

IMPLEMENTING HIGH-PERFORMANCE COMPLEX MATRIX MULTIPLICATION VIA THE 1M METHOD

FIELD G. VAN ZEE*

Abstract. Almost all efforts to optimize high-performance matrix-matrix multiplication have been focused on the case where matrices contain real elements. The community’s collective assumption appears to have been that the techniques and methods developed for the real domain carry over directly to the complex domain. As a result, implementors have mostly overlooked a class of methods that compute complex matrix multiplication using only real matrix products. This is the second in a series of articles that investigate these so-called induced methods. In the previous article, we found that algorithms based on the more generally applicable of the two methods—the 4M method—lead to implementations that, for various reasons, often underperform their real domain counterparts. To overcome these limitations, we derive a superior 1M method for expressing complex matrix multiplication, one which addresses virtually all of the shortcomings inherent in 4M. Implementations are developed within the BLIS framework, and testing on microarchitectures by three vendors confirms that the 1M method yields performance that is generally competitive with solutions based on conventionally implemented complex kernels, sometimes even outperforming vendor libraries.

Key words. high-performance, complex, matrix, multiplication, microkernel, kernel, BLAS, BLIS, 1m, 2m, 4m, induced, linear algebra, DLA

AMS subject classifications. 65Y04

1. Introduction. Over the last several decades, matrix multiplication research has resulted in methods and implementations that primarily target the real domain. Recent trends in implementation efforts have condensed virtually all matrix product computation into relatively small *kernels*—building blocks of highly optimized code (typically written in assembly language) upon which more generalized functionality is constructed via various levels of nested loops [23, 5, 3, 22, 2]. Because most effort is focused on the real domain, the complex domain is either left as an unimplemented afterthought—perhaps because the product is merely a proof-of-concept or prototype [5], or because the project primarily targets applications and uses cases that require only real computation [2]—or it is implemented in a manner that mimics the real domain down to the level of the assembly kernel [23, 3, 4].¹ Most modern microarchitectures lack machine instructions for directly computing complex arithmetic on complex numbers, and so when the effort to implement these kernels is undertaken, kernel developers encounter additional programming challenges that do not manifest in the real domain. Specifically, these kernel developers must explicitly orchestrate computation on the real and imaginary components in order to implement multiplication and addition on complex scalars, and they must do so in terms of vector instructions to ensure high performance is achievable.

Pushing the nuances and complexities of complex arithmetic down to the level of the kernel allows the higher-level loop infrastructure within the matrix multiplication to remain largely the same as its real domain counterpart. (See Figure 1.1.) However,

*Oden Institute for Computational Engineering & Sciences, The University of Texas at Austin, Austin, TX (field@cs.utexas.edu)

Funding: This research was partially sponsored by grants from Intel Corporation and the National Science Foundation (Award ACI-1550493). *Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).*

¹ Because they exhibit slightly less favorable numerical properties, we exclude Strassen-like efforts from this characterization.

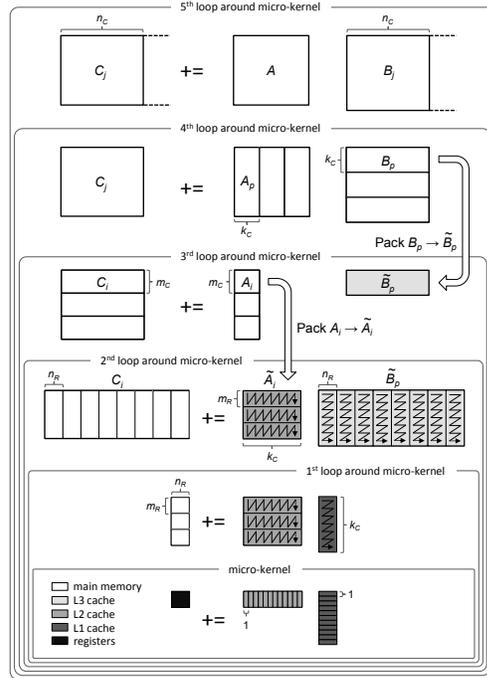


FIG. 1.1. An illustration of the algorithm for computing high-performance matrix multiplication, taken from [21], which expresses computation in terms of a so-called “block-panel” subproblem.

41 this approach also doubles the number of assembly kernels that must be written in
 42 order to fully support computation in either domain (real or complex) for the desired
 43 floating-point precisions. And while computation in the complex domain may not be
 44 of interest to all developers, it is absolutely essential for many fields and applications
 45 in part because of complex numbers’ unique ability to encode both the phase and
 46 magnitude of a wave. Thus, the maintainers of general-purpose matrix libraries—such
 47 as those that export the Basic Linear Algebra Subprograms (BLAS) [1]—are typically
 48 compelled by their diverse user bases to support general matrix multiplication (GEMM)
 49 on complex matrices despite the implementation and maintenance costs it may impose.

50 Because of how software developers have historically designed their implementa-
 51 tions, many assume that supporting complex matrix multiplication operations first
 52 requires writing complex domain kernels. To our pleasant surprise, we have discovered
 53 a new way for developers to implement high-performance complex matrix multiplica-
 54 tion *without* those kernels.

55 The predecessor to the current article investigates whether (and to what degree
 56 of effectiveness) real domain matrix multiplication kernels can be repurposed and
 57 leveraged toward the implementation of complex matrix multiplication [21]. In that
 58 article, the authors develop a new class of algorithms that implement these so-called
 59 “induced methods” for matrix multiplication in the complex domain. Instead of rely-
 60 ing on an assembly-coded complex kernel, as a conventional implementation would,
 61 these algorithms express complex matrix multiplication only in terms of real domain
 62 primitives.²

² In [21], the authors use the term “primitive” to refer to a functional abstraction that implements

63 We consider the current article a companion and follow-up to that previous
 64 work [21]. Here, we will consider a new method for emulating complex matrix multi-
 65 plication using only real domain building blocks, and we will once again show that a
 66 clever rearrangement of the real and imaginary elements within the internal “packed”
 67 matrices is key to facilitating high performance. The novelty behind this new method
 68 is that the semantics of complex arithmetic are encoded entirely within a special data
 69 layout, which allows each call to the complex matrix multiplication kernel to be re-
 70 placed with just *one* call to a real matrix multiplication kernel. This substitution
 71 is possible because a real matrix multiplication on the reorganized data mimics the
 72 computation and I/O of a comparable complex matrix multiplication on the unaltered
 73 data. Because of this one-to-one equivalence, we dub this the 1M method.

74 **1.1. Contributions.** This article makes the following contributions:

- 75 • It introduces the 1M method, which replaces each complex matrix multi-
 76 plication with only a single real matrix multiplication.³ We introduce two
 77 algorithmic variants and analyze issues germane to their high-performance
 78 implementations, including workspace, packing formats, cache behavior, mul-
 79 tithreadability, and programming effort. A detailed review shows how 1M
 80 avoids all of the major challenges observed of the 4M method.
- 81 • It promotes code reuse and portability by continuing the previous article’s
 82 focus on solutions which may be cast in terms of real matrix multiplication
 83 kernels. Such solutions have clear implications for developer productivity, as
 84 they allow kernel authors to focus their efforts on fewer and simpler kernels.
- 85 • It builds on the theme of the BLIS framework as a productivity multiplier [22],
 86 further demonstrating how complex matrix multiplication may be imple-
 87 mented with relatively minor modifications to the source code and in such a
 88 way that results in immediate instantiation of complex implementations for
 89 *all* level-3 BLAS-like operations.
- 90 • It demonstrates performance of 1M implementations that is not only superior
 91 to the previous effort based on the 4M method but also competitive with
 92 solutions based on complex matrix kernels.
- 93 • It serves as a reference guide to the 1M implementations for complex matrix
 94 multiplication found within the BLIS framework, which is available to the
 95 community under the open-source 3-clause BSD software license.

96 We believe these contributions are consequential because the 1M method effectively
 97 obviates the previous state-of-the-art established via the 4M method. Furthermore,
 98 we believe the thorough treatment of induced methods encompassed by the present
 99 article and its predecessor will have lasting archival as well as pedagogical value.

100 **1.2. Notation.** In this article, we continue the notation established in [21].
 101 Specifically, we use uppercase Roman letters (e.g. A , B , and C) to refer to ma-
 102 trices, lowercase Roman letters (e.g. x , y , and z) to refer to vectors, and lowercase
 103 Greek letters (e.g. χ , ψ , and ζ) to refer to scalars. Subscripts are used typically to
 104 denote sub-matrices within a larger matrix (e.g. $A = (A_0 \mid A_1 \mid \cdots \mid A_{n-1})$) or
 105 scalars within a larger matrix or vector.

106 We make extensive use of superscripts to denote the real and imaginary compo-
 107 nents of a scalar, vector, or (sub-)matrix. For example, $\alpha^r, \alpha^i \in \mathbb{R}$ denote the real

a single real matrix multiplication. Such primitives are often not general purpose and may come with significant prerequisites to facilitate their use.

³ This proposed 1M method was first published [19].

108 and imaginary parts, respectively, of a scalar $\alpha \in \mathbb{C}$. Similarly, A^r and A^i refer to the
 109 real and imaginary parts of a complex matrix A , where A^r and A^i are real matrices
 110 with dimensions identical to A . Note that while this notation for real, imaginary, and
 111 complex matrices encodes information about content and origin, it does not encode
 112 how the matrices are actually stored. We will explicitly address storage details as
 113 implementation issues are discussed.

114 At times we find it useful to refer to the real and imaginary elements of a com-
 115 plex object indistinguishably as *fundamental elements* (or F.E.). We also abbreviate
 116 floating-point operations as “flops” and memory operations as “memops”. We define
 117 the former to be a MULTIPLY or ADD (or SUBTRACT) operation whose operands are
 118 F.E. and the latter to be a load or store operation on a single F.E.. These definitions
 119 allow for a consistent accounting of complex computation relative to the real domain.

120 We also discuss cache and register blocksizes that are key features of the matrix
 121 multiplication algorithm discussed elsewhere [22, 20, 21]. Unless otherwise noted,
 122 blocksizes n_C , m_C , k_C , m_R , and n_R refer to those appropriate for computation in the
 123 real domain. Complex domain blocksizes will be denoted with a superscript z .

124 This article discusses and references several hypothetical algorithms and func-
 125 tions. Unless otherwise noted, a call to function FUNC that implements $C := C + AB$
 126 appears as $[C] := \text{FUNC}(A, B, C)$. We will also reference functions that access
 127 properties of matrices. For example, $M(A)$ and $N(A)$ would return the m and n
 128 dimensions of a matrix A , while $RS(B)$ and $CS(B)$ would return the row and column
 129 strides of B .

130 2. Background and review.

131 **2.1. Motivation.** In [21], the authors list three primary motivating factors be-
 132 hind their effort to seek out methods for inducing complex matrix multiplication via
 133 real domain kernels:

- 134 • **Productivity.** By inducing complex matrix multiplication from real domain
 135 kernels, the number of kernels that must be supported would be halved.
 136 This allows the DLA library developers to focus on a smaller and simpler
 137 set of real domain kernels. This benefit would manifest most obviously when
 138 instantiating BLAS-like functionality on new hardware [20].
- 139 • **Portability.** Induced methods avoid dependence on complex domain kernels
 140 because they encode the idea of complex matrix product at a higher level.
 141 This would naturally allow us to encode such methods portably within a
 142 framework such as BLIS [22]. Once integrated into the framework, developers
 143 and users would benefit from the immediate availability of complex matrix
 144 multiplication implementations whenever real matrix kernels were present.
- 145 • **Performance.** Implementations of complex matrix multiplication that rely
 146 on real domain kernels would likely inherit the high-performance properties
 147 of those kernels. Any improvement to the real kernels would benefit both real
 148 and complex domains.

149 Thus, it is clear that finding a suitable induced method would carry significant benefit
 150 to DLA library and kernel developers.

151 **2.2. The 3m and 4m methods.** The authors of [21] investigated two general
 152 ways of inducing complex matrix multiplication: the 3M method and the 4M method.
 153 These methods are then contrasted to the conventional approach, whereby a blocked
 154 matrix multiplication algorithm is executed with a complex domain kernel—one that
 155 implements complex arithmetic at the scalar level, in assembly language.

156 The 4M method begins with the classic definition of complex scalar multiplication
 157 and addition in terms of real and imaginary components of $\alpha, \beta, \gamma \in \mathbb{C}$:

$$\begin{aligned} 158 \quad & \gamma^r := \gamma^r + \alpha^r \beta^r - \alpha^i \beta^i \\ 159 \quad & \gamma^i := \gamma^i + \alpha^i \beta^r + \alpha^r \beta^i \end{aligned} \tag{2.1}$$

161 We then observe that we can apply such a definition to complex matrices $A \in \mathbb{C}^{m \times k}$,
 162 $B \in \mathbb{C}^{k \times n}$, and $C \in \mathbb{C}^{m \times n}$, provided that we can reference the real and imaginary
 163 parts as logically separate submatrices:

$$\begin{aligned} 164 \quad & C^r := C^r + A^r B^r - A^i B^i \\ 165 \quad & C^i := C^i + A^i B^r + A^r B^i \end{aligned} \tag{2.2}$$

167 This definition expresses a complex matrix multiplication in terms of four matrix
 168 products (hence the name 4M) and four matrix accumulations (i.e., additions or sub-
 169 tractions).

