA enables users to execute codes on their GPUs and it is based on a SIMT programming model. Any performance improvements in the parallelization of the revised simplex algorithm would be of great interest. Using GPU computing for solving large-scale LPs is a great challenge due to the capabilities of GPU architectures.

Some related works have been proposed on the GPU parallelization for LPs. O'Leary and Jung [17] proposed a CPU-GPU implementation of the Interior Point Method for dense LPs. Their computational results on Netlib Set [18] showed that some speedup can be gained for large dense problems. Spampinato and Elster [19] presented a GPU-based implementation of the Revised Simplex Algorithm on GPU with NVIDIA CUBLAS and NVIDIA LAPACK libraries [20]. Their implementation showed a maximum speedup of 2.5 on randomly generated LPs of at most 2000 variables and 2000 constraints. Bieling et al. [21] also proposed a parallel implementation of the Revised Simplex Algorithm on GPU. They compared their GPU-based implementation with GLPK solver and found a maximum speedup of 18 in single precision. Lalami et al. [22] proposed a parallel implementation of the standard Simplex on a CPU-GPU systems. Their computational results on randomly generated dense problems of at most 4000 variables and 4000 constraints showed a maximum speedup of 12.5. Meyer et al. [23] proposed a mono and a multi-GPU implementation of the standard Simplex algorithm and compared their implementation with the CLP solver. Their implementation outperformed CLP solver on large sparse LPs. Li et al. [24] presented a GPU-based parallel algorithm, based on Gaussian elimination, for large scale LPs that outperforms the CPU-based algorithm.

This paper presents a computational study in which the basis inverse is computed with five different updating schemes: (i) Gaussian Elimination, (ii) the built-in function `inv` of MATLAB, (iii) LU decomposition, (iv) Product Form of the Inverse and (v) a Modification of the Product Form of the Inverse; and incorporate them with the revised simplex algorithm. Then, we propose a parallel implementation of PFI and MPFI schemes, which were the fastest among the five updating methods, on a CPU-GPU System, which is based on MATLAB and CUDA.

The structure of the paper is as follows. In Section 2, a brief description of the revised simplex algorithm is presented. In Section 3, five methods that have been widely used for basis inversion are presented and analyzed. Section 4 presents the computational comparison of the updating schemes. Computational tests were carried out on randomly generated LPs of at most 5000 variables and 5000 constraints. Finally, the conclusions of this paper are outlined in section 5.

## 2 Revised Simplex Method

Using a partition  $(B, N)$  (LP.1) can be written as follows:

$$\begin{aligned} \min \quad & c_B^T x_B + c_N^T x_N \\ \text{subject to} \quad & A_B x_B + A_N x_N = b \\ & x_B, x_N \geq 0 \end{aligned} \quad (LP.2)$$

In the above problem,  $A_B$  is an  $m \times m$  non-singular submatrix of  $A$ , called basic matrix or basis. The columns of  $A$  belonging to subset  $B$  are called basic and those belonging to  $N$  are called non basic. The solution of the linear problem  $x_B = (A_B)^{-1}b, x_N = 0$  is called a basic solution. A solution  $x = (x_B, x_N)$  is feasible iff  $x \geq 0$ ; otherwise, the solution is infeasible. The solution of the (DP.1) 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  $A_{\bar{B}}$ . Any iteration of simplex type algorithms is relatively expensive. The total work of an iteration of simplex type algorithms is dictated by the computation of the basis inverse. This inverse, however, does not have to be computed from scratch during each iteration of the simplex algorithm. Simplex type algorithms maintain a factorization of basis and update this factorization in each iteration. There are several schemes for updating basis inverse. In Section 3 we present eight well-known methods for the basis inverse. A formal description of the revised simplex algorithm [25] is presented in Table 1.

## 3 Basis Inversion Updating Schemes

### 3.1 Gaussian Elimination

Gaussian elimination is a method for solving systems of linear equations, which can be used to compute the inverse of a matrix. Gaussian elimination performs the following two steps: (i) Forward Elimination: reduces the given matrix to a triangular or echelon form and (ii) Back Substitution: finds the solution of the given system. Gaussian elimination with partial pivoting requires  $O(n^3)$  time complexity.

