

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

3 FIELD G. VAN ZEE*

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

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

19 **AMS subject classifications.** 65Y04

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

38 This low-level kernel approach carries distinct benefits. Pushing the nuances and
39 complexities of complex arithmetic down to the level of the kernel allows the higher-
40 level loop infrastructure within the matrix multiplication to remain largely the same

*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.

41 as its real domain counterpart. (See Figure 1.1.) Another benefit leverages a key
 42 difference between the real and complex forms of nearly all matrix computations:
 43 arithmetic intensity. Complex matrix multiplication (regardless of how it is imple-
 44 mented) requires four times the number of floating-point operations but only twice
 45 the number of memory operations. Encoding complex arithmetic within assembly
 46 code allows the real and imaginary components of the multiplication operands to be
 47 loaded into registers and then resued at virtually no cost. The impact of the two-fold
 48 increase in memory operations is further minimized thanks to the standard format
 49 for storing complex matrices, which places an element’s real and imaginary parts
 50 adjacent to one another², allowing both to be accessed conveniently via contiguous
 51 load and store instructions. Thanks to this low-cost reuse and accommodating stor-
 52 age format, implementations based on assembly-based complex kernels are capable of
 53 achieving a somewhat larger fraction of the hardware’s peak performance relative to
 54 real domain kernels.

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

65 Because of how software developers have historically designed their implemen-
 66 tations, many assume that supporting complex matrix multiplication operations first
 67 requires writing complex domain kernels. To our pleasant surprise, we have discovered
 68 a new way for developers to implement high-performance complex matrix multiplica-
 69 tion *without* those kernels.

70 The predecessor to the current article investigates whether (and to what degree of
 71 effectiveness) real domain matrix multiplication kernels can be repurposed and lever-
 72 aged toward the implementation of complex matrix multiplication [21]. The authors
 73 develop a new class of algorithms that implement these so-called “induced methods”
 74 for matrix products in the complex domain. Instead of relying on an assembly-coded
 75 complex kernel, as a conventional implementation would, these algorithms express
 76 complex matrix multiplication only in terms of real domain primitives.³ We consider
 77 the current article a companion and follow-up to that previous work [21].

78 In this article, we will consider a new method for emulating complex matrix
 79 multiplication using only real domain building blocks, and we will once again show
 80 that a clever rearrangement of the real and imaginary elements within the internal
 81 “packed” matrices is key to facilitating high performance. The novelty behind this
 82 new method is that the semantics of complex arithmetic are encoded entirely within
 83 a special data layout, which allows each call to the complex matrix multiplication
 84 kernel to be replaced with just *one* call to a real matrix multiplication kernel. This
 85 substitution is possible because a real matrix multiplication on the reorganized data

² The widely-accepted BLAS interface requires use of this standard format.

³ In [21], the authors use the term “primitive” to refer to a functional abstraction that implements
 a single real matrix multiplication. Such primitives are often not general purpose and may come with
 significant prerequisites to facilitate their use.

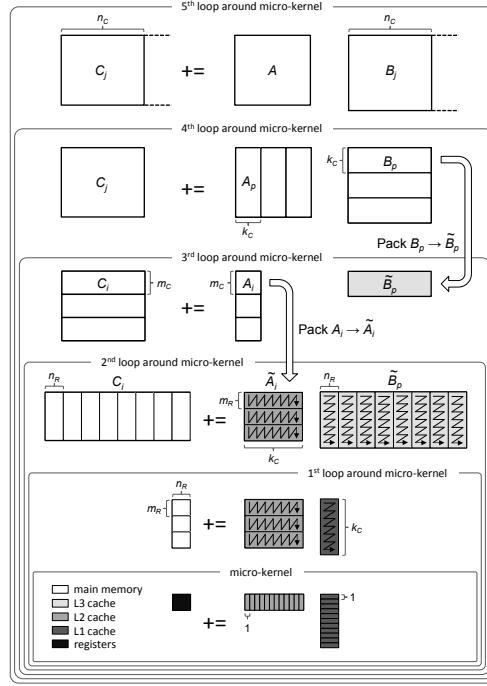


FIG. 1.1. An illustration of the algorithm for computing high-performance matrix multiplication, taken from [21], which expresses computation in terms of a so-called “block-panel” subproblem.

86 mimics the computation and I/O of a comparable complex matrix multiplication on
 87 the unaltered data. Because of this one-to-one equivalence, we call it the 1M method.

88 **1.1. Contributions.** This article makes the following contributions:

89 • It introduces⁴ the 1M method along with two algorithmic variants and an
 90 analysis of issues germane to their high-performance implementations, in-
 91 cluding workspace, packing formats, cache behavior, multithreadability, and
 92 programming effort. A detailed review shows how 1M avoids all of the major
 93 challenges observed of the 4M method.

