

# EFFICIENT PARALLEL NON-NEGATIVE LEAST SQUARES ON MULTI-CORE ARCHITECTURES

YUANCHENG LUO AND RAMANI DURAISWAMI  
UNIVERSITY OF MARYLAND, COLLEGE PARK

**Abstract.** We parallelize a version of the *active-set* iterative algorithm derived from the original works of Lawson and Hanson (1974) on multi-core architectures. This algorithm requires the solution of an unconstrained least squares problem in every step of the iteration for a matrix composed of the *passive columns* of the original system matrix. To achieve improved performance, we use parallelizable procedures to efficiently update and *downdate* the *QR* factorization of the matrix at each iteration, to account for inserted and removed columns. We use a reordering strategy of the columns in the decomposition to reduce computation and memory access costs. We consider graphics processing units (GPUs) as a new mode for efficient parallel computations and compare our implementations to that of multi-core CPUs. Both synthetic and non-synthetic data are used in the experiments.

**Key words.** Non-negative least squares, active-set, *QR* updating, parallelism, multi-core, GPU, deconvolution

**AMS subject classifications.** 15A06, 15A23, 65Y05, 65Y20

**1. Introduction.** A central problem in data-modelling is the optimization of underlying parameters specifying a linear model used to describe observed data. The underlying parameters of the model form a set  $n$  variables in a  $n \times 1$  vector  $x = \{x_1, \dots, x_n\}^T$ . The observed data is composed of  $m$  observations in a  $m \times 1$  vector  $b = \{b_1, \dots, b_m\}^T$ . Suppose that the observed data are linear functions of the underlying parameters in the model, then the function's values at data points may be expressed as a  $m \times n$  matrix  $A$  where  $Ax = b$  describes a linear mapping from the parameters in  $x$  to the observations in  $b$ .

In the general case where  $m \geq n$ , the dense overdetermined system of linear equations may be solved via a least squares approach. The usual way to solve the least squares problem is with the *QR* decomposition of the matrix  $A$  where  $A = QR$ , with  $Q$  an orthogonal  $m \times n$  matrix, and  $R$  an upper-triangular  $n \times n$  matrix. Modern implementations for general matrices use successive applications of the Householder transform to form *QR*, though variants based on Givens rotation or Gram-Schmidt orthogonalization are also viable. Such algorithms carry an associated  $O(mn^2)$  time-complexity. The resulting matrix equation may be rearranged to  $Rx = Q^T b$  and solved via back-substitution for  $x$ .

Sometimes, the underlying parameters are constrained to be non-negative in order to reflect real-world prior information. When the data is corrupted by noise, the estimated parameters may not satisfy these constraints, producing answers which are not usable. In these cases, it is necessary to explicitly enforce non-negativity, leading to the non-negative least squares (NNLS) problem considered in this paper.

The seminal work of Lawson and Hanson [19] provide the first widely used method for solving this non-negative least squares problem. This algorithm, later referred to as the *active-set method*, partitions the set of parameters or variables into the active and passive-sets. The active-set contains the variables with values forcibly set to zero and which violate the constraints in the problem. The passive-set contains the variables that do not violate the constraint. By iteratively updating a feasibility vector with components from the passive-set, each iteration is reduced to an unconstrained linear least squares sub-problem that is solvable via *QR*.

For many signal processing applications, NNLS problems in a few hundred to a thousand variables arise. In time-delay estimation for example, multiple systems are continuously stored or streamed for processing. A parallel method for solving multiple NNLS problems would enable on-line applications, in which the estimation can be performed as data is acquired. Motivated by such an application, we develop an efficient algorithm and its implementations on both multi-core CPUs and modern GPUs.

Section 1.2 summarizes alternative solutions to the NNLS problem. Section 2 establishes notation and formally describes the active-set algorithm. Section 3 presents a new method for updating the  $QR$  decompositions for the active-set algorithm. Sections 4-5 describe parallelism on multi-core CPUs and GPU like architectures. Section 6 provides a motivating application from remote estimation and section 7 compares the GPU and CPU results from experiments.

**1.1. Non-negative Least Squares.** We formally state the NNLS problem: Given a  $m \times n$  matrix  $A \in \mathbb{R}^{m \times n}$ , find a non-negative  $n \times 1$  vector  $x \in \mathbb{R}^n$  that minimizes the functional  $f(x) = \frac{1}{2}\|Ax - b\|^2$  i.e.

$$(1.1) \quad \min_x f(x) = \frac{1}{2}\|Ax - b\|^2, \quad x_i \geq 0.$$

The Karush-Kuhn-Tucker (KKT) conditions necessary for an optimal constrained solution to an objective function  $f(x)$  can be stated as follows [18]: Suppose  $\hat{x} \in \mathbb{R}^n$  is a local minimum subject to inequality constraints  $g_j(x) \leq 0$  and equality constraints  $h_k(x) = 0$ , then there exists vectors  $\mu, \lambda$  such that

$$(1.2) \quad \nabla f(\hat{x}) + \lambda^T \nabla h(\hat{x}) + \mu^T \nabla g(\hat{x}) = 0, \quad \mu \geq 0, \quad \mu^T g(\hat{x}) = 0.$$

In order to apply the KKT conditions to the minimization function (1.1), let  $\nabla f(x) = A^T(Ax - b)$ ,  $g_j(x) = -x_j$ , and  $h_k(x) = 0$ . This leads to the necessary conditions,

$$(1.3) \quad \mu = \nabla f(\hat{x}), \quad \nabla f(\hat{x})^T \hat{x} = 0, \quad \nabla f(\hat{x}) \geq 0, \quad \hat{x} \geq 0$$

that must be satisfied at the optimal solution.

**1.2. Survey of NNLS Algorithms.** A comprehensive review of the methods for solving the NNLS problem can be found in [9]. The first widely used algorithm, proposed by Lawson and Hanson in [19], is the active-set method that we implement on the GPU. Although many newer methods have since surpassed the active-set method for large and sparse matrix systems from our survey, the active-set method remains competitive for *small* to moderate sized systems with unstructured and dense matrices.

In [6], improvements to the original active-set method are developed for the Fast NNLS (FNNLS) variant. By reformulating the normal equations that appear in the pseudo-inverse for the least squares sub-problem, the cross-product matrices  $A^T A$  and  $A^T b$  can be pre-computed. This contribution leads to significant speed-ups in the presence of multiple right-hand-sides. In [2], further redundant computations are avoided by grouping similar right-hand-side observations that would lead to similar pseudo-inverses.