Gaussian elimination has been implemented on MATLAB using the *mldivide* operator. In order to find the new basis inverse using Gaussian elimination, one can use the following equation:

$$(A_B)^{-1} = A_B \setminus I \quad (1)$$

### 3.2 Built-in Function Inv of MATLAB

The basis inverse can be computed using the built-in function of MATLAB called *inv*, which uses LAPACK routines to compute the basis inverse. Due to the fact that this function is already compiled and optimized for MATLAB, its execution time is smaller compared with the other relevant methods that compute the explicit basis inverse; time-complexity, though, remains  $O(n^3)$ .

Table 1: Revised Simplex Algorithm

<p><b>Step 0.</b> (<i>Initialization</i>).  Start with a feasible partition <math>(B, N)</math>. Compute <math>(A_B)^{-1}</math> and vectors <math>x_B, w</math> and <math>s_N</math>.</p> <p><b>Step 1.</b> (<i>Test of optimality</i>).  if <math>s_N \geq 0</math>  then STOP. (LP.1) is optimal.  else  Choose the index <math>l</math> of the entering variable using a pivoting rule.  Variable <math>x_l</math> enters the basis.</p> <p><b>Step 2.</b> (<i>Minimum ratio test</i>).  Compute the pivot column <math>h_l = (A_B)^{-1}A_l</math>.  if <math>h_l \leq 0</math> then  STOP. (LP.1) is unbounded.  else  Choose the leaving variable <math>x_k = x_{B[r]}</math> using the following relation:  <math display="block">x_{B[r]} = \frac{x_{B[r]}}{h_{il}} = \min \left\{ \frac{x_{B[i]}}{h_{il}} : h_{il} &lt; 0 \right\}</math></p> <p><b>Step 3.</b> (<i>Pivoting</i>).  Swap indices <math>k</math> and <math>l</math>. Update the new basis inverse <math>(A_{\overline{B}})^{-1}</math>, using an updating scheme.  Go to Step 1.</p>
--

### 3.3 LU Decomposition

LU decomposition method factorizes a matrix as the product of an upper ( $U$ ) and a lower ( $L$ ) triangular factors, which can be used to compute the inverse of a matrix. In order to compute the  $U$  and  $L$  factors, the built-in function of MATLAB called `lu` has been used. LU decomposition can be computed in time  $O(n^3)$ .

### 3.4 MPFI

MPFI updating scheme has been presented by Benhamadou [2]. The main idea of this method 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 the following equation:

$$(\overline{A_B})^{-1} = (\overline{A_B})_r^{-1} + v \otimes (A_B)_r^{-1} \quad (2)$$

The updating scheme of the inverse is shown in equation (3).

$$(A_{\overline{B}})^{-1} = \begin{vmatrix} b_{11} & \cdots & b_{1m} \\ \vdots & \ddots & \vdots \\ 0 & 0 & 0 \\ \vdots & & \vdots \\ b_{m1} & \cdots & b_{mm} \end{vmatrix} + (A_B)^{-1} : \begin{vmatrix} b_{r1} & \cdots & b_{rr} & \cdots & b_{rm} \\ -\frac{h_{1l}}{h_{rl}} \\ \vdots \\ -\frac{1}{h_{rl}} \\ \vdots \\ -\frac{h_{ml}}{h_{rl}} \end{vmatrix} \quad (3)$$

The outer product of equation (3) requires  $m^2$  multiplications and the addition of two matrices requires  $m^2$  additions. Hence, the complexity is  $\Theta(m^2)$ .

