

Hierarchical Parallel Matrix Multiplication on Large-Scale Distributed Memory Platforms

Jean-Noël Quintin
Extrem Computing R&D
Bull
France
jean-noel.quintin@bull.net

Khalid Hasanov
University College Dublin
Dublin 4, Belfield
Ireland
khalid.hasanov@ucdconnect.ie

Alexey Lastovetsky
University College Dublin
Dublin 4, Belfield
Ireland
Alexey.Lastovetsky@ucd.ie

Abstract—Matrix multiplication is a very important computation kernel both in its own right as a building block of many scientific applications and as a popular representative for other scientific applications.

Cannon's algorithm which dates back to 1969 was the first efficient algorithm for parallel matrix multiplication providing theoretically optimal communication cost. However this algorithm requires a square number of processors. In the mid-1990s, the SUMMA algorithm was introduced. SUMMA overcomes the shortcomings of Cannon's algorithm as it can be used on a non-square number of processors as well. Since then the number of processors in HPC platforms has increased by two orders of magnitude making the contribution of communication in the overall execution time more significant. Therefore, the state of the art parallel matrix multiplication algorithms should be revisited to reduce the communication cost further.

This paper introduces a new parallel matrix multiplication algorithm, Hierarchical SUMMA (HSUMMA), which is a redesign of SUMMA. Our algorithm reduces the communication cost of SUMMA by introducing a two-level virtual hierarchy into the two-dimensional arrangement of processors. Experiments on an IBM BlueGene/P demonstrate the reduction of communication cost up to 2.08 times on 2048 cores and up to 5.89 times on 16384 cores.

I. INTRODUCTION

Matrix multiplication is a very important computation kernel both in its own right as a building block of many scientific applications and as a popular representative for other scientific applications.

Cannon's algorithm [1] which was introduced in 1967 was the first efficient algorithm for parallel matrix multiplication providing theoretically optimal communication cost. However this algorithm requires a square number of processors which makes it impossible to be used in a general purpose library. Later introduced Fox's algorithm [2] has the same restriction.

The 3D algorithm [3] which dates back to the 1990s organizes the p processors as $p^{\frac{1}{3}} \times p^{\frac{1}{3}} \times p^{\frac{1}{3}}$ 3D mesh and achieves a factor of $p^{\frac{1}{6}}$ less communication cost than 2D parallel matrix multiplication algorithms. However, in order to get this improvement the 3D algorithm requires $p^{\frac{1}{3}}$ extra copies of the matrices. That means that on one million cores the 3D algorithm will require 100 extra copies of the matrices which would be a significant problem on some platforms. Therefore, the 3D algorithm is only practical for relatively small matrices.

One implementation of Fox's algorithm on a general $P \times Q$ processor grid is PUMMA [4] which was designed for block-cyclic distributed matrices. PUMMA consists of $Q - 1$ shifts for matrix A , $LCM(P, Q)$ broadcasts for matrix B and the number of local multiplications is $LCM(P, Q)$. Here $LCM(P, Q)$ is the least common multiple of P and Q . The main shortcomings of PUMMA come from the fact that it always tries to use the largest possible matrices for both computation and communication. In this case, large memory space is required to store them temporarily, the effect of the block size is marginal and the most important it is difficult to overlap computation with communication.

In the mid-1990s SUMMA [5] was introduced. Like PUMMA, SUMMA was designed for a general $P \times Q$ processor grid. Unlike PUMMA it does not require the largest possible matrices for computation and communication and therefore allows to pipeline them. In addition, SUMMA was implemented in practice in ScaLAPACK [6]: the most popular parallel numerical linear algebra package.

Recently introduced 2.5D algorithm [7] generalizes the 3D algorithm by parametrizing the extent of the third dimension of the processor arrangement: $\frac{p}{c} \times \frac{p}{c} \times c$, $c \in [1, p^{\frac{1}{3}}]$. However, the 2.5D algorithm is efficient only if there is free amount of extra memory to store c copies of the matrices. On the other hand, it is expected that exascale systems will have a dramatically shrinking memory space per core [8]. Therefore, the 2.5D algorithm can not be scalable on the future exascale systems.

Matrix multiplication is a problem known to be very compute intensive. On the other hand, as HPC moves towards exascale, the cost of matrix multiplication will be dominated by communication cost. Therefore, the state of the art parallel matrix multiplication algorithms should be revisited to reduce the communication cost further.

The contributions of this paper are as follows:

- We introduce a new design to parallel matrix multiplication algorithm by introducing a two-level virtual hierarchy into the two-dimensional arrangement of processors. We apply our approach to SUMMA which is a state of the art algorithm. We call our algorithm hierarchical SUMMA(HSUMMA).

- We model the performance of SUMMA and HSUMMA and theoretically prove that HSUMMA reduces the communication cost of SUMMA. Then we provide experimental results on a cluster of Grid5000 and BlueGene/P which are reasonable representative and span a good spectrum of loosely and tightly coupled platforms. We use SUMMA as the only competitor to our algorithm because it is the most general and scalable parallel matrix multiplication algorithm, which decreases its per-processor memory footprint with the increase of the number of processors for a given problem size, and is used in the most famous parallel numerical linear algebra packages such as ScaLAPACK. In addition, because of its practicality SUMMA plays a starting point to implement parallel matrix multiplications on specific platforms. As a matter of fact, the most used parallel matrix multiplication algorithm for heterogeneous platforms [9] [10] was based on SUMMA as well. Therefore, despite it was introduced in the mid-1990s SUMMA is still a state of the art algorithm.

II. PREVIOUS WORK

In this section, we detail SUMMA algorithm which is the motivating work for our algorithm. Then, we outline and discuss the existing broadcast algorithms which can be used inside SUMMA to improve its communication cost.

A. SUMMA algorithm

SUMMA [5] implements the matrix multiplication $C = A \times B$ over a two-dimensional $p = s \times t$ processor grid. For simplicity, let us assume that the matrices are square $n \times n$ matrices. These matrices are distributed over the processor grid by block-distribution.

