

Improving the Accuracy of High Performance BLAS Implementations using Adaptive Blocked Algorithms

Matthew Badin*, Paolo D’Alberto†, Lubomir Bic*, Michael Dillencourt*, Alexandru Nicolau*

*University of California Irvine

Irvine, CA 92697 USA

†Yahoo! Inc.

First Ave 701, Sunnyvale CA 94089 USA

Email: mbadin@uci.edu, paolo@ics.uci.edu, bic@ics.uci.edu, dillenco@ics.uci.edu, nicolau@ics.uci.edu

Abstract—Matrix multiply is ubiquitous in scientific computing. Considerable effort has been spent on improving its performance. Once methods that make efficient use of the processor have been exhausted, methods that use less operations than the canonical matrix multiply must be explored. Combining the two methods yields a hybrid matrix multiply algorithm. Hybrid matrix multiply algorithms tend to be less accurate than the canonical matrix multiply implementation, leaving room for improvement. There are well-known techniques for improving accuracy, but they tend to be slow and it is not immediately obvious how best to apply them to hybrid algorithms without lowering performance. Previous attempts have focused on the bottom of the hybrid matrix multiply algorithm, modifying the high-performance matrix multiply implementation. In contrast, the top-down approach presented here does not require the modification of the high-performance matrix multiply implementation at the bottom, nor does it require modification of the fast asymptotic matrix multiply algorithm at the top. The three-level hybrid algorithm presented here not only has up to 10% better performance than the fastest high-performance matrix multiply, but is also more accurate.

Keywords-Recursive Matrix Multiply; Pairwise Summation; Hybrid Matrix Multiply;

I. INTRODUCTION

Matrix multiplication is a fundamental operation. It is used in everything from simulations[1] to implementation of all other level 3 Basic Linear Algebra Subroutines (BLAS) [2]. Accordingly, there has been high demand for fast matrix multiply.

High performance implementations have reached their limit in terms of efficient use of current processor architecture [3], [4]. This has prompted researchers to introduce *hybrid algorithms* that combine two different approaches [5]. At the top is a fast asymptotic algorithm, a recursive algorithm matrix for multiplication that reduces the number of recursive calls to multiply submatrices by increasing the number of additions [6], [7]. At the leaves, a highly tuned high-performance matrix multiplication implementation such as GotoBLAS[3] or Intel MKL[4] is used. Unfortunately these two-level hybrid approaches have worse accuracy than simply using a high-performance matrix multiply implementation.

There are a number of ways of improving the accuracy of matrix multiply and consequently the hybrid matrix multiply algorithms, but they tend to have poor performance and are difficult to adapt to hybrid matrix multiply algorithms. Pairwise summation [8] is a well-known technique for improving the forward error bound of matrix multiply. It reduces the forward error from $O(n)$ to $O(\log_2 n)$ by distributing the additions over a balanced tree. Adapting pairwise summation to hybrid algorithms has been difficult. Previous attempts have focused on a bottom up approach [9], applied at the kernel level [2]. The bottom up approach has two main drawbacks. First, it requires the modification of a highly tuned matrix multiply implementation [3]. Second, the improvement in accuracy comes at a direct cost in performance. This is because the additional temporary space required for each additional level of pairwise summation and the additional instructions necessary to implement pairwise summation reduce locality and hence increase the number of cache misses.

We introduce a different way of applying pairwise summation to matrix multiply. We apply recursive matrix multiply, at the outer product instead of the inner product. Our method does not require the modification of the underlying matrix multiply kernel implementation. We show how our algorithm can be applied to traditional two level hybrid matrix multiply algorithms. This not only improves the overall accuracy of the hybrid algorithm, but also achieves better accuracy than high-performance matrix multiply implementations such as GotoBLAS[3], while still improving performance by up to 10% over just using the underlying high-performance implementation.

The rest of this paper is organized as follows. Section II motivates and introduces our three-level hybrid approach. In Section III we discuss various ways in which the inner product of a matrix multiply kernel can be implemented and how it affects accuracy. We show that the bound on the relative error in Goto matrix multiply [3] is effectively $O(\sqrt{n})$. We then present a brief overview of pairwise summation (Section IV). The fact that pairwise summation has a relative error of $O(\log_2 n)$ partially explains why our

accuracy is better than that of the GotoBLAS implementation. In Section V, we describe how best to adapt pairwise summation to matrix multiply so that it can be easily adapted to traditional two-level hybrid matrix multiply algorithms, yielding recursive matrix multiply. We then develop the forward error bound for our adaptation, which shows that the placement of recursive matrix multiply on top of the high-performance matrix multiply implementation and rather than inside it is the better approach. Finally, we present experimental results (Section VI). Our results demonstrate that our three-level hybrid approach using Strassen’s algorithm [6] at the top offers better accuracy than Strassen’s hybrid algorithm alone while still offering improved performance over a typical high-performance implementation. When we use Winograd’s variant[7] of Strassen’s method at the top, the resulting algorithm is faster than the GotoBLAS general matrix multiply implementation, and it is also more accurate when the input is positive.

