

Reliable Generation of High-Performance Matrix Algebra

Geoffrey Belter,[†] Elizabeth Jessup,[‡] Thomas Nelson,[†] Boyana Norris,[§] and Jeremy G. Siek[†]

[†]Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309

Email: [belter,Thomas.Nelson,Jeremy.Siek]@Colorado.EDU

[‡]Department of Computer Science, University of Colorado, Boulder, CO 80309

Email: jessup@Colorado.EDU

[§]Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439

Email: norris@mcs.anl.gov

Abstract—Scientific programmers often turn to vendor-tuned Basic Linear Algebra Subprograms (BLAS) to obtain portable high performance. However, many numerical algorithms require several BLAS calls in sequence, and those successive calls result in suboptimal performance. The entire sequence needs to be optimized in concert. Instead of vendor-tuned BLAS, a programmer could start with source code in Fortran or C (e.g., based on the Netlib BLAS) and use a state-of-the-art optimizing compiler. However, our experiments show that optimizing compilers often attain only one-quarter the performance of hand-optimized code. In this paper we present a domain-specific compiler for matrix algebra, the Build to Order BLAS (BTO), that reliably achieves high performance using a scalable search algorithm for choosing the best combination of loop fusion, array contraction, and multithreading for data parallelism. The BTO compiler generates code that is between 16% slower and 39% faster than hand-optimized code.

I. INTRODUCTION

Traditionally, scientific programmers have used linear algebra libraries such as the Basic Linear Algebra Subprograms (BLAS) [14, 15, 23] and the Linear Algebra PACKage (LAPACK) [3] to perform their linear algebra calculations. A programmer links an application to vendor-tuned or autotuned implementations of these libraries to achieve efficient, portable applications. For programs that rely on kernels with high computational intensity, such as matrix-matrix multiply, this approach can achieve near optimal performance [35]. However, memory bandwidth, not computational capacity, limits the performance of many scientific applications [3], with data movement expected to dominate the costs in the foreseeable future [2].

Because each BLAS performs a single mathematical operation, such as matrix-vector multiplication, a tuned BLAS library has a limited scope within which to optimize memory traffic. Moreover, separately compiled BLAS limit the scope of parallelization on modern parallel architectures. Each BLAS call spawns threads and must synchronize before returning, but much of this synchronization is unnecessary when considering the entire sequence of matrix algebra operations. The BLAS Technical Forum suggested several new routines that combine sequences of routines, thereby enabling a larger scope for optimization [8, 18]. However, the number of useful BLAS

combinations is larger than is feasible to implement for each new architecture.

Instead of using vendor-optimized BLAS, a scientific programmer can start with source code in Fortran or C, perhaps based on the Netlib BLAS, and then use a state-of-the-art optimizing compiler to tune the code for the architecture of interest. However, our experiments with two industrial compilers (Intel and Portland Group) and one research compiler (Pluto [9]) show that, in many cases, these compilers achieve only one-quarter of the performance of hand-optimized code (see Section V-B). This result is surprising because the benchmark programs we tested are sequences of nested loops with affine array accesses, and the optimizations that we applied by hand (loop fusion, array contraction, and multithreading for data parallelism) are well established. Nevertheless, for some benchmarks, the compilers fail to recognize that an optimization is legal (semantics-preserving). For other benchmarks, they miscalculate the profitability of choosing one combination of optimizations over another combination.

These observations demonstrate that achieving *reliable*, automatic generation of high-performance matrix algebra is nontrivial. In particular, the three main challenges are (1) recognizing whether an optimization is legal, (2) accurately assessing the profitability of optimizations and their parameters, and (3) efficiently searching a large, discontinuous space of optimization choices and parameters. In this paper, we present the Build to Order BLAS (BTO) compiler. It is the first compiler that solves all three of these challenges specifically for the domain of *matrix algebra*.

BTO accepts as input a sequence of matrix and vector operations in a subset of MATLAB, together with a specification of the storage formats for the inputs and outputs, and produces optimized kernels in C. This choice of input language is part of our solution to the problem of determining whether an optimization is legal. The input language makes all data dependencies explicit, so there is no difficulty recognizing whether an optimization is semantics-preserving or not. Further, BTO uses a carefully designed internal representation for optimization choices that rules out many illegal choices while at the same time succinctly representing all the legal choices. To accurately assess profitability, the BTO compiler relies on a hybrid approach. Analytic modeling is used for

coarse-grained pruning whereas empirical timing is used to make the ultimate decisions. To find the best combination of optimizations within a large search space, BTO uses a genetic algorithm whose initial population is the result of a greedy, heuristic search.

We described earlier prototypes of BTO in several papers [5, 6, 20, 32]. In these papers, we considered only a subset of the optimizations considered here; moreover, at the time of their writing, we had not yet developed a search algorithm that was scalable with respect to the number of optimizations and their parameters. The following are the specific, technical contributions of this paper.

- 1) We present an internal representation for optimization choices that is *complete* (includes all legal combinations of loop fusion, array contraction, and multithreading for data parallelism) but that inherently rules out many illegal combinations (Section III).
- 2) We present a scalable and effective search strategy: a genetic algorithm with an initial population seeded by a greedy search. We describe this strategy in Section IV and show in Section V-B that it produces code that is between 16% slower and 39% faster than hand-optimized code.
- 3) We compare this genetic/greedy search strategy with several other strategies in order to reveal the rationale behind this strategy (Section V-D).

We discuss related work in Section VI and conclude the paper in Section VII with a brief discussion of future work.

II. BTO OVERVIEW

Figure 1 shows an example BTO input file for the BATAX subprogram that performs the operations $y \leftarrow \beta A^T A x$ for matrix A , vectors x and y , and scalar β . The user of BTO specifies the input types, including storage formats and a sequence of matrix, vector, and scalar operations; but the user does not specify how the operations are to be implemented. That is, the user does not identify such details as the kinds of loops or the number of threads. The BTO compiler produces a C implementation in two broad steps. It first chooses how to implement the operations in terms of loops, maximizing spatial locality by traversing memory via contiguous accesses. It then searches empirically for the combination of optimization decisions that maximizes performance. Sections III and IV describe the search process.

```

BATAX
in:
  x : vector(column), beta : scalar,
  A : matrix(row)
out:
  y : vector(column)
{
  y = beta * A' * (A * x)
}

```

Fig. 1. BTO input file for the BATAX kernel.

Throughout the compilation process, BTO uses a dataflow graph representation, illustrated in Figure 2 for the BATAX

kernel. The square boxes correspond to the input and output matrices and vectors, and the circles correspond to the operations (operations are labeled with numbers, which are used to identify the operations in the remainder of the paper).




Fig. 2. Dataflow graph for $y \leftarrow \beta A^T A x$.

The BTO compiler uses a type system based on a container abstraction, which describes the iteration space of matrices and vectors. Containers may be oriented horizontally or vertically and can be nested. We assume that moving from one element to the next in a container is a constant-time operation and good for spatial locality, but we place no other restrictions on what memory layouts can be viewed as containers. The types are defined by the following grammar, in which R designates row, C designates column, and S designates scalar.

orientations $O ::= C \mid R$
 types $T ::= O \langle T \rangle \mid S$

Figure 3 shows several types with a corresponding diagram depicting the container shapes: a row container with scalar elements (upper left), a nested container for a row-major matrix (right), and a partitioned row container (lower left). During the creation of the dataflow graph, each node is assigned a type. The input and outputs are assigned types derived from the input file specification, whereas the types associated with intermediate results are inferred by the BTO compiler.