We can see the size of the matrices as $\frac{n}{b} \times \frac{n}{b}$ by introducing a block of size b . Then each element in A , B , and C is a square $b \times b$ block, and the unit of computation is the updating of one block, that is, a matrix multiplication of size b . For simplicity, we assume that n is a multiple of b . The algorithm can be formulated as follows: The algorithm consists of $\frac{n}{b}$ steps. At each step

- Each processor holding part of the pivot column of the matrix A horizontally broadcasts its part of the pivot column along processor row.
- Each processor holding part of the pivot row of the matrix B vertically broadcasts its part of the pivot row along processor column.
- Each processor updates each block in its C rectangle with one block from the pivot column and one block from the pivot row, so that each block c_{ij} , $(i, j) \in (1, \dots, \frac{n}{b})$ of matrix C will be updated as $c_{ij} = c_{ij} + a_{ik} \times b_{kj}$.
- After $\frac{n}{b}$ steps of the algorithm, each block c_{ij} of matrix

$$C \text{ will be equal to } \sum_{k=1}^{\frac{n}{b}} a_{ik} \times b_{kj}$$

Figure 1 shows the communication pattern at the third step with one block per processor.

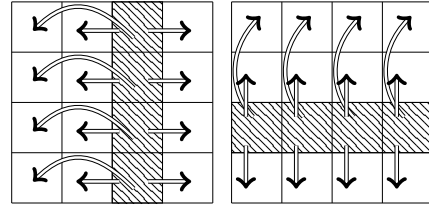


Fig. 1. Communication Pattern of SUMMA for the third step with four processors and one block per processor

B. MPI Broadcast Algorithms

Collective communications are fundamental tools to parallelize matrix multiplication algorithms. We have already seen that the communication pattern of SUMMA is based on broadcast and an improvement in the broadcast algorithm can improve the communication cost of SUMMA as well. Therefore it is worth to outline existing broadcast algorithms.

A lot of research have been done into MPI [11] collective communications and especially into MPI broadcast algorithms [12] [13] [14]. Early implementations of broadcast algorithms assumed homogeneous and fully connected networks. They are based on simple binary or binomial trees. On the other hand, some algorithms have been introduced in order to be more effective for large message sizes and use the benefits of hierarchical networks by using pipelined trees or recursive halving algorithms [12]. However, neither the broadcast APIs nor numerous broadcast algorithms are application specific, and most of the time improvements come from platform parameters and very often they are for specific architectures, such as mesh, hypercube and fat tree [15].

In the next section we introduce hierarchical SUMMA algorithm. Our algorithm is neither an improvement of an existing broadcast algorithm nor a new broadcast algorithm. HSUMMA is an application specific but platform independent hierarchical matrix multiplication algorithm which reduces communication cost of SUMMA and allows better overlapping of communications and computation. Indeed, HSUMMA can use any of the existing optimized broadcast algorithms and still reduce the communication cost of SUMMA as demonstrated in Section IV-C.

III. HIERARCHICAL SUMMA

Let us assume we have $p = s \times t$ processors distributed over the same two-dimensional virtual processor grid as in SUMMA, the matrices are square $n \times n$ matrices, b is the block size. The distribution of the matrices is the same as in SUMMA. HSUMMA partitions the virtual $s \times t$ processor grid into a higher level $I \times J$ arrangement of rectangular groups of processors, so that inside each group there is a two-dimensional $\frac{s}{I} \times \frac{t}{J}$ grid of processors. Figure 2 gives an example of such two-level hierarchical arrangement of

processors. In this example a 6×6 grid of processors is arranged into two-level 3×3 grids of groups and 2×2 grid of processors inside a group.

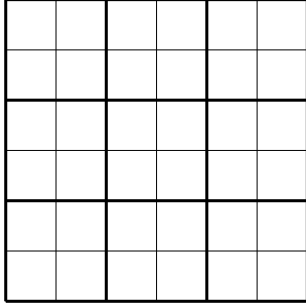


Fig. 2. Hierarchical platform as nine groups, four processors per group

Let $P_{(x,y)(i,j)}$ denote the processor (i,j) inside the group (x,y) . HSUMMA splits the communication phase of the SUMMA algorithm into two phases and consists of $\frac{n}{b}$ steps. The algorithm can be summarized as follows:

- Horizontal broadcast of the pivot column of the matrix A is performed as follows:
 - First, each processor $P_{(k,y)(i,j)}$, $k \in (1, \dots, I)$ holding part of the pivot column of the matrix A horizontally broadcasts its part of the pivot column to the processors $P_{(k,z)(i,j)}$, $z \neq y$, $z \in (1, \dots, I)$ in the other groups.
 - Now, inside each group (x,y) processor $P_{(x,y)(i,j)}$ has the required part of the pivot column of the matrix A and it further horizontally broadcasts it to the processors $P_{(x,y)(i,c)}$, $c \neq j$, $c \in (1, \dots, \frac{s}{I})$ inside the group.
- Vertical broadcast of the pivot row of the matrix B is performed as follows:
 - First, each processor $P_{(x,k)(i,j)}$, $k \in (1, \dots, I)$ holding part of the pivot row of the matrix B vertically broadcasts its part of the pivot row to the processors $P_{(z,k)(i,j)}$, $z \neq x$, $z \in (1, \dots, I)$ in the other groups.
 - Now, inside each group (x,y) processor $P_{(x,y)(i,j)}$ has the required part of the pivot row of the matrix B and it further vertically broadcast it to the processors $P_{(x,y)(r,j)}$, $r \neq i$, $r \in (1, \dots, \frac{t}{J})$ inside the group.
- Each processor inside a group updates each block in its C rectangle with one block from the pivot column and one block from the pivot row, so that each block c_{ij} , $(i,j) \in (1, \dots, \frac{n}{b})$ of matrix C will be updated as $c_{ij} = c_{ij} + a_{ik} \times b_{kj}$.
- After $\frac{n}{b}$ steps of the algorithm, each block c_{ij} of matrix

$$C \text{ will be equal to } \sum_{k=1}^{\frac{n}{b}} a_{ik} \times b_{kj}$$

The communication phases described above are illustrated by Figure 3. In general it is possible to use one block size