170 The 3M method relies on a Strassen-like algebraic equivalent of Eq. 2.2:

$$\begin{aligned} 171 \quad & C^r := C^r + A^r B^r - A^i B^i \\ 172 \quad & C^i := C^i + (A^r + A^i)(B^r + B^i) - A^r B^r - A^i B^i \end{aligned}$$

174 This re-expression reduces the number of matrix products to three at the expense of
 175 increasing the number of accumulations from four to seven. However, when the cost
 176 of a matrix product greatly exceeds that of an accumulation, this trade-off can result
 177 in a net reduction in computational runtime.

178 The authors of [21] observe that both methods may be applied to any particular
 179 level of a blocked matrix multiplication algorithm, resulting in several algorithms,
 180 each exhibiting somewhat different properties. Furthermore, they show how either
 181 method’s implementation is facilitated by reordering real and imaginary elements
 182 within the internal storage format used when making packed copies of the current
 183 matrix blocks.⁴ The blocked algorithm used in that article is shown in Figure 1.1 and
 184 revisited in Section 2.4 of the present article.

185 Algorithms that implement the 3M method were found to yield “effective flops
 186 per second” performance that not only exceeded that of 4M, but also approached or
 187 exceeded the theoretical peak rate of the hardware.⁵ Unfortunately, these compelling
 188 results come at a cost: the numerical properties of implementations based on 3M
 189 are slightly less robust than that of algorithms based on the conventional approach
 190 or 4M. And although the author of [6] found that 3M was stable enough for most
 191 practical purposes, many applications will be unwilling to stray from the numerical
 192 expectations implicit in conventional matrix multiplication. Thus, going forward, we
 193 will turn our attention away from 3M and instead focus on the 4M as the standard
 194 reference method against which we will compare.

⁴ Others have exploited the careful design of packing and computational primitives in an effort to improve performance, including in the context of Strassen’s algorithm [7, 9, 10, 11], the computation of the K-Nearest Neighbors [24], tensor contraction [8], and Fast Fourier Transform [17].

⁵ Note that 3M and other Strassen-like algorithms are able to exceed the hardware’s theoretical peak performance when measured in *effective* flops per second: that is, the 3M implementation’s wall clock time—now shorter because of avoided matrix products—divided into the flop count of a *conventional* algorithm.

195 **2.3. Previous findings.** For the reader’s convenience, we will now summarize
 196 the key findings, observations, and other highlights from the previous article regarding
 197 algorithms and implementations based on the 4M method [21].

- 198 • Since all algorithms in the 4M family execute the same number of flops, the
 199 algorithms’ relative performance depends entirely on (1) the number of mem-
 200 ops executed and (2) the level of cache from which F.E. of the packed matrices
 201 \tilde{A}_i and \tilde{B}_p are reused⁶. The number of memops is affected only by a halving
 202 of certain cache blocksize needed in order to leave cache footprints of \tilde{A}_i and
 203 \tilde{B}_p unchanged. The level of cache from which F.E. are reused is determined
 204 by the level of the GEMM algorithm to which the 4M method was applied.
- 205 • The lowest-level application, algorithm 4M_1A, efficiently moves F.E. of A , B ,
 206 and C from main memory to the L1 cache only once per rank- k_C update and
 207 reuses F.E. from the L1 cache. It relies on a relatively simple packing format
 208 and requires negligible, fixed-size workspace, is well-suited for multithreading,
 209 and is minimally disruptive to the BLIS framework. Algorithm 4M_1A can
 210 also be extended relatively easily to all other level-3 operations.
- 211 • The conventional assembly-based approach to complex matrix multiplication
 212 can be viewed as a special case of 4M in which F.E. are reused from registers
 213 rather than cache. In this way, a conventional implementation embodies the
 214 lowest-level application of 4M possible, in which the method is applied to
 215 individual scalars (and then optimally encoded via vector instructions).
- 216 • The way complex numbers are stored has a significant effect on performance.
 217 The standard format adopted by the community (and required by the BLAS),
 218 which uses an interleaved pair-wise storage of real and imaginary values,
 219 naturally favors conventional implementations because they can reuse F.E.
 220 from vector registers. However, this storage is awkward for algorithms based
 221 on 4M (and 3M) because it stymies the use of vector instructions for loading
 222 and storing F.E. of C^r and C^i . The 4M_1A algorithm already suffers from a
 223 *quadrupling*⁷ of the number of memops on C in addition to being forced to
 224 access these F.E. in a non-contiguous manner.
- 225 • While the performance of 4M_1A exceeds an unoptimized reference imple-
 226 mentation, it not only falls short of a comparable conventional solution, it
 227 also falls short of its real domain “benchmark”—that is, the performance of a
 228 similar problem size in the real domain computed by an optimized algorithm
 229 using the same real domain kernel.

230 **2.4. Revisiting the matrix multiplication algorithm.** In this section, we
 231 review a common algorithm for high-performance matrix multiplication on conven-
 232 tional microprocessor architectures. This algorithm was first reported on in [3] and
 233 further refined in [22]. Figure 1.1 illustrates the key features of this algorithm.

234 The current state-of-the-art formulation of the matrix multiplication algorithm
 235 consists of six loops, the last of which resides within a microkernel that is typically
 236 highly optimized for the target hardware. These loops partition the matrix operands
 237 using carefully chosen cache (n_C , k_C , and m_C) and register (m_R and n_R) blocksizes

⁶ Here, the term “reuse” refers to the reuse of F.E. that corresponds to the recurrence of A^r , A^i , B^r , and B^i in Eq. 2.2, not the reuse of whole (complex) elements that naturally occurs in the execution of the GEMM algorithm in Figure 1.1.

⁷ A factor of two comes from the fact that, as shown in Eq. 2.2, 4M touches C^r and C^i twice each, while another factor of two comes from the cache blocksize scaling required on k_C in order to maintain the cache footprints of micropanels of \tilde{A}_i and \tilde{B}_p .

238 that result in submatrices residing favorably at various levels of the cache hierarchy
 239 so as to allow data to be reused many times. In addition, submatrices of A and B are
 240 copied (“packed”) to temporary workspace matrices (\tilde{A}_i and \tilde{B}_p , respectively) in such
 241 a way that allows the microkernel to subsequently access matrix elements contiguously
 242 in memory, which improves cache and TLB performance. The cost of this packing is
 243 amortized over enough computation that its impact on overall performance is negli-
 244 gible for all but the smallest problems. At the lowest level, within the microkernel
 245 loop, an $m_R \times 1$ micro-column and a $1 \times n_R$ micro-row are loaded from the current
 246 micropanels of \tilde{A}_i and \tilde{B}_p , respectively, so that the outer product of these vectors
 247 may be computed to update the corresponding $m_R \times n_R$ submatrix, or micro-tile, of
 248 C . The individual floating-point operations that constitute these tiny rank-1 updates
 249 are oftentimes executed via vector instructions (if the architecture supports them) in
 250 order to maximize utilization of the floating-point unit(s).

251 The algorithm captured by Figure 1.1 forms the basis for all level-3 implementa-
 252 tions found in the BLIS framework (as of this writing). This algorithm is based on a
 253 so-called block-panel matrix multiplication.⁸ The register (m_R , n_R) and cache (m_C ,
 254 k_C , n_C) blocksizes labeled in the algorithmic diagram are typically chosen by the
 255 kernel developer as a function of hardware characteristics, such as the vector register
 256 set, cache sizes, and cache associativity. The authors of [15] present an analytical
 257 model for identifying suitable (if not optimal) values for these blocksizes.

258 **3. 1m method.** The primary motivation for seeking a better induced method
 259 comes from the observation that 4M inherently must update real and imaginary F.E.
 260 of C : (1) in separate steps, and may not use vector instructions to do so (due to the
 261 standard interleaved storage format); and (2) twice as frequently, in the case of 4M_1A,
 262 due to the algorithm’s half-of-optimal cache blocksize k_C . As reviewed in Section 2.3,
 263 this imposes a significant drag on performance. If there existed an induced method
 264 that could update real and imaginary elements in one step, it may conveniently avoid
 265 both issues.

266 **3.1. Derivation.** Consider the classic definition of complex scalar multiplication
 267 and accumulation, shown in Eq. 2.1, refactored and expressed in terms of matrix and
 268 vector notation:

$$269 \quad (3.1) \quad \begin{pmatrix} \gamma^r \\ \gamma^i \end{pmatrix} += \begin{pmatrix} \alpha^r & -\alpha^i \\ \alpha^i & \alpha^r \end{pmatrix} \begin{pmatrix} \beta^r \\ \beta^i \end{pmatrix}$$

271 Here, we have a singleton complex matrix multiplication problem that can naturally
 272 be expressed as a tiny real matrix multiplication where $m = k = 2$ and $n = 1$.
 273 Let us assume we implement this very small matrix multiplication according to the
 274 high-performance algorithm discussed in Section 2.4.

275 From this, we make the following key observation: If we pack α to \tilde{A}_i in such a
 276 way that duplicates α^r and α^i to the second column of the micropanel (while also
 277 swapping the placement of the duplicates and negating the duplicated α^i), and if
 278 we pack β to \tilde{B}_p such that β^i is stored to the second row of the micropanel (which,
 279 granted, only has one column), then a real domain GEMM microkernel executed on
 280 those micropanels will compute the correct result in the complex domain and do so
 281 with a *single* invocation of that microkernel.

⁸ This terminology describes the shape of the typical problem computed by the macro-kernel, i.e. the second loop around the microkernel. An alternative algorithm that casts its largest cache-bound subproblem in terms of panel-block matrix multiplication is discussed in [19].

282 Thus, Eq. 3.1 serves as a packing template that hints at how the data must be
 283 stored. Furthermore, this template can be generalized. We augment α, β, γ with
 284 conventional row and column indices to denote the complex elements of matrices A ,
 285 B , and C , respectively. Also, let us apply the Eq. 3.1 to the special case of $m = 3$,
 286 $n = 4$, and $k = 2$ to better observe the general pattern.

$$287 \quad (3.2) \quad \begin{pmatrix} \gamma_{00}^r & \gamma_{01}^r & \gamma_{02}^r & \gamma_{03}^r \\ \gamma_{00}^i & \gamma_{01}^i & \gamma_{02}^i & \gamma_{03}^i \\ \gamma_{10}^r & \gamma_{11}^r & \gamma_{12}^r & \gamma_{13}^r \\ \gamma_{10}^i & \gamma_{11}^i & \gamma_{12}^i & \gamma_{13}^i \\ \gamma_{20}^r & \gamma_{21}^r & \gamma_{22}^r & \gamma_{23}^r \\ \gamma_{20}^i & \gamma_{21}^i & \gamma_{22}^i & \gamma_{23}^i \end{pmatrix} += \begin{pmatrix} \alpha_{00}^r & -\alpha_{00}^i & \alpha_{01}^r & -\alpha_{01}^i \\ \alpha_{00}^i & \alpha_{00}^r & \alpha_{01}^i & \alpha_{01}^r \\ \alpha_{10}^r & -\alpha_{10}^i & \alpha_{11}^r & -\alpha_{11}^i \\ \alpha_{10}^i & \alpha_{10}^r & \alpha_{11}^i & \alpha_{11}^r \\ \alpha_{20}^r & -\alpha_{20}^i & \alpha_{21}^r & -\alpha_{21}^i \\ \alpha_{20}^i & \alpha_{20}^r & \alpha_{21}^i & \alpha_{21}^r \end{pmatrix} \begin{pmatrix} \beta_{00}^r & \beta_{01}^r & \beta_{02}^r & \beta_{03}^r \\ \beta_{00}^i & \beta_{01}^i & \beta_{02}^i & \beta_{03}^i \\ \beta_{10}^r & \beta_{11}^r & \beta_{12}^r & \beta_{13}^r \\ \beta_{10}^i & \beta_{11}^i & \beta_{12}^i & \beta_{13}^i \end{pmatrix}$$

289 From this, we can make the following observations:

- 290 • The complex matrix multiplication $C := C + AB$ with $m = 3$, $n = 4$, and
 291 $k = 2$ becomes a real matrix multiplication with $m = 6$, $n = 4$, and $k = 4$.
 292 In other words, the m and k dimensions are doubled for the purposes of the
 293 real GEMM primitive.
- 294 • If the primitive is the real GEMM microkernel, and we assume that matrices
 295 A and B above represent column-stored and row-stored micropanels from
 296 \hat{A}_i and \hat{B}_p , respectively, and also that the dimensions are conformal to the
 297 register blocksizes of this microkernel (i.e., $m = m_R$ and $n = n_R$) then the
 298 micropanels of \hat{A}_i are packed from a $\frac{1}{2}m_R \times \frac{1}{2}k_C$ submatrix of A , which, when
 299 expanded in the special packing format, appears as the $m_R \times k_C$ micropanel
 300 that the real GEMM microkernel expects.
- 301 • Similarly, the micropanels of \hat{B}_p are packed from a $\frac{1}{2}k_C \times n_R$ submatrix of
 302 B , which, when reordered into a second special packing format, appears as
 303 the $k_C \times n_R$ micropanel that the real GEMM microkernel expects.