II. THREE LEVEL HYBRID

High performance matrix multiply implementations such as Atlas[2], Intel MKL[4] and GotoBLAS[3] attempt to maximize the utilization of the processor. These implementations have largely reached their limit. As a result, two-level hybrid matrix have been heavily explored as a mean of further improving the performance of matrix multiply [5], [10], [11]. These algorithms combine a fast asymptotic algorithm for matrix multiplication at the top levels with high performance matrix multiply kernels at the leaves. A major problem with this approach is that fast asymptotic matrix multiply algorithms trade off fewer matrix multiplications against more matrix additions. Since addition is a fundamental source of rounding error, this trade off increases the error, even if only a few applications of the fast asymptotic matrix multiply algorithm are used[8].

Rounding errors in additions can be reduced using accurate summation algorithms, such as distillation [12] and pairwise summation [8]. Since accurate summation algorithms can be computationally expensive, it is important to decide where in the hybrid algorithm they are best applied. One possibility is to apply an accurate summation algorithm within the fast asymptotic matrix multiply algorithm, replacing all ordinary additions with the accurate summation algorithm. This fails to have any affect on the overall accuracy of the hybrid matrix multiply algorithm because even small errors are amplified by the multiplication inside the fast asymptotic matrix multiply algorithm [8]. A second possibility is to deploy the accurate summation algorithm inside the matrix multiplication kernel itself, improving the accuracy of the leaf computations of the hybrid matrix multiply algorithm. This approach degrades performance due to cache constraints. Since there must be enough cache for both the accurate summation algorithm and the high-performance matrix multiply algorithm, the accurate sum-

mation algorithm can only be used in a very limited manner before it impacts performance[9]. This is discussed further in Section V.

If we want to avoid the above difficulties, then we must deploy the accurate summation algorithm outside of either the matrix multiply kernel or the fast asymptotic matrix multiply algorithm. This will result in a three level hybrid matrix multiply algorithm. There are three possible places where the accurate summation algorithm may be placed: (1) beneath the high-performance kernel, (2) above the fast asymptotic matrix multiply algorithm, or (3) between the two. The first option is not a reasonable choice. The second option has a similar disadvantage: if we deploy the accurate summation algorithm above the fast asymptotic matrix multiply, then this significantly diminishes the performance advantage gained by using fast asymptotic algorithm, defeating the purpose of using it. Thus we are left with the third option, placing the accurate summation algorithm between the fast asymptotic matrix multiply algorithm at the top and the high performance matrix multiply-kernel at the bottom.

We expand upon this notion further in Sections IV and VI, below. To keep the paper concise, we focus only on one family of fast asymptotic matrix multiply algorithms, namely Strassen’s algorithm [6] and Winograd’s variant of Strassen’s algorithm[7]. But first we need to discuss the source of rounding errors in matrix multiply implementations and, in particular, high performance matrix multiply implementations. Then we can discuss how to reduce this error in the context of hybrid matrix multiply algorithms.

III. INNER PRODUCT VARIATIONS

There are many different approaches to computing matrix products and to computing individual inner products. While some of these choices are largely dictated by architecture and performance considerations [2], [3], they impact the overall accuracy of the computation. This becomes important when dealing with high performance matrix multiply kernels. It is possible for the high performance kernel to be more accurate than the canonical matrix multiply algorithm, which improves the accuracy of the overall hybrid algorithm.