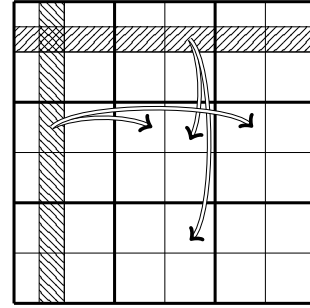


Fig. 3. HSUMMA between groups

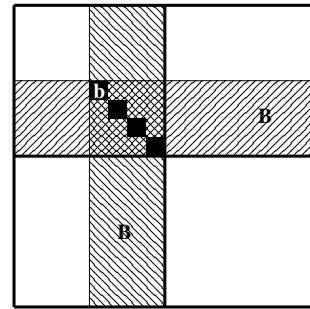


Fig. 4. HSUMMA inside group

between groups and another block size inside a group. In this case the size of sent data between the groups is at least the same as the size of data sent inside a group. Therefore, the block size inside a group should be less than or equal to the block size between groups. Let us assume the block size between groups is B and inside a group is b . Then, the number of steps in the higher level will be equal to the number of blocks between groups: $\frac{n}{B}$. In each iteration between the

groups, the number of steps inside a group is $\frac{B}{b}$, so the total number of steps of HSUMMA, $\frac{n}{B} \times \frac{B}{b}$, will be the same as the number of steps of SUMMA. The amount of data sent is the same as in SUMMA. The steps inside a group are shown in Figure 4. It is clear that SUMMA is a special case of HSUMMA when the number of groups equals to one or to the total number of processors.

One may ask why to introduce a new hierarchical algorithm if MPI implementations already provide us with high-performance broadcast algorithms. As we have already mentioned before, MPI optimizations of broadcast are platform specific and do not depend on the application. On the other hand, HSUMMA is an application specific hierarchical algorithm which optimizes communication cost at the application level and is platform independent. A general purpose broadcast algorithm can not replace the communication pattern of HSUMMA as in each level it calculates pivot row and pivot column before broadcasting and it is very application specific.

The pseudocode for HSUMMA is Algorithm 1.

Algorithm 1: Hierarchical SUMMA algorithm

```

/*The A,B,C matrices are distributed on a
virtual 2-D grid of  $p = s \times t$  processors.
Here are the instructions executed by the
processor  $P_{(x,y)}(i,j)$  (this is the processor
 $(i,j)$  inside the group  $(x,y)$ ).*/
Data:  $NB_{Block\_Group}$ : Number of steps in the higher level
Data:  $NB_{Block\_Inside}$ : Number of steps in the lower level
Data:  $(M, L, N)$ : Matrix dimensions
Data:  $A, B$ : two input sub-matrices of size  $(\frac{M}{s} \times \frac{L}{t}, \frac{L}{s} \times \frac{N}{t})$ 
Result:  $C$ : result sub-matrix of size  $\frac{M}{s} \times \frac{N}{t}$ 
begin
  /* Broadcast A and B */
  MPI_Comm group_col_comm /* communicator
between  $P_{(*,y)}(i,j)$  processors */
  MPI_Comm group_row_comm /* communicator
between  $P_{(x,*)}(i,j)$  processors */
  MPI_Comm col_comm /* communicator between
 $P_{(x,y)}(*,j)$  processors */
  MPI_Comm row_comm /* communicator between
 $P_{(x,y)}(i,*)$  processors */

  for iter_group = 0; iter_group <  $NB_{Block\_Group}$ ; iter_group ++ do
    if i == Pivot_inside_group_col(iter_group) then
      if x == Pivot_group_col(iter_group) then
        /* Get direct access to the
iter_groupth group block of A */
        Copy_Block_group( Block_group_A, A, iter_group )
        MPI_Bcast( Block_group_A, TypeBlock_group_A,
Pivot_group_col(iter_group), group_row_comm )

      if j == Pivot_inside_group_row(iter_group) then
        if y == Pivot_group_row(iter_group) then
          /* Get direct access to the
iter_groupth group block of B */
          Copy_Block_group( Block_group_B, B, iter_group )
          MPI_Bcast( Block_group_B, TypeBlock_group_B,
Pivot_group_row(iter_group), group_col_comm )

    for iter = 0; iter <  $NB_{Block\_Inside}$ ; iter ++ do
      if i == Pivot_inside_group_col(iter) then
        /* Get access to the iterth
block of Block_group_A on this
processor */
        Copy_Block_A( Block_A, Block_group_A, iter )
        MPI_Bcast( Block_A, TypeBlock_A,
Pivot_col(iter), row_comm )

      if j == Pivot_inside_group_row(iter) then
        /* Get access to the iterth
block of Block_group_B on this
processor */
        Copy_Block_B( Block_B, Block_group_B, iter )
        MPI_Bcast( Block_B, TypeBlock_B,
Pivot_row(iter), col_comm )

      DGemm( Block_A, Block_B, C )

```

IV. THEORETICAL ANALYSIS

In this section SUMMA and HSUMMA are theoretically analysed and compared. First of all, for simplicity we assume that the matrices are $n \times n$ square matrices. Let b be block size inside one group and B be block size between groups. The execution time depends on the communication time (i.e. the broadcast algorithm and the communication model). As a communication model we use Hockney's model [16] which represents the time of sending of a message of size m between two processors as $\alpha + m\beta$. Here, α is the latency, β is the reciprocal of network bandwidth. In addition, let us assume that a combined floating point computation (for one addition and multiplication) time is γ . Binomial tree and Van de Geijn broadcast algorithms [13] [17] are used to analyse both our algorithm and SUMMA. It is known that the costs of these broadcast algorithms are as follows:

- Binomial tree algorithm: $\log_2(p) \times (\alpha + m \times \beta)$
- Van de Geijn algorithm: $(\log_2(p) + p - 1)\alpha + 2^{\frac{p-1}{p}}m\beta$

A. Analysis of SUMMA

For simplicity let us assume the $n \times n$ matrices are distributed over a two-dimensional $\sqrt{p} \times \sqrt{p}$ grid of processors and the block size is b . This algorithm has $\frac{n}{b}$ steps. In each step, processors broadcast pivot row and pivot column. So the computation cost in each step is $O(2 \times \frac{n^2}{p} \times b)$. Hence, the overall computation cost will be $O(\frac{2n^3}{p})$.