A second class of algorithms is iterative optimization methods. Unlike the active-set approach, these methods are not limited to a single active constraint at each iteration. In [17], a Projective Quasi-Newton NNLS approach uses gradient projections to avoid pre-computing  $A^T A$  and non-diagonal gradient scaling to improve convergence

and reduce zigzagging. Another approach in [10] produces a sequence of vectors optimized at a single coordinate with all other coordinates fixed. These vectors have an efficiently computable analytical solution that converge to the solution.

Other methods outside the scope of this review include the Principal Block Pivoting method for large sparse NNLS in [7], and the Interior Point Newton-like method in [1], [16] for moderate and large problems.

**2. Active-set Method.** Given a set of  $m$  linear equations in  $n$  unknowns which are constrained to be non-negative, let the active-set  $Z$  be the subset of variables which violate the non-negativity constraint or are zero and the passive-set  $P$  be the variables with positive values. Lawson and Hanson observe that only a small subset of variables remains in the candidate active-set  $Z$  at the solution. If the true active-set  $Z$  is known, then the NNLS problem is solved by an unconstrained least squares problem using the variables from the passive-set.

---

**Algorithm 1** Active-set method for non-negative least squares [19]

---

**Require:**  $A \in \mathbb{R}^{m \times n}$ ,  $x = 0 \in \mathbb{R}^n$ ,  $b \in \mathbb{R}^m$ , set  $Z = \{1, 2, \dots, n\}$ ,  $P = \emptyset$

**Ensure:** Solution  $\hat{x} \geq 0$  s.t.  $\hat{x} = \arg \min_{\frac{1}{2}} \|Ax - b\|^2$

```

1: while true do
2:   Compute negative gradient  $w = A^T(b - Ax)$ 
3:   if  $Z \neq \emptyset$  and  $\max_{i \in Z}(w_i) > 0$  then
4:     Let  $j = \arg \max_{i \in Z}(w_i)$ 
5:     Move  $j$  from set  $Z$  to  $P$ 
6:     while true do
7:       Let matrix  $A^P \in \mathbb{R}^{m \times *}$  s.t.  $A^P = \{\text{columns } A_i \text{ s.t. } i \in P\}$ 
8:       Compute least squares solution  $y$  for  $A^P y = b$ 
9:       if  $\min(y_i) \leq 0$  then
10:        Let  $\alpha = -\min_{i \in P}(\frac{x_i}{x_i - y_j})$  s.t. (column  $j \in A^P$ ) = (column  $i \in A$ )
11:        Update feasibility vector  $x = x + \alpha(y - x)$ 
12:        Move from  $P$  to  $Z$ , all  $i \in P$  s.t.  $x_i = 0$ 
13:      else
14:        Update  $x = y$ 
15:        break
16:      end if
17:    end while
18:  else
19:    return  $x$ 
20:  end if
21: end while

```

---

In algorithm 1, the candidate active-set  $Z$  is updated by first moving the largest positive component variable in the negative gradient  $w$  to the passive-set (line 5). This selects the component with the most negative gradient that reduces the residual 2-norm. The variables in the passive-set form a candidate linear least squares system  $A^P y = b$  where matrix  $A^P$  contain the column vectors in matrix  $A$  that correspond to indices in the passive-set (lines 7, 8). At each iteration, the feasibility vector  $x$  moves towards the solution vector  $y$  while preserving non-negativity (line 11). Convergence to the optimal solution is proven in [19].

The termination condition (line 3) checks if the gradient is strictly positive or

if the residual can no longer be minimized. At termination, the following relations satisfy the optimality conditions in (1.3):

1.  $w_i \leq 0 \quad i \in Z$  termination condition (line 3).
2.  $w_i = 0 \quad i \in P$  solving least squares sub-problem (line 8).
3.  $x_i = 0 \quad i \in Z$  updating sets (line 12).
4.  $x_i > 0 \quad i \in P$  updating  $x$  (lines 10-11).

The variables in the passive-set form the corresponding columns of the matrix  $A^P$  in the unconstrained least squares sub-problem  $A^P y = b$ . As discussed previously, the cost of solving the unconstrained least squares sub-problem is  $O(mn^2)$  via  $QR$ . If there are  $k$  iterations, then the cost of  $k$  independent decompositions is  $O(kmn^2)$ . However, the decompositions at each iteration share a similar structure in matrix  $A^P$ , and this can be taken advantage of. We observe the following properties of matrix  $A^P$  as the iterations proceed:

1. The active and passive-sets generally exchange a single variable per iteration; one column is added or removed from matrix  $A^P$ .
2. Most exchanges move variables from the active-set into the passive-set; early iterations add variables to an empty passive-set to build the feasible solution, while later iterations add and remove variables to refine the solution.

Hence, we develop a general method for  $QR$  column updating and downdating that takes advantage of the pattern of movement between variables in the active and passive-sets. To achieve real-time and on-line processing, the method must be parallelizable on GPUs or other multi-core architectures. We note that the improvements made to the active-set NNLS proposed in [6], [2] do not apply to our problem, and moreover do not account for possible efficiencies suggested by the observations above.

**3. Proposed Algorithm.** The first property of matrix  $A^P$  suggests that a full  $A^P = QR$  decomposition is unnecessary. Instead, we consider an efficient  $QR$  column updating and downdating method.

1.  $QR$  Updating: A new variable added to set  $P$  expands matrix  $A^P$  by a single column. Update previous matrices  $Q, R$  with this column insertion.
2.  $QR$  Downdating: The removal of a variable from set  $P$  shrinks matrix  $A^P$  by a single column. Downdate previous matrices  $Q, R$  with this column deletion.

The second property of matrix  $A^P$  suggests that we can optimize the cost for  $QR$  updating in terms of floating point operations (flops) and column or row memory accesses. We observe that many  $QR$  updating methods minimize computations when inserting columns at the right-most index. Our method takes advantage of this by maintaining a separate ordering for the columns of matrix  $A^P$  by the relative times of insertions and deletions across iterations. That is, a column insertion always appends to the end of a reordered matrix  $\hat{A}^P$ . We describe the effects of the reordering strategy for various updating methods in sections 3.1-3.3. We also show that the modified Gram-Schmidt and Givens rotation methods are the most cost efficient with respect to the reordering strategy for overdetermined and square systems.