304 It is easy to see by inspection that the real matrix multiplication implied by
 305 Eq. 3.2 induces the desired complex matrix multiplication. We will refer to the packing
 306 format used on matrix A above as the 1E format, since the F.E. are “expanded”
 307 (i.e., duplicated to the next column, with the duplicates swapped and the imaginary
 308 duplicate negated). Similarly, we will refer to the packing format used on matrix B
 309 above as the 1R format, since the F.E. are merely reordered (i.e., imaginary elements
 310 moved to the next row). Thus, the 1M method is fundamentally about reordering the
 311 matrix data so that a subsequent real matrix multiplication on that reordered data is
 312 equivalent to a complex matrix multiplication on the original data.⁹

313 **3.2. Two variants.** Notice that implicit in the 1M method suggested by Eq. 3.2
 314 is the fact that matrix C is stored by columns. This assumption is important; when A
 315 and B are packed according to the 1E and 1R formats, respectively, C must be stored
 316 by columns in order to allow the real domain primitive (or microkernel) to correctly
 317 update the individual real and imaginary F.E. of C with the corresponding F.E. from
 318 the matrix product AB .

319 Suppose that we instead refactored and expressed Eq. 2.1 as follows:

$$320 \quad (3.3) \quad (\gamma^r \ \gamma^i) += (\alpha^r \ \alpha^i) \begin{pmatrix} \beta^r & \beta^i \\ -\beta^i & \beta^r \end{pmatrix}$$

321

⁹ The authors of [17] also investigated the use of transforming the data layout during packing to facilitate complex matrix multiplication. And while they employ techniques similar to those of the 1M method, their approach differs in that it does not recycle the existing real domain microkernel.

TABLE 3.1
1M complex domain blocksizes as a function of real domain blocksizes

Variant	Blocksizes, in terms of real domain values, required for . . .						
	k_C^z	m_C^z	n_C^z	m_R^z	m_P^z	n_R^z	n_P^z
1M_C	$\frac{1}{2}k_C$	$\frac{1}{2}m_C$	n_C	$\frac{1}{2}m_R$	m_P	n_R	n_P
1M_R	$\frac{1}{2}k_C$	m_C	$\frac{1}{2}n_C$	m_R	m_P	$\frac{1}{2}n_R$	n_P

Note: Blocksizes m_P and n_P represent the so-called “packing dimensions” for the micro-panels of \tilde{A}_i and \tilde{B}_p , respectively. These values are analogous to the leading dimensions of matrices stored by columns or rows. In BLIS microkernels, typically $m_R = m_P$ and $n_R = n_P$, but sometimes the kernel author may find it useful for $m_R < m_P$ or $n_R < n_P$.

322 This gives us a different template, one that implies different packing formats for A
 323 and B . Applying Eq. 3.3 to the special case of $m = 4$, $n = 3$, and $k = 2$ yields:

(3.4)

$$\begin{pmatrix} \gamma_{00}^r & \gamma_{00}^i & \gamma_{01}^r & \gamma_{01}^i & \gamma_{02}^r & \gamma_{02}^i \\ \gamma_{10}^r & \gamma_{10}^i & \gamma_{11}^r & \gamma_{11}^i & \gamma_{12}^r & \gamma_{12}^i \\ \gamma_{20}^r & \gamma_{20}^i & \gamma_{21}^r & \gamma_{21}^i & \gamma_{22}^r & \gamma_{22}^i \\ \gamma_{30}^r & \gamma_{30}^i & \gamma_{31}^r & \gamma_{31}^i & \gamma_{32}^r & \gamma_{32}^i \end{pmatrix} + = \begin{pmatrix} \alpha_{00}^r & \alpha_{00}^i & \alpha_{01}^r & \alpha_{01}^i \\ \alpha_{10}^r & \alpha_{10}^i & \alpha_{11}^r & \alpha_{11}^i \\ \alpha_{20}^r & \alpha_{20}^i & \alpha_{21}^r & \alpha_{21}^i \\ \alpha_{30}^r & \alpha_{30}^i & \alpha_{31}^r & \alpha_{31}^i \end{pmatrix} \begin{pmatrix} \beta_{00}^r & \beta_{00}^i & \beta_{01}^r & \beta_{01}^i & \beta_{02}^r & \beta_{02}^i \\ -\beta_{00}^i & \beta_{00}^r & -\beta_{01}^i & \beta_{01}^r & -\beta_{02}^i & \beta_{02}^r \\ \beta_{10}^r & \beta_{10}^i & \beta_{11}^r & \beta_{11}^i & \beta_{12}^r & \beta_{12}^i \\ -\beta_{10}^i & \beta_{10}^r & -\beta_{11}^i & \beta_{11}^r & -\beta_{12}^i & \beta_{12}^r \end{pmatrix}$$

326 In this variant, we see that matrix B , not A , is stored according to the 1E format
 327 (where columns become rows), while matrix A is stored according to 1R (where rows
 328 become columns). Also, we can see that matrix C must be stored by rows in order to
 329 allow the real GEMM microkernel to correctly update its F.E. with the corresponding
 330 values from the matrix product AB .

331 Henceforth, we will refer to the 1M variant exemplified in Eq. 3.2 as 1M_C since
 332 it is predicated on column storage of the output matrix C , and we will refer to the
 333 variant depicted in Eq. 3.4 as 1M_R since it assumes C is stored by rows.

334 **3.3. Determining complex blocksizes.** As we alluded in Section 3.1, the
 335 appropriate blocksizes to use with 1M are a function of the real domain blocksizes.
 336 This makes sense because the idea is to fool the real GEMM microkernel, and the
 337 various loops for register and cache blocking around the microkernel, into thinking
 338 that it is computing a real domain matrix multiplication. Which blocksizes must be
 339 modified (halved) and which are used unchanged depends on the variant of 1M being
 340 executed—or, more specifically, which matrix is packed according to the 1E format.

341 Table 3.1 summarizes the complex domain blocksizes prescribed for 1M_C and
 342 1M_R as a function of the real domain values.

343 Those familiar with the matrix multiplication algorithm implemented by the BLIS
 344 framework, as depicted in Figure 1.1, may be unfamiliar with m_P and n_P , the so-
 345 called packing dimensions. These values are the leading dimensions of the micropanel.
 346 On most architectures, $m_P = m_R$ and $n_P = n_R$, but in some situations it may be
 347 convenient (or necessary) to use $m_R < m_P$ or $n_R < n_P$. In any case, these packing
 348 dimensions are never scaled, even when their corresponding register blocksizes *are*
 349 scaled to accommodate the 1E format, because the halving that would otherwise be
 350 called for is cancelled out by the doubling of F.E. that manifests in the 1E format.

Algorithm: $[C] := \text{RMMBP}(A, B, C)$ for ($j = 0 : n - 1 : n_C$) Identify B_j, C_j from B, C for ($p = 0 : k - 1 : k_C$) Identify A_p, B_{jp} from A, B_j PACK $B_{jp} \rightarrow \tilde{B}_p$ for ($i = 0 : m - 1 : m_C$) Identify A_{pi}, C_{ji} from A_p, C_j PACK $A_{pi} \rightarrow \tilde{A}_i$ for ($h = 0 : n_C - 1 : n_R$) Identify \tilde{B}_{ph}, C_{jih} from \tilde{B}_p, C_{ji} for ($l = 0 : m_C - 1 : m_R$) Identify \tilde{A}_{il}, C_{jihl} from \tilde{A}_i, C_{jih} $C_{jihl} := \text{RKERN}(\tilde{A}_{il}, \tilde{B}_{ph}, C_{jihl})$
--

FIG. 3.1. Abbreviated pseudo-code for implementing the general matrix multiplication algorithm depicted in Figure 1.1. Here, RKERN calls a real domain GEMM microkernel.

3.4. Algorithms.

352 **3.4.1. General algorithm.** Before investigating 1M method algorithms, we will
 353 first provide algorithms for computing real matrix multiplication to serve as a reference
 354 for the reader. Specifically, in Figure 3.1 we provide pseudo-code for RMMBP, which
 355 depicts a real domain instance of the block-panel algorithm shown in Figure 1.1.

356 **3.4.2. 1m-specific algorithm.** Applying 1M_C and 1M_R to the block-panel
 357 algorithm depicted in Figure 1.1 yields two nearly identical algorithms, 1M_C_BP and
 358 1M_R_BP, respectively. Their differences can be encoded within a few conditional
 359 statements within key parts of the high and low levels of code. Figure 3.2 shows a
 360 hybrid algorithm that encompasses both, supporting row- and column-stored C .

361 In Figure 3.2 (right), we illustrate the 1M *virtual* microkernel. This function,
 362 VK1M, consists largely of a call to the real domain microkernel RKERN with some
 363 additional logic needed to properly induce complex matrix multiplication in all cases.
 364 Some of the details of the virtual microkernel will be addressed later.

365 **3.5. Performance properties.** Table 3.2 tallies the total number of F.E. mem-
 366 ops required by 1M_C_BP and 1M_R_BP. For comparison, we also include the corre-
 367 sponding memop counts for a selection of 4M algorithms as well as a conventional
 368 assembly-based solution, as first published in Table III in [21].

369 Notice that 1M_C_BP and 1M_R_BP incur additional memops relative to a conven-
 370 tional assembly-based solution because, unlike the latter, 1M implementations cannot
 371 reuse¹⁰ all real and imaginary F.E. from vector registers.

372 We can hypothesize that the observed performance signatures of 1M_C_BP and
 373 1M_R_BP may be slightly different because each places the additional memop overhead
 374 that is unique to 1M on different parts of the computation. This stems from the fact
 375 that there exists an asymmetry in the assignment of packing formats to matrices in
 376 each 1M variant. Specifically, 50% more memops—relative to a conventional assembly
 377 solution—are required during the initial packing and the movement between caches

¹⁰ Here, the term “reuse” refers to the same reuse described in Footnote 6.

Algorithm: $[C] := 1M_?_BP(A, B, C)$	$[C] := vk1M(A, B, C)$
<pre> Set bool COLSTORE if RS(C) = 1 for (j = 0 : n - 1 : n_C) Identify B_j, C_j from B, C for (p = 0 : k - 1 : k_C) Identify A_p, B_jp from A, B_j if COLSTORE PACK1R B_jp → B̃_p else PACK1E B_jp → B̃_p for (i = 0 : m - 1 : m_C) Identify A_pi, C_ji from A_p, C_j if COLSTORE PACK1E A_pi → Ã_i else PACK1R A_pi → Ã_i for (h = 0 : n_C - 1 : n_R) Identify B_ph, C_jih from B̃_p, C_ji for (l = 0 : m_C - 1 : m_R) Identify A_il, C_jihl from Ã_i, C_jih C_jihl := vk1M(Ã_il, B̃_ph, C_jihl) </pre>	<pre> Acquire workspace W Determine if using W; set USEW if (USEW) Alias C_use ← W, C_in ← 0 else Alias C_use ← C, C_in ← C Set bool COLSTORE if RS(C_use) = 1 if (COLSTORE) CS(C_use) × = 2 else RS(C_use) × = 2 N(A) × = 2; M(B) × = 2 C_use := RKERN(A, B, C_in) if (USEW) C := W </pre>

FIG. 3.2. Left: Pseudo-code for Algorithms 1M_C_BP and 1M_R_BP, which result from applying 1M_C and 1M_R algorithmic variants to the block-panel algorithm depicted in Figure 1.1. Here, PACK1E and PACK1R pack matrices into the 1E and 1R formats, respectively. Right: Pseudo-code for a virtual microkernel used by all 1M algorithms.

378 for the matrix packed according to 1E since that format writes four F.E. for every
 379 two that it reads from the source operand. (Packing to 1R incurs the same number
 380 of memops as an assembly-based solution.) Also, if 1M_C_BP and 1M_R_BP use real
 381 microkernels with different micro-tile shapes (i.e., different values of m_R and n_R),
 382 those microkernels' differing performance properties will likely cause the performance
 383 signatures of 1M_C_BP and 1M_R_BP to deviate further.

384 Table 3.3 summarizes Table 3.2 and adds: (1) the level of the memory hierarchy
 385 from which each matrix is reused; and (2) a measure of memory movement efficiency.

386 **3.6. Algorithm details.** This section lays out important details that must be
 387 handled when implementing the 1M method.

388 **3.6.1. Microkernel I/O preference.** Within the BLIS framework, microker-
 389 nels are registered with a property that describes their input/output *preference*. The
 390 I/O preference describes whether the microkernel is set up to ideally use vector in-
 391 structions to load and store elements of the micro-tile by rows or by columns. This
 392 property typically originates from the semantic orientation of vector registers used to
 393 accumulate the $m_R \times n_R$ micropanel product. Whenever possible, the BLIS frame-
 394 work will perform logical transpositions¹¹ so that the apparent storage of C matches
 395 the preference property of the microkernel being used. This guarantees that the mi-
 396 crokernel will be able to load and store F.E. of C using vector instructions.