Tiling is a common strategy within high performance matrix multiply kernels. The primary goal of tiling is to maximize reuse of the data in the cache and hence improve performance. Subtleties in how the individual tiles are added to the whole greatly impact the forward error bound on the inner product [8]. The difference is sometimes referred to as *preload* vs. *postload* [9]. We adopt this terminology in this paper. The following discussion will demonstrate that in practice these differences allow some high performance matrix multiply kernels to achieve an error bound that is effectively $O(\sqrt{n})$ for common problem sizes [9], [3]. We will argue (in Section V, below) that not only does our three level hybrid algorithm beat this bound, it continues to do so

```

ALGORITHM: Tiled Matrix Multiply
procedure tiled( $A, B, C, M, K, N, bSize$ )
for ( $i = 0, i < M, i += bSize$ )
  for ( $j = 0, j < N, j += bSize$ )
    for ( $k = 0, k < K, k += bSize$ )
       $AOffset = i * K + k$ 
       $BOffset = k * N + j$ 
       $COffset = i * N + j$ 
      multiply( $A+AOffset, K, B+BOffset, N,$ 
               $C+COffset, N, bSize, bSize, bSize$ )

```

Figure 1. Tiled Matrix Multiply

```

ALGORITHM: Direct Matrix Multiply (Preload)
procedure multiply( $A, ldA, B, ldB, C, ldC, M, K, N$ )
begin
for ( $i = 0, i < M, i ++$ )
  for ( $j = 0, < N, j ++$ )
    for ( $k = 0, k < K, k ++$ )
       $C[i * ldC + j] += A[i * ldA + k] * B[k * ldB + j]$ 

```

Figure 2. Preload Matrix Multiply

even when using a fast asymptotic matrix multiply algorithm that is less accurate than the canonical matrix multiply [8].

A. Preload and Postload

Preload and postload are best discussed in the context of a specific tiling strategy. Consider the code shown in Figure 1, which multiplies the $M \times K$ matrix A by the $K \times N$ matrix B to produce the $M \times N$ matrix C . The matrices are stored in row-major order. The tiles are square, of size $bSize \times bSize$. For simplicity, we assume that the matrix dimensions are all even multiples of the tile size. The strategy in Figure 1 is to compute C one column of tiles at a time, always keeping in memory a $bSize \times k$ strip of A along with the tile of C currently being computed. The code repeatedly multiplies a tile of A by a tile of B and adds the result to a tile of C . Each partial tile computation is performed by calling the procedure

```
multiply( $A, ldA, B, ldB, C, ldC, M, K, N$ )
```

where A, B , and C are the addresses of the start of the tiles; the tiles are of sizes $M \times K$, $K \times N$, and $M \times N$ respectively; and ldA , ldB , and ldC are the *strides* (i.e., the row lengths) of the three matrices.

The error of the algorithm shown in Figure 1 depends on how the inner kernel (i.e., the `multiply()` procedure) is implemented. If the implementation is the preload code shown in Figure 2, the summation of the inner product is recursive (standard) summation of K terms. In contrast, the postload code shown in Figure 3 computes the contribution of the tile to the inner product in a temporary variable

```

ALGORITHM: Direct Matrix Multiply (Postload)
procedure multiply( $A, ldA, B, ldB, C, ldC, M, K, N$ )
begin
for ( $i = 0, i < M, i ++$ )
  for ( $j = 0, < N, j ++$ )
     $sum = 0$ 
    for ( $k = 0, k < K, k ++$ )
       $sum += A[i * ldA + k] * B[k * ldB + j]$ 
    end for
     $C[i * ldC + j] += sum$ 
  end for
end for

```

Figure 3. Postload Matrix Multiply

and adds this sum to the inner product. When the code of Figures 1 and 3 is combined, the net effect is to compute the dot product as the sum of $K/bSize$ partial sums, each of which is the sum of $bSize$ terms. This has a significant effect on the forward error bound, as will be discussed in the next subsection.

B. Forward Error Bound

In order to do error analysis of the tiling strategies presented in the previous section, a model must be used to account for the rounding errors, particularly the rounding errors in addition. We use the standard floating point model for estimating error bounds presented in [8], as it widely used and holds for IEEE standard arithmetic. In this model, the result of an individual floating point operations is $fl(x + y) = (x + y)(1 + \delta)$, where $|\delta| \leq u$ and δ is the relative error of an individual operation, u is the “unit in last place” (commonly known as *ulp*), and $fl(\cdot)$ is the value contained inside the parenthesis rounded to machine precision. This means the previous statement says that the real answer, rounded to machine precision, is equal to the real answer plus some small relative error, less than or equal to one *ulp*. Using this model, the following forward error bound for the inner product of matrix multiply is derived in [8]:

$$|x^T y - fl(x^T y)| \leq \gamma_n \sum_{i=1}^n |x_i y_i| = \gamma_n |x^T y|, \quad (\text{III.1})$$

where $\gamma_n = nu/(1 - nu)$. If the inner product is broken up into n/k strips, each containing k summands, the forward error bound becomes:

$$|s_n - \widehat{s}_n| \leq \gamma_{\frac{n}{k} + k - 1} |x^T y|. \quad (\text{III.2})$$