**3.1. QR Updating by Modified Gram-Schmidt.** The reordering strategy allows a new column  $a_i$  from the matrix  $A^P = [a_1, \dots, a_i, \dots, a_n]$  to be treated as the right-most column in the decomposition. We define list  $\hat{P}$  as an ordered list of column indices from set  $P$  such that the associated column  $\hat{p}_{i-1}$  is added in a prior

iteration to column  $\hat{p}_i$ . The reordered decomposition  $\hat{A}^P = \hat{Q}\hat{R}$  is

$$(3.1) \quad \begin{aligned} \hat{A}^P &= [a_{\hat{p}_1}, \dots, a_{\hat{p}_{i-1}}, a_{\hat{p}_i}], \\ \hat{Q} &= [q_{\hat{p}_1}, \dots, q_{\hat{p}_{i-1}}, q_{\hat{p}_i}], \\ \hat{R} &= [r_{\hat{p}_1}, \dots, r_{\hat{p}_{i-1}}, r_{\hat{p}_i}] \end{aligned}$$

where  $\hat{Q}$  is a  $m \times i$  matrix and  $\hat{R}$  is an  $i \times i$  matrix. To compute column  $q_{\hat{p}_i}$ , we orthogonalize the inserted column  $a_i$  with all the previous columns in matrix  $\hat{Q}$  via vector projections. To compute column  $r_{\hat{p}_i}$ , we take the inner products between column  $a_i$  and columns in  $\hat{Q}$ , or the equivalent matrix-vector product  $\hat{Q}^T a_i$ . Both quantities are found using the Modified Gram-Schmidt (MGS) procedure in algorithm 2.

---

**Algorithm 2** Reordered MGS QR Column Updating

---

**Require:** Reordered list  $\hat{P}$  contains the elements in set  $P$ , index  $i$  the variable added to set  $P$ , column  $a_i$  the new column in  $A^P$ , columns  $q_j \in Q$

**Ensure:**  $\hat{A}^P = \hat{Q}\hat{R}$ , update vector  $\hat{Q}^T b$ , list  $\hat{P}$

- 1: Let vector  $u = a_i$
  - 2: **for all** column index  $k \in \text{list } \hat{P}$  **do**
  - 3:    $u = u - \langle q_k, u \rangle q_k$
  - 4:    $\hat{R}_{ki} = \langle a_i, q_k \rangle$
  - 5: **end for**
  - 6:  $q_i = \frac{u}{\|u\|}$
  - 7:  $\hat{R}_{ii} = \|u\|$
  - 8:  $\hat{Q}^T b_i = \langle q_i, b \rangle$
  - 9: Add  $i$  to list  $\hat{P}$
- 

With the reordering strategy in algorithm 2, a new column  $a_i$  is always inserted in the right-most position of matrix  $\hat{A}^P$ . The number of columns read from memory in matrix  $\hat{Q}$  is the size of set  $P$ , denoted as  $\ell \leq n$  and is used to form column  $q_i$ . The number of column memory writes per step is two, as column  $q_i$  appends to matrix  $\hat{Q}$  and the projection step writes a single column to matrix  $\hat{R}$ . Updating matrices  $\hat{Q}$  and  $\hat{R}$  requires  $6m\ell + 3m + 1$  flops. The asymptotic complexity is  $O(mn)$ .

Without the reordering strategy, column  $a_i$  can be inserted into the middle of matrix  $\hat{A}^P$ . This requires computing column  $q_i$  the re-orthogonalization of the  $\ell - i$  columns to its right. The memory access costs of computing  $q_i$  is  $i$  number of columns reads from matrix  $\hat{Q}$  and two columns writes to matrices  $\hat{Q}$  and  $\hat{R}$ . The re-orthogonalization costs of column  $q_j$  where  $j > i$  is equivalent to a new column insertion into matrix  $A^P$ . This is because the MGS method does not compute the null-space of the basis vectors in matrix  $\hat{Q}$ . Orthogonalizing columns  $q_j$  and  $q_{j+1}$  with respect to column  $q_i$  does not preserve the orthogonality between  $q_j$  and  $q_{j+1}$ . Thus, each of the  $\ell - i + 1$  columns must be reinserted with an additional  $\ell(\ell - i + 1)$  column reads and  $2(\ell - i + 1)$  column writes. Updating matrices  $\hat{R}$  and  $\hat{Q}$  requires a total of  $(3m\ell + 3mi + 1)(\ell - i + 2)$  flops. The asymptotic complexity is  $O(mn^2)$ .

**3.2. Alternative QR Updating by Rotations.** Rotation based methods for updating  $QR$  are possible. In [12],  $\hat{Q}$  and  $\hat{R}$  are treated as  $m \times m$  matrices where matrix  $\hat{Q}$  is initially the identity. When inserting column  $a_i$ , the method appends  $m \times 1$  column vector  $r_{\hat{p}_i} = \hat{Q}^T a_i$  to matrix  $\hat{R}$ . A series of rotation transformations

introduces zeros to rows  $\{i + 1, i + 2, \dots, m\}$  of column  $r_{\hat{p}_i}$  to preserve the upper-triangular property. The rotation transformations then update the columns to the right of index  $i$  in matrices  $\hat{R}$  and  $\hat{Q}$ . A similar step follows updating the right-hand side  $\hat{Q}^T b_i$ .

Without the reordering strategy, the costs of this rotation method depend on index  $i$ . Column  $r_{\hat{p}_i}$  requires  $m - i$  rotation transformations. Each transformation requires two row memory reads and writes to matrix  $\hat{R}$  and two column memory reads and writes to matrix  $\hat{Q}$  for a total of  $2(m - i)$ . This is disadvantageous as the number of column and row accesses is bound by  $m$  and multiple columns and rows of matrices  $\hat{R}$  and  $\hat{Q}$  are modified. Updating matrix  $\hat{Q}$  and  $\hat{R}$  requires  $6m(m - i)$  and  $2m^2 + 8(m - i) + 6(\ell - i + 1)(\ell/2 - i/2 - 1)$  flops respectively. The asymptotic complexity is  $O(m^2 + n^2)$ .

With the reordering strategy, index  $i = \ell + 1$  and so many of the costs are reduced. There are no columns to the right of index  $i$  in matrix  $\hat{R}$  so updating is limited to single column memory write of column  $r_{\hat{p}_i}$ . Updating matrix  $\hat{Q}$  now requires  $m - \ell - 1$  column reads and writes each while applying the transformations. Updating matrices  $\hat{Q}$  and  $\hat{R}$  requires  $6m(m - \ell - 1)$  and  $2m^2 + 8(m - \ell - 1)$  flops respectively. The asymptotic complexity is  $O(m^2)$ .