94 • It promotes code reuse and portability by continuing the previous article’s
 95 focus on solutions which may be cast in terms of real matrix multiplication
 96 kernels. Such solutions have clear implications for developer productivity, as
 97 they allow kernel authors to focus their efforts on fewer and simpler kernels.

98 • It builds on the theme of the BLIS framework as a productivity multiplier [22],
 99 further demonstrating how complex matrix multiplication may be imple-
 100 mented with relatively minor modifications to the source code and in such a
 101 way that results in immediate instantiation of complex implementations for
 102 *all* level-3 BLAS-like operations.

103 • It demonstrates performance of 1M implementations that is not only superior
 104 to the previous effort based on the 4M method but also competitive with
 105 solutions based on complex matrix kernels.

106 • It serves as a reference guide to the 1M implementations for complex matrix

⁴ This proposed 1M method was first published [19].

107 multiplication found within the BLIS framework, which is available to the
 108 community under the open-source 3-clause BSD software license.

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

113 **1.2. Notation.** In this article, we continue the notation established in [21].
 114 Specifically, we use uppercase Roman letters (e.g. A , B , and C) to refer to ma-
 115 trices, lowercase Roman letters (e.g. x , y , and z) to refer to vectors, and lowercase
 116 Greek letters (e.g. χ , ψ , and ζ) to refer to scalars. Subscripts are used typically to
 117 denote sub-matrices within a larger matrix (e.g. $A = (A_0 | A_1 | \cdots | A_{n-1})$) or
 118 scalars within a larger matrix or vector.

119 We make extensive use of superscripts to denote the real and imaginary compo-
 120 nents of a scalar, vector, or (sub-)matrix. For example, $\alpha^r, \alpha^i \in \mathbb{R}$ denote the real
 121 and imaginary parts, respectively, of a scalar $\alpha \in \mathbb{C}$. Similarly, A^r and A^i refer to the
 122 real and imaginary parts of a complex matrix A , where A^r and A^i are real matrices
 123 with dimensions identical to A . Note that while this notation for real, imaginary, and
 124 complex matrices encodes information about content and origin, it does not encode
 125 how the matrices are actually stored. We will explicitly address storage details as
 126 implementation issues are discussed.

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

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

137 This article discusses several hypothetical algorithms and functions. Unless oth-
 138 erwise noted, a call to function FUNC that implements $C := C + AB$ appears as [
 139 C] := FUNC(A , B , C). We will also reference functions that access properties of
 140 matrices. For example, M(A) and N(A) would return the m and n dimensions of a
 141 matrix A , while RS(B) and CS(B) would return the row and column strides of B .

142 2. Background and review.

143 **2.1. Motivation.** In [21], the authors list three primary motivating factors be-
 144 hind their effort to seek out methods for inducing complex matrix multiplication via
 145 real domain kernels:

- 146 **• Productivity.** By inducing complex matrix multiplication from real domain
 147 kernels, the number of kernels that must be supported would be halved.
 148 This allows the DLA library developers to focus on a smaller and simpler
 149 set of real domain kernels. This benefit would manifest most obviously when
 150 instantiating BLAS-like functionality on new hardware [20].
- 151 **• Portability.** Induced methods avoid dependence on complex domain kernels
 152 because they encode the idea of complex matrix product at a higher level.
 153 This would naturally allow us to encode such methods portably within a
 154 framework such as BLIS [22]. Once integrated into the framework, developers

155 and users would benefit from the immediate availability of complex matrix
 156 multiplication implementations whenever real matrix kernels were present.

157 • **Performance.** Implementations of complex matrix multiplication that rely
 158 on real domain kernels would likely inherit the high-performance properties
 159 of those kernels. Any improvement to the real kernels would benefit both real
 160 and complex domains.

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

163 **2.2. The 3m and 4m methods.** The authors of [21] investigated two general
 164 ways of inducing complex matrix multiplication: the 3M method and the 4M method.
 165 These methods are then contrasted to the conventional approach, whereby a blocked
 166 matrix multiplication algorithm is executed with a complex domain kernel—one that
 167 implements complex arithmetic at the scalar level, in assembly language.

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

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

170 We then observe that we can apply such a definition to complex matrices $A \in \mathbb{C}^{m \times k}$,
 171 $B \in \mathbb{C}^{k \times n}$, and $C \in \mathbb{C}^{m \times n}$, provided that we can reference the real and imaginary
 172 parts as logically separate submatrices:

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

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

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

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

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

181 The authors of [21] observe that both methods may be applied to any particular
 182 level of a blocked matrix multiplication algorithm, resulting in several algorithms,
 183 each exhibiting somewhat different properties. Furthermore, they show how either
 184 method's implementation is facilitated by reordering real and imaginary elements
 185 within the internal storage format used when making packed copies of the current
 186 matrix blocks.⁵ The blocked algorithm used in that article is shown in Figure 1.1 and
 187 revisited in Section 2.4 of the present article.