Here, s_n is the true value of the inner product, \widehat{s}_n is the computed value, and k is equal to $bSize$. This means that $bSize$, which is usually chosen based upon the architecture of the machine, has a dramatic effect on the forward error

bound. An interesting side effect of (III.2) is when $k = c\sqrt{n}$, as proven in [9]:

$$\gamma_{(c\sqrt{n} + \frac{n}{c\sqrt{n}} - 1)} = \gamma_{(c\sqrt{n} + \frac{1}{c}\sqrt{n} - 1)} = \gamma_{((c + \frac{1}{c})\sqrt{n} - 1)}. \quad (\text{III.3})$$

As illustrated by Castaldo, Whaley and Chronopoulos[9], when $k = 60$, you achieve roughly the same bound for $n \in [900, 14400]$ as $60 = 2\sqrt{900}$ and $60 = \frac{1}{2}\sqrt{14400}$. This means for most practical problem sizes, you end up with a forward error bound that is effectively $O(\sqrt{n})$ in practice. This effect is illustrated by comparing GotoBLAS [3] to the canonical $O(n^3)$ matrix multiply algorithm in Figure 5. In the next section we will describe how it is possible to achieve a lower bound by further changing the order in which addition is carried out. This will give us a mechanism, that when combined with a traditionally less accurate two level hybrid matrix multiply algorithm, will allow us to produce an algorithm more accurate and faster than just the high performance matrix multiply algorithm alone, which we demonstrate in Section VI.

IV. PAIRWISE SUMMATION

As discussed in the previous section, if the addition is broken up, the longest path from the summand to the total is shortened, it is possible to reduce the error. Where standard summation is simply linear, pairwise summation breaks apart the addition into a balanced tree, where the leaves are the summands and the root is the total. Pairwise summation is usually defined as a special case of recursive (standard) summation and this time will be no different. Recursive summation is traditionally defined in terms of a set of summands, by removing two elements from the set, adding them together, and placing the new sum back into the set. This recursion repeats, hence the name, until only one element remains, the final sum. In contrast, pairwise summation adds adjacent elements together creating new summands, then again adds adjacent summands together. The process of adding adjacent summands repeats until a final sum is reached. A recursive algorithm to illustrate pairwise summation can be found in Figure 4, where n is the number of elements and the elements are numbered x_1 to x_n . The following forward error bound for pairwise summation is derived in [8]:

$$|s_n - \widehat{s}_n| \leq \gamma_{\log_2 n} \sum_{i=1}^n |x_i| \quad (\text{IV.1})$$

In the next section we will describe some of the pitfalls of applying pairwise summation directly to the inner product along with describing a better method of adapting pairwise summation to matrix multiplication, at the outer product level.

ALGORITHM: Pairwise Summation

```

procedure pwsun( $n, x_1, \dots, x_n$ )
begin
if( $n == 1$ )
return  $x_1$ 
else if( $n == 2$ )
return  $x_1 + x_2$ 
else
 $k = \lfloor n \div 2 \rfloor$ 
return pwsun( $k, x_1, \dots, x_k$ ) + pwsun( $n - k, x_{k+1}, \dots, x_n$ )
end if
end procedure

```

Figure 4. Pairwise Summation




Figure 5. Goto Absolute Error

V. DIRECT RECURSIVE MATRIX MULTIPLY

As suggested by Higham, the forward error bound can be greatly improved by applying pairwise summation to the entire inner product, producing a forward relative error bound of $\gamma_{\lceil \log_2 n \rceil + 1} |x|^T |y|$ [8]. However, the full implementation of this matrix multiplication variation is not practical as the performance tends to be abysmal. Previous attempts to apply pairwise summation to dense matrix multiply kernels have therefore focused on a bottom up approach, attempting to increase the accuracy of the individual inner product (or products as tiling is usually used) while using as little additional temporary space as possible [9]. This is done by applying pairwise summation at the tile level and only applying a few levels, usually only three. The amount of temporary space required for this approach is a multiple of the tile size and grows linearly with the number of levels of pairwise summation that are applied.