### 3.5 PFI

The PFI scheme, in order to update the new basis  $(A_{\overline{B}})^{-1}$ , uses information only about the entering and leaving variables along with the current basis  $(A_B)^{-1}$ . The new basis inverse can be updated at any iteration using the (4).

$$(\overline{A_B})^{-1} = (A_B E)^{-1} = E^{-1} (A_B)^{-1} \quad (4)$$

where  $E^{-1}$  is the inverse of the eta-matrix and can be computed by the following equation:

$$E^{-1} = I - \frac{1}{h_{rl}} (h_l - e_l) e_l^T = \begin{bmatrix} 1 & & -h_{1l} & & \\ & \ddots & \vdots & & \\ & & 1/h_{rl} & & \\ & & \vdots & \ddots & \\ & & -h_{ml}/h_{rl} & & 1 \end{bmatrix} \quad (5)$$

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

## 4 Computational Results of Serial Implementations

Computational studies have been widely used, in order to examine the practical efficiency of an algorithm or even compare algorithms. The computational comparison of the aforementioned five updating schemes has been performed on a quad-processor Intel Core i7 3.4 GHz with 32 Gbyte of main memory running under Microsoft Windows 7 64-bit. The algorithms have been implemented using MATLAB Professional R2012a. MATLAB (MATrix LABoratory) is a powerful programming environment and is especially designed for matrix computations in general. All times in the following tables are measured in seconds.

The test set used in the computational study was randomly generated. Problem instances have the same number of constraints and variables. The largest problem tested has 5,000 constraints and

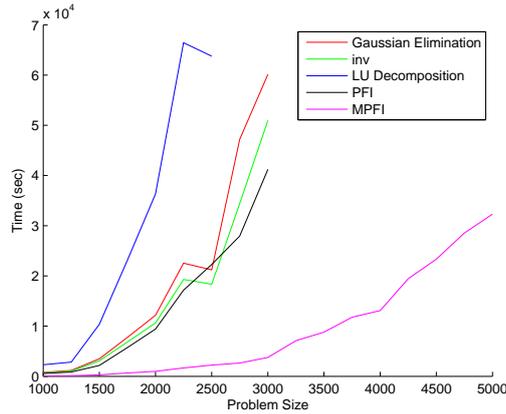


Figure 1: Basis Inverse Time Comparison

5,000 variables. All instances are dense. For each instance, we averaged times over 10 runs. A time limit of 20 hours was set that explains why there are no measurements for some updating methods on large instances. It should be noted that in MATLAB R2012a multithreading is enabled by default thus our implementations are automatically parallelized and executed using the available multicore CPU.

Table 3 presents the results from the execution of the above mentioned updating schemes. We have also included the execution time from MATLAB's linprog built-in function, a function for solving linear programming problems. MATLAB's linprog function includes two algorithms for large-scale and medium-scale optimization. The large-scale algorithm is based on Interior Point Solver [26], a primal-dual interior-point algorithm. LIPSOL used a Cholesky-infinity factorization that causes overhead during the factorization of dense matrices and as a result it cannot solve problems with more than 1,500 variables and constraints. Due to this restriction, we have used in our comparison the medium-scale algorithm, which is a variation of the simplex method. Table 2 includes the execution time for the basis inverse of each updating scheme, while Table 3 presents the total execution time. The execution time of the basis inverse and the whole algorithm for each updating scheme is also graphically illustrated in Figures 1 and 2, respectively.

The MPFI updating scheme has the best performance. On the other hand, LU updating method has the worst performance. Another significant issue is the performance of Gaussian elimination, PFI, function inv and linprog of MATLAB which are close to each other and the results are not quite satisfactory.

Table 2: Basis Inverse Time (secs)

Problem size	Gaussian	LU			MPFI
	Elimination	inv	Decomposition	PFI	
1,000	793.30	702.85	2,314.62	555.65	83.36
1,250	1,173.10	1,036.74	2,820.17	845.49	121.18
1,500	3,487.13	3,075.34	10,351.52	2,160.71	292.43
1,750	7,746.69	6,843.46	23,153.03	5,743.58	680.61
2,000	12,157.33	10,622.34	36,338.81	9,411.65	981.08
2,250	22,547.04	19,288.01	66,437.50	17,156.07	1,667.49
2,500	21,209.41	18,336.32	63,775.88	22,266.25	2,228.82
2,750	47,205.29	34,422.09	-	27,952.71	2,648.20
3,000	60,134.08	50,964.37	-	41,204.74	3,762.74
3,250	-	-	-	-	7,100.16
3,500	-	-	-	-	8,793.45
3,750	-	-	-	-	11,785.46
4,000	-	-	-	-	13,087.66
4,250	-	-	-	-	19,432.13
4,500	-	-	-	-	23,344.92
4,750	-	-	-	-	28,534.90
5,000	-	-	-	-	32,317.85