Fig. 3. Vector, partitioned vector, and matrix with their corresponding types.

Note on the polyhedral model: The type system used by BTO and the polyhedral model [22] share a common goal: both describe a schedule to traverse an iteration space. Much of BTO's functionality can be accomplished by using polyhedral-based tools. There are two motivations for using a domain-specific type system as BTO does: (1) ability to seamlessly perform additional optimizations (array contraction), and (2) extensibility with regard to sparse matrix storage formats.

III. SEARCH SPACE

This section describes the search space and challenges with regard to efficiently representing the space. We present a domain specific representation that enables BTO to eliminate many illegal points without spending any search time on them.

This section also sets up the discussion for specific search strategies in Section IV.

A. Description of Search Space

The optimization search space we consider here has three dimensions: (1) loop fusion, (2) dimension or direction of data partitioning, and (3) number of threads. Even considering only these three dimensions, there is a combinatorial explosion of optimization combinations that BTO considers. This search space is sparse, first consisting of a high ratio of illegal compared to legal programs. Within the legal programs, only a handful achieve good performance. The search space is also discrete because performance tends to cluster with no continuity between clusters. Efficiently searching this space is the goal, and doing so requires a well-designed representation.

Early versions of BTO's representation were too specific and therefore limited the performance. For example, they applied heuristics such as fusing loops at every opportunity. Experimental data show that in some cases it is best to back off from full fusion, and the representation needs to become more generic to accommodate that.

At the other extreme, we discovered that an overly generic representation leads to the evaluation of an intractable number of illegal versions. For example, we tried a string-of-digits representation that we describe in the next subsection. With this approach, search time was dominated by the identification and discarding of illegal programs.

Figure 4 shows a graphical representation of an overly general search space and what area of that search space BTO currently searches. The gray areas represent illegal programs. This area is large and, spending time in it makes search times intractable. This section describes a representation that allows BTO to spend time only on the section labeled *BTO Considered Search Space*, which contains many fewer illegal points. Finally to further improve search times, within the legal space, BTO prunes points it deems unlikely to be unprofitable. The rest of this section walks through the findings that led to our current approach, as well as the representation that enables a scalable search.




Fig. 4. Representation of a typical search space showing how BTO avoids spending time searching a large portion of illegal points.

B. Features of Search Space

In an effort to interface with existing search tools, we initially represented every fusion and partitioning decision in an easy to manipulate set of digits. For fusion we used an

adjacency matrix that created $((n-1) * n)/2$ digits, where n is the number of operations; partitions were represented with $2n$ digits, where each operation had a direction of partition and a thread count. As an example, a three-operation program would be represented as follows.

$$f, f, f, w, t, w, t, w, t$$

$$\begin{array}{lcl} 0 & \leq & f \leq \text{max_loop_nest_depth} \\ 1 & \leq & w \leq 3 \\ 0 & \leq & t \leq \text{max_thread_count} \end{array}$$

In the presence of one level of data partitioning and for a matrix-vector type operation with a maximum thread count of 8, there are over 1.2 million combinations of loop fusion and thread parallelism.

The primary advantage to this approach was that a search tool could easily manipulate these strings of digits with no domain knowledge. Unfortunately, search time was dominated by discarding illegal points. Many of the illegal points were caused by a disrespect of interaction between digits. We now summarize two important features that this representation does not encode.

Fusion is an equivalence relation. If an operation a is fused with operation b and b is fused with c , then c must be fused with a . Consider a three-operation program and the representation of fusion with an adjacency matrix M , where $M[a, b]$ shows the depth of fusion between the loop nests of a and b . Below, we show a valid fusion choice on the left and an invalid fusion choice on the right. Each value in the matrix specifies fusing up to two levels of nested loops. The matrix on the left describes fusing the outer loop of all three operations, but only b and c have the inner loop fused. The matrix on the right indicates fusing the inner loop of a with b and b with c , but not a and c , which of course is impossible. We can describe these constraints as forcing the relation specified in the adjacency matrix to be an equivalence relation at every depth.

	a	b	c
a		1	1
b			2
c			

	a	b	c
a		2	1
b			2
c			

Fused operations must use the same number of threads.

Consider a fuse graph that specifies a fusion of operations a and b but then a partition that specifies a use 4 threads and b use 6 threads. Partitioning the two operations with different thread counts guarantees that fusion of these two operations is not possible.

Given the previous example program of three matrix-vector operation with a maximum thread count of 8, respecting these features will bring the possible points to consider down to a little over 1,000, or less than one-tenth of a percent of points to consider without respecting these features.

C. Domain Specific Representation

Designing a representation that respects the previously discussed features requires domain knowledge. At the expense

of having to custom-build search tools, we designed a representation to disallow, with no search time required, a large number of illegal points.

Loop fusion is represented by *fuse sets*. Each operation is given a unique identifier, and loops are represented with $\{\}$. A single loop operation (e.g., dot product) is represented as $\{ID\}$, where ID is a number identifying an operation node in the dataflow graph. A two-loop operation, such as a matrix-vector product, is represented as $\{\{ID\}\}$. When discussing a specific $\{\}$, we annotate it using $\{i\}$, where i describes an axis of the iteration space. We use i to describe the iteration over rows of a matrix and j for columns of a matrix. A complete iteration space for a matrix can be described as $\{i\{j\}\}$ or $\{j\{i\}\}$.

Fusion is described by putting two operations within the $\{\}$. For example, outer-loop fusion of two matrix-vector products is described by $\{j\{i1\}\{i2\}\}$, and fusion including the inner loops is described by $\{j\{i1\ 2\}\}$. This notation encodes the equivalence relation of loop fusion, disallowing a huge number of illegal fusion combinations.

In BTO, fuse sets are actually more general than described in the previous paragraph. In addition to representing loops, fuse sets can represent iterations over tiles, spawning threads for data parallelism, or loop unrolling. We refer to increasing the dimensionality of the iteration space in this way as “partitioning” since it conceptually cuts a matrix or vector into smaller parts. A matrix-vector operation of $\{i\{j1\}\}$ can be partitioned as $\{p(i)\{i\{j1\}\}\}$ or $\{p(j)\{i\{j1\}\}\}$, where the $\{\}$ s annotated with $p(i)$ and $p(j)$ describe the new iteration dimension and the existing i or j loop variables that the partition affects. The search tool must specify which existing loop is being modified and how many threads should be used. The important point here is that we can represent any level of nesting and describe both C loops and data parallelism. By extending the fuse set representation to partitioning, thread counts can be assigned to each set, eliminating the consideration of points with mismatched thread counts within a fused operation.

BTO uses this representation to enumerate or manipulate the fuse sets and to generate the search space. This approach allows BTO to never touch the majority of the illegal points it encountered with more general-purpose search tools.

D. Discarding Remaining Illegal Points

Recall Figure 4 where the representation applied by BTO reduces the search space to the area labeled *BTO Considered Search Space*. In this search space, a significant number of illegal points remain. Identifying them as early as possible is key to a fast search. This section describes how BTO discards the remaining illegal points. Figure 2 shows the dataflow graph for the BATAX operation $y \leftarrow \beta A^T A x$ first described in Section II. Figure 5 shows each operation in BATAX numbered according to its corresponding number in the dataflow graph. Let us assume for simplicity that subgraphs are fixed. Thus, although the scaling by β could be located differently in the graph, in this example it cannot.