This approach has two main drawbacks. The first is that

performance suffers as the levels of pairwise summation increase, because as more temporary space is required to store the temporary sums. This temporary space reduces the amount of space available for the actual matrix multiply kernel, reducing matrix reuse and harming performance. This effect can be mitigated by only applying a few levels of pairwise summation, usually only three [9]. But this reduces the accuracy, partially defeating the purpose of applying pairwise summation. The second drawback is that this implementation must be done at the kernel level of the high performance matrix multiply implementation, increasing the complexity of the implementation. This complexity issue is compounded in implementations such as GotoBLAS which have a different kernel for each specific architecture [3].

There is a much simpler way of adapting pairwise summation to matrix multiply via a top down approach, directly using the definition of matrix multiply. Given matrices

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix}, \quad B = \begin{bmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{bmatrix}, \quad \text{and} \\ C = \begin{bmatrix} C_{1,1} & C_{1,2} \\ C_{2,1} & C_{2,2} \end{bmatrix}$$

the entries of C are given by

$$\begin{aligned} C_{1,1} &= A_{1,1} \times B_{1,1} + A_{1,2} \times B_{2,1}, \\ C_{1,2} &= A_{1,1} \times B_{1,2} + A_{1,2} \times B_{2,2}, \\ C_{2,1} &= A_{2,1} \times B_{1,1} + A_{2,2} \times B_{2,1}, \quad \text{and} \\ C_{2,2} &= A_{2,1} \times B_{1,2} + A_{2,2} \times B_{2,2}. \end{aligned}$$

Each submatrix of C can be computed using the inner products shown above. If we further decompose the submatrices of A , B , and C into submatrices, this process can be repeated, yielding a balanced tree computation. This algorithm is called recursive matrix multiply.

In the next subsection we will discuss the effect of computing a matrix product by running direct recursion until the submatrices are down to a certain size, and then passing these leaves to a high performance matrix multiply implementation. We compare this top-down approach to pairwise summation, with a bottom-up approach implemented within the high performance matrix multiply implementation.

A. Forward Error Bound

Using the same standard model as described in Section V-A, let k be the leaf size at which direct recursion is terminated. From (III.1), we know that the forward error in the computation of each leaf is bounded above by

$$\gamma_k \sum_{i=1}^k |x_i y_i| = \gamma_k |x|^T |y|.$$

Since there are n/k leaves, it follows from (IV.1) that the forward error of adding up the leaves after they are computed satisfies

$$|s_n - \widehat{s}_n| \leq \gamma_{\lceil \log_2(\frac{n}{k}) \rceil + 1} |x|^T |y|.$$

Combining the two bounds yields

$$|s_n - \widehat{s}_n| \leq |x|^T |y| (1 + \delta)^k (1 + \delta)^{\log_2(\frac{n}{k})}, \quad (\text{V.1})$$

which when simplified becomes

$$|s_n - \widehat{s}_n| \leq \gamma_{(\lceil \log_2(\frac{n}{k}) \rceil + k)} |x|^T |y|. \quad (\text{V.2})$$

In contrast, it is shown in [9] that the forward error in a bottom-up approach with three levels of pairwise summation and a leaf size of k satisfies:

$$|s_n - \widehat{s}_n| \leq \gamma_{k+2(\sqrt{\frac{n}{k}}-1)} |x|^T |y| \quad (\text{V.3})$$

Increasing the number of levels degrades performance significantly and marginally improves the forward error bound, but for any fixed number of levels the error bound is still asymptotically worse than $\log(n/k)$. Hence our approach using direct recursion yields a smaller forward error bound than a bottom up approach applying pairwise summation.

When the leaf size k is fixed, direct recursion is more accurate for large problems than GotoBLAS [3], since the bound on the relative error in the former grows in proportion to $\log_2 n$, while the bound on the relative error in the latter grows in proportion to \sqrt{n} .

It can be seen from (V.2) that the leaf size dominates the relative error when recursive matrix multiply is placed in the middle. This means that if instead of fixing the leaf size we fix the number of levels of recursion, accuracy will be degraded. We will see in the results section that for the largest problem sizes, if we fix the number of levels of recursion and let the leaf size vary, the error approaches the error of not inserting the intermediate incursion at all. Empirically, the crossover point appears to be about $8k$, as seen in Figure 7 and Table I. This is discussed in more detail in Section VI.

Our approach allows a large problem size to have an error very close to that of a much smaller problem size, when k is fixed. By placing direct recursive matrix multiply in the middle of the traditional hybrid matrix multiply implementation, we offer a path between the fast asymptotic matrix multiply algorithm at the top and the high performance matrix multiply implementation at the bottom without greatly increasing the forward error in between.