**3.3. Alternative QR Updating by Semi-normal Equations.** The corrected semi-normal equations (CSNE) can be used to update a  $\ell \times \ell$  matrix  $\hat{R}$  without the construction of matrix  $\hat{Q}$ . The stability analysis of this method is provided by [3]. With the reordering strategy, the problem treats  $\hat{R}^T \hat{R}x = \hat{A}^P b$  where column  $r_{\hat{p}_i}$  is computed by

$$(3.2) \quad \begin{aligned} \hat{R}^T \hat{R}z &= \hat{A}^P a_i, \\ s &= a_i - \hat{A}^P z, \\ \hat{R}^T \hat{R}\delta z &= s, \\ z &= z + \delta z, \\ r_{\hat{p}_i} &= \begin{bmatrix} \hat{R}z \\ \|\hat{A}^P z - a_i\| \end{bmatrix}. \end{aligned}$$

Although the method does not compute and store matrix  $\hat{Q}$ , it requires both row and column access to matrix  $\hat{R}$  and more operations to produce column  $r_{\hat{p}_i}$ . Computing and correcting for vector  $z$  entails four back-substitutions using matrices  $\hat{R}^T$ ,  $\hat{R}$ , and  $\hat{A}^P$ . All four back-substitutions requires  $\ell$  row or column memory reads from matrix  $\hat{R}$  each. Two of the back-substitutions require  $m$  row memory reads from matrix  $\hat{A}^P$ . The total number of column and row memory reads in the method is  $3m + 2\ell$  and one column memory write to update matrix  $\hat{R}$ . The entire procedure requires  $6m^2 + 4m + 1 + 3\ell^2 + 3\ell$  flops. The asymptotic complexity is  $O(m^2 + n^2)$ .

Without the reordering strategy, the same CSNE method computes column  $r_{\hat{p}_i}$  and a series of rotations introduces zeros below index  $i$ . The rotation transformations are then applied to the columns to the right of index  $i$  in matrix  $\hat{R}$ . This requires an additional  $\ell - i$  row memory reads and writes to matrix  $\hat{R}$  each and  $6(\ell - i + 1)(\ell/2 - i/2 - 1) + 3(\ell - i)$  flops. The asymptotic complexity is  $O(n^2)$ . The costs for the updates are summarized in Table 3.1.

**3.4. QR DOWNDATING BY ROTATIONS.** The reordering strategy is less applicable to the downdating scheme as deleted columns may not be in the right-most index. Suppose that column  $a_i$  is removed from matrix  $A^P = [a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n]$ . Let



We refer to [11] for precautions when computing the rotation coefficients  $c$ ,  $r$  in algorithm 3. When updating matrices  $\tilde{R}$  and  $\tilde{Q}$ , a row or column is fixed so the transformation requires  $2(\ell - i + 1)$  row and column memory reads and writes each. Updating matrices  $\tilde{Q}$  and  $\tilde{R}$  requires  $6m(m - i)$  and  $6(\ell - i) + 6(\ell - i + 1)(\ell/2 - i/2 - 1)$  flops respectively. The asymptotic complexity is  $O(m^2 + n^2)$ . The costs for the downdates are summarized in Table 3.1.

Algorithm	Col/row accesses	Up/down $Q$ flops	Up/down $R$ flops
MGS/up/reorder	$\ell + 2$	$6m\ell + 3m + 1$	included in $Q$
MGS/up/unorder	$\ell(\ell - i + 2) + 2(\ell - i + 2)$	$(3m\ell + 3mi + 1)(\ell - i + 2)$	included in $Q$
Rot/up/reorder	$2(m - \ell - 1)$	$6m(m - \ell - 1)$	$2m^2 + 8(m - \ell - 1)$
Rot/up/unorder	$4(m - \ell - 1)$	$6m(m - i)$	$2m^2 + 8(m - i) + 6(\ell - i + 1)(\ell/2 - i/2 - 1)$
CSNE/up/reorder	$3m + 2\ell + 1$	NA	$6m^2 + 4m + 1 + 3\ell^2 + 3\ell$
CSNE/up/unorder	$3m + 4\ell - 2i + 1$	NA	$6m^2 + 4m + 1 + 3\ell^2 + 6(\ell - i + 1)(\ell/2 - i/2 - 1) + 6\ell - 3i$
Rot/down/NA	$4(\ell - i + 1)$	$6m(m - i)$	$6(\ell - i) + 6(\ell - i + 1)(\ell/2 - i/2 - 1)$

TABLE 3.1

*Costs for QR updating/downdating methods with respect to the reordering strategy. The rotation and CSNE methods have flops of order  $m^2$ . For overdetermined and square systems where  $\ell \leq n \leq m$ , this quantity is minimized for the modified Gram-Schmidt method.*

**4. Multi-core CPU Architectures.** The multi-core trend began as a response to the slowdown of Moore’s Law while manufactures approached the limitations in single-core clock speeds. With additional cores added on chip, individual CPU threads can be assigned and processed by their own units in hardware. Thus, a single problem is decomposed and solved by several threads without over-utilizing a single core. This gave multi-threading an edge over traditional single-core processors as data and instruction level caches could be dedicated to a smaller sub-set of operations.

Such Multiple-Instruction-Multiple-Data (MIMD) architectures support task-level parallelism where each core can asynchronously execute separate threads on separate data regions. The individual cores are often super-scalar and thus capable of processing out-of-order instructions in their pipeline. This allows multi-core architectures to simulate data-level parallelism from Single-Instruction-Multiple-Data (SIMD) like architectures such as the GPU with added proficiency. Furthermore, multi-core architectures have access to a common pool of main memory off-die and capable of multi-level caching per core and per processor on-die. For both data and task-level parallelism, this allows memory to be decomposed and cached on a per-core basis for efficient reuse.

Several application programming interfaces (APIs) and libraries take advantage of these shared memory multiprocessing environments for high performance computing. Open Message Passing (OpenMP) is an API based on fork-join operations where the program enters into a designated parallel region [22]. Each thread exhibits both task and data-level parallelism as it independently executes code within a same parallel region. The Intel Math Kernel Library (MKL) is a set of optimized math routines with calls to Basic Linear Algebra Sub-programs (BLAS) and Linear Algebra PACKage (LAPACK) libraries [14]. Many of its fundamental matrix and vector routines are blocked and solved across multiple threads.

**4.1. CPU Implementation.** To exploit the advantages of multi-threading, we adopt both the OpenMP API and the Intel MKL in the CPU implementation. One

way to map each linear system to a thread is to declare the entire NNLS algorithm within a parallel OpenMP region. That is, a specified fraction of threads execute NNLS on a mutually exclusive set of linear system of equations. The remaining threads are dedicated to the MKL library in order to accelerate common matrix-vector and vector-vector operations used to solve the unconstrained least squares sub-problem.