The broadcasts of each row and column are independent at each step and they can be done in parallel. For this analysis the network congestion is neglected. The amount of data transferred by each broadcast is $\frac{n}{\sqrt{p}} \times b$. The total communication cost of SUMMA can be computed by multiplying the communication cost of each step by the number of steps depending on the broadcast algorithm. The results are:

- Binomial Tree:
 $\log_2(p) \times \left(\alpha \frac{n}{b} + \beta \times \frac{n^2}{\sqrt{p}} \right)$
- Van de Geijn broadcast:
 $(\log_2(p) + 2(\sqrt{p} - 1))\alpha \frac{n}{b} + 4\left(1 - \frac{1}{\sqrt{p}}\right) \frac{n^2}{\sqrt{p}}\beta$

B. Analysis of HSUMMA

To simplify the analysis, let us assume there are G groups arranged as $\sqrt{G} \times \sqrt{G}$ grid of processors groups. Let B denote the block size between groups (we also call such a block an outer block), b be block size inside a group, and $n \times n$ be the size of the matrices.

There is one step per outer block, thus there will be $\frac{n}{B}$ steps in the highest level called outer steps. Each outer block belongs to \sqrt{p} processors. These processors broadcast the part of the outer block along the row or column of processor in parallel. The size of sent data is $2 \frac{n \times B}{\sqrt{p}}$ per processor which has a part of the outer block.

Inside one group, processors are arranged in a grid of size: $\frac{\sqrt{p}}{\sqrt{G}} \times \frac{\sqrt{p}}{\sqrt{G}}$. After the reception of the outer block each group multiplies the outer block using the SUMMA algorithm inside the group. Thus there are $\frac{B}{b}$ inner steps to execute. The inner block belongs to $\frac{\sqrt{p}}{\sqrt{G}}$ processors. These processors send $2 \frac{n \times b}{\sqrt{p}}$ amount of data per inner step.

The overall communication time is equal to the sum of the communication times between the groups and inside the groups.

- Inner Communication cost (inside each group):
 - Binomial Tree: $\log_2 \left(\frac{p}{G} \right) \times \left(\alpha \times \frac{n}{b} + \beta \times \frac{n^2}{\sqrt{p}} \right)$
 - Van de Geijn broadcast: $\left(\log_2 \left(\frac{p}{G} \right) + 2 \left(\frac{\sqrt{p}}{\sqrt{G}} - 1 \right) \right) \times \alpha \times \frac{n}{b} + 4 \left(1 - \frac{\sqrt{G}}{\sqrt{p}} \right) \times \frac{n^2}{\sqrt{p}} \beta$
- Outer Communication cost (between groups):
 - Binomial Tree: $\log_2(G) \times \left(\alpha \times \frac{n}{B} + \beta \times \frac{n^2}{\sqrt{p}} \right)$
 - Van de Geijn broadcast: $\left(\log_2(G) + 2 \left(\sqrt{G} - 1 \right) \right) \times \alpha \times \frac{n}{B} + 4 \left(1 - \frac{1}{\sqrt{G}} \right) \times \frac{n^2}{\sqrt{p}} \beta$

We can summarize the cost of HSUMMA and SUMMA as in Table I and Table II.

TABLE I
COMPARISON WITH BINOMIAL TREE BROADCAST

Algorithm	Comp. Cost	Latency Factor		Bandwidth Factor	
		inside groups	between groups	inside groups	between groups
SUMMA	$\frac{2n^3}{p}$	$\log_2(p) \times \frac{n}{b}$		$n^2 \times \frac{\log_2(p)}{\sqrt{p}}$	
HSUMMA	$\frac{2n^3}{p}$	$\log_2 \left(\frac{p}{G} \right) \times \frac{n}{b}$	$\log_2(G) \times \frac{n}{B}$	$\log_2 \left(\frac{p}{G} \right) \times \frac{n^2}{\sqrt{p}}$	$\log_2(G) \times \frac{n^2}{\sqrt{p}}$

C. Theoretical Prediction

One of the goals of this section is to demonstrate that independent on the broadcast algorithm employed by SUMMA, HSUMMA will either outperform SUMMA or be at least equally fast. In this section we introduce a general model for broadcast algorithms and theoretically predict SUMMA and HSUMMA. In the model we assume no contention and assume all the links are homogeneous. We prove that even with this simple model the extremums of the communication cost function can be predicted.

Again we assume that the time taken to send a message of size m between any two processors is modeled as $T(m) = \alpha + m \times \beta$, where α is the latency and β is the reciprocal bandwidth.

We model a broadcast time for a message of size m among p processors by formula (1). This model generalizes all homogeneous broadcast algorithms such as pipelined/non-pipelined flat, binary, binomial, linear, scatter/gather broadcast algorithms [18] [19] which are used inside state of the art broadcast implementations like MPICH and Open MPI.

$$T_{bcast}(m, p) = L(p) \times \alpha + m \times W(p) \times \beta \quad (1)$$

In (1) we assume that $L(1) = 0$ and $W(1) = 0$. It is also assumed that $L(p)$ and $W(p)$ are monotonic and differentiable functions in the interval $(1, p)$ and their first derivatives are constants or monotonic in the interval $(1, p)$.