VI. RESULTS

The high performance matrix multiply implementation chosen to illustrate the effects of postload was GotoBLAS [3]. Atlas [2] was also considered, as was Intel MKL [4], however, Atlas was found to have worse performance and Intel MKL had nearly identical performance to that of GotoBLAS.




Figure 6. Strassen 3-Level Performance




Figure 7. Strassen 3-Level Hybrid Accuracy

A. Testing Methodology

Unless otherwise stated, the matrices tested were square matrices. The reference was generated with GotoBLAS [3], using double precision general matrix multiply (dgemm) which is sufficient when the test are conducted in single precision. The input for all tests was a uniform distribution between [0,1], generated using GNU R, the seeds provided by random.org. The machine used was a Q9450 Penryn Intel Quad processor running at 3GHz on a 64-bit Kubuntu 10.04 desktop installation. The metrics used to compare various implementations, in terms of accuracy, are those of relative error and absolute error [8]. Though several other metrics exist [13], [14], [9], they do not appear to be widely used. For performance, we compare implementations by how quickly

each algorithm can perform the essential $O(n^3)$ operations, whether or not the implementation actually does them, measured in flops (floating point operations per second). The high performance matrix multiply implementation used was GotoBLAS [3] where not only the implementation, but also the entire hybrid matrix multiply algorithm, are single threaded. Implementation details and the impact of multi-threading on accuracy and performance will be explored in future works. The implementation used as the performance reference for Winograd’s variant in Table III is described in [5]. In both the tables and graphs, “2x SR, 2x DR” translates into applying Strassen’s algorithm recursively twice, followed by applying direct recursive matrix multiply twice before switching to a high performance implementation, in this case GotoBLAS [3].

B. Measuring Error

High performance implementations such as GotoBLAS have traditionally enjoyed better accuracy than the canonical matrix multiply, as illustrated in Figure 5, which compares the absolute max error of the standard matrix multiply as compared to GotoBLAS’s single precision general matrix multiply (sgemm), both computed in single precision. Without this attribute of additional accuracy, hybrid matrix multiply algorithms would not be useful as the error of the asymptotically fast matrix multiply algorithm would be too large. In practice though, the additional accuracy offered by high performance implementations such as GotoBLAS makes hybrid algorithms possible[5]. The additional accuracy provided by GotoBLAS however has been traditionally attributed to a feature of the architecture, namely, 80 bit floating point registers that are available on Intel x86 processors. Unfortunately the effect of extended precision registers has not been true for several years, at least in modern implementations, and most importantly, it is not true in the current version of Goto BLAS [3]. The problem dates back to the introduction of the SSE instruction set, which mapped MMX registers to FP registers, which are 80 bit registers on Intel processors[15]. When the SSE2 SIMD instruction set was introduced (single instruction, multiple data), the MMX registers were replaced with XMM registers, which are 128 bit registers [15]. These registers operate on packed double precision or packed single precision data, two double precision values or four single precision values, respectively. These registers no longer offer additional accuracy [15]. This was tested empirically by comparing the results of vector dot products produced by the $O(n^3)$ matrix multiply algorithm in single precision against that produced using inline assembly using the SSE2 instruction set (four vectors each), the results were identical. By inspecting Goto’s implementation for the Penryn processor, as the inner matrix multiply kernel in Goto is tailored to specific processors, it is easy to see that XMM registers are used and not MMX. This means that the error observed in Figure 5 is entirely because of the

preload effect described in Section III-A and not something specific to the architecture. It may appear in Figure 5 that Goto’s implementation is logarithmic, however as discussed in Section V-A, Goto’s implementation is effectively square root[9]. As can be clearly seen from the graph, a simple tiling strategy combined with postload matrix multiply offers a very low error for practical problem sizes.