**5. GPU Architectures.** Recent advances in general purpose graphics processing units (GPUs) have given rise to highly programmable architectures designed with parallel applications in mind. Moreover, GPUs are considered to be typical of future generations of highly parallel, multi-threaded, multi-core processors with tremendous computational horsepower. They are well-suited for algorithms that map to a Single-Instruction-Multiple-Thread (SIMT) architecture. Hence, GPUs achieve a high arithmetic intensity (ratio of arithmetic operation to memory operations) when performing the same operations across multiple threads on a multi-processor.

GPUs are often designed as a set of multiprocessors, each containing a smaller set of scalar-processors (SP) with a Single-Instruction-Multiple-Data (SIMD) architecture. Hardware multi-threading under a SIMT architecture maps multiple threads to a single SP. A single SP handles the instruction address and register states of multiple threads so that they may execute independently. The multiprocessor's SIMT unit schedules batches of threads to execute a common instruction. If threads of the same batch diverge via a data-dependent conditional branch, then all the threads along the separate branches are serialized until they converge back to the same execution path.

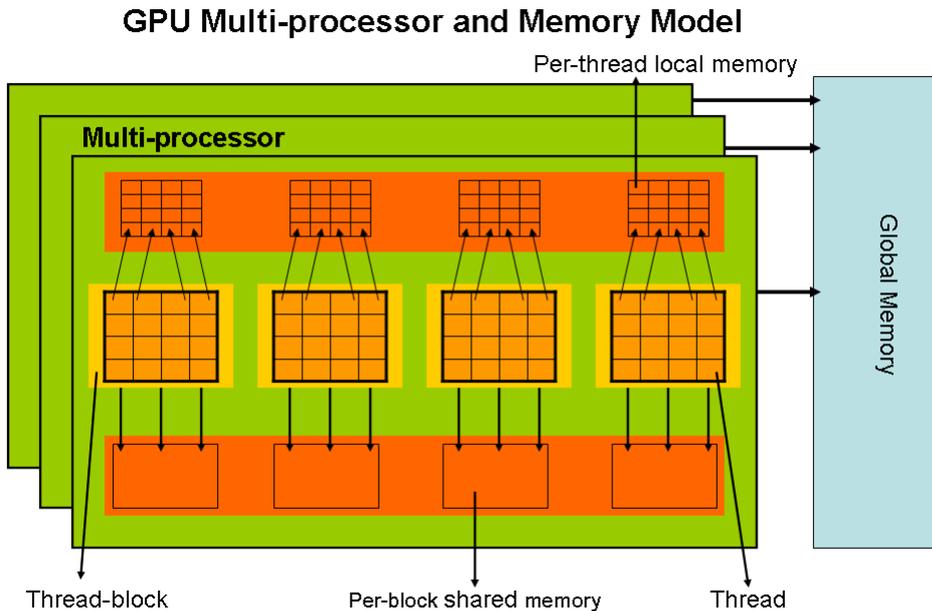


FIG. 5.1. GPU multiprocessor and memory model

GPUs have a hierarchical memory model with significantly different access times

to each level. At the top, all multiprocessors may access a global memory pool on the device. This is the common space where input data is generally copied and stored from main memory by the host. It is also the slowest memory to access as a single query from a multi-processor has a 400 to 600 clock cycles latency on a cache-miss. See [20] for a discussion on coalesced global memory accesses which reads or writes to a continuous chunk of memory at a cost of one query and implicit caching on the Fermi architecture. On the same level, texture memory is also located on the device but can only be written to from hosts. However, it is faster than global memory when access patterns are spatially local. On the next level, SPs on the same multi-processor have access to a fast shared memory space. This enables explicit inter-thread communication and temporary storage for frequently accessed data. Constant memory, located on each multi-processor, are cached and optimized for broadcasting to multiple threads. On the lowest level, a SP has its own private set of registers distributed amongst its assigned threads. The latency for accessing both shared and per-processor registers normally adds zero extra clock cycles to the instruction time.

Programming models such as NVIDIA’s Compute Unified Device Architecture (CUDA) [20] and OpenCL [21] organize threads into thread-blocks, which in turn are arranged in a 2D grid. A thread-block refers to a 1-D or 2-D patch of threads that are executed on a single multiprocessor. These threads efficiently synchronize their instructions and pass data via shared memory. Instructions are generally executed in parallel until a conditional branch or an explicit synchronization barrier is declared. The synchronization barrier ensures that the thread-block waits for all its threads to complete its last instruction. Thus, two levels of data parallelism are achieved. The threads belonging to the same thread-block execute in lock-step as they process a set of data. Individual thread-blocks execute asynchronously but generally with the same set of instructions on a different set of data.

While efficient algorithms on sequential processors must reduce the number of computations and cache-misses, parallel algorithms on GPUs are more concerned with minimizing data dependencies and optimizing accesses to the memory hierarchy. Data dependency increases the number of barrier synchronizations amongst threads and is often subject to the choice of the algorithm. Memory access patterns present a difficult bottleneck on multiple levels. While latency is the first concern for smaller problems, we run into a larger issue with memory availability as the problem size grows. That is, the shared memory and register availability are hard limits that bound the size and efficiency of thread-blocks. A register memory bound per SP limits the number of threads assigned to each SP and so decreases the maximum number of threads and thread-blocks running per multi-processor. A shared memory bound per multi-processor limits the number of thread-blocks assigned to a multi-processor and so decreases the total number of threads processed per multi-processor.

**5.1. GPU Implementation.** One way to map each linear system onto a GPU is to consider every thread-block as an independent vector processor. Each thread-block of size  $m \times 1$  maps to the elements in a column vector and asynchronously solves for a mutually exclusive set of linear systems. The number of thread-blocks that fit onto a single multi-processor depends on the column size  $m$  or the number of equations in the linear system. This poses a restriction on the size of linear systems that our GPU implementation can solve as the maximum size  $m$  is constrained to a fraction of the amount of shared memory available per multi-processor. Fortunately, this is not an issue for applications where  $m$  is small (500-1000) and the number of linear systems to be solved is large. However for arbitrarily sized linear systems of

equations, our GPU implementation is not generalizable. We note that this is not an algorithmic constraint but rather a design choice for our application. Our multi-core CPU implementation of the same algorithm can solve for arbitrarily sized linear systems. We discuss the details of the GPU implementation in sections 5.2-5.3.

**5.2. Parallelizing QR Methods.** Full  $QR$  decompositions on the GPU via blocked MGS, Givens rotations, and Householder reflections are implemented in [23], [15]. While [15] cites that the blocked MGS and Givens rotation methods are ill-suited for large systems on GPUs as they suffer from instability and synchronization overhead, we are interested in only the  $QR$  updating and downdating schemes for a large number of *small* systems. We show that it is possible for  $m$  threaded multi-processors to efficiently perform the MGS updating and Givens rotation downdating steps.