With this general model the theoretical communication cost of SUMMA will be as follows:

$$T_S(n, p) = 2 \left(\frac{n}{b} \times L(\sqrt{p}) \alpha + \frac{n^2}{\sqrt{p}} \times W(\sqrt{p}) \beta \right) \quad (2)$$

In the same way we can express the communication cost of HSUMMA as the sum of the latency cost and the bandwidth cost:

$$T_{HS}(n, p, G) = T_{HS_l}(n, p, G) + T_{HS_b}(n, p, G) \quad (3)$$

Here $G \in [1, p]$ and we take $b = B$ for simplicity. The latency cost $T_{HS_l}(n, p, G)$ and the bandwidth cost $T_{HS_b}(n, p, G)$ will be given by the following formulas:

$$T_{HS_l}(n, p, G) = 2 \frac{n}{b} \times \left(L(\sqrt{G}) + L\left(\frac{\sqrt{p}}{\sqrt{G}}\right) \right) \alpha \quad (4)$$

$$T_{HS_b}(n, p, G) = 2 \frac{n^2}{\sqrt{p}} \times \left(W(\sqrt{G}) + W\left(\frac{\sqrt{p}}{\sqrt{G}}\right) \right) \beta \quad (5)$$

It is clear that $T_S(n, p)$ is a special case of $T_{HS}(n, p, G)$ when $G = 1$ or $G = p$.

Let us investigate extremums of T_{HS} as a function of G for a fixed p and n . We have $b = B$.

$$\frac{\partial T_{HS}}{\partial G} = \frac{n}{b} \times L_1(p, G) \alpha + \frac{n^2}{\sqrt{p}} \times W_1(p, G) \beta \quad (6)$$

Here, $L_1(p, G)$ and $W_1(p, G)$ are defined as follows:

$$L_1(p, G) = \left(\frac{\partial L(\sqrt{G})}{\partial \sqrt{G}} \times \frac{1}{\sqrt{G}} - \frac{\partial L\left(\frac{\sqrt{p}}{\sqrt{G}}\right)}{\partial \frac{\sqrt{p}}{\sqrt{G}}} \times \frac{\sqrt{p}}{G\sqrt{G}} \right) \quad (7)$$

$$W_1(p, G) = \left(\frac{\partial W(\sqrt{G})}{\partial \sqrt{G}} \times \frac{1}{\sqrt{G}} - \frac{\partial W\left(\frac{\sqrt{p}}{\sqrt{G}}\right)}{\partial \frac{\sqrt{p}}{\sqrt{G}}} \times \frac{\sqrt{p}}{G\sqrt{G}} \right) \quad (8)$$

It can be easily shown that, if $G = \sqrt{p}$ then $L_1(p, G) = 0$ and $W_1(p, G) = 0$, thus, $\frac{\partial T_{HS}}{\partial G} = 0$. In addition, $\frac{\partial T_{HS}}{\partial G}$ changes the sign in the interval $(1, p)$ depending on the value of G . That means $T_{HS}(n, p, G)$ has extremum at $G = \sqrt{p}$ for fixed n and p . The expression of $\frac{\partial T_{HS}}{\partial G}$ shows that, depending on the ratio of α and β the extremum can be either minimum or maximum in the interval $(1, p)$. If $G = \sqrt{p}$ is the minimum point it means that with $G = \sqrt{p}$ HSUMMA will outperform

TABLE II
COMPARISON WITH VAN DE GEIJN BROADCAST

Algorithm	Comp. Cost	Latency Factor		Bandwidth Factor	
		inside groups	between groups	inside groups	between groups
SUMMA	$\frac{2n^3}{p}$	$(\log_2(p) + 2(\sqrt{p} - 1)) \times \frac{n}{b}$		$4\left(1 - \frac{1}{\sqrt{p}}\right) \times \frac{n^2}{\sqrt{p}}$	
HSUMMA	$\frac{2n^3}{p}$	$\left(\log_2\left(\frac{p}{G}\right) + 2\left(\frac{\sqrt{p}}{\sqrt{G}} - 1\right)\right) \times \frac{n}{b}$	$(\log_2(G) + 2(\sqrt{G} - 1)) \times \frac{n}{B}$	$4\left(1 - \frac{\sqrt{G}}{\sqrt{p}}\right) \times \frac{n^2}{\sqrt{p}}$	$4\left(1 - \frac{1}{\sqrt{G}}\right) \times \frac{n^2}{\sqrt{p}}$
HSUMMA($G = \sqrt{p}, b = B$)	$\frac{2n^3}{p}$	$(\log_2(p) + 4(\sqrt[4]{p} - 1)) \times \frac{n}{b}$		$8\left(1 - \frac{1}{\sqrt[4]{p}}\right) \times \frac{n^2}{\sqrt{p}}$	

SUMMA, otherwise HSUMMA with $G = 1$ or $G = p$ will have the same performance as SUMMA.

Now lets apply this analysis to the HSUMMA communication cost function obtained for Van de Geijn broadcast algorithm (see Table II) and again assuming $b = B$ for simplicity. We will have:

$$\frac{\partial T_{HSV}}{\partial G} = \frac{G - \sqrt{p}}{G\sqrt{G}} \times \left(\frac{n\alpha}{b} - 2\frac{n^2}{p} \times \beta\right) \quad (9)$$

It is clear that if $G = \sqrt{p}$ then $\frac{\partial T_{HSV}}{\partial G} = 0$. Depending on the ratio of α and β , the communication cost as a function of G has either minimum or maximum in the interval $(1, p)$.

- If

$$\frac{\alpha}{\beta} > 2\frac{nb}{p} \quad (10)$$

then $\frac{\partial T_{HSV}}{\partial G} < 0$ in the interval $(1, \sqrt{p})$ and $\frac{\partial T_{HSV}}{\partial G} > 0$ in (\sqrt{p}, p) . Thus T_{HS} has the minimum in the interval $(1, p)$ and the minimum point is $G = \sqrt{p}$.

- If

$$\frac{\alpha}{\beta} < 2\frac{nb}{p} \quad (11)$$

then T_{HS} has the maximum in the interval $(1, p)$ and the maximum point is $G = \sqrt{p}$. The function gets its minimum at either $G = 1$ or $G = p$.