C. Strassen’s Algorithm

Before discussing the data further, it is important to emphasize a key purpose of hybrid algorithms and ultimately the limitations of matrix multiply and the machines the algorithms run on, namely, the reason for the existence of hybrid matrix multiply algorithms. Goto’s implementation of matrix multiply and his resulting BLAS implementation is widely considered to be the fastest implementation of the canonical matrix multiply [3]. This is achieved through a key insight into current x86 processor architecture, resulting in a highly efficient implementation. This means, that in order to improve performance any further, a asymptotically fast matrix multiply implementation must be used, something that reduces the total number of multiplications. Though this detail may seem minor, it can be quickly illustrated in Figure 6 and 7. Figure 6 is the overall speed of the implementations, as described in Section VI-A, comparing Goto’s matrix multiply, a standard hybrid implementation of Strassen’s algorithm on top of Goto’s matrix multiply which applies Strassen’s algorithm until a leaf size of 1k or less before switching to GotoBLAS, and that of our adaptation by placing direct recursive matrix multiply in between. To be explicit, our algorithm runs Strassen’s algorithm twice, unlike the standard hybrid implementation which is variable with a fixed leaf of 1k or less, we then run direct recursive matrix multiply twice before finally switching to a high performance matrix multiply implementation. This means the leaf size passed to the high performance matrix multiply implementation is not fixed, that the levels of recursion are, the consequences of which are discussed in Section V-A. By comparing the implementations it is evident that our algorithm is faster than the high performance matrix multiply implementation, but not as fast as a standard Strassen hybrid. By also looking at Figure 7, which compares the algorithms in terms of max absolute error, it is clear that our algorithm is more accurate than the standard Strassen hybrid matrix multiply and that neither are more accurate than just using Goto’s matrix multiply. The importance of the discussion at the beginning of the subsection now becomes apparent, namely, though both implementations are not as accurate as the high performance matrix multiply implementation by itself, a hybrid matrix multiply algorithm is the only way to improve performance past the high performance canonical matrix multiply implementations. This means that if one wants an algorithm that is faster than the fastest implementation of the canonical matrix multiply, they must

use a hybrid matrix multiply algorithm, however, if they want more accuracy than what is offered by a standard hybrid matrix multiply algorithm and still better performance than the canonical high performance matrix multiply implementation, our algorithm offers a solution. Fortunately this can be improved further by taking advantage of a peculiarity of certain asymptotically fast matrix multiply algorithms, in particular, that Winograd’s variant is more accurate in practice than Strassen’s when the input matrices are positive, as we will see in the next subsection.

D. Winograd’s Algorithm

Winograd luckily has an interesting property in that it does not have “true” subtraction, which means that when the input is positive, it is actually more accurate than Strassen’s algorithm [5]. This is easily demonstrated in Table I where the first column represents the problem size and the rows are the accuracy of individual algorithms for that problem size. When Winograd is combined with a fixed two levels of direct recursive matrix multiply, the error grows linearly with the size of the leaf, as described in Section V-A, which can be clearly seen in Table I. This means two applications of direct recursive matrix multiply offers better accuracy up until the size of the leaf overwhelms the benefit of two levels of direct recursive matrix multiply, resulting in no benefit in terms of accuracy at around 8k. Fortunately we already know this can be further improved by using a fixed leaf size and variable levels of direct recursive matrix multiply. By drawing attention to Table I we can clearly see the benefit of a fixed leaf of 256. Instead of the error growing linearly in terms of the leaf size, the leaf error becomes a constant and instead the error grows linearly in terms of the level of direct recursive matrix multiply, which grows slower, as discussed in Section V-A. Table II contains the absolute error using a slightly larger fixed leaf size of 320. Though not as accurate as a fixed leaf size of 256, it has the added benefit of being faster with this particular high performance matrix multiply kernel implementation, as seen in Table III. By comparing both Tables II and III, a fixed leaf size of 320 appears to be a good choice in terms of accuracy and speed and that a larger leaf size would not be desirable. The reason for this is though performance may vary by architecture and implementation, accuracy largely does not. The only benefit of a different leaf size would be better accuracy if the leaf was smaller and if the high performance implementation offered it, better performance for a smaller leaf size. Unfortunately GotoBLAS [3] appears to be implemented with larger problem sizes in mind. This is revealed in Table III where Goto’s implementation appears to plateau starting at a problem size of 2k, far too large to be used as a leaf computation when accuracy is important. However, even with this limitation, we have still demonstrated what we believe to be the first three level hybrid matrix multiply algorithm that is not only more accurate than the fastest

Winograd Fixed Leaf Variants				
Size	SGEMM	2xWin	2xWin2xDR	WinDR256
1024	0.000331	0.000225	0.000108	0.000136
2048	0.000509	0.000602	0.000228	0.000231
4096	0.000733	0.000848	0.000562	0.000450
8192	0.001194	0.001413	0.001415	0.000834

Table I
WINOGRAD HYBRIDS, ABSOLUTE ERROR

Winograd Fixed Leaf Variants				
Size	SGEMM	2xWin	2xWin2xDR	WinDR320
1024	0.000331	0.000225	0.000108	0.000216
2048	0.000509	0.000602	0.000228	0.000334
4096	0.000733	0.000848	0.000562	0.000579
8192	0.001194	0.001413	0.001415	0.001001