BTO performs type inference on the initial dataflow graph to check whether the input program makes sense, assigning types

```

t0 = A * x           1
t1 = A' * t0         2
y = t1 * beta        3

```

Fig. 5. Operation listing for $y \leftarrow \beta A^T A x$.

to all operations in the process. As BTO considers different optimization choices, it incrementally updates the types to determine quickly whether an optimization choice results in incompatible types.

In particular, illegal data dependency chains can be created with the fuse set representation and therefore must be checked against the data flow graph for correctness. The following is a partial list of the possible fuse sets for the running example.

```

a :  {{{1}}}{{{2}}}{{{3}}}
b :  {{{1}}}{{{2}}}{{{3}}}
c :  {{{12}}}{{{3}}}
d :  {{{1}}}{{{3}}}{{{2}}}
e :  {{{1}}}{{{2}}}{{{3}}}
f :  {{{123}}}

```

Fuse set d says to fuse operations 1 and 3. However, referring to the dataflow graph in Figure 2, one can see that there is a data dependency (operation 2) between 1 and 3.

A more subtle data dependency is caused by reduction operations. Figure 6 shows the pseudocode for the example. Examination of the outer loops (lines 1 and 4) show that the iterations are compatible and are legal to fuse. Looking at the inner loops (lines 2 and 5) we see compatible loops and assume fusion is possible. However, on line 3, $t0[i]$ is the destination of an accumulation and is not available for use until the inner loop is complete. The next operation consumes this result and so the inner loops cannot be fused.

```

for i in 1 to M           1
  for j in 1 to N         2
    t0[i] += A[i,j] * x[j] 3
  for i in 1 to M         4
    for j in 1 to N       5
      t1[j] += A[i,j] * t0[i] 6
    for j in 1 to N       7
      y[j] = t1[j] * beta   8

```

Fig. 6. Pseudocode for unfused operations as shown in Figure 5.

The introduction of loops, the type inference, and the legality of partition introduction are all based on the underlying type system employed by BTO. This system is described in detail in previous papers [5]. Briefly, a set of rules describes legal linear algebra operations based on the types involved in the operation. Certain rules cause a reduction, so an examination of the types involved in an operation provides the loop nests and flags any loops as performing a reduction. In order to catch the reduction data dependency, data flow analysis is combined with the result of examining the type to determine that results are the destination of a reduction and that fusion cannot occur.

The legality of every partitioning must also be checked for each operation. In the absence of fusion, doing so is simply of a matter of checking the type of each operand and the result

of a given operation. The challenge is in identifying the set of partitions for each operation such that fusion remains possible. The first operation of the BATAX example, $t0 = A \times x$, can be partitioned in the following ways.

$$\begin{array}{ll} (1) & t0(p) = A(p, :) \times x \quad \left| \quad \{p(i)\{i\{j1\}\}\} \right. \\ (2) & t0 += A(:, p) \times x(p) \quad \left| \quad \{p(j)\{i\{j1\}\}\} \right. \end{array}$$

Here we show the slicing of the matrix using the colon notation for a complete iteration and p for a subblock on which to operate in parallel (borrowing notation from MATLAB). On the right is the representation as a fuse set. Partitioning (1) cuts the rows of A and vector $t0$ while the second cuts the columns of A and the vector x . Partitioning (2) leads to a reduction at the parallel level, so $t0$ is not available for use until after a join. The second operation of the example, $t1 = A \times t0$ can be partitioned in the following ways.

$$\begin{array}{ll} (3) & t1(p) = A(p, :) \times t0 \quad \left| \quad \{p(j)\{i\{j2\}\}\} \right. \\ (4) & t1 += A(:, p) \times t0(p) \quad \left| \quad \{p(i)\{i\{j2\}\}\} \right. \end{array}$$

The question is how to partition operations 1 and 2 so that they can achieve fusion. Data dependence analysis says that the partition of operation 1, which introduces a reduction, will cause fusion to fail, so operation 1 must be partitioned by using method (1) thus limiting the options for operation 2. The matrix A is shared so, to achieve fusion after partitioning, A needs to be accessed the same way in both partition loops. From partitioning (1) we see that A is accessed as $A(p, :)$. Because operation 2 accesses the transpose of A , we must select partitioning (4), accessing A as $A(:, p)$. By examining the $\{ \}$ notation we similarly observe that the partitions introduced in (1) and (4) both generate $\{p(i)\}$. In large fuse sets, the likelihood of finding a correct set of operation partitions randomly is small.

BTO uses a similar approach to that used in the BATAX example. At the start of a search, BTO enumerates the possible ways of partitioning for each individual operation. Then, when given a set of operations to fuse in the presence of partitioning, a list of operation partitionings that will allow fusion is found efficiently by comparing the shared data structures in the operation (e.g., the matrix A in BATAX). This list may consist of zero to many combinations that work for a fuse set, but all will be legal. This approach quickly rules out the illegal combinations, leaving only the legal points to consider.

E. Discarding Unprofitable Points

We again refer back to Figure 4, this time considering the *BTO Legal Points*, a small section labeled *BTO Pruned* represents legal points that typically exhibit poor performance. BTO uses a handful of heuristics to prune these poorly performing points. The first heuristic is to perform fusion only on operations that share an operand. For example, if one loop writes to a temporary matrix and another loop reads from the temporary, then fusing the two loops reduces memory traffic. Similarly, if two loops read from the same matrix, then fusion is likely to be profitable. On the other hand, fusing loops that do not share an operand is unlikely to reduce memory traffic.

The next heuristic is that array contraction is always applied to temporary data structures in the presence of fusion. Again, reducing memory traffic almost always improves performance.

The second two heuristics eliminate points without having to spend any time on those that are unprofitable. The array contraction is always performed while the contiguous traversal is encoded in the type system exploited by BTO.

IV. GENETIC/GREEDY SEARCH STRATEGY

This section describes the BTO search strategy based on a genetic algorithm whose initial population is determined by a greedy search that tries to maximally fuse loops. We refer to this search strategy as MFGA, for Maximal Fusion followed by Genetic Algorithm. Section V explores why this search is used, and the value of heuristics and alternatives.

We explain MFGA using the $y \leftarrow \beta A^T A x$ BATAX example from the previous section. Genetic algorithms are a broad category of global optimization techniques inspired by biological evolution [25]. In genetic algorithms, each code version is called an organism. A genetic algorithm uses a population of organisms. At each generation, the worst organisms are removed from the population and are replaced with newly generated organisms.