If we take $G = \sqrt{p}$ in the HSUMMA communication cost function (see Table II) and assume the above conditions the optimal communication cost function will be as follows:

$$(\log_2(p) + 4(\sqrt[4]{p} - 1)) \times \frac{n}{b} \times \alpha + 8\left(1 - \frac{1}{\sqrt[4]{p}}\right) \times \frac{n^2}{\sqrt{p}} \times \beta \quad (12)$$

Thus, we have proved that depending on the ratio of α and β HSUMMA will either reduce the communication cost of SUMMA or in the worst case have the same performance as SUMMA.

We will use this model to predict the performance of HSUMMA on Grid5000, BlueGene/P and future exascale platforms.

V. EXPERIMENTS

Our experiments were carried out on a cluster of Grid5000 and a BlueGene/P (BG/P) platform which are fairly representative and span a good spectrum of loosely and tightly coupled platforms. The details of the platforms are given in the appropriate sections. The times in our experimental results are the mean times of 30 experiments.

A. Experiments on Grid5000

Some of our experiments were carried out on the Graphene cluster of Nancy site of Grid5000 platform. We have used Intel MKL BLAS for sequential operations, MPICH-2 for MPI implementation and our implementations of the matrix multiplication algorithms. In addition to MPICH we also did some experiments with Open MPI on Grid5000 and got similar results. Thus in this paper we just present the experiments with MPICH implementation of MPI. The size of matrices in our experiments on Grid5000 is 8192×8192 . Figure 5 compares SUMMA and HSUMMA with block size 64. It is clear that smaller block sizes lead to a larger number of steps and this in turn will affect the latency cost. It can be seen that in this case HSUMMA outperforms SUMMA with huge difference.

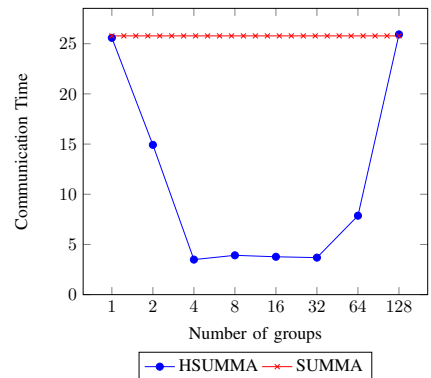


Fig. 5. HSUMMA on Grid5000. Communication time in seconds. $b = B = 64, n = 8192$ and $p = 128$

Figure 6 represents the same comparison but with block size 512. This block size is the maximum possible one with this configuration. In this case the improvement is 1.6 times as the minimum communication time of HSUMMA and SUMMA are 2.81 and 4.53 seconds respectively. In addition, theoretically HSUMMA should have the same performance as SUMMA when $G = 1$ or $G = p$ and the figures verify that in practice. That means HSUMMA can never be worse than SUMMA. In the worst case it will have the same performance as SUMMA.

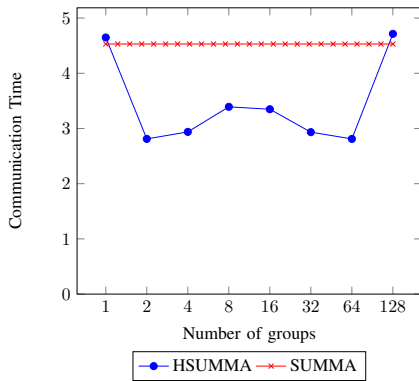


Fig. 6. HSUMMA on Grid5000. Communication time in seconds. $b = B = 512$, $n = 8192$ and $p = 128$

Figure 7 shows experimental results from scalability point of view. Here we use the largest possible block size for both algorithms. If we used block size 64 for scalability plot we would see the significant difference between HSUMMA and SUMMA. However, even with this configuration which is optimal for SUMMA it can be seen that on small platforms both SUMMA and HSUMMA have the same performance, however, the trend shows that on larger platforms HSUMMA will outperform SUMMA and therefore is more scalable.

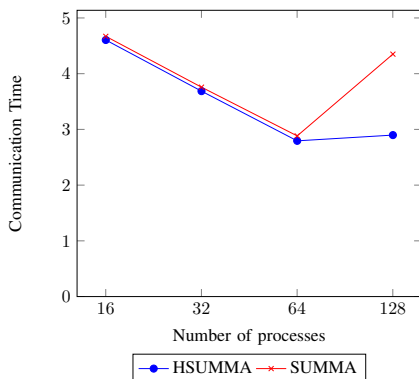


Fig. 7. HSUMMA and SUMMA on Grid5000. Communication time in seconds. $b = B = 512$ and $n = 8192$

The experiments show that with any number of groups HSUMMA outperforms SUMMA on Grid5000.

1) *Validation of the Analytical Model on Grid5000:* We take the following approximately real parameters for Graphene cluster of Grid5000:

- Latency: 1E-4
- Reciprocal bandwidth: 1E-9
- p : 8192
- n : 8192
- b : 64

The algorithmic parameters are the same as in our experiments on Grid5000. It is clear that $\frac{\alpha}{\beta} > 2 * \frac{8192 * 64}{128} = 8192$ and therefore according to our theoretical analysis HSUMMA has minimum in the interval $(1, p)$. We do not have experimental minimum exactly at $G = \sqrt{p}$ as predicted by our theoretical results. However, this does not downgrade the importance of our analytical model because the main goal of our analytical analysis is to predict if HSUMMA will be more efficient than SUMMA on the target platform or not. If this is the case, the optimal number of groups can be easily found experimentally by using only few iterations of HSUMMA with different values of G and thus can be incorporated into the algorithm.

B. Experiments on BlueGene/P

Some of our experiments were carried out on Shaheen BlueGene/P at Supercomputing Laboratory at King Abdulah University of Science&Technology (KAUST) in Thuwal, Saudi Arabia. Shaheen is a 16-rack BlueGene/P. Each node is equipped with four 32-bit, 850 Mhz PowerPC 450 cores and 4GB DDR memory. VN (Virtual Node) mode with torus connection was used for the experiments on the BG/P. The Blue Gene/P architecture provides a three-dimensional point-to-point Blue Gene/P torus network which interconnects all compute nodes and global networks for collective and interrupt operations. Use of this network is integrated into the BlueGene/P MPI implementation.