For the MGS update step, most of the operations are formulated as vector inner products, scalar-vector products, and vector-vector summations. These operations lead to an one-to-one mapping between the  $m \times 1$  column vector coordinates and the  $m$  threaded thread-block. Such operations are computable via parallel reduction techniques from [4]. In algorithm 2, we parallelize all four inner products (lines 3, 4, 6, 8) in  $\log m$  parallel time each. The inner loop iterates for  $\ell$  or at most  $n$  times. Thus, we obtain an order reduction in parallel time-complexity to  $O(n \log m)$ .

For the Givens rotation downdate step, we obtain an one-to-many mapping between the  $n \times 1$  row vector elements and the  $m$  threads in a thread-block for matrix  $R$ . We obtain an one-to-one mapping for the  $m \times 1$  column vector elements in matrix  $Q$ . Computing vector  $Q^T b$  follows a similar relation. For obtaining the rotation coefficients  $c, s$ , a single thread computes and broadcasts to the rest of the thread-block. In algorithm 3, the inner loop (lines 3-4) updates both matrices  $\tilde{R}$  and  $\tilde{Q}$  in parallel  $O(1)$  time. Writing row and column data and updating vector  $\tilde{Q}^T b$  (line 7) are thread-independent and computable in  $O(1)$  parallel time. Thus, we obtain an order reduction in parallel time-complexity to  $O(n)$ .

Parallel reductions are often performed on the GPU in place of common vector-vector operations using prefix sum discussed in [13]. Algorithm 4 sums 512 elements in 9 parallel flops, 5 thread-synchronizations, and 18 parallel shared memory accesses. Each of the 512 threads reserves a memory slot in shared memory. The unique thread ID or tID denotes the corresponding data index in the shared memory array. At each step, half the threads from the previous step sum up the data entries stored in the other half of shared memory. The process continues until index 0 in the shared memory array contains the total summation.

**5.3. Memory Usage.** To take advantage of different access times on the GPU memory hierarchy, the input and intermediate data can be stored and accessed on different levels for efficient reuse. Local intermediate vectors can either be stored in shared memory or alternatively in dedicated registers spanning all threads in a thread-block. List  $\hat{P}$  is stored in shared memory as multiple threads require synchronization to update and downdate the same column. The right-hand-side vector  $\hat{Q}^T b$  is stored in registers since no thread accesses elements outside its one-to-one mapping in the update and downdate steps.

Global memory accesses on the GPU are unavoidable for updating large matrices  $\hat{Q}$  and  $\hat{R}$ . We store matrix  $\hat{Q}^T$  so that column vector accesses are coalesced in row-oriented programming models and matrix  $\hat{R}$  as the Givens rotations update the rows. Matrices  $\hat{Q}$  and  $\hat{R}$  are stored in-place unlike the compact format in (3.1), (3.3). We allocate  $m \times n$  blocks of global memory and use the reordered list  $\hat{P}$  to associate

**Algorithm 4** CUDA parallel floating-point summation routine [13]

---

```

__device__ float reduce512( float smem512[], unsigned short tID){
    __syncthreads();
    if(tID < 256) smem512[tID] += smem512[tID + 256];
    __syncthreads();
    if(tID < 128) smem512[tID] += smem512[tID + 128];
    __syncthreads();
    if(tID < 64) smem512[tID] += smem512[tID + 64];
    __syncthreads();
    if(tID < 32){
        smem512[tID] += smem512[tID + 32];
        smem512[tID] += smem512[tID + 16];
        smem512[tID] += smem512[tID + 8];
        smem512[tID] += smem512[tID + 4];
        smem512[tID] += smem512[tID + 2];
        smem512[tID] += smem512[tID + 1];
    }
    __syncthreads();
    return smem512[0];
}

```

---

column and row indices for the update and downdate steps. This is to avoid any physical shifts of column vectors in global memory. Rather, we parallel shift the list  $\hat{P}$  when a variable is removed from the passive-set.

The MGS update step reads  $\ell$  number of columns in matrix  $\hat{Q}$  from global memory into registers. Computing inner products and vector norms during the projections requires an intermediate shared memory vector for the parallel reduction function. The new column for matrix  $\hat{R}$  is locally stored in registers before updated to global memory. A single element for vector  $\hat{Q}^T b$  is updated and written to shared memory. The total number of parallel shared memory accesses is  $39\ell + 2$ . The total number of parallel global memory accesses is  $\ell + 2$ .

The Givens rotations downdate step accesses two columns of matrix  $\tilde{Q}$  and two rows of matrix  $\tilde{R}$  for each of the  $\ell - i$  transformation. Since row  $i$  of matrix  $\tilde{R}$  and column  $i$  of matrix  $\tilde{Q}$  are fixed across transformations, they are stored and updated in shared memory. The other row and column are directly updated in global memory. Updating vector  $\tilde{Q}^T b$  requires two shared memory reads and writes. The total number of parallel shared memory accesses is  $2(\ell - i) + 2$ . The total number of parallel global memory accesses is  $2(\ell - i + 1)$ .

**6. Application.** In remote sensing, a discrete-time deconvolution recovers a signal  $x$  that has been convolved with a transfer function  $s$ . The known signal  $s$  is often convolved with an unknown signal  $x$  that satisfies certain characteristics.

$$\begin{aligned}
 (6.1) \quad y(t) = s(t) * x(t) &= \int_{-\infty}^{\infty} s(\tau)x(t - \tau) d\tau = \int_{-\infty}^{\infty} x(\tau)s(t - \tau) d\tau \\
 &= \sum_{\tau=-\infty}^{\infty} x(\tau)s(t - \tau) d\tau = \sum_{\tau=1}^n x(\tau)s(t - \tau) d\tau
 \end{aligned}$$

where  $t$  is the sample's time,  $y(t)$  the observed signal, and  $n$  the number of samples over time. To solve for the unknown signal  $x$ , we rewrite (6.1) as the following linear

system of equations  $Ax = b$ , where  $A$  is a Toeplitz matrix:

$$(6.2) \quad \begin{aligned} A &= \begin{bmatrix} s(0) & s(-1) & \cdots & s(-(n-1)) \\ s(1) & s(0) & \cdots & s(-(n-2)) \\ \vdots & \vdots & \ddots & \vdots \\ s(n-1) & s(n-2) & \cdots & s(0) \end{bmatrix}, \\ x &= [x(1) \ \cdots \ x(n)]^T, \\ b &= [y(1) \ \cdots \ y(n)]^T, \\ Ax &= b. \end{aligned}$$

Efficient algorithms for the deconvolution problem, which either exploit the simple structure of the convolution in Fourier space, or which exploit the Toeplitz structure of the matrix, are available in [5], [8]. However, if signal  $x$  is known to be non-negative and the data  $y(t)$  is corrupted by noise, then we may treat the deconvolution as a NNLS problem to ensure non-negativity.