397 This preference property is merely an interesting performance detail for conven-
 398 tional implementations (real and complex). However, in the case of 1M, it becomes
 399 crucial for constructing a correctly-functioning implementation. More specifically, the

¹¹ This amounts to swapping the row and column strides and swapping the m and n dimensions.

TABLE 3.2
F.E. memops incurred by various algorithms, broken down by stage of computation

Algorithm	F.E. memops required to ... ^a				
	update micro-tiles ^b C^r, C^i	pack \tilde{A}_i	move \tilde{A}_i from L2 to L1 cache	pack \tilde{B}_p	move \tilde{B}_p from L3 to L1 cache
4M_H	$8mn \frac{k}{k_C}$	$8mk \frac{n}{n_C}$	$4mk \frac{n}{n_R}$	$8kn$	$4kn \frac{m}{m_C}$
4M_1B	$8mn \frac{k}{k_C}$	$8mk \frac{2n}{n_C}$	$4mk \frac{n}{n_R}$	$8kn$	$4kn \frac{2m}{m_C}$
4M_1A	$8mn \frac{2k}{k_C}$	$8mk \frac{n}{n_C}$	$4mk \frac{n}{n_R}$	$8kn$	$4kn \frac{m}{m_C}$
assembly	$4mn \frac{k}{k_C}$	$4mk \frac{n}{n_C}$	$2mk \frac{n}{n_R}$	$4kn$	$2kn \frac{m}{m_C}$
1M_C_BP	$4mn \frac{2k}{k_C}$	$6mk \frac{n}{n_C}$	$4mk \frac{n}{n_R}$	$4kn$	$2kn \frac{2m}{m_C}$
1M_R_BP		$4mk \frac{n}{n_C}$	$2mk \frac{2n}{n_R}$	$6kn$	$4kn \frac{m}{m_C}$

^a We express the number of iterations executed in the 5th, 4th, 3rd, and 2nd loops as $\frac{n}{n_C}, \frac{k}{k_C}, \frac{m}{m_C}$, and $\frac{n}{n_R}$. The precise number of iterations along a dimension x using a cache blocksize x_C would actually be $\lceil \frac{x}{x_C} \rceil$. Similarly, when blocksize scaling of $\frac{1}{2}$ is required, the precise value $\lceil \frac{x}{x_C/2} \rceil$ is expressed as $\frac{2x}{x_C}$. These simplifications allow easier comparison between algorithms while still providing meaningful approximations.

^b As described in Section 3.6.2, $m_R \times n_R$ workspace sometimes becomes mandatory, such as when $\beta^i \neq 0$. When workspace is employed in a 4M-based algorithm, the number of F.E. memops incurred updating the micro-tile typically doubles from the values shown here.

TABLE 3.3
Performance properties of various algorithms

Algorithm	Total F.E. memops required (Sum of columns of Table 3.2)	Level from which F.E. of matrix X are reused, and l_{L1} : # of times each cache line is moved into the L1 cache (per rank- k_C update).					
		C	l_{L1}^C	A	l_{L1}^A	B	l_{L1}^B
4M_H	$8mn \left(\frac{k}{k_C} \right) + 4mk \left(\frac{2n}{n_C} + \frac{n}{n_R} \right) + 2kn \left(4 + \frac{2m}{m_C} \right)$	MEM	4	MEM	4	MEM	4
4M_1B	$8mn \left(\frac{k}{k_C} \right) + 4mk \left(\frac{4n}{n_C} + \frac{n}{n_R} \right) + 2kn \left(4 + \frac{4m}{m_C} \right)$	L2	2 ^a	L2	1	L1	1
4M_1A	$8mn \left(\frac{2k}{k_C} \right) + 4mk \left(\frac{2n}{n_C} + \frac{n}{n_R} \right) + 2kn \left(4 + \frac{2m}{m_C} \right)$	L1	1 ^a	L1	1	L1	1
assembly	$4mn \left(\frac{k}{k_C} \right) + 2mk \left(\frac{2n}{n_C} + \frac{n}{n_R} \right) + 2kn \left(2 + \frac{m}{m_C} \right)$	REG	1	REG	1	REG	1
1M_C_BP	$4mn \left(\frac{2k}{k_C} \right) + 2mk \left(\frac{3n}{n_C} + \frac{2n}{n_R} \right) + 2kn \left(2 + \frac{2m}{m_C} \right)$	REG	1	L2 ^b	1	REG	1
1M_R_BP	$4mn \left(\frac{2k}{k_C} \right) + 2mk \left(\frac{2n}{n_C} + \frac{2n}{n_R} \right) + 2kn \left(3 + \frac{2m}{m_C} \right)$	REG	1	REG	1	L1 ^b	1

^a This assumes that the micro-tile is not evicted from the L1 cache during the next call to RKERN.

^b In the case of 1M algorithms, we consider F.E. of A and B to be “reused” from the level of cache in which the 1E-formatted matrix resides.

400 microkernel’s I/O preference determines whether the 1M_C or 1M_R algorithm is pre-
 401 scribed. Generally speaking, a 1M_C algorithmic variant must employ a microkernel
 402 that prefers to access C by columns, while a 1M_R algorithmic variant must use a
 403 microkernel that prefers to access C by rows.

404 **3.6.2. Workspace.** In some cases, a small amount of $m_R \times n_R$ workspace is
 405 needed. These cases fall into one of four scenarios: (1) C is row-stored and the real
 406 microkernel RKERN has a column preference; (2) C is column-stored and RKERN has
 407 a row preference; (3) C is general-stored (i.e., neither $\text{RS}(C)$ nor $\text{CS}(C)$ is unit); and
 408 (4) $\beta^i \neq 0$. If any of these conditions hold, then the 1M virtual microkernel will need
 409 to use workspace. This corresponds to the setting of USEW in VK1M (in Figure 3.2),
 410 which causes RKERN to compute the micropanel product normally but store it to the
 411 workspace W . Subsequently, the result in W is then accumulated back to C .

412 Cases (1) and (2), while supported, actually never occur in practice because BLIS
 413 will perform a logical transposition of the operation, when necessary, so that the
 414 storage of C will always appear to match the I/O preference of the microkernel.

415 Case (3) is needed because the real microkernel is programmed to support the up-
 416 dating of *real* matrices stored with general stride, which cannot emulate the updating
 417 of *complex* matrices stored with general stride. The reason is even when stored with
 418 general stride, complex matrices use the standard storage format, which interleaves
 419 real and imaginary F.E. in contiguous pairs. There is no way to coax this pattern
 420 of data access from a real domain microkernel, given its existing API. Thus, general
 421 stride support must be implemented outside RKERN, within VK1M.

422 Case (4) is needed because real domain microkernels are not capable of scaling C
 423 by complex scalars β when $\beta^i \neq 0$.

424 **3.6.3. Handling alpha and beta scalars.** As in the previous article, we have
 425 simplified the general matrix multiplication to $C := C + AB$. In practice, the operation
 426 is implemented as $C := \beta C + \alpha AB$, where $\alpha, \beta \in \mathbb{C}$. Let us use Algorithm 1M_C_BP
 427 in Figure 3.2 to consider how to support arbitrary values of α and β .

428 If no workspace is needed (because none of the four situations described in Sec-
 429 tion 3.6.2 apply), we can simply pass β^r into the RKERN call. However, if workspace *is*
 430 needed, then we must pass in a local $\beta_{\text{use}} = 0$ to RKERN, compute to local workspace
 431 W , and then apply β at the end of VK1M when W is accumulated to C .

432 When α is real, the scaling may be performed directly by RKERN. This situation
 433 is ideal since it usually incurs no additional costs.¹² Scaling by α with non-zero
 434 imaginary components can still be performed by the packing function when either \tilde{A}_i
 435 or \tilde{B}_p are packed. Though somewhat less than ideal, the overhead incurred by this
 436 treatment of α is minimal in practice since packing is a memory-bound operation.

437 **3.6.4. Multithreading.** As with Algorithm 4M_1A in the previous article, Al-
 438 gorithms 1M_C_BP and 1M_R_BP parallelize in a straightforward manner for multicore
 439 and many-core environments. Because these algorithms encode the 1M method en-
 440 tirely within the packing functions and the virtual microkernel, all other levels of code
 441 are completely oblivious to, and therefore unaffected by, the specifics of the new al-
 442 gorithms. Therefore, we expect that 1M_C_BP and 1M_R_BP will yield multithreaded
 443 performance that is on-par with that of RMMBP.

444 **3.6.5. Bypassing the virtual microkernel.** Because the 1M virtual microker-
 445 nel serves as a function wrapper to the real domain microkernel, it incurs additional

¹² This is because many microkernels multiply their intermediate AB product by α unconditionally

TABLE 4.1

Register and cache block sizes used by various BLIS implementations of matrix multiplication, as configured for an Intel Xeon E5-2690 v3 “Haswell” processor

Precision/Domain	Implementation	m_R^z	n_R^z	m_C^z	k_C^z	n_C^z
single complex	BLIS 1M_C	16/2	6	144/2	256/2	4080
	BLIS 1M_R	6	16/2	144	256/2	4080/2
	BLIS assembly (c)	8	3	56	256	4080
	BLIS assembly (r)	3	8	75	256	4080
double complex	BLIS 1M_C	8/2	6	72/2	256/2	4080
	BLIS 1M_R	6	8/2	72	256/2	4080/2
	BLIS assembly (c)	4	3	44	256	4080
	BLIS assembly (r)	3	4	192	256	4080

Note: For 1M implementations, division by 2 is made explicit to allow the reader to quickly see both the complex blocksize values as well as the values that would be used by the underlying real domain microkernels when performing real matrix multiplication. The I/O preference of the assembly-based implementations is indicated by a “(c)” or “(r)” (for column- or row-preferring).

446 overhead. Thankfully, there exists a simple workaround, one that is viable as long as
 447 $\beta^i = 0$ and C is either row- or column-stored (but not general-stored). If these con-
 448 ditions are met, the real domain *macrokernel* can be called with modified parameters
 449 to induce the equivalent complex domain subproblem. This optimization allows the
 450 virtual microkernel (and its associated overhead) to be avoided entirely.

451 Because this optimization relies only on $\beta \in \mathbb{R}$ and row- or column storage of C ,
 452 it may be applied automatically at runtime to the vast majority of use cases.

453 **3.7. Other complex storage formats.** The 1M method was developed specifi-
 454 cally to facilitate performance on complex matrices stored using the standard storage
 455 format required by the BLAS. This interleaved, pair-wise storage convention is ubiqui-
 456 tous within the community and therefore implicitly assumed. However, some current
 457 and future applications may be willing to tolerate the API changes that would allow
 458 storing a complex matrix X as two separate real matrices, X^r and X^i . For those ap-
 459 plications, the best an induced method may hope to do is implement each specialized
 460 complex matrix multiplication in terms of *two* real domain matrix multiplications—
 461 since there are two real matrices that must be updated. Indeed, there exists a variant
 462 of the 1M method, which we call the 2M method, that targets updating a matrix C
 463 that separates (entirely or by blocks) its real and imaginary F.E. [19].

464 And while treatment of non-standard storage formats is beyond the scope of this
 465 article, motivated readers may use the presentation of 2M in [19] to extend the insights
 466 presented here to develop 2M algorithms.

467 **4. Performance.** In this section we present performance results for implemen-
 468 tations of 1M algorithms on a recent Intel architecture. For comparison, we include
 469 results for a key 4M algorithm as well as those of conventional assembly-based ap-
 470 proaches in the real and complex domains.

471 **4.1. Platform and implementation details.** Results presented in this section
 472 were gathered on a single Cray XC40 compute node consisting of two 12-core Intel
 473 Xeon E5-2690 v3 processors featuring the “Haswell” microarchitecture. Each core,

474 running at a clock rate of 3.2 GHz¹³, provides a single-core peak performance of 51.2
 475 gigaflops (GFLOPS) in double precision and 102.4 GFLOPS in single precision.¹⁴
 476 Each socket has a 30MB L3 cache that is shared among cores, and each core has a
 477 private 256KB L2 cache and 32KB L1 (data) cache. Performance experiments were
 478 gathered under the Cray Linux Environment 6 operating system running the Linux
 479 4.4.103 (x86_64) kernel. Source code was compiled by the GNU C compiler (`gcc`)
 480 version 7.3.0.¹⁵ The version of BLIS used in these tests was not officially released at
 481 the time of this writing, and was adapted from version 0.6.0-11.¹⁶

482 Algorithms `1M_C_BP` and `1M_R_BP` were implemented in the BLIS framework as
 483 described in Section 3.4. We also refer to results based on existing conventional
 484 assembly-based microkernels written by hand (via GNU extended inline assembly
 485 syntax) for the Haswell microarchitecture.