A. Max Fuse

The search begins with a greedy Max-Fuse (MF) heuristic: we attempt to fuse as many of the loops as possible to the greatest depth possible, subject to the constraints described in Section III. The MF search starts from unfused but partitioned versions of the kernel in which the axis of partitioning has not yet been decided. Continuing with the BATAX example from the previous section, the following represents the unfused partitioned kernel. The X , Y , and Z are unknowns determined during the MF search.

$$\{x\{i\{j1\}\}\}\{y\{i\{j2\}\}\}\{z\{j3\}\}$$

To fuse the X and Y iterations, we need $X = Y$, so we proceed with the fusion and constrain ourselves to $X = Y$.

$$\begin{aligned} & \{x\{i\{j1\}\}\}\{x\{i\{j2\}\}\}\{z\{j3\}\} \\ \Rightarrow & \{x\{i\{j1\}\}\{i\{j2\}\}\}\{z\{j3\}\} \end{aligned}$$

At this point, X has to be $p(i)$ because the alternative, $p(j)$, would mean that the necessary results from operation 1 would not be available for operation 2. Next, we can also fuse the i iteration of operations 1 and 2.

$$\begin{aligned} & \{p(i)\{i\{j1\}\}\{i\{j2\}\}\}\{z\{j3\}\} \\ \Rightarrow & \{p(i)\{i\{j1\}\{j2\}\}\}\{z\{j3\}\} \end{aligned}$$

Because of the reduction in the matrix-vector product (operation 1), the j iteration of operations 1 and 2 cannot be fused.

Next, we consider whether the $p(i)$ iteration can be fused with Z . The $p(i)$ iteration requires a reduction before the final vector scaling of operation 3, so 3 must reside in its own thread. Finally, there is only one axis of iteration in operation 3, so Z must be $p(j)$. Therefore, the MF search produces the following organism: $\{p(i)\{i\{j1\}\{j2\}\}\}\{p(j)\{j3\}\}$.

B. Generation of Initial Population by Mutation

The initial population for the genetic algorithm is created by applying random mutation to the Max-Fuse point. Each mutation performs *one* of the following four changes: (1) add or remove fusion level, (2) add or remove partition level, (3) change the partition axis, or (4) change the number of threads. Mutations are constrained to the set of legal organisms; for example, attempting to add a mutation to the already maximally fused point from our previous example will fail, resulting in no change. However, mutations might randomly remove the partition from operation 3:

$$\{_{p(i)}\{i\{j1\}\{j2\}\}\}_{p(j)}\{j3\}\} \Rightarrow \{_{p(i)}\{i\{j1\}\{j2\}\}\}_{j3}\}$$

The Random search described in Section V-D2 consists of repeatedly applying random mutations to the organism without any further search.

C. Selection and Crossover

After the initial generation of organisms (and for every following generation), we compile and test every organism and record its runtime, which is the value the search tries to minimize. We then select $2N$ of the fittest organisms to be parents for the next generation, where the population size N can be user specified, but defaults to 20.

a) *Parent Selection Method*: The population evolves via tournament selection [25]: k random organisms are chosen to be potential parents, and the potential parent with the best fitness becomes an actual parent. This process balances hill climbing with exploration, allowing less fit organisms to sometimes become parents, and thus helping the algorithm escape locally optimal solutions that are not globally optimal. Larger values of k cause the algorithm to converge more quickly on a solution, while smaller values of k converge more slowly but increases exploration. BTO uses $k = 2$ to favor exploration.

b) *Crossover*: Crossover is a function that takes two parent organisms and randomly chooses features of the two parents to create a child organism. The key strength of genetic algorithms is that crossover can sometimes combine the strengths of two versions. Our crossover function generates the child recursively from the two parents, making fusion decisions at each level and making sure those decisions remain valid for inner levels.

Our crossover function uses the type-based representation of each expression as described in Section II and performs crossover by comparing the two types. Continuing with the $y \leftarrow \beta A^T A x$ example, consider the following two organisms A and B.

$$\begin{aligned} A &: \{_{p(i)}\{i\{j1\}\{j2\}\}\}_{j3}\} \\ B &: \{_{p(i)}\{i\{j1\}\}\}\}_{p(i)}\{i\{j2\}\}\}_{p(j)}\{j3\}\} \end{aligned}$$

Parent A partitions and partly fuses operations 1 and 2 but does not partition operation 3. Parent B has all partitions turned on but has not fused operations 1 and 2.

Crossover chooses which parent to emulate for each operation, working from the outermost fuse level inwards. Each step constrains the possibilities for the other operations. In our

example, crossover might choose parent A for the outermost level of operation 1, meaning 1 and 2 exist in the same thread (also using Parent A's partitioning axis $p(i)$ and thread number data). It then might choose Parent B for the next level, iteration i . This mechanism forces operation 1 and 2 not to be fused, resulting in $\{i\{j1\}\}\{i\{j2\}\}$. Then the crossover moves to operation 3 and the process continues. If B is chosen, the final child becomes $\{_{p(i)}\{i\{j1\}\}\{i\{j2\}\}\}_{p(j)}\{j3\}\}$.

The tournament selection process is repeated N times, creating a new generation of organisms. Fitness values are cached. If crossover ever produces an organism that was already tested in a previous generation, the old value is used to save search time.

D. Search for Number of Threads

BTO uses a fixed number of threads to execute all of the data-parallel partitions in a kernel. We refer to this as the *global thread number* heuristic. An alternative is to allow different numbers of threads for each partition, which we refer to as the *exhaustive thread* search. In Section V-D3, we present data that show that the exhaustive approach takes much more time but does not lead to significantly better results.

BTO includes the search for the best number of threads in the MFGA algorithm. The initial number is set to the number of cores in the target computer architecture. The mutation function either increments or decrements the thread number of 2. The crossover function simply picks the thread number from one of the parents. After the genetic algorithm completes, MFGA performs an additional search for the best number of threads by testing the performance when using thread counts between 2 and the number of cores, incrementing by 2.

V. RESULTS

We begin this section with a comparison of the performance of BTO-generated routines and several state-of-the-art tools and libraries that perform similar sets of optimizations, as well as hand-optimized code. The BTO compiler generates code that is between 16% slower and 39% faster than hand-optimized code. The other automated tools and libraries achieve comparable performance to BTO and hand-optimized code for only a few of the kernels. For the smaller kernels in which we can exhaustively enumerate all possible combinations of optimizations, we show that the MFGA search method finds a routine that performs within 2% of the best found in the exhaustive search. We present empirical data that motivates our choice of the MFGA search strategy, comparing MFGA to several alternative strategies and analyzing the orthogonality of fusion choices versus number of threads.

A. Test Environment and Kernels

The results in this section rely on the kernels shown in Table I. Some of these kernels respond well to loop fusion and data parallelism while others do not. Some of these kernels are in the updated BLAS [8] but have not been adopted by vendor-tuned BLAS libraries. These kernels also represent various uses of the BLAS. For example, the DGEMV kernel maps

```
// q = A * p
dgemv('N',A_nrows,A_ncols,1.0,A,lda,p,1,0.0,q,1);
// s = A' * r
dgemv('T',A_nrows,A_ncols,1.0,A,lda,r,1,0.0,s,1);
```

Fig. 7. Example sequence of BLAS calls that implement BICGK.

directly to a BLAS call while others are equivalent to multiple BLAS calls. As an example, Listing 7 shows the sequence of BLAS calls that implement the BICGK kernel.

TABLE I
KERNEL SPECIFICATIONS.