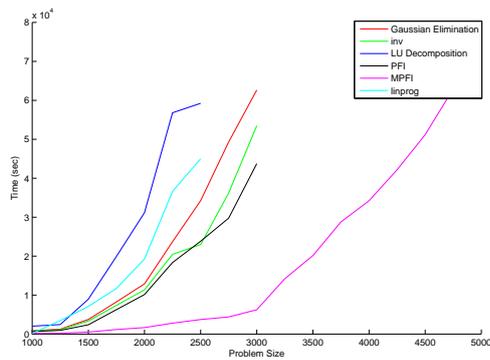


Figure 2: Total Time Comparison

Table 3: Total Time (secs)

Problem size	Gaussian Elimination	LU			MPFI	linprog
		inv	Decomposition	PFI		
1,000	871.81	774.15	2,025.42	631.22	155.51	171.68
1,250	1,269.92	1,134.50	2,446.59	940.03	216.97	3,453.09
1,500	3,746.93	3,330.55	8,946.89	2,384.46	509.71	7,097.17
1,750	8,245.07	7,347.78	19,905.57	6,255.89	1161.56	11,682.54
2,000	12,875.85	11,340.48	31,160.31	10,140.09	1670.73	19,267.61
2,250	23,709.13	20,453.34	56,804.33	18,345.68	2798.21	36,614.81
2,500	34,249.22	22,975.68	59,233.99	23,828.22	3730.49	44,998.08
2,750	49,312.02	36,202.77	-	29,762.26	4,384.83	-
3,000	62,646.43	53,472.82	-	43,740.40	6,242.79	-
3,250	-	-	-	-	14,200.80	-
3,500	-	-	-	-	20,147.04	-
3,750	-	-	-	-	28,776.84	-
4,000	-	-	-	-	34,235.18	-
4,250	-	-	-	-	42,196.15	-
4,500	-	-	-	-	51,210.67	-
4,750	-	-	-	-	62,919.07	-
5,000	-	-	-	-	70,058.59	-

## 5 Parallel Implementation of PFI and MPFI Updating Schemes on a CPU-GPU system

PFI and MPFI were the fastest updating schemes. In this section, we present the GPU-based implementations of these updating methods taking advantage of the power that modern GPUs offer. The parallel implementations of these updating methods are implemented on MATLAB and CUDA. The updating methods are built using both native MATLAB code and CUDA MEX files. 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 ( $m$ ).

### 5.1 GPU-based PFI

Let us assume that we have  $t$  gpu cores. Table 4 shows the steps that we used to compute the new basis inverse  $(A_{\bar{B}})^{-1}$  with the PFI scheme on the GPU.

Table 4: GPU-based PFI

<p><b>Step 0.</b>          Compute the column vector:</p> $v = \left[ -\frac{h_{1l}}{h_{rl}} \quad \dots \quad \frac{1}{h_{rl}} \quad \dots \quad -\frac{h_{ml}}{h_{rl}} \right]^T$ <p>Each core computes in parallel <math>m/t</math> elements of <math>v</math>. The pivot element is shared between all <math>t</math> cores.</p> <p><b>Step 1.</b>          Replace the <math>r^{th}</math> column of an identity matrix with the column vector <math>v</math>. Each core assigns in parallel <math>m/t</math> elements to the identity matrix. This matrix is the inverse of the Eta-matrix.</p> <p><b>Step 2.</b>          Perform a matrix multiplication according to (4). Each core will compute a block of the new basis.</p>
---

### 5.2 GPU-based MPFI

Table 5 shows the steps that we used to compute the new basis inverse  $(A_{\bar{B}})^{-1}$  with the MPFI scheme on the GPU.