486 All experiments were performed on randomized, column-stored matrices with
 487 GEMM scalars held constant: $\alpha = \beta = 1$. In all performance graphs, each data
 488 point represents the best of three trials.

489 Blocksizes for each of the BLIS implementations tested are provided in Table 4.1.

490 In all graphs presented in this section, the x -axes denote the problem size, the
 491 y -axes show observed floating-point performance in units of GFLOPS per core, and
 492 the theoretical peak performance coincides with the top of each graph.

493 **4.2. Sequential results.** Figure 4.1 reports performance results for various im-
 494 plementations of double- and single-precision complex matrix multiplication on a
 495 single core of the Haswell processor. For these results, all matrix dimensions were
 496 equal (e.g. $m = n = k$). Results for `1M_C_BP` (which uses a column-preferring mi-
 497 crokernel) appears on the left of Figure 4.1 while those of `1M_R_BP` (which uses a
 498 row-preferring microkernel) appears on the right.

499 Each graph in Figure 4.1 also contains three reference implementations: BLIS’s
 500 complex GEMM based on conventional assembly-coded kernels (e.g. “`cgemm` assem-
 501 bly”); BLIS’s real GEMM (e.g. “`sgemm` assembly”); and the `4M_1A` implementation
 502 found in BLIS.¹⁷ We configured all three of these reference codes to use column-
 503 preferential microkernels on the left and row-preferential microkernels on the right,
 504 as indicated by a “(c)” or “(r)” in the legends, in order to provide consistency with
 505 the 1M results.

506 As predicted in Section 3.5, we find that the performance signatures of the
 507 `1M_C_BP` and `1M_R_BP` algorithms differ slightly. This was expected given that the
 508 1E and 1R packing formats place different memory access burdens on different packed
 509 matrices, \tilde{A}_i and \tilde{B}_p , which reside in different levels of cache. It was not previously
 510 clear, however, which would be superior over the other. It seems that, at least in
 511 the sequential case, the difference is somewhat more noticeable in double-precision,
 512 though even there it is quite subtle. This difference is almost certainly due to the

¹³ This system uses Intel’s Turbo Boost 2.0 dynamic frequency throttling technology. According to [14], the maximum the clock frequency when executing AVX instructions is 3.2 GHz when utilizing one or two cores, and 3.0 GHz when utilizing three or more cores.

¹⁴ Accounting for the reduced AVX clock frequency, the peak performance when utilizing 24 cores is 48 GFLOPS/core in double precision and 96 GFLOPS/core in single precision.

¹⁵ The following optimization flags were used during compilation of BLIS and its test drivers: `-O3 -mavx2 -mfma -mfpmath=sse -march=haswell`.

¹⁶ Despite not yet having an official version number, this version of BLIS may be uniquely identified, with high probability, by the first 10 digits of its git “commit” (SHA1 hash) number: `ceee2f973e`.

¹⁷ Within any given graph of Figures 4.1 and 4.2, the 1M and `4M_1A` implementations use the same real-domain microkernel as that of the real GEMM (e.g. “`sgemm` assembly” or “`dgemm` assembly”).

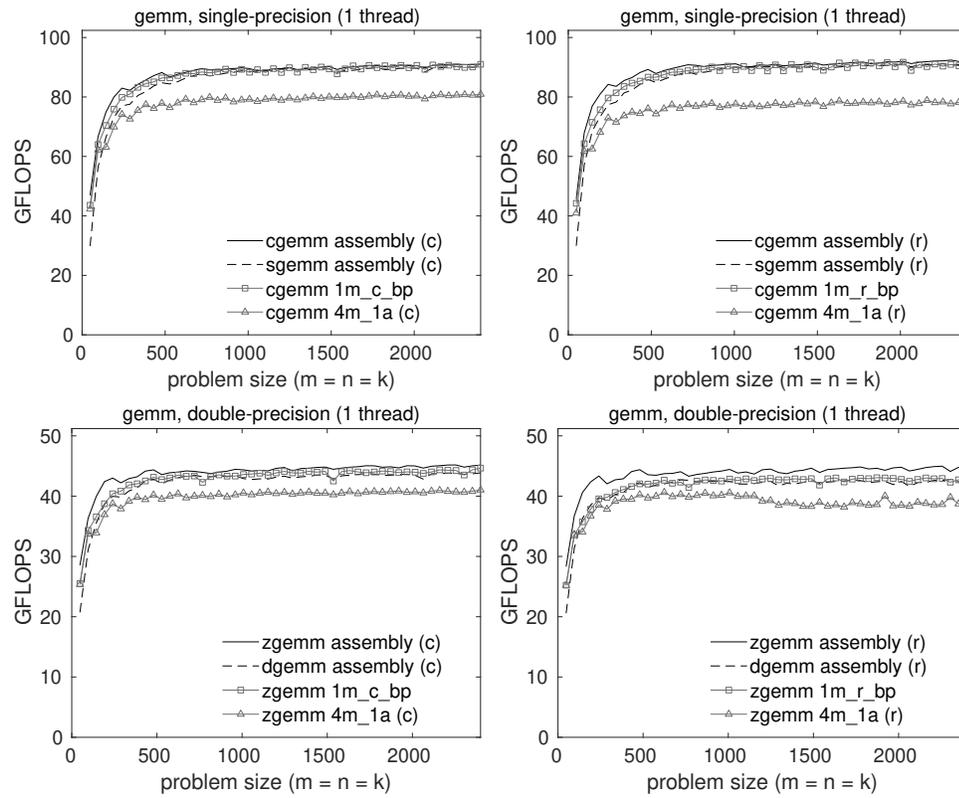


FIG. 4.1. Single-threaded performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core of an Intel Xeon E5-2690 v3 “Haswell” processor. The left and right graphs differ in which 1M implementation they report, with the left graphs reporting 1M_C_BP (which employs a column-preferring microkernel) and the right graphs reporting 1M_R_BP (which employs a row-preferring microkernel). The graphs also contain three reference curves for comparison: an assembly-coded complex GEMM, an assembly-coded real GEMM, and the 4M_1A implementation found in BLIS (with the latter two using the same microkernel as the 1M implementation shown in the same graph). For consistency with the 1M curves, these reference implementations differ from left to right graphs in the I/O preference of their underlying microkernel, indicated by a “(c)” or “(r)” (for column- or row-preferring) in the legends. The theoretical peak performance coincides with the top of each graph.

513 individual performance characteristics of the underlying row- and column-preferential
 514 microkernels. We find evidence of this in the 4M_1A results, which was also affected
 515 by the change in microkernel I/O preference.

516 In all cases, the 1M implementations outperform 4M_1A, with the margin some-
 517 what larger in single-precision.

518 The 1M implementations match or exceed the performance of their real domain
 519 GEMM benchmarks (the dotted lines in each graph) and are quite competitive with
 520 assembly-coded complex GEMM (the solid lines) regardless of the algorithm employed.

521 **4.3. Multithreaded results.** Figure 4.2 shows single- and double-precision per-
 522 formance using 24 threads, with one thread bound to each physical core of the proces-
 523 sor. Performance is presented in units of gigaflops per core to facilitate visual assess-
 524 ment of scalability. For all BLIS implementations, we employed 4-way parallelism
 525 within the 5th loop, 3-way parallelism within the 3rd loop, and 2-way parallelism in

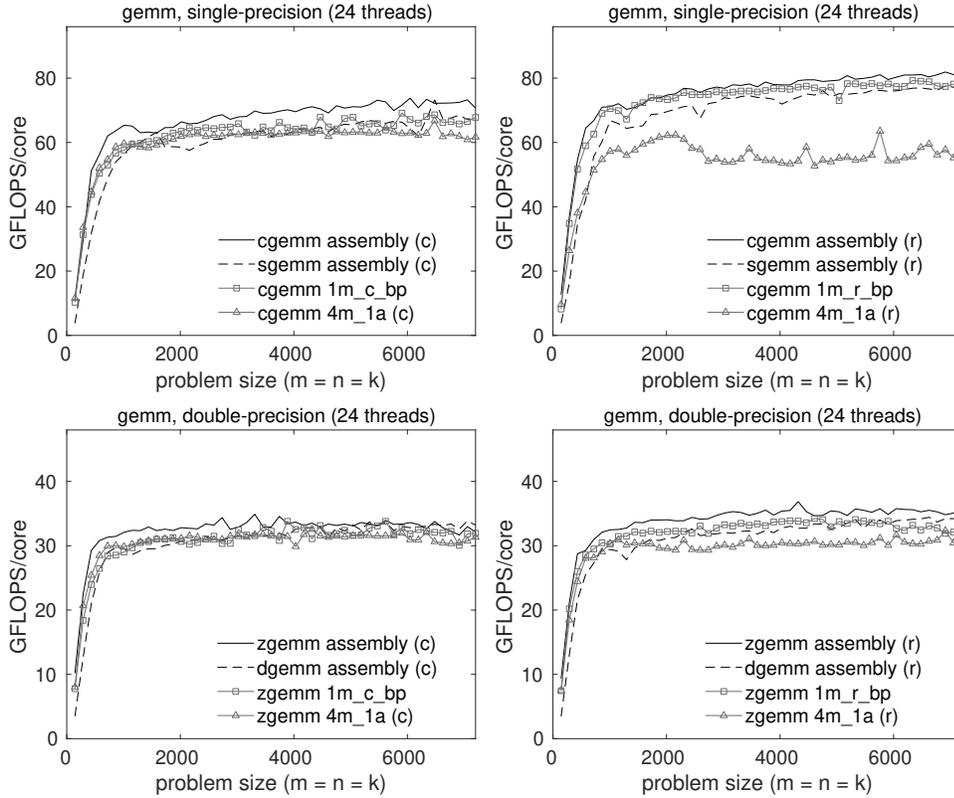


FIG. 4.2. Multithreaded performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on two Intel Xeon E5-2690 v3 “Haswell” processors, each with 12 cores. All data points reflect the use of 24 threads. The left and right graphs differ in which 1M implementation they report, with the left graphs reporting 1M_C_BP (which employs a column-prefering microkernel) and the right graphs reporting 1M_R_BP (which employs a row-prefering microkernel). The graphs also contain three reference curves for comparison: an assembly-coded complex GEMM, an assembly-coded real GEMM, and the 4M_1A implementation found in BLIS (with the latter two using the same microkernel as the 1M implementation shown in the same graph). For consistency with the 1M curves, these reference implementations differ from left to right graphs in the I/O preference of their underlying microkernel, indicated by a “(c)” or “(r)” (for column- or row-prefering) in the legends. The theoretical peak performance coincides with the top of each graph.

526 the 2nd loop for a total of 24 threads. This parallelization scheme was chosen in a
 527 manner consistent with that of the previous article using a strategy set forth in [18].

528 Compared to the single-threaded case, we find a more noticeable difference in
 529 multithreaded performance between the 1M algorithms. Specifically, the 1M_R_BP im-
 530 plementation (based on a row-prefering microkernel) outperforms that of 1M_C_BP
 531 (based on a column-prefering microkernel), with the difference more pronounced
 532 in single-precision. We suspect this is rooted not in the algorithms *per se* but in
 533 the differing microkernel implementations used by each 1M algorithm. The 1M_R_BP
 534 algorithm uses a real microkernel that is 6×16 and 6×8 in the single- and double-
 535 precision cases, respectively, while 1M_C_BP uses 16×6 and 8×6 microkernels for
 536 single- and double-precision implementations, respectively. The observed difference
 537 in performance between the 1M algorithms is likely attributable to the fact that the

538 microkernels’ different values for m_R and n_R place different latency and bandwidth
 539 requirements when reading F.E. from the caches (primarily L1 and L2). More specif-
 540 ically, larger values of m_R place a heavier burden on loading elements from the L2
 541 cache, which is usually disadvantageous since that cache may exhibit higher latency
 542 and/or lower bandwidth. By contrast, a microkernel with larger n_R loads more ele-
 543 ments (per $m_R \times n_R$ rank-1 update) from the L1 cache, which resides closer to the
 544 processor and offers lower latency and/or higher bandwidth than the L2 cache.

545 The multithreaded 1M implementation approximately matches or exceeds its real
 546 domain counterpart in all cases.

547 The 1M algorithm based on a row-preferential microkernel, 1M_R_BP, outper-
 548 forms 4M_1A, especially in single-precision where the margin is quite wide. The 1M
 549 algorithm based on column-preferential microkernels, 1M_C_BP, performs more poorly,
 550 barely edging out 4M_1A in single precision and tracking closely with 4M_1A in double
 551 precision. We suspect that 4M_1A is more resilient to the lower-performing column-
 552 preferential microkernel by virtue of the fact that the algorithm’s virtual microkernel
 553 leans heavily on the L1 cache, which on this architecture is capable of being read
 554 from and written to at relatively high bandwidth (64 bytes/cycle and 32 bytes/cycle,
 555 respectively) [13].