Kernel	Operation
AXPYDOT	$z \leftarrow w - \alpha v$ $\beta \leftarrow z^T u$
VADD	$x \leftarrow w + y + z$
WAXPBY	$w \leftarrow \alpha x + \beta y$
ATAX	$y \leftarrow A^T A x$
BICGK	$q \leftarrow A p$ $s \leftarrow A^T r$
DGEMV	$z \leftarrow \alpha A x + \beta y$
DGEMVT	$x \leftarrow \beta A^T y + z$ $w \leftarrow \alpha A x$
GEMVER	$B \leftarrow A + u_1 v_1^T + u_2 v_2^T$ $x \leftarrow \beta B^T y + z$ $w \leftarrow \alpha B x$
GESUMMV	$y \leftarrow \alpha A x + \beta B x$

The computers used for testing include recent AMD and Intel multicore architectures which we describe in Table II.

TABLE II
SPECIFICATIONS OF THE TEST MACHINES.

Processor	Cores	Speed (GHz)	L1 (KB)	L2 (KB)	L3 (MB)
Intel Westmere	24	2.66	12 x 32	12 x 256	2 x 12
AMD Phenom II X6	6	3.3	6 x 64	6 x 512	1 x 6
AMD Interlagos	64	2.2	64 x 16	16 x 2048	8 x 8

B. Comparison to Similar Tools

We begin by placing BTO performance results in context by comparing them with several state-of-the-art tools and libraries. BTO performs loop fusion and array contraction and makes use of data parallelism. BTO relies on the native compiler for loop unrolling and vectorization. The two compilers used to gather this data are the Intel C Compiler (ICC) [19] and the PGI C Compiler (PGCC) [27]. With the exception of the explicitly labeled PGCC results, all kernels are compiled using ICC. Both ICC and PGCC unroll loops and vectorize. They also identify and exploit data parallelism and perform loop fusion.

We begin with a detailed comparison of BTO and five other approaches for generating high performance code on the Intel Westmere. We then give a brief summary of similar results on the AMD Phenom and Interlagos.

The first approaches are using ICC and PGCC, which represent the best commercial compilers. The third approach is using Pluto [9], a source-to-source translator capable of performing loop fusion and identifying data parallelism. The fourth

approach is using Intel’s Math Kernel Library (MKL) [19] which is a vendor-tuned BLAS implementation targeting Intel CPUs. The fifth is a hand-tuned implementation (applying loop fusion, array contraction, and data parallelism) created by an expert in performance tuning who works in the performance library group at Apple, Inc. The input for ICC, PGCC and Pluto is a translation of the BLAS calls to C loops. The compiler flags used with ICC were “-O3 -mkl -fno-alias” and the flags for PGCC were “-O4 -fast -Mipa=fast -Mconcur -Mvect=fuse -Msafepr” (“-Msafepr” not used on Interlagos). Data parallelism is exploited by ICC, PGCC, Pluto, and MKL by using OpenMP [12]. BTO and the hand-tuned versions use Pthreads [26].

Figure 8 shows the speedup relative to ICC on the y-axis for several linear algebra kernels. (ICC performance is 1.) On the left are the three vector-vector kernels, and on the right are the six matrix-vector kernels, all from Table I.

PGCC tends to do slightly better than ICC, with speedups ranging from 1.1 to 1.5 times faster. Examining the output of PGCC shows that all but GESUMMV and GEMVER were parallelized. However, PGCC’s ability to perform loop fusion was mixed; it fused the appropriate loops in AXPYDOT, VADD, and WAXPBY but complained of a “complex flow graph” on the remaining kernels and only achieved limited fusion.

The MKL BLAS outperforms ICC by factors ranging from 1.4 to 4.2. The calls to BLAS routines prevent loop fusion, so significant speedups, such as those observed in AXPYDOT and GESUMMV, can instead be attributed to parallelism and well tuned vector implementations of the individual operations. We were unable to determine why the BLAS perform so well for AXPYDOT. Surprisingly, the BLAS DGEMV does not perform as well as Pluto and BTO. Given the lack of fusion potential in this kernel, we speculate that differences in the parallel implementations are the cause.

The Pluto results show speedups ranging from 0.7 to 5.7 times faster than ICC. The worst-performing kernels are AXPYDOT, ATAX, and DGEMVT. These three kernels represent the only cases where Pluto did not introduce data parallelism. For the remaining two vector-vector kernels, VADD and WAXPBY, Pluto created the best-performing result, slightly better than the BTO and hand-tuned versions. Inspection shows that the only difference between Pluto, hand-tuned, and BTO in these cases was the use of OpenMP for Pluto and Pthreads for hand-tuned and BTO. The fusion was otherwise identical and the difference in thread count had little effect. For the matrix-vector operations, if we enabled fusion but not parallelization with Pluto’s flags, then Pluto matched BTO with respect to fusion. However, with both fusion and parallelization enabled, Pluto sometimes misses fusion and/or parallelization opportunities. For example BICGK was parallelized but not fused. The GEMVER results depend on the loop ordering in the input file. For GEMVER, Pluto performed either complete fusion with no parallelism or incomplete fusion with parallelism; the latter provided the best performance and is shown in Figure 8.

The hand-tuned implementation is intended as a sanity check. For the vector-vector operations, the hand-tuned version

TABLE III

PERFORMANCE DATA FOR AMD PHENOM. BLAS NUMBERS FROM AMD'S ACML. SPEEDUPS RELATIVE TO UNFUSED LOOPS COMPILED WITH PGCC (PGCC PERFORMANCE IS 1 AND NOT SHOWN). BEST PERFORMING VERSION IN BOLD.

Kernel	BLAS	Pluto	HAND	BTO
AXPYDOT	0.97	1.81	1.58	1.86
VADD	0.84	1.33	1.50	1.83
WAXPBY	0.79	1.40	1.68	1.91
ATAX	1.27	0.69	2.92	2.92
BICGK	1.27	0.80	2.80	2.84
DGEMV	1.67	0.71	1.85	1.89
DGEMVT	1.67	0.71	1.85	1.89
GEMVER	1.04	1.61	2.61	2.34
GESUMMV	1.63	0.63	1.74	1.75

is within a few percent of the best implementation. Typically the fusion in both the hand tuned and the best tool based version are identical with the primary difference being either thread count or what appears to be a difference between Pthreads and OpenMP performance. In the case of the matrix-vector operations, the hand-tuned version is the best for all but DGEMV and GESUMMV, where it is equal to the best.

The BTO performance results show speedups ranging from 3.2 to 6.9 times faster than ICC. For the vector-vector operations, the performance is similar to the hand-tuned version in all cases. Inspection shows that for AXPYDOT, BTO was slightly faster than the hand-tuned version because BTO did not fuse the inner loop while the hand-tuned version did. BTO performed slightly worse than hand-tuned on WAXPBY because of a difference in thread count. Similarly, the performance of the matrix-vector operations is close but slightly lower than that of the hand-tuned version. BTO fused the same as hand-tuned for BICGK, GEMVER and DGEMVT with the only difference being in thread count. For ATAX, both BTO and hand-tuned fused the same and selected the same number of threads, but BTO was slightly slower because of where it zeroed out a data structure. In the hand-tuned version the zeroing occurred in the threads, while in BTO's case it occurred in the main thread.

Similar results on AMD Phenom and AMD Interlagos are shown in Table III and Table IV, respectively. The Pluto-generated code for the matrix-vector operations tended to perform worse than that produced for the other methods evaluated. On this computer, achieving full fusion while maintaining parallelism is of great importance. As previously discussed, Pluto tended to achieve fusion or parallelism but struggled with the combination. These results demonstrate the difficulty of portable high-performance code generation even under autotuning scenarios.