Table 5: GPU-based MPFI

**Step 0.**

Compute the column vector:

$$v = \left[ -\frac{h_{1l}}{h_{rl}} \quad \cdots \quad \frac{1}{h_{rl}} \quad \cdots \quad -\frac{h_{ml}}{h_{rl}} \right]^T$$

Each core computes in parallel  $m/t$  elements of  $v$ . The pivot element is shared between all  $t$  cores.

**Step 1.**

Compute the outer product  $v \otimes (A_{B_r})^{-1}$  with matrix multiplication. Each core will compute a block of the new matrix.

**Step 2.**

Set the  $r^{th}$  row of  $(A_{\overline{B}})^{-1}$  equal to zero. Each core computes in parallel  $t/p$  rows of  $(A_{\overline{B}})^{-1}$ .

**Step 3.**

Add matrix  $(A_{\overline{B}})^{-1}$  with the resulted matrix from step 1. Each core will compute a block of the new basis inverse.

## 6 Computational Results of Parallel Implementations

The same randomly generated test set is also used in order to test the performance of the GPU-based implementations. The computational comparison of the parallel implementations has been also performed on a quad-processor Intel Core i7 3.4 GHz with 32 Gbyte of main memory running under Microsoft Windows 7 64-bit and NVIDIA Quadro 6000 with 6 Gbyte of memory and 448 CUDA cores. The mex files have been implemented using CUDA 4.2 and Microsoft Visual Studio 2012. Table 6 presents the results from the execution of the GPU-based implementations of PFI and MPFI updating schemes. For each implementation, the table shows the CPU time for the basis inverse and the total time.

Table 7 presents the speedup obtained by the GPU-based implementations regarding the CPU time for the basis inverse and the total time. We now plot the ratios taken from Table 7 in Fig. 3. The total time is in logarithmic scale.

From the above results, we observe: (i) the MPFI scheme is much faster than PFI both in serial and in GPU-based implementation, (ii) using PFI scheme, the speedup gained from the GPU implementation is around 7 for the time of basis inverse and 5.5 for total time when the problem size reaches to 3,000x3,000, and (iii) using MPFI scheme, the speedup gained from the GPU implementation is around 19 for the time of basis inverse and 5 for total time when the problem size reaches to 5,000x5,000.

Table 6: Basis Inverse and Total Time of the GPU-based Implementations (secs)

<b>Problem</b>	<b>PFI</b>		<b>MPFI</b>	
	<b>Time of basis inverse</b>	<b>Total time</b>	<b>Time of basis inverse</b>	<b>Total time</b>
1,000	149.56	208.57	41.56	101.42
1,250	198.98	275.93	37.50	113.80
1,500	455.32	625.03	65.55	239.37
1,750	1,122.77	1,509.01	123.63	510.37
2,000	1,611.04	2,166.60	147.14	704.66
2,250	2,868.02	3,765.82	230.96	1,147.32
2,500	3,838.33	5,043.59	293.38	1,512.32
2,750	4,480.71	5,867.20	309.58	1,712.80
3,000	5,846.40	7,820.77	405.98	2,372.87
3,250	-	-	702.99	5,108.20
3,500	-	-	775.44	6,806.43
3,750	-	-	938.33	8,964.75
4,000	-	-	1,005.20	9,923.24
4,250	-	-	1,305.05	11,163.00
4,500	-	-	1,409.72	12,738.97
4,750	-	-	1,659.01	14,139.12
5,000	-	-	1,709.94	14,067.99