556 **4.4. Comparing to other implementations.** While our primary goal is not
 557 to compare the performance of the newly developed 1M implementations with that
 558 of other established BLAS solutions, some basic comparison is merited and thus we
 559 have included Figure 4.3 (left). These graphs are similar to those in Figure 4.1,
 560 except that: we show only implementations based on row-preferential microkernels;
 561 we omit 4M_1A; and we include results for complex GEMM implementations provided
 562 by OpenBLAS 0.3.6 [16] and Intel MKL 2019 Update 4 [12].

563 Figure 4.3 (right) shows multithreaded performance of the same implementations
 564 running with 24 threads.

565 These graphs show that BLIS’s complex assembly-based and 1M implementations
 566 typically outperform OpenBLAS while falling short in most (but not all) cases when
 567 compared to Intel’s MKL library.

568 **4.5. Additional results.** Additional performance results were gathered on a
 569 Marvell ThunderX2 compute server as well as an AMD EPYC (Zen) system. For
 570 brevity, we present and discuss that data in the appendix available online as sup-
 571 plementary materials. Those results reinforce the narrative provided here, lending
 572 even more evidence that the 1M method is capable of yielding high-performance im-
 573 plementations of complex matrix multiplication that are competitive with (and often
 574 outperform) other leading library solutions.

575 5. Observations.

576 **5.1. 4m limitations circumvented.** The previous article concluded by iden-
 577 tifying a number of limitations inherent in the 4M method. We now revisit this list
 578 and briefly discuss whether, to what degree, and how those limitations are overcome
 579 by algorithms based on the 1M method.

580 **Number of calls to primitive.** The most versatile 4M algorithm, 4M_1A, incurs
 581 up to a four-fold increase in function call overhead over a comparable assembly-based
 582 implementation. By comparison, 1M algorithms require at most a doubling of micro-
 583 kernel function call overhead, and in certain common cases (e.g., when $\beta \in \mathbb{R}$ and C
 584 is row- or column-stored), this overhead can be avoided completely. The 1M method

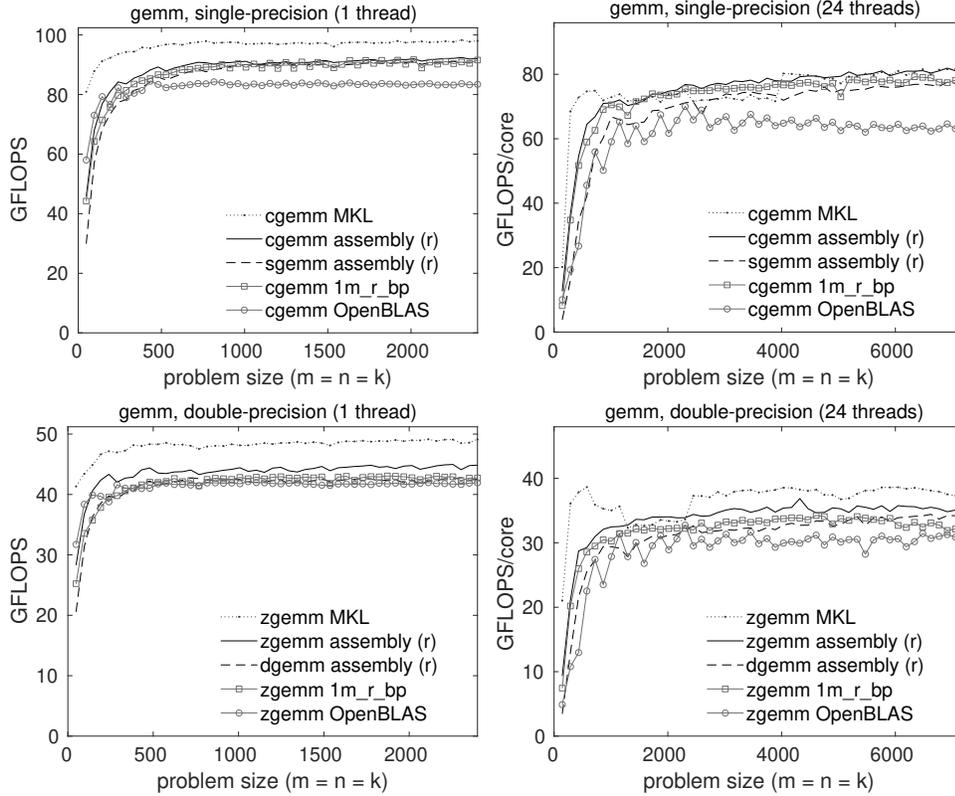


FIG. 4.3. Single-threaded (left) and multithreaded (right) performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core (left) or 12 cores (right) of an Intel Xeon E5-2690 v3 “Haswell” processor. All multithreaded data points reflect the use of 24 threads. The 1M curves are identical from those shown in Figures 4.1 and 4.2. The theoretical peak performance coincides with the top of each graph.

585 is a clear improvement over 4M due to its one-to-one substitution of the matrix mul-
 586 tiplication primitive.

587 **Inefficient reuse of input data from A , B , and C .** The most cache-efficient
 588 application of 4M is the lowest level algorithm, 4M_1A, which reuses F.E. of A , B ,
 589 and C from the L1 cache. But, as shown in Table 3.3, both 1M_R and 1M_C variants
 590 reuse F.E. of two of the three matrices from registers, with 1M_R_BP reusing F.E. of
 591 the third matrix from the L1 cache.

592 **Non-contiguous output to C .** Algorithms based on the 4M method must up-
 593 date only the real and then only the imaginary parts of the output matrix, twice
 594 each. When C is stored (by rows or columns) in the standard format, with real and
 595 imaginary F.E. interleaved, this piecemeal approach prevents the real microkernel
 596 from using vector load and store instructions on C during those four updates. The
 597 1M method avoids this issue altogether by packing A and B to formats that allow the
 598 real microkernel to update contiguous real and imaginary F.E. of C simultaneously.

599 **Reduction of k_C .** Algorithm 4M_1A requires that the real microkernel’s pre-
 600 ferred k_C blocksize be halved in the complex algorithm in order to maintain proper

601 cache footprints of \tilde{A}_i and \tilde{B}_p as well the footprints of their constituent micropanels.¹⁸
 602 Using these sub-optimally sized micropanels can noticeably hobble the performance
 603 of 4M_1A. Looking back at Table 3.1, it may seem like 1M suffers a similar handicap;
 604 however, the reason for halving k_C and its effect are both completely different. In the
 605 case of 1M, the use of $k_C^z = \frac{1}{2}k_C$ is simply a conversion of units (complex elements
 606 to real F.E.) for the purposes of identifying the size of the complex submatrices to
 607 be packed that will induce the optimal k_C value from the perspective of the real mi-
 608 crokernel, *not* a reduction in the F.E. footprint of the micropanels operated upon by
 609 that real microkernel. The ability of 1M to achieve high performance when $k = \frac{1}{2}k_C$
 610 is actually a strength for certain higher-level applications, such as Cholesky, LU, and
 611 QR factorizations based on rank- k update. Those operations tend to perform better
 612 when the algorithmic blocksize (corresponding to k_C) is as narrow as possible in order
 613 to limit the amount of computation in the lower-performing unblocked subproblem.

614 **Framework accommodation.** The 1M algorithms are no more disruptive to
 615 the BLIS framework than the most accommodating of 4M algorithms, 4M_1A. This is
 616 because, like with 4M_1A, almost all of the 1M implementation details are sequestered
 617 within the packing routines and the virtual microkernel.

618 **Interference with multithreading.** Because the 1M algorithms are imple-
 619 mented entirely within the packing routines and virtual microkernel, they parallelize
 620 just as easily as the most thread-friendly of the 4M algorithms, 4M_1A, and entirely
 621 avoid the threading difficulties of higher-level 4M algorithms.¹⁹

622 **Non-applicability to two-operand operations.** Certain higher-level appli-
 623 cations of 4M are inherently incompatible with two-operand operations because they
 624 would overwrite the original contents of the input/output operand even though subse-
 625 quent stages of computation depend on that original input. 1M avoids this limitation
 626 entirely. Like 4M_1A, 1M can easily be applied to two-operand level-3 operations such
 627 as TRMM and TRSM.²⁰

628 **5.2. Summary.** The analysis above suggests that the 1M method solves or
 629 avoids most of the performance-degrading weaknesses of 4M and in the remaining
 630 cases is no worse off than the best 4M algorithm.

631 **5.3. Limitations of 1m.** Although the 1M method avoids most of the weak-
 632 nesses inherent to the 4M method, a few notable caveats remain.

633 **Non-real values of beta.** In the most common cases where $\beta^i = 0$, the 1M
 634 implementation may employ the optimization described in Section 3.6.5. However,
 635 when $\beta^i \neq 0$, the virtual microkernel must be called. In such cases, 1M yields slightly
 636 lower performance due to extra memops.²¹

637 **Algorithmic dependence on I/O preference.** If the real domain microkernel
 638 is row-preferential (and thus performs row-oriented I/O on C), then the 1M implemen-
 639 tation must choose an algorithm based on the 1M_R variant. But (in this scenario),

¹⁸ Recall that the halving of k_C for 4M_1A was motivated by the desire to keep not just two, but four real micropanels in the L1 cache simultaneously. These correspond to the real and imaginary parts of the current micropanels of \tilde{A}_i and \tilde{B}_p .

¹⁹ This thread-friendly property holds even when the virtual microkernel is bypassed altogether as discussed in Section 3.6.5

²⁰ As with 4M_1A, 1M support for TRSM requires a separate pair of virtual microkernels that fuse a matrix multiplication with a triangular solve with n_R right-hand sides.

²¹ The 4M method suffers lower performance when $\beta^i \neq 0$ for similar reasons.

640 if 1M_C is instead preferred for some reason, then either the underlying microkernel
 641 needs to be updated to handle both row- and column-oriented I/O, or a new column-
 642 preferential microkernel must be written. A similar caveat holds if the real domain
 643 microkernel is column-preferential and the 1M_R variant is preferred.

644 **Higher bandwidth on \tilde{A}_i and \tilde{B}_p .** Compared to a conventional, assembly-
 645 based GEMM, implementations based on the 1M method require twice as much mem-
 646 ory bandwidth when reading packed matrices \tilde{A}_i and \tilde{B}_p . Microkernels that encode
 647 complex arithmetic at the assembly level are able to load real and imaginary F.E.
 648 and then reuse those F.E. from registers, thus increasing the microkernel’s arithmetic
 649 intensity. By contrast, the 1M method’s reliance on real domain microkernels means
 650 that it must reuse real and imaginary F.E. from some level of cache and thus incur
 651 additional memory traffic.²² The relative benefit of the conventional approach is likely
 652 to be most visible when parallelizing GEMM across all cores of a many-core system
 653 since that situation tends to saturate memory bandwidth.

654 **5.4. Further discussion.** Before concluding, we offer some final thoughts on
 655 the 1M method and its place in the larger spectrum of approaches to implementing
 656 complex matrix multiplication.

657 **5.4.1. Geometric interpretation.** Matrix multiplication is sometimes thought
 658 of as a three-dimensional operation with a contraction (accumulation) over the k -
 659 dimension. This interpretation carries into the complex domain as well. However, when
 660 each complex element is viewed in terms of its real and imaginary components, we
 661 find that a fourth pseudo-dimension of computation (of fixed size 2) emerges, one
 662 which also involves a contraction. The 1M method reorders and duplicates elements
 663 of A and B in such a way that exposes and “flattens” this extra dimension of com-
 664 putation. This, combined with the exposed treatment of real and imaginary F.E.,
 665 causes the resulting floating-point operations to appear indistinguishable from a real
 666 domain matrix multiplication with m and k dimensions (for column-stored C) or k
 667 and n dimensions (for row-stored C) that are twice as large.

668 **5.4.2. Data reuse: efficiency vs. programmability.** Both the conventional
 669 approach and 1M move data efficiently through the memory hierarchy.²³ However,
 670 once in registers, a conventional complex microkernel reuses those loaded values to
 671 perform twice as many flops as 1M. The previous article observes that all 4M algo-
 672 rithms make different variations of the same tradeoff: by forgoing the reuse of F.E.
 673 from registers and instead reusing those data from some level of cache, the algorithms
 674 avoid the need to explicitly encode complex arithmetic at the assembly level. As it
 675 turns out, 1M makes a similar tradeoff, but gives up less while gaining more: it is
 676 able to effectively reuse F.E. from two of the three matrix operands from registers
 677 while still avoiding the need for a complex microkernel, and it manages to replace
 678 that kernel operation with a single real matrix multiplication. And we would argue
 679 that increasing programmability and productivity by forfeiting a modest performance
 680 advantage is a good trade to make under almost any circumstance.

681 **5.4.3. Storage.** The supremacy of the 1M method is closely tied to the inter-
 682 leaved storage of real and imaginary values—specifically, of the output matrix C . If
 683 applications are motivated to instead store complex matrices in non-standard formats,

²² The 4M method suffers the same “bandwidth penalty” as 1M for the same reason.

²³ This is in contrast to, for example, Algorithm 4M_HW, which the previous article showed makes rather inefficient use of cache lines as they travel through the L3, L2, and L1 caches.

684 such as two real matrices, (one each for real and imaginary components) the 2M ap-
 685 proach (for numerically sensitive settings) as well as low-level applications of 3M (for
 686 numerically insensitive settings) may become more appropriate [19, 21].