Summary: Compared with the best alternative approach for a given kernel, BTO performance ranges from 16% slower to 39% faster. Excluding hand-written comparison points, BTO performs between 14% worse and 229% better. Pluto, ICC, PGCC, and BLAS all achieve near-best performance for only a few points; however, BTO's performance is most consistent across kernels and computers. Excluding the hand-optimized results, BTO finds the best version for 7 of 9 kernels on the Intel Westmere, all 9 kernels on the AMD Phenom, and 7 of

TABLE IV

PERFORMANCE DATA FOR AMD INTERLAGOS. BLAS NUMBERS FROM AMD'S ACML. SPEEDUPS RELATIVE TO UNFUSED LOOPS COMPILED WITH PGCC (PGCC PERFORMANCE IS 1 AND NOT SHOWN). BEST PERFORMING VERSION IN BOLD.

Kernel	BLAS	Pluto	HAND	BTO
AXPYDOT	0.82	1.60	1.73	1.61
VADD	0.43	1.05	1.14	1.15
WAXPBY	0.34	1.06	1.16	1.11
ATAX	2.49	0.43	4.09	4.28
BICGK	2.35	1.60	3.03	4.22
DGEMV	2.45	0.89	1.66	2.07
DGEMVT	2.43	0.43	4.08	4.03
GEMVER	1.70	2.00	4.15	4.05
GESUMMV	2.36	0.37	1.65	2.03

9 kernels on the AMD Interlagos. Surprisingly, on the AMD Phenom, BTO surpassed the hand-optimized code for 7 of the 9 kernels and tied for one kernel.

C. MFGA Compared to Exhaustive Searches

In Section V-B, we presented results showing that BTO's MFGA search strategy finds high-performing versions for a range of kernels. In this section, we show how the performance of the MFGA search strategy compares with the best version that can be produced using exhaustive or nearly exhaustive search strategies on Intel Westmere. These strategies require long-running searches that can take days to complete. For the smaller kernels, a completely exhaustive search is possible. For larger kernels, exhaustive search was not feasible, so we instead use a strategy that is exhaustive with respect to each optimization, but orthogonal between optimizations. For the largest kernels, GEMVER and DGEMV, even the orthogonal approach took too much time, not completing even after weeks of running.

We compared the performance of kernels produced by MFGA as percentage of the exhaustive search for smaller kernels or as a percentage of the orthogonal search for larger kernels such as DGEMVT and GESUMMV. The results show that scalable search produces kernel performance within 1-2% of the best performance.

D. Evaluation of Search Methods

In the previous sections, we demonstrated that BTO is capable of generating high-performance routines. In this section, we examine the data that led to creating the MFGA search strategy. All of the experiments in this section were performed on the Intel Westmere.

1) *Orthogonality of Fusion and Thread Search:* The MFGA strategy, for the most part, treats decisions regarding fusion and thread count orthogonally, which significantly reduces the size of the search space. However, before we could employ this search method, we first had to show that it would lead to no degradation in performance.

We define *orthogonal search* as first searching only the fusion dimension, then using only the best candidate, searching every viable thread count. We evaluated the effectiveness and




Fig. 8. Performance data for Intel Westmere. Speedups relative to unfused loops compiled with ICC (ICC performance is 1 and not shown). The left three kernels are vector-vector while the right six are matrix-vector operations. In all cases, BTO generates code that is between 16% slower and 39% faster than hand-optimized code and significantly faster than library and compiler-optimized versions.

search time of the orthogonal search as compared to an exhaustive search using the smaller kernels: ATAX, AXPYDOT, BICGK, VADD, and WAXPBY. For all kernels, orthogonal search found the best-performing version while taking 1-8% of the time of exhaustive search, demonstrating that searching the space orthogonally dramatically reduces search time without sacrificing performance. This reduction in search time results in part from the chosen orthogonal ordering. By searching the fusion space first, we often dramatically reduce the number of data-parallel loops and hence the size of the subsequent thread-count search space.

Thus, we see that fusion and thread search can be conducted orthogonally without a significant loss of kernel performance.

2) *Fusion Search*: Next we focus on fusion strategies. In this section we analyze our choice of using a combination of a genetic algorithm and the max-fuse heuristic.

We compare four search strategies on our most challenging kernel, GEMVER. In particular, we test random search, our genetic algorithm without the max-fuse heuristic, the max-fuse heuristic by itself, and the combination of the max-fuse heuristic with the genetic algorithm (MFGA). As described in Section IV, the random search strategy and the genetic algorithm use the same mutation schemes, and thus their comparison shows the benefit of the crossover and selection methods.

Figure 9 shows the performance over time of each of the search methods. (MF is a single point near 3 GFLOPS.) Because the search is stochastic, each of the lines in the chart is the average of two runs. MFGA finds the optimal point in less than 10 minutes on average. Without the MF heuristic, GA alone eventually reaches 90% of MFGA but requires over an hour of search time. The Random search plateaus without ever finding the optimal value. The MF heuristic by itself achieves 40% of MFGA.

In conclusion, a combination of GA and MF is the best strategy for the fusion portion of the search.

3) *Thread Search*: Using the MFGA heuristic described in the previous section, we explore several possible thread search strategies, including the *global* thread number and the *exhaustive* strategies discussed in Section IV-D. The baseline test is the MFGA search with number of threads set equal

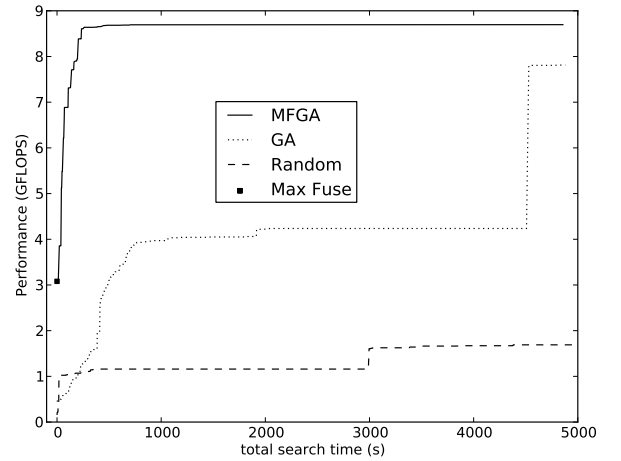


Fig. 9. GEMVER performance over time for different search strategies on Intel Westmere. MFGA finds the best version more quickly and consistently than either search individually.

to the number of cores (24 for these experiments), which we refer to as the *const* strategy. Recall that the *global* strategy starts with MFGA and then searches over a single parameter for all loop nests for the number of threads. Recall that the *exhaustive* search replaces the single thread parameter with the full space of possible thread counts, i.e., considering the number of threads for each loop nest individually.

The results for seven kernels are in Figure 10. The top chart shows the final performance of the best version found in each case.

Searching over the thread space improves the final performance compared with using a constant number of threads (e.g., equal to the number of cores), with negligible difference in kernel performance between the global thread count (fixed count for all threads) and fully exhaustive approaches (varying thread counts for different operations). The bottom chart in Figure 10 shows the total search cost of the different thread search approaches, demonstrating that global thread search improves scalability without sacrificing performance.