188 Algorithms that implement the 3M method were found to yield “effective flops
 189 per second” performance that not only exceeded that of 4M, but also approached or

190 ⁵ Others have exploited the careful design of packing and computational primitives in an effort to
 191 improve performance, including in the context of Strassen's algorithm [7, 9, 10, 11], the computation
 192 of the K-Nearest Neighbors [24], tensor contraction [8], and Fast Fourier Transform [17].

199 exceeded the theoretical peak rate of the hardware.⁶ Unfortunately, these compelling
 200 results come at a cost: the numerical properties of implementations based on 3M
 201 are slightly less robust than that of algorithms based on the conventional approach
 202 or 4M. And although the author of [6] found that 3M was stable enough for most
 203 practical purposes, many applications will be unwilling to stray from the numerical
 204 expectations implicit in conventional matrix multiplication. Thus, going forward, we
 205 will focus on 4M as the standard reference method against which we will compare.

206 It is worth briefly considering the simplest approach to implementing the 4M
 207 method, which is hinted at by Eq. 2.2. This straightforward algorithm would invoke
 208 the real domain GEMM four times, computing $A^r B^r$ and $A^i B^i$ to update C^r and $A^i B^r$
 209 and $A^r B^i$ to update C^i . This is possible as long as the matrices’ real and imaginary
 210 parts are separately addressable, as they are in modern BLAS-like frameworks such
 211 as BLIS [22]. The previous article studied this high-level instance of the 4M method,
 212 dubbed Algorithm 4M_HW, and classified it into a family of related algorithms. Unfor-
 213 tunately, Algorithm 4M_HW does not perform well with the standard format for storing
 214 complex matrices. The reason: Algorithm 4M_HW computes with A^r , A^i , B^r , B^i , C^r ,
 215 and C^i as separate logical matrices, but since their F.E. are stored non-contiguously,
 216 those F.E. cannot be accessed efficiently on modern hardware. If Algorithm 4M_HW
 217 instead computed upon matrices that split the storage of their real and imaginary
 218 parts into two separate matrices, each with contiguous rows or columns, then its ex-
 219 pected performance would rise to match that of real matrix multiplication. However,
 220 even with this somewhat exotic “split” complex storage format, Algorithm 4M_HW
 221 would carry some disadvantages relative to the new induced method discussed later
 222 in this article.⁷

223 Thus, our aim is to develop an induced method that (1) yields performance that is
 224 at least as high as that of a corresponding real domain GEMM while also (2) allowing
 225 applications to continue using the standard storage format⁸ and (3) avoiding key
 226 disadvantages inherent in the various 4M algorithms, including 4M_HW.

227 **2.3. Previous findings.** For the reader’s convenience, we will now summarize
 228 the key findings, observations, and other highlights from the previous article regarding
 229 algorithms and implementations based on the 4M method [21].

230 • Since all algorithms in the 4M family execute the same number of flops, the
 231 algorithms’ relative performance depends entirely on (1) the number of mem-
 232 ops executed and (2) the level of cache from which F.E. of the packed matrices
 233 \tilde{A}_i and \tilde{B}_p are reused⁹. The number of memops is affected only by a halving
 234 of certain cache blocksize needed in order to leave cache footprints of \tilde{A}_i and

⁶ 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.

⁷ For example, parallelizing Algorithm 4M_HW may be limited by the three implicit synchronization points that would occur between the four invocations of real domain GEMM. Also, it was shown in the previous article that Algorithm 4M_HW inherently can only be applied to two-operand level-3 operations such as TRMM and TRSM by using $m \times n$ workspace [21].

⁸ An implementation may use the split format internally while still requiring the standard storage format at the user level. However, this technique, which is employed on more granular scale by Algorithm 4M_1A in the previous article, would incur a noticeable increase in memory operations and serve as a net drag on performance [21].

⁹ 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.

235 \tilde{B}_p unchanged. The level of cache from which F.E. are reused is determined
 236 by the level of the GEMM algorithm to which the 4M method was applied.
 237 • The lowest-level application, algorithm 4M_1A, efficiently moves F.E. of A , B ,
 238 and C from main memory to the L1 cache only once per rank- k_C update and
 239 reuses F.E. from the L1 cache. It relies on a relatively simple packing format
 240 and requires negligible, fixed-size workspace, is well-suited for multithreading,
 241 and is minimally disruptive to the BLIS framework. Algorithm 4M_1A can
 242 also be extended relatively easily to all other level-3 operations.
 243 • The conventional assembly-based approach to complex matrix multiplication
 244 can be viewed as a special case of 4M in which F.E. are reused from registers
 245 rather than cache. In this way, a conventional implementation embodies the
 246 lowest-level application of 4M possible, in which the method is applied to
 247 individual scalars (and then optimally encoded via vector instructions).
 248 • The way complex numbers are stored has a significant effect on performance.
 249 The standard format adopted by the community (and required by the BLAS),
 250 which uses an interleaved pair-wise storage of real and imaginary values,
 251 naturally favors conventional implementations because they can reuse F.E.
 252 from vector registers. However, this storage is awkward for algorithms based
 253 on 4M (and 3M) because it stymies the use of vector instructions for loading
 254 and storing F.E. of C^r and C^i . The 4M_1A algorithm already suffers from a
 255 *quadrupling*¹⁰ of the number of memops on C in addition to being forced to
 256 access these F.E. in a non-contiguous manner.
 257 • While the performance of Algorithm 4M_1A exceeds that of its simpler sibling,
 258 4M_HW, it not only falls short of a comparable conventional solution, it also
 259 falls short of its real domain “benchmark”—that is, the performance of a
 260 similar problem size in the real domain computed by an optimized algorithm
 261 using the same real domain kernel.