All the sequential computations in our experiments were performed by using DGEMM routine from IBM ESSL library. The size of the matrices for all our experiments on the BG/P is 65536×65536 . We use our implementation of SUMMA for comparison with HSUMMA as the performance of ScaLAPACK implementation lingers behind our implementation.

The benefit of HSUMMA comes from the optimal number of groups. Therefore, it is interesting to see how different numbers of groups affect the communication cost of HSUMMA on a large platform. Figure 8 shows HSUMMA on 16384 cores. In order to have a fair comparison we use the same block size inside a group and between the groups. The figure shows that the execution time of SUMMA is 50.2 seconds and the communication time is 36.46 seconds. On the other hand, the minimum execution time of HSUMMA is 21.26 and the minimum communication time is 6.19 when $G = 512$. Thus the execution time of HSUMMA is 2.36 times and the communication time is 5.89 times less than that of SUMMA on 16384 cores. On the other hand, HSUMMA achieves 2.08 times less communication time and 1.2 times less overall execution time than SUMMA on 2048 cores. We also did

experiments on BlueGene/P cores smaller than 2048 and the results showed that on smaller numbers of cores the performance of HSUMMA and SUMMA was almost the same.

The "zigzags" on the figure can be explained by the fact that mapping communication layouts to network hardware on BlueGene/P impacts the communication performance and it was observed by P. Balaji et al. [20] as well. When we group processors we do not take into account the platform parameters. However, according to our preliminary observations these "zigzags" can be eliminated by taking platform parameters into account while grouping. In addition, the effects of square versus non-square meshes also a reason for that.

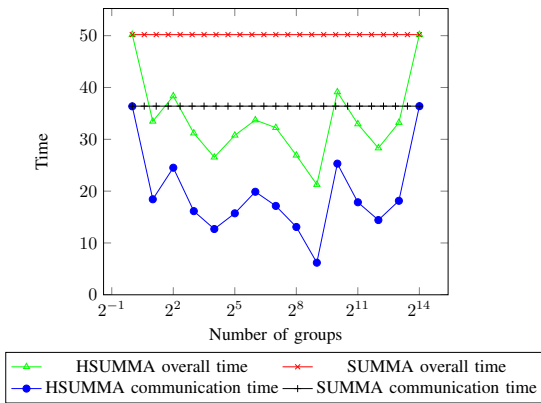


Fig. 8. SUMMA and HSUMMA on 16384 cores on BG/P. Execution and communication time. $b = B = 256$ and $n = 65536$

Figure 9 represents scalability analysis of SUMMA and HSUMMA from communication point of view. It can be seen that HSUMMA is more scalable than SUMMA and this pattern suggests that the communication performance of HSUMMA gets much better than that of SUMMA while the number of cores increases.

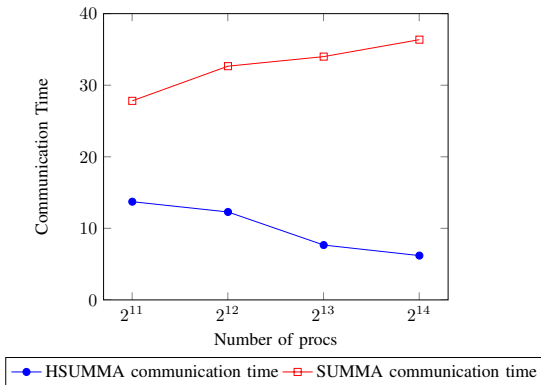


Fig. 9. HSUMMA and SUMMA on BlueGene/P with VN mode. Communication time in seconds. $b = B = 256$ and $n = 65536$

1) *Validation of the Analytical Model on BlueGene/P:* The parameters of the BlueGene/P are as follows:

- Latency: 3E-6

- Bandwidth: 1E-9
- p : 16384
- n : 65536
- b : 256

Here again we use the same values of the algorithmic parameters as in our experiments. By using these values it can be shown that $\frac{\alpha}{\beta} > 2\frac{nb}{p}$ which proves the communication function of HSUMMA has the minimum in the interval $(1, p)$. For some ratios of n and p the above condition may not hold. However, in this case the cost of matrix multiplication will be dominated by computation cost and even in this case HSUMMA can be used just by using one or p group.

C. Prediction on Exascale

We use the following parameters to predict performance of HSUMMA on exascale platforms. These platform parameters are obtained from a recent report on exascale architecture roadmap [8].

- Total flop rate (γ): 1E18 flops
- Latency: 500 ns
- Bandwidth: 100 GB/s
- Problem size: $n = 2^{22}$
- Number of processors: $p = 2^{20}$
- Block size: $b = 256$

Again we have $\frac{\alpha}{\beta} > 2\frac{nb}{p}$ which means HSUMMA can be efficient and outperform SUMMA on exascale platforms and the theoretical plot is shown in Figure 10.

These analyses show that with any realistic platform parameters HSUMMA reduces the communication cost of SUMMA. However, one of the useful features of HSUMMA is that in the worst case it can use just one or p group and have exactly the same performance as SUMMA.

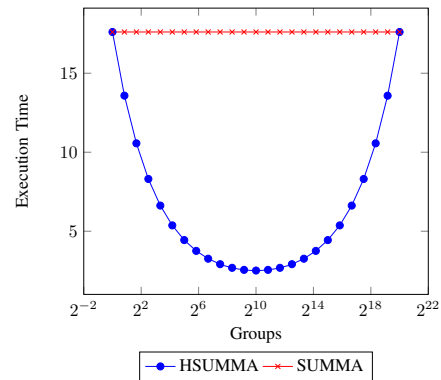


Fig. 10. Prediction of SUMMA and HSUMMA on Exascale. Execution time in seconds. $p = 1048576$

VI. CONCLUSIONS

We can conclude that our two-level hierarchical approach to parallel matrix multiplication significantly reduces the communication cost on large platforms such as BlueGene/P. Our experiments show that HSUMMA achieves 2.08 times less

communication time than SUMMA on 2048 cores and 5.89 times less communication cost on 16384 cores. Moreover, the overall execution time of HSUMMA is 1.2 times less than the overall execution time of SUMMA on 2048 cores, and 2.36 times less on 16384 cores. This trend suggests that, while the number of processors increases our algorithm will be more scalable than SUMMA. In addition, our experiments on Grid5000 show that our algorithm can be effective on small platforms as well. All these results prove that whatever stand-alone application-oblivious optimized broadcast algorithms are made available for exascale platforms, they cannot replace application specific optimizations of communication cost.