687 **6. Conclusions.** We began the article by reviewing the general motivations for
 688 induced methods for complex matrix multiplication as well as the specific methods,
 689 3M and 4M, studied in the previous article. Then, we recast complex scalar multipli-
 690 cation (and accumulation) in such a way that revealed a template that could be used
 691 to fashion a new induced method, one that casts complex matrix multiplication in
 692 terms of a single real matrix product. The key is the application of two new packing
 693 formats on the left- and right-hand matrix product operands that allows us to dis-
 694 guise the complex matrix multiplication as a real matrix multiplication with slightly
 695 modified input parameters. This 1M method is shown to have two variants, one each
 696 favoring row-stored and column-stored output matrices. When implemented in the
 697 BLIS framework, competitive performance was observed for 1M algorithms on three
 698 modern microarchitectures. Finally, we reviewed the limitations of the 4M method
 699 that are overcome by 1M and concluded by discussing a few high-level observations.

700 The key takeaway from our study of induced methods is that the real and imag-
 701 inary elements of complex matrices can always be reordered to accommodate the
 702 desired fundamental primitives, whether those primitives are defined to be various
 703 forms of real matrix multiplication (as is the case for the 4M, 3M, 2M, and 1M meth-
 704 ods), or vector instructions (as is the case for microkernels that implement complex
 705 arithmetic in assembly code). Indeed, even in the real domain, the classic matrix
 706 multiplication algorithm’s packing format is simply a reordering of data that targets
 707 the fundamental primitive implicit in the microkernel—namely, an $m_R \times n_R$ rank-1
 708 update. The family of induced methods presented here and in the previous article ex-
 709 pand upon this basic reordering so that the mathematics of complex arithmetic can be
 710 expressed at different levels of the algorithm and of its corresponding implementation,
 711 each yielding different benefits, costs, and performance.

712 **Acknowledgements.** We thank the Texas Advanced Computing Center for pro-
 713 viding access to the the Intel Xeon “Lonestar5” (Haswell) compute node on which
 714 the performance data presented in Section 4 were gathered. We also kindly thank
 715 Marvell and Oracle Corporation for arranging access to the Marvell ThunderX2 and
 716 AMD EPYC (Zen) systems, respectively, on which the performance data presented in
 717 Appendix A were gathered. Finally, we thank Devangi Parikh for helpfully gathering
 718 the results on the ThunderX2 system.

719

REFERENCES

- 720 [1] J. J. DONGARRA, J. DU CROZ, S. HAMMARLING, AND I. DUFF, *A set of level 3 basic linear*
 721 *algebra subprograms*, ACM Trans. Math. Soft., 16 (1990), pp. 1–17.
 722 [2] G. FRISON, D. KOUZOUPIS, T. SARTOR, A. ZANELLI, AND M. DIEHL, *BLASFEO: Basic lin-*
 723 *ear algebra subroutines for embedded optimization*, ACM Trans. Math. Soft., 44 (2018),
 724 pp. 42:1–42:30, <https://doi.org/10.1145/3210754>, <http://doi.acm.org/10.1145/3210754>.
 725 [3] K. GOTO AND R. A. VAN DE GEIJN, *Anatomy of high-performance matrix multiplication*, ACM
 726 Trans. Math. Soft., 34 (2008), pp. 12:1–12:25, <https://doi.org/10.1145/1356052.1356053>,
 727 <http://doi.acm.org/10.1145/1356052.1356053>.
 728 [4] K. GOTO AND R. A. VAN DE GEIJN, *High-performance implementation of the level-3 BLAS*,
 729 ACM Trans. Math. Soft., 35 (2008), pp. 4:1–4:14, <https://doi.org/10.1145/1377603.1377607>,
 730 <http://doi.acm.org/10.1145/1377603.1377607>.
 731 [5] J. A. GUNNELS, G. M. HENRY, AND R. A. VAN DE GEIJN, *A family of high-performance matrix*
 732 *multiplication algorithms*, in Proceedings of the International Conference on Computational

- 733 Sciences-Part I, ICCS '01, Berlin, Heidelberg, 2001, Springer-Verlag, pp. 51–60, <http://dl.acm.org/citation.cfm?id=645455.653765>.
- 734
- 735 [6] N. J. HIGHAM, *Stability of a method for multiplying complex matrices with three real matrix*
- 736 *multiplications*, SIAM J. Matrix Anal. App., 13 (1992), pp. 681–687, [https://doi.org/10.](https://doi.org/10.1137/0613043)
- 737 [1137/0613043](https://doi.org/10.1137/0613043), <https://doi.org/10.1137/0613043>.
- 738 [7] J. HUANG, *Practical fast matrix multiplication algorithms*, (2018). PhD thesis, The University
- 739 of Texas at Austin.
- 740 [8] J. HUANG, D. A. MATTHEWS, AND R. A. VAN DE GEIJN, *Strassen's algorithm for tensor contrac-*
- 741 *tion*, SIAM Journal on Scientific Computing, 40 (2018), pp. C305–C326, [https://doi.org/](https://doi.org/10.1137/17M1135578)
- 742 [10.1137/17M1135578](https://doi.org/10.1137/17M1135578), <https://doi.org/10.1137/17M1135578>, <https://arxiv.org/abs/https://doi.org/10.1137/17M1135578>.
- 743
- 744 [9] J. HUANG, L. RICE, D. A. MATTHEWS, AND R. A. VAN DE GEIJN, *Generating families of*
- 745 *practical fast matrix multiplication algorithms*, in 31th IEEE International Parallel and
- 746 Distributed Processing Symposium (IPDPS 2017), May 2017, pp. 656–667, [https://doi.](https://doi.org/10.1109/IPDPS.2017.56)
- 747 [org/10.1109/IPDPS.2017.56](https://doi.org/10.1109/IPDPS.2017.56).
- 748 [10] J. HUANG, T. M. SMITH, G. M. HENRY, AND R. A. VAN DE GEIJN, *Strassen's algorithm reloaded*,
- 749 in Proceedings of the International Conference for High Performance Computing, Network-
- 750 ing, Storage and Analysis, SC '16, Piscataway, NJ, USA, 2016, IEEE Press, pp. 59:1–59:12,
- 751 <http://dl.acm.org/citation.cfm?id=3014904.3014983>.
- 752 [11] J. HUANG, C. D. YU, AND R. A. VAN DE GEIJN, *Implementing Strassen's algorithm with CUT-*
- 753 *LASS on NVIDIA Volta GPUs*, FLAME Working Note #88, TR-18-08, The University
- 754 of Texas at Austin, Department of Computer Science, 2018, [https://apps.cs.utexas.edu/](https://apps.cs.utexas.edu/apps/sites/default/files/tech_reports/GPUStrassen.pdf)
- 755 apps/sites/default/files/tech_reports/GPUStrassen.pdf.
- 756 [12] INTEL, *Math Kernel Library*. <https://software.intel.com/en-us/mkl>, 2019.
- 757 [13] INTEL CORPORATION, *Intel[®] 64 and IA-32 Architectures Optimization Reference Manual*,
- 758 no. 248966-033, June 2016.
- 759 [14] INTEL CORPORATION, *Intel[®] Xeon[®] Processor E5 v3 Product Family: Processor Specification*
- 760 *Update*, no. 330785-010US, September 2016.
- 761 [15] T. M. LOW, F. D. IGUAL, T. M. SMITH, AND E. S. QUINTANA-ORTÍ, *Analytical modeling is*
- 762 *enough for high-performance BLIS*, ACM Trans. Math. Soft., 43 (2016), pp. 12:1–12:18,
- 763 <https://doi.org/10.1145/2925987>, <http://doi.acm.org/10.1145/2925987>.
- 764 [16] *OpenBLAS*. <http://xianyi.github.com/OpenBLAS/>, 2019.
- 765 [17] D. T. POPOVICI, F. FRANCHETTI, AND T. M. LOW, *Mixed data layout kernels for vector-*
- 766 *ized complex arithmetic*, in 2017 IEEE High Performance Extreme Computing Conference
- 767 (HPEC), Sep. 2017, pp. 1–7, <https://doi.org/10.1109/HPEC.2017.8091024>.
- 768 [18] T. M. SMITH, R. A. VAN DE GEIJN, M. SMELYANSKIY, J. R. HAMMOND, AND F. G. VAN
- 769 ZEE, *Anatomy of high-performance many-threaded matrix multiplication*, in Proceedings
- 770 of the 28th IEEE International Parallel & Distributed Processing Symposium (IPDPS),
- 771 IPDPS '14, Washington, DC, USA, 2014, IEEE Computer Society, pp. 1049–1059, <https://doi.org/10.1109/IPDPS.2014.110>, <https://doi.org/10.1109/IPDPS.2014.110>.
- 772
- 773 [19] F. G. VAN ZEE, *Inducing complex matrix multiplication via the 1m method*, FLAME Working
- 774 Note #85 TR-17-03, The University of Texas at Austin, Department of Computer Sciences,
- 775 February 2017.
- 776 [20] F. G. VAN ZEE, T. SMITH, F. D. IGUAL, M. SMELYANSKIY, X. ZHANG, M. KISTLER, V. AUSTEL,
- 777 J. GUNNELS, T. M. LOW, B. MARKER, L. KILLOUGH, AND R. A. VAN DE GEIJN, *The BLIS*
- 778 *framework: Experiments in portability*, ACM Trans. Math. Soft., 42 (2016), pp. 12:1–12:19,
- 779 <http://doi.acm.org/10.1145/2755561>.
- 780 [21] F. G. VAN ZEE AND T. M. SMITH, *Implementing high-performance complex matrix multiplica-*
- 781 *tion via the 3m and 4m methods*, ACM Trans. Math. Soft., 44 (2017), pp. 7:1–7:36.
- 782 [22] F. G. VAN ZEE AND R. A. VAN DE GEIJN, *BLIS: A framework for rapidly instantiating BLAS*
- 783 *functionality*, ACM Trans. Math. Soft., 41 (2015), pp. 14:1–14:33, [http://doi.acm.org/10.](http://doi.acm.org/10.1145/2764454)
- 784 [1145/2764454](http://doi.acm.org/10.1145/2764454).
- 785 [23] R. C. WHALEY, A. PETITET, AND J. J. DONGARRA, *Automated empirical optimization of soft-*
- 786 *ware and the ATLAS project*, Parallel Computing, 27 (2001), pp. 3–35, [https://doi.org/](https://doi.org/https://doi.org/10.1016/S0167-8191(00)00087-9)
- 787 [https://doi.org/https://doi.org/10.1016/S0167-8191\(00\)00087-9](https://doi.org/https://doi.org/10.1016/S0167-8191(00)00087-9), [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0167819100000879)
- 788 [article/pii/S0167819100000879](http://www.sciencedirect.com/science/article/pii/S0167819100000879). New Trends in High Performance Computing.
- 789 [24] C. D. YU, J. HUANG, W. AUSTIN, B. XIAO, AND G. BIROS, *Performance optimization for*
- 790 *the k-nearest neighbors kernel on x86 architectures*, in Proceedings of the International
- 791 Conference for High Performance Computing, Networking, Storage and Analysis, SC '15,
- 792 New York, NY, USA, 2015, ACM, pp. 7:1–7:12, <https://doi.org/10.1145/2807591.2807601>,
- 793 <http://doi.acm.org/10.1145/2807591.2807601>.

794 Appendix A. Additional Performance Results.

795 In this section we present performance results for implementations of 1M method
 796 on two additional types of hardware. The primary purpose of gathering these results
 797 was to confirm 1M performance on additional architectures beyond the Intel Haswell
 798 system reported on in the main article.

799 **A.1. Marvell ThunderX2.** In this section, we report the performance of the
 800 1M method on the Marvell ThunderX2, a high-performance ARMv8 microarchitec-
 801 ture.

802 **A.1.1. Platform and implementation details.** Results presented in this sec-
 803 tion were gathered on a single compute node consisting of two 28-core Marvell Thun-
 804 derX2 CN9975 processors.²⁴ Each core, running at a clock rate of 2.2 GHz, provides
 805 a single-core peak performance of 17.6 gigaflops (GFLOPS) in double precision and
 806 35.2 GFLOPS in single precision. Each socket has a 32MB L3 cache that is shared
 807 among cores, and each core has a private 256KB L2 cache and 32KB L1 (data) cache.
 808 Performance experiments were gathered under the Ubuntu 16.04 operating system
 809 running the Linux 4.15.0 kernel. Source code was compiled by the GNU C compiler
 810 (`gcc`) version 7.3.0.²⁵ The version of BLIS used in these tests was version 0.5.0-1.²⁶

811 In this section, we show 1M results for only Algorithm 1M_C_BP. Unlike the
 812 results shown in the main article, we did not develop conventional assembly-based
 813 microkernels and thus cannot compare against a complex domain solution based on
 814 those kernels. For further comparison, we measured performance for the complex
 815 GEMM implementations found in OpenBLAS²⁷ and ARMPL 18.4.0.