262 **2.4. Revisiting the matrix multiplication algorithm.** In this section, we
 263 review a common algorithm for high-performance matrix multiplication on conven-
 264 tional microprocessor architectures. This algorithm was first reported on in [3] and
 265 further refined in [22]. Figure 1.1 illustrates the key features of this algorithm.

266 The current state-of-the-art formulation of the matrix multiplication algorithm
 267 consists of six loops, the last of which resides within a microkernel that is typically
 268 highly optimized for the target hardware. These loops partition the matrix operands
 269 using carefully chosen cache (n_C , k_C , and m_C) and register (m_R and n_R) blocksizes
 270 that result in submatrices residing favorably at various levels of the cache hierarchy
 271 so as to allow data to be reused many times. In addition, submatrices of A and B are
 272 copied (“packed”) to temporary workspace matrices (\tilde{A}_i and \tilde{B}_p , respectively) in such
 273 a way that allows the microkernel to subsequently access matrix elements contiguously
 274 in memory, which improves cache and TLB performance. The cost of this packing is
 275 amortized over enough computation that its impact on overall performance is negli-
 276 gible for all but the smallest problems. At the lowest level, within the microkernel
 277 loop, an $m_R \times 1$ micro-column and a $1 \times n_R$ micro-row are loaded from the current
 278 micropanels of \tilde{A}_i and \tilde{B}_p , respectively, so that the outer product of these vectors
 279 may be computed to update the corresponding $m_R \times n_R$ submatrix, or micro-tile, of
 280 C . The individual floating-point operations that constitute these tiny rank-1 updates

¹⁰ 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 .

281 are oftentimes executed via vector instructions (if the architecture supports them) in
 282 order to maximize utilization of the floating-point unit(s).

283 The algorithm captured by Figure 1.1 forms the basis for all level-3 implementations
 284 found in the BLIS framework (as of this writing). This algorithm is based on a
 285 so-called block-panel matrix multiplication.¹¹ The register (m_R, n_R) and cache ($m_C,$
 286 k_C, n_C) blocksizes labeled in the algorithmic diagram are typically chosen by the
 287 kernel developer as a function of hardware characteristics, such as the vector register
 288 set, cache sizes, and cache associativity. The authors of [15] present an analytical
 289 model for identifying suitable (if not optimal) values for these blocksizes.

290 **3. 1m method.** The primary motivation for seeking a better induced method
 291 comes from the observation that 4M inherently must update real and imaginary F.E.
 292 of C : (1) in separate steps, and may not use vector instructions to do so (due to the
 293 standard interleaved storage format); and (2) twice as frequently, in the case of 4M-1A,
 294 due to the algorithm's half-of-optimal cache blocksize k_C . As reviewed in Section 2.3,
 295 this imposes a significant drag on performance. If there existed an induced method
 296 that could update real and imaginary elements in one step, it may conveniently avoid
 297 both issues.

298 **3.1. Derivation.** Consider the classic definition of complex scalar multiplication
 299 and accumulation, shown in Eq. 2.1, refactored and expressed in terms of matrix and
 300 vector notation:

$$301 \quad (3.1) \quad \begin{pmatrix} \gamma^r \\ \gamma^i \end{pmatrix} += \begin{pmatrix} \alpha^r & -\alpha^i \\ \alpha^i & \alpha^r \end{pmatrix} \begin{pmatrix} \beta^r \\ \beta^i \end{pmatrix}$$

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

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

314 Thus, Eq. 3.1 serves as a packing template that hints at how the data must be
 315 stored. Furthermore, this template can be generalized. We augment α, β, γ with
 316 conventional row and column indices to denote the complex elements of matrices A ,
 317 B , and C , respectively. Also, let us apply the Eq. 3.1 to the special case of $m = 3$,
 318 $n = 4$, and $k = 2$ to better observe the general pattern.