**7. Experiments.** As a baseline, we note that Matlab’s `lsqnonneg` function implements the same active-set algorithm but with a full  $QR$  decomposition for the least squares sub-problem. Matlab 2009b and later versions use Intel’s Math Kernel Library (MKL) with multi-threading to resolve the least squares sub-problems. For a better comparison, we first port the `lsqnonneg` function into native C-code with calls to multi-threaded MKL BLAS and LAPACK functions. The results from this implementation (CPU `lsqnonneg`) show a 1.5-3x speed-up over the Matlab `lsqnonneg` function in our experiments. Next, we apply our updating and downdating strategies with column reordering using MKL BLAS functions to the CPU version. The results from this second implementation (CPU NNLS) show a 1-3x speed-up that depends on the number of column updates and downdates. Last, we compare the `lsqnonneg` variants to alternative NNLS algorithms from literature.

To compare GPU implementation with the multi-threaded CPU variants, we begin timing the point of entry and exit out of the GPU kernel function. Memory transfer and pre-processing times in the case of non-synthetic data are omitted. Both the GPU and CPU variants also obtain identical solutions subject to rounding error within the same number of iterations for all data sets. We find that for a fewer number of linear systems, the CPU implementations outperforms the GPU as only a fraction of the GPU cores are utilized. When the number of linear systems surpasses the number of multi-processors, the GPU scales better on an order of 1-3x than our fastest CPU implementation.

For reference, we use a Dual Quad-Core Intel(R) Xeon(R) X5560 CPU @ 2.80GHz (8 cores) for testing our CPU implementations. The CPU codes compiled under both Intel `icc 11.1`, `gcc 4.5.1`, and linked to MKL 10.1.2.024 yield similar results for 8 run-time threads. The codes tested between Matlab 2010b and 2009b also yield comparable results. Mixing the number of threads assigned between OpenMP and MKL did not have a large impact on our system. We use a NVIDIA Tesla C2050 (448 cores across 14 multi-processors) and codes compiled under CUDA 3.2 for the GPU implementation and testing.

**7.1. Non-Synthetic Test Data.** For real-world data, we use terrain laser imaging sets obtained from the NASA Laser Vegetation Imaging Sensor (LVIS)<sup>1</sup>. Each

data set contains multiple 1-D Gaussian-like signals  $s$  and observations of total return energy  $b$  of size  $m = n = 432$ . In this deconvolution problem, the transfer functions  $s$  represent the single impulse energy fired over time on ground terrain and the observed signals  $b$  produces a waveform that indicate the reflected energy over time. The signals  $s$  are generally 15-25 samples wide so the computed matrices  $A$  are toeplitz banded and sparse. NNLS solves for corresponding pairs of matrix  $A$  and vector  $b$  to obtain the sparse non-negative solutions  $x$  that represent the times of arrival for a series of a fired impulses. This estimates the ranges or distances to a surface target.

For comparing the NNLS methods, we record the run-times in relation to the number of column updates and dwnupdates for the least squares sub-problem. The results from CPU NNLS show a 11x speed-up over the GPU implementation when solving for a single system. This is due to the underutilization of cores in all but the multi-processor currently assigned to the linear system of interest. For a larger number of systems, the GPU results show a 1-2x speed-up over CPU NNLS due in part to the larger number of processing units suited for vector operations in the algorithm. The results between CPU NNLS and CPU lsqnonneg show the performance gains from fewer flops and memory accesses attained by the column reordering, updating, and downdating strategies.

Number of systems	1	24	48	96	192
Number of updates	108	1477	2806	5695	12067
Number of dwnupdates	14	220	406	834	1839
GPU NNLS	0.2172	0.2238	0.2356	0.4508	0.7203
CPU NNLS	0.0186	0.1636	0.2990	0.5989	1.2489
CPU lsqnonneg	0.0770	0.7163	1.2908	2.6225	5.5460
Matlab lsqnonneg	0.1342	1.1236	2.0152	4.1268	8.7176
Matlab FNNLS [6]	0.0493	0.5936	1.1135	2.2639	4.6148
Matlab interior-points [16]	8.3642	135.7896	261.3665	528.4468	1082.8714
Matlab PQN-NNLS [17]	1.4161	138.2953	217.6929	479.9867	862.9674

TABLE 7.1

*Runtime (seconds) comparisons of NNLS and lsqnonneg variants for signal deconvolution. Signal and observation data taken from LVIS Sierra Nevada, USA (California, New Mexico), 2008.*

**7.2. Synthetic Test Data.** For the first set of synthetic data, we generate mean shifted 1-D Gaussians with  $\sigma = 4.32$  to store as columns in matrix  $A$  of size  $m = n = 512$ . In this Gaussian fitting problem, each system uses the same matrix  $A$  but with non-negative random vectors  $b$ . The choice of the  $\sigma$  parameter ensures that the mean shifted Gaussians are not too wide as to allow early convergence and not too narrow as to locally affect only a few variables. Furthermore, matrix  $A$  is now considered dense and vectors  $b$  no longer reflect real-world values. We expect the average number of iterations or column updates and dwnupdates to exceed that of the real-world data cases.

The total speed-up of GPU NNLS over the CPU variants are more pronounced (3x compared to CPU NNLS, 23x compared to CPU lsqnonneg). The larger ratio of column dwnupdate to update steps suggests that our reordering strategy and fast Givens rotation method in the downdating step outperforms the lsqnonneg variants.

<sup>1</sup><https://lvis.gsfc.nasa.gov/index.php>

Number of systems	1	24	48	96	192
Number of updates	165	3907	7792	15541	30966
Number of downdates	92	2109	4196	8362	16599
GPU NNLS	0.4257	0.5094	0.9546	1.7862	3.2672
CPU NNLS	0.0654	1.2238	2.4141	4.8067	9.6250
CPU lsqnonneg	0.4791	8.7611	17.2483	34.5396	69.4889
Matlab lsqnonneg	0.9437	19.5904	38.6469	77.1072	155.7700
Matlab FNNLS [6]	0.4937	11.7176	23.9502	47.4635	91.4500
Matlab interior-points [16]	1.6317	40.7017	83.9788	164.4839	328.9569
Matlab PQN-NNLS [17]	3.1106	128.6616	253.3796	504.3153	989.2051

TABLE 7.2

*Runtime (seconds) comparisons of NNLS and lsqnonneg variants on multiple systems of mean shifted Gaussian matrix  $A$  and random vectors  $b$ .*

For the second set of synthetic data, we generate both random matrices  $A$  of size  $m = n = 512$  and non-negative random vectors  $b$ . The number of column updates and downdates is less than that of the two previous experiments. Furthermore, the total number of column updates dominates the number of column downdates. The results between GPU and CPU NNLS show that both implementations have similar run-time scaling as the number of systems increases. This suggests that the most of the performance gains in prior experiments are from the GPU downdating steps. The results between CPU NNLS and CPU lsqnonneg leads to a similar conclusion as the performance gain (1.7x) is minimum compared to the prior two experiments.