At the moment, we select the optimal number of groups sampling over valid values. However, it can be easily automated and incorporated into the implementation by using few iterations of HSUMMA.

Our algorithm does not change the distribution of the matrices in SUMMA. Currently, our algorithm was designed for block-checkerboard distribution of the matrices. However, we believe that by using block-cyclic distribution the communication can be better overlapped and parallelized and thus the communication cost can be reduced even further. Thus, theoretical and practical analysis of our algorithm with block-cyclic distribution is one of our main future works. In addition, until now we got all these improvements without overlapping the communications on the virtual hierarchies.

We also plan to investigate the algorithm with more than two levels of hierarchy as we believe that in this case it is possible to get even better performance. In addition, we plan to apply the same approach to other numerical linear algebra kernels such as QR/LU factorization.

ACKNOWLEDGMENTS

The research in this paper was supported by IRCSET(Irish Research Council for Science, Engineering and Technology) and IBM, grant numbers EPSG/2011/188 and EP-SPD/2011/207.

Some of the experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see <https://www.grid5000.fr>)

Another part of the experiments in this research were carried out using the resources of the Supercomputing Laboratory at King Abdullah University of Science&Technology (KAUST) in Thuwal, Saudi Arabia.

REFERENCES

- [1] L. Cannon, "A cellular computer to implement the kalman filter algorithm," Ph.D. dissertation, Bozeman, MT, USA, 1969.
- [2] G. C. Fox, S. W. Otto, and A. J. G. Hey, "Matrix algorithms on a hypercube i: Matrix multiplication," *Parallel Computing*, vol. 4, pp. 17–31, April 1987.
- [3] R. Agarwal, S. M. Balle, F. G. Gustavson, M. Joshi, and P. Palkar, "A three-dimensional approach to parallel matrix multiplication," *IBM Journal of Research and Development*, vol. 39, 1995.

- [4] J. Choi, D. W. Walker, and J. Dongarra, "Pumma: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers," *Concurrency - Practice and Experience*, vol. 6, no. 7, pp. 543–570, 1994.
- [5] R. van de Geijn and J. Watts, "Summa: Scalable universal matrix multiplication algorithm," *Concurrency: Practice and Experience*, vol. 9, no. 4, pp. 255–274, April 1997.
- [6] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley, *ScalAPACK user's guide*, J. J. Dongarra, Ed. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 1997.
- [7] E. Solomonik and J. Demmel, "Communication-optimal parallel 2.5d matrix multiplication and lu factorization algorithms," in *Proceedings of the 17th international conference on Parallel processing - Volume Part II*, ser. Euro-Par'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 90–109. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2033408.2033420>
- [8] M. Kondo, "Report on exascale architecture roadmap in Japan," 2012.
- [9] O. Beaumont, V. Boudet, F. Rastello, and Y. Robert, "Matrix multiplication on heterogeneous platforms," 2001.
- [10] A. Lastovetsky and J. Dongarra, *High Performance Heterogeneous Computing*. Wiley, 2009.
- [11] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra, *MPI-The Complete Reference, Volume 1: The MPI Core*, 2nd ed. Cambridge, MA, USA: MIT Press, 1998.
- [12] R. Thakur, "Improving the performance of collective operations in mpich," in *Recent Advances in Parallel Virtual Machine and Message Passing Interface. Number 2840 in LNCS, Springer Verlag (2003) 257267 10th European PVM/MPI Users Group Meeting*. Springer Verlag, 2003, pp. 257–267.
- [13] M. Barnett, S. Gupta, D. G. Payne, L. Shuler, R. Geijn, and J. Watts, "Interprocessor collective communication library (intercom)," in *In Proceedings of the Scalable High Performance Computing Conference*. IEEE Computer Society Press, 1994, pp. 357–364.
- [14] R. Thakur and R. Rabenseifner, "Optimization of collective communication operations in mpich," *International Journal of High Performance Computing Applications*, vol. 19, pp. 49–66, 2005.
- [15] D. Scott, "Efficient all-to-all communication patterns in hypercube and mesh topologies," in *Distributed Memory Computing Conference, 1991. Proceedings., The Sixth*, apr-1 may 1991, pp. 398–403.
- [16] R. W. Hockney, "The communication challenge for mpp: Intel paragon and meiko cs-2," *Parallel Comput.*, vol. 20, no. 3, pp. 389–398, Mar. 1994. [Online]. Available: [http://dx.doi.org/10.1016/S0167-8191\(06\)80021-9](http://dx.doi.org/10.1016/S0167-8191(06)80021-9)
- [17] M. Shroff and R. A. V. D. Geijn, "Collmark: Mpi collective communication benchmark," Tech. Rep., 2000.
- [18] J. Pješivac-Grbović, "Towards automatic and adaptive optimizations of MPI collective operations," Ph.D. dissertation, University of Tennessee, Knoxville, December, 2007.
- [19] P. Patarasuk, X. Yuan, and A. Faraj, "Techniques for pipelined broadcast on ethernet switched clusters," *J. Parallel Distrib. Comput.*, vol. 68, no. 6, pp. 809–824, Jun. 2008. [Online]. Available: <http://dx.doi.org/10.1016/j.jpdc.2007.11.003>
- [20] P. Balaji, R. Gupta, A. Vishnu, and P. Beckman, "Mapping communication layouts to network hardware characteristics on massive-scale blue gene systems," *Comput. Sci.*, vol. 26, no. 3-4, pp. 247–256, Jun. 2011. [Online]. Available: <http://dx.doi.org/10.1007/s00450-011-0168-y>