$$319 \quad (3.2) \quad \begin{pmatrix} \gamma_{00}^r & \gamma_{01}^r & \gamma_{02}^r & \gamma_{03}^r \\ \gamma_{00}^i & \gamma_{01}^i & \gamma_{02}^i & \gamma_{03}^i \\ \gamma_{10}^r & \gamma_{11}^r & \gamma_{12}^r & \gamma_{13}^r \\ \gamma_{10}^i & \gamma_{11}^i & \gamma_{12}^i & \gamma_{13}^i \\ \gamma_{20}^r & \gamma_{21}^r & \gamma_{22}^r & \gamma_{23}^r \\ \gamma_{20}^i & \gamma_{21}^i & \gamma_{22}^i & \gamma_{23}^i \end{pmatrix} += \begin{pmatrix} \alpha_{00}^r & -\alpha_{00}^i & \alpha_{01}^r & -\alpha_{01}^i \\ \alpha_{00}^i & \alpha_{00}^r & \alpha_{01}^i & \alpha_{01}^r \\ \alpha_{10}^r & -\alpha_{10}^i & \alpha_{11}^r & -\alpha_{11}^i \\ \alpha_{10}^i & \alpha_{10}^r & \alpha_{11}^i & \alpha_{11}^r \\ \alpha_{20}^r & -\alpha_{20}^i & \alpha_{21}^r & -\alpha_{21}^i \\ \alpha_{20}^i & \alpha_{20}^r & \alpha_{21}^i & \alpha_{21}^r \end{pmatrix} \begin{pmatrix} \beta_{00}^r & \beta_{01}^r & \beta_{02}^r & \beta_{03}^r \\ \beta_{00}^i & \beta_{01}^i & \beta_{02}^i & \beta_{03}^i \\ \beta_{10}^r & \beta_{11}^r & \beta_{12}^r & \beta_{13}^r \\ \beta_{10}^i & \beta_{11}^i & \beta_{12}^i & \beta_{13}^i \end{pmatrix}$$

¹¹ 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].

321 From this, we can make the following observations:

- 322 • The complex matrix multiplication $C := C + AB$ with $m = 3$, $n = 4$, and
323 $k = 2$ becomes a real matrix multiplication with $m = 6$, $n = 4$, and $k = 4$.
324 In other words, the m and k dimensions are doubled for the purposes of the
325 real GEMM primitive.
- 326 • If the primitive is the real GEMM microkernel, and we assume that matrices
327 A and B above represent column-stored and row-stored micropans from
328 \tilde{A}_i and \tilde{B}_p , respectively, and also that the dimensions are conformal to the
329 register blocksizes of this microkernel (i.e., $m = m_R$ and $n = n_R$) then the
330 micropans of \tilde{A}_i are packed from a $\frac{1}{2}m_R \times \frac{1}{2}k_C$ submatrix of A , which, when
331 expanded in the special packing format, appears as the $m_R \times k_C$ micropans
332 that the real GEMM microkernel expects.
- 333 • Similarly, the micropans of \tilde{B}_p are packed from a $\frac{1}{2}k_C \times n_R$ submatrix of
334 B , which, when reordered into a second special packing format, appears as
335 the $k_C \times n_R$ micropans that the real GEMM microkernel expects.

336 It is easy to see by inspection that the real matrix multiplication implied by
337 Eq. 3.2 induces the desired complex matrix multiplication. We will refer to the packing
338 format used on matrix A above as the 1E format, since the F.E. are “expanded”
339 (i.e., duplicated to the next column, with the duplicates swapped and the imaginary
340 duplicate negated). Similarly, we will refer to the packing format used on matrix B
341 above as the 1R format, since the F.E. are merely reordered (i.e., imaginary elements
342 moved to the next row). Thus, the 1M method is fundamentally about reordering the
343 matrix data so that a subsequent real matrix multiplication on that reordered data is
344 equivalent to a complex matrix multiplication on the original data.¹²

345 **3.2. Two variants.** Notice that implicit in the 1M method suggested by Eq. 3.2
346 is the fact that matrix C is stored by columns. This assumption is important; when A
347 and B are packed according to the 1E and 1R formats, respectively, C must be stored
348 by columns in order to allow the real domain primitive (or microkernel) to correctly
349 update the individual real and imaginary F.E. of C with the corresponding F.E. from
350 the matrix product AB .

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

$$352 \quad (3.3) \quad (\gamma^r \ \gamma^i) += (\alpha^r \ \alpha^i) \begin{pmatrix} \beta^r & \beta^i \\ -\beta^i & \beta^r \end{pmatrix}$$