Table II
WINOGRAD HYBRIDS, ABSOLUTE ERROR

Winograd Fixed Leaf Variants					
Algorithm	1024	2048	4096	9024	
SGEMM	18.83	22.64	22.99	23.30	
Win	21.83	23.66	26.15	29.49	
2xWin 2xDR	14.31	18.47	22.71	25.71	
2xWin 256DR	17.89	18.57	19.80	20.42	
2xWin 320DR	20.45	21.07	22.77	24.07	

Table III
WINOGRAD HYBRIDS, SPEED IN GFLOPS

high performance matrix multiply implementation [3], but is also more accurate.

VII. CONCLUSION

In this paper we have presented what we believe to be the first three level hybrid matrix multiply algorithm. We have demonstrated not only the usefulness of a three level hybrid matrix multiply algorithm, but identified a clear need for one, a need for something that improves both performance and accuracy over a traditional high performance matrix multiply implementation. We have proven that our ordering of the individual levels is the correct ordering for the hybrid matrix multiply algorithm. We have also built a model that can be used for determining the proper number of applications of direct recursive matrix multiply based upon the problem size, which will allow any application programmer to tailor the three level hybrid algorithm to their needs. We have also demonstrated empirically that this model is correct and that it is possible to out perform a traditional high performance matrix multiply implementation by up to 10% while still offering superior accuracy. Furthermore, we have demonstrated that the performance and accuracy of our three level hybrid matrix multiply algorithm will only improve in the future with high performance matrix multiply implementations specifically tailored for hybrid algorithms.

REFERENCES

- [1] M. Yokokawa, K. Itakura, A. Uno, T. Ishihara, and Y. Kaneda, "16.4 Tflops direct numerical simulation of turbulence by a Fourier spectral method on the Earth Simulator," in *Proceedings of The 2002 ACM/IEEE Conference on Supercomputing*, 2002.
- [2] R. C. Whaley and J. Dongarra, "Automatically tuned linear algebra software," in *PPSC*, 1999.
- [3] K. Goto and R. van de Geijn, "On reducing TLB misses in matrix multiplication," The University of Texas at Austin, Working Note 9, November 2002.
- [4] Intel, "Intel math kernel library (intel mkl) 10.2," 2010. [Online]. Available: <http://software.intel.com/en-us/intel-mkl/>
- [5] P. D'Alberto and A. Nicolau, "Adaptive Winograd's matrix multiplications," *ACM Trans. Math. Softw.*, vol. 36, pp. 3:1–3:23, March 2009. [Online]. Available: <http://doi.acm.org/10.1145/1486525.1486528>
- [6] V. Strassen, "Gaussian elimination is not optimal." *Numerische Mathematik*, vol. 14, no. 3, pp. 354–356, 1969.
- [7] S. Winograd, "A new algorithm for inner product," *IEEE Trans. Comput.*, vol. 17, pp. 693–694, July 1968. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1311073.1311294>
- [8] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed. SIAM, 2002.
- [9] A. M. Castaldo, R. C. Whaley, and A. T. Chronopoulos, "Reducing floating point error in dot product using the superblock family of algorithms," *SIAM J. Sci. Comput.*, vol. 31, no. 2, pp. 1156–1174, 2008.
- [10] S. Huss-Lederman, E. M. Jacobson, J. R. Johnson, A. Tsao, and T. Turnbull, "Implementation of Strassen's algorithm for matrix multiplication," in *Proceedings of Supercomputing '96*, 1996, pp. 9–6.
- [11] C. C. Douglas, M. Heroux, G. Sliselman, R. M. Smith, and R. M., "Gemmw: A portable level 3 Blas Winograd variant of Strassen's matrix-matrix multiply algorithm," 1994.
- [12] Y. Nievergelt, "Analysis and applications of Priest's distillation," *ACM Trans. Math. Softw.*, vol. 30, no. 4, pp. 402–433, 2004.
- [13] A. Ziv, "Relative distance - an error measure in round-off error analysis," *Mathematics of Computation*, no. 160, pp. 563–569, October 1982.
- [14] F. W. J. Olver, "A new approach to error arithmetic," *SIAM Journal on Numerical Analysis*, vol. 15, no. 2, pp. 368–393, April 1978.
- [15] Intel, *Intel 64 and IA-32 Architectures Software Developer's Manual - Volume 1: Basic Architecture*, Intel Corporation, June 2010.