816 All other parameters, such as values of α and β , and the number of trials per-
 817 formed for each problem size, as well as graphing conventions, such as scaling of the
 818 y -axis, remain identical to those of the main article.

819 **A.1.2. Analysis.** Figure A.1 contains single-threaded (left) and multithreaded
 820 (right) performance of single-precision (top) and double-precision (bottom) complex
 821 GEMM implementations. In addition to the 1M_C_BP implementation within BLIS,
 822 we also show the corresponding real domain GEMM implementation and the `cgemm` or
 823 `zgemm` found in OpenBLAS and ARMPL. For all BLIS implementations, we employed
 824 4-way parallelism within the 5th loop and 14-way parallelism within the 3rd loop for
 825 a total of 56 threads.

826 In Figure A.1 (top-left), single-precision 1M and its corresponding real domain
 827 benchmark track each other closely in the single-threaded configurations tested, as we
 828 would have expected. Somewhat surprisingly, the vendor library, ARMPL, does not
 829 appear to scale well at 56 threads, as shown in Figure A.1 (top-right). Also somewhat
 830 surprisingly, OpenBLAS performance is consistently low, even for sequential execu-
 831 tion. This suggests that while parallelism may be well-configured, their kernel is likely
 832 underperforming.

²⁴ While four-way symmetric multithreading is available on this hardware, the feature was disabled at boot-time so that the operating system detects only one logical core per physical core and schedules threads accordingly.

²⁵ The following optimization flags were used during compilation of BLIS and its test drivers: `-O3 -ftree-vectorize -mtune=cortex-a57`. In addition to those flags, the following flags were also used when compiling assembly kernels: `-march=armv8-a+fp+simd -mcpu=cortex-a57`.

²⁶ This version of BLIS may be uniquely identified, with high probability, by the first 10 digits of its `git` “commit” (SHA1 hash) number: e90e7f309b.

²⁷ This version of OpenBLAS may be uniquely identified, with high probability, by the first 10 digits of its `git` commit number: 52d3f7af50.

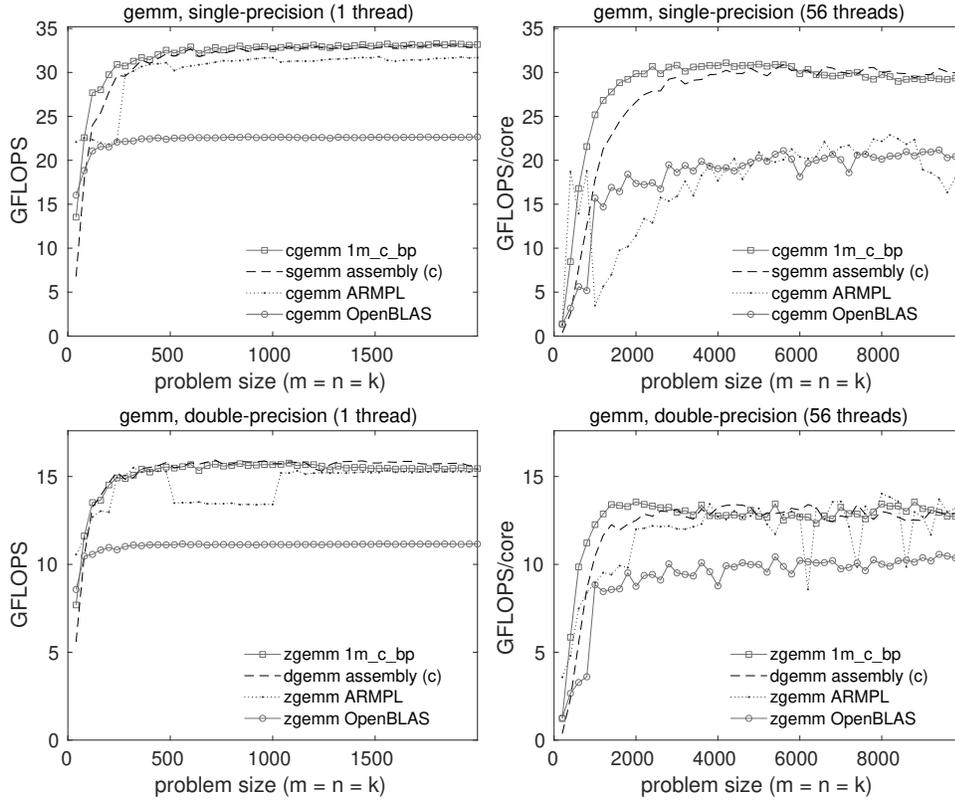


FIG. A.1. Single-threaded (left) and multithreaded (right) performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core (left) or 56 cores (right) of a Marvell ThunderX2 CN9975 processor. All multithreaded data points reflect the use of 56 threads. The real domain GEMM implementation from BLIS uses a column-preferential microkernel, as indicated the a “(c)” in the legends. (The 1M_C_BP implementation uses the same column-preferential microkernel as the real domain GEMM implementation.) The theoretical peak performance coincides with the top of each graph.

833 Figure A.1 (bottom) tells a similar story of performance among double-precision
 834 implementations, except that all BLIS implementations are, for reasons not immedi-
 835 ately obvious, somewhat less efficient relative to peak performance than their single-
 836 precision counterparts. ARMPL performance is more competitive for both one and 56
 837 threads, though the single-core graph exposes evidence of a “crossover point” strat-
 838 egy gone awry. ARMPL also seems to exhibit large swings in performance for certain
 839 large, multithreaded problem sizes. Once again, OpenBLAS performance is much
 840 lower, but consistently so.

841 In summary, BLIS’s 1M implementation performs extremely well on the Marvell
 842 CN9975 when computing in single precision. Performance and scalability in double
 843 precision, while not quite as impressive, is still highly competitive, especially when
 844 compared to OpenBLAS and the ARM Performance Library.

845 **A.2. AMD Zen.** In this section, we report the performance of the 1M method
 846 on the AMD Zen microarchitecture.

847 **A.2.1. Platform and implementation details.** Results presented in this sec-
 848 tion were gathered on a single compute node consisting of two 32-core AMD EPYC
 849 7551 (Zen) processors.²⁸ Each core runs at a clock rate of 3.0 GHz when using a
 850 single core and 2.55 GHz when utilizing all cores simultaneously. The former clock
 851 rate yields a single-core peak performance of 24.0 GFLOPS in double precision and
 852 48.0 GFLOPS in single precision, and the latter clock rate yields a multicore peak
 853 performance of 20.4 GFLOPS/core and 40.8 GFLOPS/core for single- and double-
 854 precision computation, respectively. Each socket has a 64MB of L3 cache (distributed
 855 as 8MB for each four-core complex) that is shared among cores, and each core has a
 856 private 512KB L2 cache and 32KB L1 (data) cache. Performance experiments were
 857 gathered under the Ubuntu 18.04 operating system running the Linux 4.15.0 kernel.
 858 Source code was compiled by the GNU C compiler (`gcc`) version 7.4.0.²⁹ The version
 859 of BLIS used in these tests was version 0.6.0-1266.³⁰

860 In this section, we show 1M results for only Algorithm 1M_R_BP. For reference, we
 861 also measured performance for the complex GEMM implementations found in Open-
 862 BLAS 0.3.7 and Intel MKL 2020 (initial release).

863 All other parameters, such as values of α and β , and the number of trials per-
 864 formed for each problem size, as well as graphing conventions, such as scaling of the
 865 y -axis, remain identical to those of the main article.

866 **A.2.2. Analysis.** Figure A.2 contains single-threaded (left) and multithreaded
 867 (right) performance of single-precision (top) and double-precision (bottom) complex
 868 GEMM implementations. In addition to the 1M_R_BP implementation within BLIS, we
 869 also show the corresponding real and complex domain GEMM implementations based
 870 on conventional assembly-coded kernels. We also show the `cgemm` or `zgemm` found in
 871 OpenBLAS and MKL. For all BLIS implementations, we employed 2-way parallelism
 872 within the 5th loop, 8-way parallelism within the 3rd loop, and 4-way parallelism
 873 within the 2nd loop for a total of 64 threads.

874 In Figure A.2 (top-left), all implementations track closely together except for
 875 MKL.³¹ We see a similar pattern for single-threaded double precision in Figure A.2
 876 (bottom-left).

877 In Figure A.2 (top-right) and (bottom-right), we see multithreaded performance
 878 when utilizing all 64 cores of the AMD EPYC system. The relative performance of
 879 1M_R_BP is consistent with the results seen previously on Haswell. That is, the 1M
 880 method facilitates performance that meets or exceeds the performance of an optimized
 881 real domain implementation of GEMM (i.e., one that uses the same microkernels as 1M),
 882 but falls slightly short of the performance of a conventional assembly-coded complex
 883 domain GEMM. Once again, MKL performance suffers noticeably on AMD hardware.
 884 OpenBLAS lags somewhat behind the BLIS-based implementations, but performance
 885 unexpectedly drops for very large problem sizes. This behavior was reproducible,
 886 though the exact problem size at which the drop-off occurred shifted across repeated

²⁸ While two-way symmetric multithreading is available on this hardware, a maximum of one logical core per physical core was utilized during our tests.

²⁹ The following optimization flags were used during compilation of BLIS and its test drivers: `-O3 -march=znver1`. Furthermore, all test drivers were run via `numactl -i all`.

³⁰ This version of BLIS may be uniquely identified, with high probability, by the first 10 digits of its `git` “commit” (SHA1 hash) number: `f391b3e2e7`.

³¹ We hypothesize that as MKL parses the results of the `CPUID` instruction, it detects an unexpected CPU vendor (AMD instead of Intel) and therefore selects a “fallback” (safe but low-performing) kernel. If this is the case, then the fix would be trivial, which suggests that MKL’s underperformance on AMD hardware is deliberate.

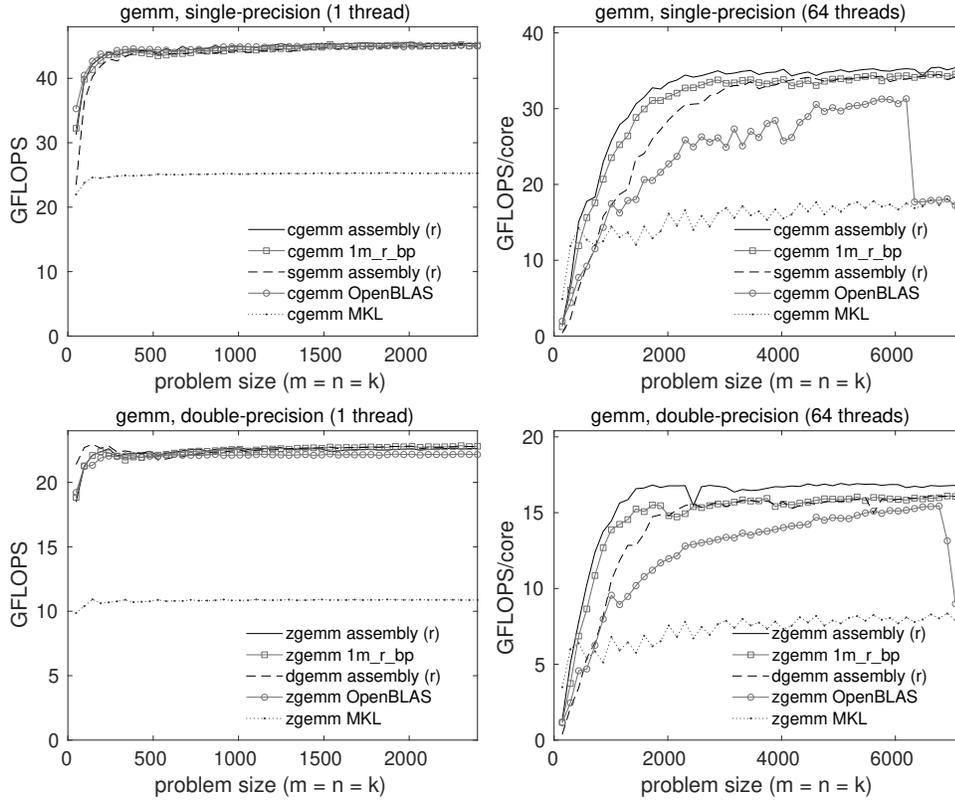


FIG. A.2. Single-threaded (left) and multithreaded (right) performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core (left) or 64 cores (right) of an AMD EPYC 7551 (Zen) processor. All multithreaded data points reflect the use of 64 threads. The real and complex domain GEMM implementations from BLIS use row-preferential microkernels, as indicated the a “(r)” in the legends. (The 1M_R_BP implementation uses the same row-preferential microkernel as the real domain GEMM implementation.) The theoretical peak performance coincides with the top of each graph.

887 experiments.

888 In summary, BLIS’s 1M implementation performs very well on the AMD EPYC
 889 7551 when computing in single and double precision, exceeding the performance of
 890 both OpenBLAS and MKL. Scalability (relative to theoretical peak) is also quite
 891 good in both precisions considering the challenges that NUMA-based architectures
 892 sometimes pose to parallelization efforts.