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

$$356 \quad (3.4) \quad \begin{pmatrix} \gamma_{00}^r \ \gamma_{00}^i \ \gamma_{01}^r \ \gamma_{01}^i \ \gamma_{02}^r \ \gamma_{02}^i \\ \gamma_{10}^r \ \gamma_{10}^i \ \gamma_{11}^r \ \gamma_{11}^i \ \gamma_{12}^r \ \gamma_{12}^i \\ \gamma_{20}^r \ \gamma_{20}^i \ \gamma_{21}^r \ \gamma_{21}^i \ \gamma_{22}^r \ \gamma_{22}^i \\ \gamma_{30}^r \ \gamma_{30}^i \ \gamma_{31}^r \ \gamma_{31}^i \ \gamma_{32}^r \ \gamma_{32}^i \end{pmatrix} += \begin{pmatrix} \alpha_{00}^r \ \alpha_{00}^i \ \alpha_{01}^r \ \alpha_{01}^i \\ \alpha_{10}^r \ \alpha_{10}^i \ \alpha_{11}^r \ \alpha_{11}^i \\ \alpha_{20}^r \ \alpha_{20}^i \ \alpha_{21}^r \ \alpha_{21}^i \\ \alpha_{30}^r \ \alpha_{30}^i \ \alpha_{31}^r \ \alpha_{31}^i \end{pmatrix} \begin{pmatrix} \beta_{00}^r \ \beta_{00}^i \ \beta_{01}^r \ \beta_{01}^i \ \beta_{02}^r \ \beta_{02}^i \\ -\beta_{00}^i \ \beta_{00}^r \ -\beta_{01}^i \ \beta_{01}^r \ -\beta_{02}^i \ \beta_{02}^r \\ \beta_{10}^r \ \beta_{10}^i \ \beta_{11}^r \ \beta_{11}^i \ \beta_{12}^r \ \beta_{12}^i \\ -\beta_{10}^i \ \beta_{10}^r \ -\beta_{11}^i \ \beta_{11}^r \ -\beta_{12}^i \ \beta_{12}^r \end{pmatrix}$$

358 In this variant, we see that matrix B , not A , is stored according to the 1E format
359 (where columns become rows), while matrix A is stored according to 1R (where rows

¹² 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 $n_R < n_P$.

360 become columns). Also, we can see that matrix C must be stored by rows in order to
361 allow the real GEMM microkernel to correctly update its F.E. with the corresponding
362 values from the matrix product AB .

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

366 **3.3. Determining complex blocksizes.** As we alluded in Section 3.1, the
367 appropriate blocksizes to use with 1M are a function of the real domain blocksizes.
368 This makes sense because the idea is to fool the real GEMM microkernel, and the
369 various loops for register and cache blocking around the microkernel, into thinking
370 that it is computing a real domain matrix multiplication. Which blocksizes must be
371 modified (halved) and which are used unchanged depends on the variant of 1M being
372 executed—or, more specifically, which matrix is packed according to the 1E format.

373 Table 3.1 summarizes the complex domain blocksizes prescribed for 1M_C and
374 1M_R as a function of the real domain values.

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

383 **3.4. Algorithms.**

384 **3.4.1. General algorithm.** Before investigating 1M method algorithms, we will
385 first provide algorithms for computing real matrix multiplication to serve as a reference
386 for the reader. Specifically, in Figure 3.1 we provide pseudo-code for RMMBP, which
387 depicts a real domain instance of the block-panel algorithm shown in Figure 1.1.

388 **3.4.2. 1m-specific algorithm.** Applying 1M_C and 1M_R to the block-panel
389 algorithm depicted in Figure 1.1 yields two nearly identical algorithms, 1M_C_BP and
390 1M_R_BP, respectively. Their differences can be encoded within a few conditional
391 statements within key parts of the high and low levels of code. Figure 3.2 shows a
392 hybrid algorithm that encompasses both, supporting row- and column-stored C .

Algorithm: $[C] := \text{RMMPB}(A, B, C)$
<pre> 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} $C_{jihl} := \text{RKERN}(\tilde{A}_{il}, \tilde{B}_{ph}, C_{jihl})$ </pre>

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. The algorithm is left-justified to facilitate comparison with Algorithms 1M_C_BP and 1M_R_BP in Figure 3.2 (left).

Algorithm: $[C] := 1M_?_BP(A, B, C)$	$[C] := \text{VK1M}(A, B, C)$
<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 $B_{jp} \rightarrow \tilde{B}_p$ else PACK1E $B_{jp} \rightarrow \tilde{B}_p$ for ($i = 0 : m - 1 : m_C$) Identify A_{pi}, C_{ji} from A_p, C_j if COLSTORE PACK1E $A_{pi} \rightarrow \tilde{A}_i$ else PACK1R $A_{pi} \rightarrow \tilde{A}_i$ 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} $C_{jihl} := \text{VK1M}(\tilde{A}_{il}, \tilde{B}_{ph}, C_{jihl})$ </pre>	<pre> Acquire workspace W Determine if using W; set USEW if (USEW) Alias $C_{\text{use}} \leftarrow W$, $C_{\text{in}} \leftarrow 0$ else Alias $C_{\text{use}} \leftarrow C$, $C_{\text{in}} \leftarrow C$ Set bool COLSTORE if RS(C_{use}) = 1 if (COLSTORE) CS(C_{use}) $\times= 2$ else RS(C_{use}) $\times= 2$ $N(A) \times= 2$; $M(B) \times= 2$ $C_{\text{use}} := \text{RKERN}(A, B, C_{\text{in}})$ 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.

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

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.