Table 7: Speedup Obtained by the GPU-based Implementations

<b>Problem</b>	<b>PFI</b>		<b>MPFI</b>	
	<b>Basis inverse</b>	<b>Total</b>	<b>Basis inverse</b>	<b>Total</b>
1,000	3.72	3.03	2.01	1.53
1,250	4.25	3.41	3.23	1.91
1,500	4.75	3.81	4.46	2.13
1,750	5.12	4.15	5.51	2.28
2,000	5.84	4.68	6.67	2.37
2,250	5.98	4.87	7.22	2.44
2,500	5.80	4.72	7.60	2.47
2,750	6.24	5.07	8.55	2.56
3,000	7.05	5.59	9.27	2.63
3,250	-	-	10.10	2.78
3,500	-	-	11.34	2.96
3,750	-	-	12.56	3.21
4,000	-	-	13.02	3.45
4,250	-	-	14.89	3.78
4,500	-	-	16.56	4.02
4,750	-	-	17.20	4.45
5,000	-	-	18.90	4.98

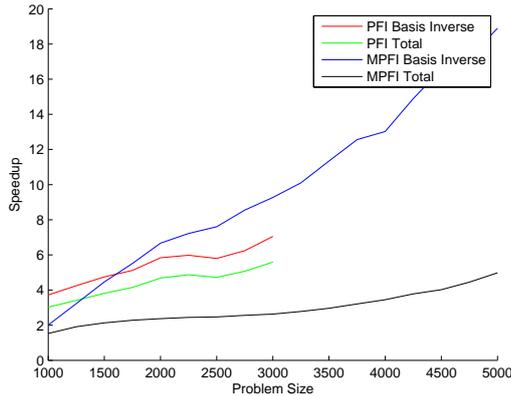


Figure 3: Speedup Comparison

## 7 Conclusions

The basis inverse is the most time-consuming step in simplex type algorithms, so it is essential to implement a fast and numerically stable updating method. In this paper, we performed a computational comparison of five updating schemes and incorporate them to the revised simplex algorithm. The results of the computational study showed that MPFI updating scheme is the fastest when solving large dense LPs.

We proposed a GPU-based implementation for PFI and MPFI updating schemes, which were the fastest serial implementations, and implemented them on MATLAB and CUDA. We performed again a computational study and find that GPU-based implementations of PFI and MPFI outperform the serial ones. More specifically, the speedup for PFI method is up to 7 for the time of basis inverse and 5.5 for the total time and the speedup for MPFI method is up to 19 for the time of basis inverse and 5 for total time. Our approach allow us to solve problems of size 10,000x10,000.

In future work, we plan to implement all the steps of the algorithm in order to fully parallelize the revised simplex method for GPUs. Moreover, we plan also to test sparse LPs and also port our application to a multi-GPU architecture.

## 8 Acknowledgements

The authors thank NVIDIA for support through their Academic Partnership Program.

## References

- [1] G. B. Dantzig and W. Orchard-Hays, "The product form of the inverse in the simplex method", *Math. Comp.*, Vol. 8, 1954, pp. 64–67.

- [2] M. Benhamadou, "On the simplex algorithm 'revised form'", *Advances in Engineering Software*, Vol. 33, 2002, pp. 769–777.
- [3] R. K. Brayton, F. G. Gustavson and R. A. Willoughby, "Some results on sparse matrices", *Mathematics of Computation*, Vol. 24, 1970, pp. 937–954.
- [4] H. Markowitz, "The elimination form of the inverse and its applications to linear programming", *Management Science*, Vol. 3, 1957, pp. 255–269.
- [5] R. H. Bartels and G. H. Golub, "The simplex method of linear programming using LU decomposition", *Communications of the ACM*, Vol. 12, 1969, pp. 266–268.
- [6] J. J. H. Forrest and J. A. Tomlin, "Updated triangular factors of the basis to maintain sparsity in the product form simplex method", *Mathematical Programming*, Vol. 2, 1972, pp. 263–278.
- [7] J. Reid, "A sparsity-exploiting variant of the Bartels-Golub decomposition for linear programming bases", *Mathematical Programming*, Vol. 24, 1982, pp. 55–69.
- [8] L. M. Suhl and U. H. Suhl, "A fast LU update for linear programming", *Annals of Operations Research*, Vol. 43, No. 1, 1993, pp. 33–47.
- [9] M. Saunders, "A fast and stable implementation of the simplex method using Bartels-Golub updating", In: J. Bunch and S.T. Rachev, Eds., *Sparse Matrix Computation*, Academic Press, New York, 1976, pp. 213–226.
- [10] D. Goldfarb, "On the Bartels-Golub decomposition for linear programming bases", *Mathematical Programming*, Vol. 13, 1977, pp. 272–279.
- [11] J. L. Nazareth, "Computer Solution of Linear Programs", Oxford University Press, Oxford, UK, 1987.
- [12] V. Chvatal, "Linear Programming", W. H. Freeman and Company, New York, USA, 1983.
- [13] P. F. McCoy and J. A. Tomlin, "Some experiments on the accuracy of three methods of updating the inverse in the simplex method", Technical Report, Stanford University, 1974.
- [14] S. Lim, G. Kim and S. Park, "A comparative study between various LU update methods in the simplex method", *Journal of the Military Operations Research Society of Korea*, Vol. 29, No. 1, 2003.
- [15] E. S. Badr, K. Paparrizos, N. Samaras and A. Sifaleras, "On the basis inverse of the exterior point simplex algorithm", *Proceedings of the 17th National Conference of Hellenic Operational Research Society (HELORS)*, Rio, Greece, 16–18 June 2005, pp. 677–687.