Fig. 10. Best runtime (top) and search time (bottom) for exhaustive and global searches. A constant thread number (e.g., equal to the number of cores) cannot achieve the runtime performance of either global or exhaustive thread search. Searching over a global thread count results in a much shorter search time without significantly worsening kernel performance.

VI. RELATED WORK

We describe the relationship between our contributions in this paper and related work in four areas of the literature: loop restructuring compilers, search strategies for autotuning, partitioning matrix computations, and empirical search.

Loop Fusion and Parallelization: Megiddo and Sarkar [24] study the problem of deciding which loops to fuse in a context where parallelization choices have already been made (such as an OpenMP program). They model this problem as a weighted graph whose nodes are loops and whose edges are labeled with the runtime cost savings resulting from loop fusion. Because the parallelization choices are fixed prior to the fusion choices, their approach sometimes misses the optimal combination of parallelization and fusion decisions.

Darte and Huard [13], on the other hand, study the space of all fusion decisions followed by parallelization decisions. Pouchet et al. [28] take a similar approach, they use an orthogonal approach that exhaustively searches over fusion decisions, then uses the polyhedral model with analytic models to make tiling and parallelization decisions. These approaches roughly correspond to the orthogonal search technique described in Section V-D1.

Bondhugula et al. [9] employs the heuristic of maximally fusing loops. Loop fusion is generally beneficial, but too much can be detrimental as it can put too much pressure on registers and cache [21].

Bondhugula et al. [10] develop an analytic model for predicting the profitability of fusion and parallelization and show speedups relative to other heuristics such as always fuse and never fuse. However, they do not validate their model against the entire search space as we do where possible here.

Search for Autotuning: Vuduc et al. [34] study the optimization space of applying register tiling, loop unrolling, software pipelining, and software prefetching to matrix multiplication. They show that this search space is difficult (a very

small number of combinations achieve good performance), and present a statistical method for determining when a search has found a point that is close enough to the best.

Balaprakash et al. [4] study the effectiveness of several search algorithms (random search, genetic algorithms, Nelder-Mead simplex) to find the best combination of optimization decisions from among loop unrolling, scalar replacement, loop parallelization, vectorization, and register tiling as implemented in the Orio autotuning framework [17]. They conclude that the modified Nelder-Mead method is effective for their search problem.

Chen et al. [11] develop a framework for empirical search over many loop optimizations such as permutation, tiling, unroll-and-jam, data copying, and fusion. They employ an orthogonal search strategy, first searching over unrolling factors, then tiling sizes, and so on. Tiwari et al. [33] describe an autotuning framework that combines ActiveHarmony’s parallel search backend with the CHiLL transformation framework.

Looptool [29] and AutoLoopTune [30] support loop fusion, unroll-and-jam and array contraction. AutoLoopTune also supports tiling. POET [37] also supports a number of loop transformations.

Partitioning Matrix Computations: The approach to partitioning matrix computations described in this paper is inspired by the notion of a blocked matrix view in the Matrix Template Library [31]. Several researchers have subsequently proposed similar abstractions, such as the hierarchically tiled arrays of Almasi et al. [1] and the support for matrix partitioning in FLAME [16].

Search with Empirical Evaluation: Bilmes et al. [7] and Whaley and Dongarra [35] autotune matrix multiplication using empirical evaluation to determine the profitability of optimizations. Zhao et al. [38] use exhaustive search and empirical testing to select the best combination of loop fusion decisions. Yi and Qasem [36] apply empirical search to determine the profitability of optimizations for register reuse, SSE vectorization, strength reduction, loop unrolling, and prefetching. Their framework is parameterized with respect to the search algorithm and includes numerous search strategies.

VII. CONCLUSIONS AND FUTURE WORK

For many problems in high-performance computing, the best solutions require extensive testing and tuning. We present an empirical autotuning approach for dense matrix algebra that is reliable and scalable. Our tool considers loop fusion, array contraction, and shared memory parallelism.

Our experiments have shown that the BTO autotuning system outperforms standard optimizing compilers and a vendor-optimized BLAS library in most cases, and our results are competitive with hand-tuned code. We also describe how we developed our search strategies and tested the usefulness of each part of the search.

Two big expansions of functionality are planned: distributed memory support and extension of matrix formats to include triangular, banded, and sparse. These extensions will improve the usefulness of BTO, while also providing an important stress test for the scalability of the search algorithms and code generation.