397 **3.5. Performance properties.** Table 3.2 tallies the total number of F.E. mem-
 398 ops required by 1M_C_BP and 1M_R_BP. For comparison, we also include the corre-
 399 sponding memop counts for a selection of 4M algorithms as well as a conventional
 400 assembly-based solution, as first published in Table III in [21].

401 Notice that 1M_C_BP and 1M_R_BP incur additional memops relative to a conven-
 402 tional assembly-based solution because, unlike the latter, 1M implementations cannot
 403 reuse¹³ all real and imaginary F.E. from vector registers.

404 We can hypothesize that the observed performance signatures of 1M_C_BP and
 405 1M_R_BP may be slightly different because each places the additional memop overhead
 406 that is unique to 1M on different parts of the computation. This stems from the fact
 407 that there exists an asymmetry in the assignment of packing formats to matrices in
 408 each 1M variant. Specifically, 50% more memops—relative to a conventional assem-
 409 bly solution—are required during the initial packing and the movement between caches
 410 for the matrix packed according to 1E since that format writes four F.E. for every
 411 two that it reads from the source operand. (Packing to 1R incurs the same number
 412 of memops as an assembly-based solution.) Also, if 1M_C_BP and 1M_R_BP use real
 413 microkernels with different micro-tile shapes (i.e., different values of m_R and n_R),
 414 those microkernels’ differing performance properties will likely cause the performance
 415 signatures of 1M_C_BP and 1M_R_BP to deviate further.

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

¹³ Here, the term “reuse” refers to the same reuse described in Footnote 9.

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.

418 **3.6. Algorithm details.** This section lays out important details that must be
419 handled when implementing the 1M method.

420 **3.6.1. Microkernel I/O preference.** Within the BLIS framework, microker-
421 nels are registered with a property that describes their input/output *preference*. The
422 I/O preference describes whether the microkernel is set up to ideally use vector in-
423 structions to load and store elements of the micro-tile by rows or by columns. This
424 property typically originates from the semantic orientation of vector registers used to
425 accumulate the $m_R \times n_R$ micropanel product. Whenever possible, the BLIS frame-
426 work will perform logical transpositions¹⁴ so that the apparent storage of C matches
427 the preference property of the microkernel being used. This guarantees that the mi-
428 crokernel will be able to load and store F.E. of C using vector instructions.

429 This preference property is merely an interesting performance detail for conven-
430 tional implementations (real and complex). However, in the case of 1M, it becomes
431 crucial for constructing a correctly-functioning implementation. More specifically, the
432 microkernel’s I/O preference determines whether the 1M_C or 1M_R algorithm is pre-
433 scribed. Generally speaking, a 1M_C algorithmic variant must employ a microkernel
434 that prefers to access C by columns, while a 1M_R algorithmic variant must use a
435 microkernel that prefers to access C by rows.

436 **3.6.2. Workspace.** In some cases, a small amount of $m_R \times n_R$ workspace is
437 needed. These cases fall into one of four scenarios: (1) C is row-stored and the real
438 microkernel RKERN has a column preference; (2) C is column-stored and RKERN has
439 a row preference; (3) C is general-stored (i.e., neither RS(C) nor CS(C) is unit); and
440 (4) $\beta^i \neq 0$. If any of these conditions hold, then the 1M virtual microkernel will need
441 to use workspace. This corresponds to the setting of USEW in VK1M (in Figure 3.2),

¹⁴ This amounts to swapping the row and column strides and swapping the m and n dimensions.

442 which causes RKERN to compute the micropanel product normally but store it to the
 443 workspace W . Subsequently, the result in W is then accumulated back to C .

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

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

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

456 **3.6.3. Handling alpha and beta scalars.** As in the previous article, we have
 457 simplified the general matrix multiplication to $C := C + AB$. In practice, the operation
 458 is implemented as $C := \beta C + \alpha AB$, where $\alpha, \beta \in \mathbb{C}$. Let us use Algorithm 1M_C_BP
 459 in Figure 3.2 to consider how to support arbitrary values of α and β .

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

464 When α is real, the scaling may be performed directly by RKERN. This situation
 465 is ideal since it usually incurs no additional costs.¹⁵ Scaling by α with non-zero
 466 imaginary components can still be performed by the packing function when either \tilde{A}_i
 467 or \tilde{B}_p are packed. Though somewhat less than ideal, the overhead incurred by this
 468 treatment of α is minimal in practice since packing is a memory-bound operation.

469 **3.6.4. Multithreading.** As with Algorithm 4M_1A in the previous article, Al-
 470 gorithms 1M_C_BP and 1M_R_BP parallelize in a straightforward manner for multicore
 471 and many-core environments. Because these algorithms encode the 1M method en-
 472 tirely within the packing functions and the virtual microkernel, all other levels of code
 473 are completely oblivious to, and therefore unaffected by, the specifics of the new al-
 474 gorithms. Therefore, we expect that 1M_C_BP and 1M_R_BP will yield multithreaded
 475 performance that is on-par with that of RMMBP.