Number of systems	1	24	48	96	192
Number of updates	52	1169	2361	4766	9525
Number of downdates	0	13	28	58	124
GPU NNLS	0.1194	0.1397	0.2526	0.4875	0.8667
CPU NNLS	0.0068	0.1157	0.2245	0.4352	0.8794
CPU lsqnonneg	0.0223	0.4665	0.8959	1.7269	3.5243
Matlab lsqnonneg	0.0330	0.8096	1.5583	3.0097	6.1627
Matlab FNNLS [6]	0.0248	0.5902	1.1602	2.2920	4.6022
Matlab interior-points [16]	0.4480	10.7729	21.1767	41.9441	83.6695
Matlab PQN-NNLS [17]	1.5935	55.8196	110.0892	221.7277	441.7190

TABLE 7.3

*Runtime (seconds) comparisons of NNLS and lsqnonneg variants on multiple systems of random matrices  $A$  and random vectors  $b$ .*

**8. Concluding Remarks.** In this paper, we have presented an efficient procedure for solving least squares sub-problems in the active-set algorithm. We have shown that prior  $QR$  decompositions may be used to update and solve similar least squares sub-problems. Furthermore, a reordering of variables in the passive-set yielded fewer computations in the update step. This has lead to substantial speed-ups over existing methods in both the GPU and CPU implementations. Applications to satellite based terrain mapping and in microphone array signal processing are being worked on. Both GPU and CPU source codes are available on-line at [http:](http://)

[//www.cs.umd.edu/~yluo1/Projects/NNLS.html](http://www.cs.umd.edu/~yluo1/Projects/NNLS.html).

**9. Acknowledgements.** We would like to acknowledge James Blair and Michelle Hofton at NASA for providing us with data for the deconvolution problem, NVIDIA and NSF award 0403313 for facilities, and ONR award N00014-08-10638 for support.

#### REFERENCES

- [1] S. BELLAVIA, M. MACCONI, AND B. MORINI, *An interior point Newton-like method for non-negative least squares problems with degenerate solution*, Numerical Linear Algebra with Applications, 13 (2006), pp. 825–846.
- [2] M. H. VAN BENTHEM, M. R. KEENAN, *Fast algorithm for the solution of large-scale non-negativity-constrained least squares problems*. Journal of Chemometrics, Vol. 18, (2004), pp. 441–450.
- [3] A. BJÖRCK, *Stability Analysis of the method of semi-normal equations for least squares problems*, Linear Algebra Appl., 88/89 (1987), pp. 31–48.
- [4] G. E. BLELLOCH, *Prefix Sums and Their Applications*, Technical Report CMU-CS-90-190, Carnegie Mellon University, Pittsburgh, PA, Nov. (1990).
- [5] A. W. BOJANCZYK, R. P. BRENT AND F. R. DE HOOG, *QR factorization of toeplitz matrices*, Numerische Mathematik, 49 (1986), pp. 81–94.
- [6] R. BRO, S. D. JONG, *A fast non-negativity-constrained least squares algorithm*, Journal of Chemometrics, Vol. 11, No. 5, (1997), pp. 393–401.
- [7] M. CATRAL, L. HAN, M. NEUMANN, AND R. PLEMMONS, *On reduced rank nonnegative matrix factorization for symmetric nonnegative matrices*, Lin. Alg. Appl., 393 (2004), pp. 107–126.
- [8] R. H. CHAN, J. G. NAGY, AND R. J. PLEMMONS, *FFT-based preconditioners for toeplitz-block least squares problems*. SIAM J. Numer. Anal. 30 (1993), pp. 1740–1768.
- [9] D. CHEN, R. J. PLEMMONS, *Nonnegativity Constraints in Numerical Analysis*, Symp on the Birth of Numerical Analysis, Leuven, Belgium, (2007).
- [10] V. FRANC, V. HLAVC, AND M. NAVARA, *Sequential coordinate-wise algorithm for non-negative least squares problem*, Research report CTU-CMP-2005-06, Center for Machine Perception, Czech Technical University, Prague, Czech Republic, (2005).
- [11] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Third ed., The Johns Hopkins University Press, Baltimore, MD, 1996, pp. 223–236.
- [12] S. HAMMARLING, C. LUCAS, *Updating the QR factorization and the least squares problem*, MIMS EPrint, Manchester Institute for Mathematical Sciences, University of Manchester, Manchester, (2008).
- [13] M. HARRIS, *Optimizing parallel reduction in CUDA*, (2007).
- [14] INTEL, *Math kernel library reference manual*, (2010).
- [15] A. KERR, D. CAMPBELL, AND M. RICHARDS, *QR decomposition on GPUs, GPGPU-2: Proceedings of 2nd Workshop on General Purpose Processing on Graphics Processing Units*. New York, NY, USA: ACM, (2009), pp. 71–78.
- [16] S. J. KIM, K. KOH, M. LUSTIG, S. BOYD, AND D. GORINEVSKY, *An interior-point method for large-scale  $l_1$ -regularized least squares*, IEEE Journal of Selected Topics in Signal Processing 1, no. 4, 606617 (2007).
- [17] D. KIM, S. SRA, AND I. S. DHILLON, *A new projected quasi-Newton Approach for the non-negative least squares problem*, Technical Report TR-06-54, Computer Sciences, The University of Texas at Austin, (2006).
- [18] H. W. KUHN AND A. W. TUCKER, *Nonlinear programming*, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, Berkeley, CA, (1951), pp. 481–492.
- [19] C. L. LAWSON AND R. J. HANSON, *Solving least squares Problems*, PrenticeHall, 1987.
- [20] NVIDIA, *CUDA programming guide 3.2*, (2011).
- [21] NVIDIA, *OpenCL programming guide for CUDA architectures 3.1*, (2010).
- [22] SUN MICROSYSTEMS, INC, *OpenMP API user guide*, (2003).
- [23] V. VOLKOV AND J. W. DEMMEL, *Benchmarking GPUs to tune dense linear algebra*, SC08, (2008).