REFERENCES

- [1] Gheorghe Almasi, Luiz De Rose, Jose Moreira, and David Padua. Programming for locality and parallelism with hierarchically tiled arrays. In *The 16th International Workshop on Languages and Compilers for Parallel Computing*, pages 162–176, College Station, TX, 2003.
- [2] S. Amarasinghe, D. Campbell, W. Carlson, A. Chien, W. Dally, E. Elnohazy, M. Hall, R. Harrison, W. Harrod, K. Hill, et al. Exascale software study: Software challenges in extreme scale systems. *DARPA IPTO, Air Force Research Labs, Tech. Rep.*, 2009.
- [3] W. K. Anderson, W. D. Gropp, D. K. Kaushik, D. E. Keyes, and B. F. Smith. Achieving high sustained performance in an unstructured mesh CFD application. In *Proceedings of the 1999 ACM/IEEE Conference on Supercomputing (CDROM)*, Supercomputing '99, Portland, Oregon, 1999. ACM.
- [4] P. Balaprakash, S. Wild, and P. Hovland. Can search algorithms save large-scale automatic performance tuning? *Procedia CS*, 4:2136–2145, 2011.
- [5] Geoffrey Belter, E. R. Jessup, Ian Karlin, and Jeremy G. Siek. Automating the generation of composed linear algebra kernels. In *SC '09: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, pages 1–12, New York, 2009. ACM. ISBN 978-1-60558-744-8. doi: <http://doi.acm.org/10.1145/1654059.1654119>.
- [6] Geoffrey Belter, Jeremy G. Siek, Ian Karlin, and E. R. Jessup. Automatic generation of tiled and parallel linear algebra routines. In *Fifth International Workshop on Automatic Performance Tuning (iWAPT 2010)*, pages 1–15, Berkeley, CA, June 2010.
- [7] Jeff Bilmes, Krste Asanovic, Chee-Whye Chin, and Jim Demmel. Optimizing matrix multiply using PhiPAC: A portable, high-performance, ANSI C coding methodology. In *ICS '97: Proceedings of the 11th International Conference on Supercomputing*, pages 340–347, New York, 1997. ACM Press. ISBN 0-89791-902-5. doi: <http://doi.acm.org/10.1145/263580.263662>.
- [8] L. Susan Blackford, James Demmel, Jack Dongarra, Iain Duff, Sven Hammarling, Greg Henry, Michael Heroux, Linda Kaufman, Andrew Lumsdaine, Antoine Petit, Roldan Pozo, Karin Remington, and R. Clint Whaley. An updated set of Basic Linear Algebra Subprograms (BLAS). *ACM Transactions on Mathematical Software*, 28(2):135–151, June 2002.
- [9] U. Bondhugula, A. Hartono, J. Ramanujam, and P. Sadayappan. Pluto: A practical and fully automatic polyhedral program optimization system. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation (PLDI 08)*, pages 101–113, Tucson, AZ, June 2008.
- [10] Uday Bondhugula, Oktay Gunluk, Sanjeeb Dash, and Lakshminarayanan Renganarayanan. A model for fusion and code motion in an automatic parallelizing compiler. In *Proceedings of the 19th International Conference on Parallel Architectures and Compilation Techniques*, PACT '10, pages 343–352, New York, 2010. ACM.
- [11] C. Chen, J. Chame, and M. Hall. CHILL: A framework for composing high-level loop transformations. Technical Report 08-897, Department of Computer Science, University of Southern California, 2008.
- [12] Leonardo Dagum and Ramesh Menon. Openmp: An industry-standard API for shared-memory programming. *IEEE Comput. Sci. Eng.*, 5(1):46–55, January 1998. ISSN 1070-9924. doi: 10.1109/99.660313. URL <http://dx.doi.org/10.1109/99.660313>.
- [13] Alain Darté and Guillaume Huard. Loop shifting for loop parallelization. Technical Report 2000-22, Ecole Normale Supérieure de Lyon, May 2000.
- [14] Jack J. Dongarra, Jeremy De Croz, Sven Hammarling, and Richard J. Hanson. An extended set of FORTRAN Basic Linear Algebra Subprograms. *ACM Transactions on Mathematical Software*, 14(1):1–17, March 1988.
- [15] Jack J. Dongarra, Jeremy Du Croz, Sven Hammarling, and Iain Duff. A set of level 3 Basic Linear Algebra Subprograms. *ACM Transactions on Mathematical Software*, 16(1):1–17, March 1990.
- [16] John A. Gunnels, Fred G. Gustavson, Greg M. Henry, and Robert A. van de Geijn. FLAME: Formal linear algebra methods environment. *ACM Trans. Math. Softw.*, 27(4):422–455, 2001.
- [17] Albert Hartono, Boyana Norris, and Ponnuswamy Sadayappan. Annotation-based empirical performance tuning using Orio. In *IPDPS '09: Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing*, pages 1–11, Washington, DC, 2009. IEEE Computer Society. ISBN 978-1-4244-3751-1. doi: <http://dx.doi.org/10.1109/IPDPS.2009.5161004>. URL <http://www.mcs.anl.gov/uploads/cels/papers/P1556.pdf>. Also available as Preprint ANL/MCS-P1556-1008.
- [18] Gary W. Howell, James W. Demmel, Charles T. Fulton, Sven Hammarling, and Karen Marmol. Cache efficient bidiagonalization using BLAS 2.5 operators. *ACM Trans. Math. Softw.*, 34:14:1–14:33, May 2008.
- [19] Intel. Intel Composer. <http://software.intel.com/en-us/articles/intel-compilers>, April 2012.
- [20] Ian Karlin, Elizabeth Jessup, Geoffrey Belter, and Jeremy G. Siek. Parallel memory prediction for fused linear algebra kernels. *SIGMETRICS Perform. Eval. Rev.*, 38:43–49, March 2011. ISSN 0163-5999. doi: <http://doi.acm.org/10.1145/1964218.1964226>. URL <http://doi.acm.org/10.1145/1964218.1964226>.
- [21] Ian Karlin, Elizabeth Jessup, and Erik Silken. Modeling the memory and performance impacts of loop fusion. *Journal of Computational Science*, In press, 2011. ISSN 1877-7503. doi: DOI:10.1016/j.jocs.2011.03.002.
- [22] Richard M. Karp, Raymond E. Miller, and Shmuel Winograd. The organization of computations for uniform recurrence equations. *J. ACM*, 14(3):563–590, July 1967. ISSN 0004-5411. doi: 10.1145/321406.321418. URL <http://doi.acm.org/10.1145/321406.321418>.
- [23] C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh. Basic Linear Algebra Subprograms for Fortran usage. *ACM Transactions on Mathematical Software*, 5

- (3):308–323, September 1979.
- [24] Nimrod Megiddo and Vivek Sarkar. Optimal weighted loop fusion for parallel programs. In *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '97, pages 282–291, New York, 1997. ACM.
 - [25] M. Mitchell. *An introduction to genetic algorithms*. The MIT Press, 1998. ISBN 0262631857.
 - [26] Frank Mueller. Pthreads library interface. Technical report, Florida State University, 1999.
 - [27] Portland Group. Portland group compiler. <http://www.pggroup.com>, April 2012.
 - [28] Louis-Noël Pouchet, Uday Bondhugula, Cédric Bastoul, Albert Cohen, J. Ramanujam, and P. Sadayappan. Combined iterative and model-driven optimization in an automatic parallelization framework. In *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '10, pages 1–11, Washington, DC, November 2010. IEEE Computer Society.
 - [29] Apan Qasem, Guohua Jin, and John Mellor-Crummey. Improving performance with integrated program transformations. Technical Report TR03-419, Department of Computer Science, Rice University, October 2003.
 - [30] Apan Qasem, Ken Kennedy, and John Mellor-Crummey. Automatic tuning of whole applications using direct search and a performance-based transformation system. *The Journal of Supercomputing: Special Issue on Computer Science Research Supporting High-Performance Applications*, 36(9):183–196, May 2006.
 - [31] Jeremy G. Siek. A modern framework for portable high performance numerical linear algebra. Master's thesis, University of Notre Dame, 1999.
 - [32] Jeremy G. Siek, Ian Karlin, and E. R. Jessup. Build to order linear algebra kernels. In *Workshop on Performance Optimization for High-Level Languages and Libraries (POHLL 2008)*, pages 1–8, Miami, FL, April 2008.
 - [33] Ananta Tiwari, Chun Chen, Jacqueline Chame, Mary Hall, and Jeffrey K. Hollingsworth. A scalable autotuning framework for compiler optimization. In *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium*, Rome, Italy, May 2009.
 - [34] Richard Vuduc, James W. Demmel, and Jeff A. Bilmes. Statistical models for empirical search-based performance tuning. *International Journal of High Performance Computing Applications*, 18(1):65–94, 2004. doi: 10.1177/1094342004041293. URL <http://hpc.sagepub.com/content/18/1/65.abstract>.
 - [35] R. Clint Whaley and Jack J. Dongarra. Automatically tuned linear algebra software. In *Supercomputing '98: Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)*, pages 1–27, Washington, DC, 1998. IEEE Computer Society. ISBN 0-89791-984-X.
 - [36] Qing Yi and Apan Qasem. Exploring the optimization space of dense linear algebra kernels. In *Languages and Compilers for Parallel Computing: 21th International Workshop, LCPC 2008, Edmonton, Canada, July 31 - August 2, 2008, Revised Selected Papers*, pages 343–355, Berlin, 2008. Springer-Verlag. ISBN 978-3-540-89739-2. doi: http://dx.doi.org/10.1007/978-3-540-89740-8_24.
 - [37] Qing Yi, Keith Seymour, Haihang You, Richard Vuduc, and Dan Quinlan. POET: Parameterized optimizations for empirical tuning. In *Proceedings of the Parallel and Distributed Processing Symposium, 2007*, pages 1–8, Long Beach, CA, March 2007. IEEE. doi: 10.1109/IPDPS.2007.370637.
 - [38] Y. Zhao, Q. Yi, K. Kennedy, D. Quinlan, and R. Vuduc. Parameterizing loop fusion for automated empirical tuning. Technical Report UCRL-TR-217808, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, December 2005.