- [16] N. Ploskas, N. Samaras and K. Margaritis, "A parallel implementation of the revised simplex algorithm using OpenMP: some preliminary results", *Optimization Theory, Decision Making, and Operational Research Applications*, Springer Proceedings in Mathematics & Statistics, Vol. 31, 2013, pp. 163–175.
- [17] D. P. O’Leary and J. H. Jung, "Implementing an interior point method for linear programs on a CPU-GPU system", *Electronic Transactions on Numerical Analysis*, Vol. 28, 2008, pp. 879-899.
- [18] D. M. Gay, "Electronic mail distribution of linear programming test problems", *Mathematical Programming Society COAL Newsletter*, Vol. 13, 1985, pp. 10–12.
- [19] D. G. Spampinato and A. C. Elster, "Linear optimization on modern GPUs", *Proceedings of the 23rd IEEE International Parallel and Distributed Processing Symposium, IPDPS 2009*, 2009, Rome, Italy.
- [20] NVIDIA Corporation, "CUDA-CUBLAS Library 2.0", 2013.  
<https://developer.nvidia.com/cublas>
- [21] J. Bieling, P. Peschlow and P. Martini, "An efficient GPU implementation of the revised Simplex method", *Proceedings of the 24th IEEE International Parallel and Distributed Processing Symposium, IPDPS 2010*, April 2010, Atlanta, USA.
- [22] M. E. Lalami, V. Boyer and D. El-Baz, "Efficient Implementation of the Simplex Method on a CPU-GPU System", *Proceedings of the 2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and PhD Forum, IPDPSW 2011*, 2011, Washington, USA, pp. 1999–2006.
- [23] X. Meyer, P. Albuquerque and B. Chopard, "A multi-GPU implementation and performance model for the standard simplex method", *Proceedings of the 1st International Symposium and 10th Balkan Conference on Operational Research*, 22 - 25 September 2011, Thessaloniki, Greece, pp. 312–319.
- [24] J. Li, R. Lv, X. Hu and Z. Jiang, "A GPU-Based Parallel Algorithm for Large Scale Linear Programming Problem", In: J. Watada, G. Phillips-Wren, L. Jai and R. J. Howlett, Eds., *Intelligent Decision Technologies, SIST 10*, Springer Berlin Heidelberg, Springer-Verlag, Berlin, 2011, pp. 37-46.
- [25] G. B. Dantzig, A. Orden and P. Wolfe, "The Generalized Simplex Method", *RAND P-392-1*, 1953.
- [26] Y. Zhang, "Solving Large-Scale Linear Programs by Interior-Point Methods Under the MATLAB Environment", *Technical Report TR96-01*, Department of Mathematics and Statistics, University of Maryland, July 1995, Baltimore, USA.