476 **3.6.5. Bypassing the virtual microkernel.** Because the 1M virtual microker-
 477 nel serves as a function wrapper to the real domain microkernel, it incurs additional
 478 overhead. Thankfully, there exists a simple workaround, one that is viable as long as
 479 $\beta^i = 0$ and C is either row- or column-stored (but not general-stored). If these con-
 480 ditions are met, the real domain *macrokernel* can be called with modified parameters
 481 to induce the equivalent complex domain subproblem. This optimization allows the
 482 virtual microkernel (and its associated overhead) to be avoided entirely.

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

485 **3.7. Other complex storage formats.** The 1M method was developed specif-
 486 ically to facilitate performance on complex matrices stored using the standard storage
 487 format required by the BLAS. This interleaved storage convention for real and imag-

¹⁵ This is because many microkernels multiply their intermediate AB product by α unconditionally

TABLE 4.1

Register and cache blocksizes 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).

488 binary values is ubiquitous within the community and therefore implicitly assumed.
 489 However, some applications may be willing to tolerate API changes that would allow
 490 storing a complex matrix X as two separate real matrices X^r and X^i . For those ap-
 491 plications, the best an induced method may hope to do is implement each specialized
 492 complex matrix multiplication in terms of *two* real domain matrix multiplications—
 493 since there are two real matrices that must be updated. Indeed, there exists a variant
 494 of the 1M method, which we call the 2M method, that targets updating a matrix C
 495 that separates (entirely or by blocks) its real and imaginary F.E. [19].

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

500 **4.1. Platform and implementation details.** Results presented in this section
 501 were gathered on a single Cray XC40 compute node consisting of two 12-core Intel
 502 Xeon E5-2690 v3 processors featuring the “Haswell” microarchitecture. Each core,
 503 running at a clock rate of 3.2 GHz¹⁶, provides a single-core peak performance of 51.2
 504 gigaflops (GFLOPS) in double precision and 102.4 GFLOPS in single precision.¹⁷
 505 Each socket has a 30MB L3 cache that is shared among cores, and each core has a
 506 private 256KB L2 cache and 32KB L1 (data) cache. Performance experiments were
 507 gathered under the Cray Linux Environment 6 operating system running the Linux
 508 4.4.103 (x86_64) kernel. Source code was compiled by the GNU C compiler (gcc)
 509 version 7.3.0.¹⁸ The version of BLIS used in these tests was not officially released at

¹⁶ 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.

¹⁷ 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.

¹⁸ The following optimization flags were used during compilation of BLIS and its test drivers: -O3 -mavx2 -mfma -mfpmath=sse -march=haswell.

510 the time of this writing, and was adapted from version 0.6.0-11.¹⁹

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

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

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

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

522 **4.2. Sequential results.** Figure 4.1 reports performance results for various im-
 523 plementations of double- and single-precision complex matrix multiplication on a
 524 single core of the Haswell processor. For these results, all matrix dimensions were
 525 equal (e.g. $m = n = k$). Results for 1M_C_BP (which uses a column-preferring mi-
 526 crokernel) appears on the left of Figure 4.1 while those of 1M_R_BP (which uses a
 527 row-preferring microkernel) appears on the right.

528 Each graph in Figure 4.1 also contains three reference implementations: BLIS’s
 529 complex GEMM based on conventional assembly-coded kernels (e.g. “cgemm assem-
 530 bly”); BLIS’s real GEMM (e.g. “sgemm assembly”); and the 4M_1A implementation
 531 found in BLIS.²⁰ We configured all three of these reference codes to use column-
 532 preferential microkernels on the left and row-preferential microkernels on the right,
 533 as indicated by a “(c)” or “(r)” in the legends, in order to provide consistency with
 534 the 1M results.

535 As predicted in Section 3.5, we find that the performance signatures of the
 536 1M_C_BP and 1M_R_BP algorithms differ slightly. This was expected given that the
 537 1E and 1R packing formats place different memory access burdens on different packed
 538 matrices, \tilde{A}_i and \tilde{B}_p , which reside in different levels of cache. It was not previously
 539 clear, however, which would be superior over the other. It seems that, at least in
 540 the sequential case, the difference is somewhat more noticeable in double-precision,
 541 though even there it is quite subtle. This difference is almost certainly due to the
 542 individual performance characteristics of the underlying row- and column-preferential
 543 microkernels. We find evidence of this in the 4M_1A results, which was also affected
 544 by the change in microkernel I/O preference.

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

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

550 Finally, the curious reader may recall our brief hypothetical discussion of execut-
 551 ing Algorithm 4M_HW on a split complex storage format from Section 2.2 and wonder
 552 where such an implementation would fall relative to the measured performance data.
 553 Since Algorithm 4M_HW on a split format would mimic the execution of four unrelated

¹⁹ 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`.

²⁰ 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”).

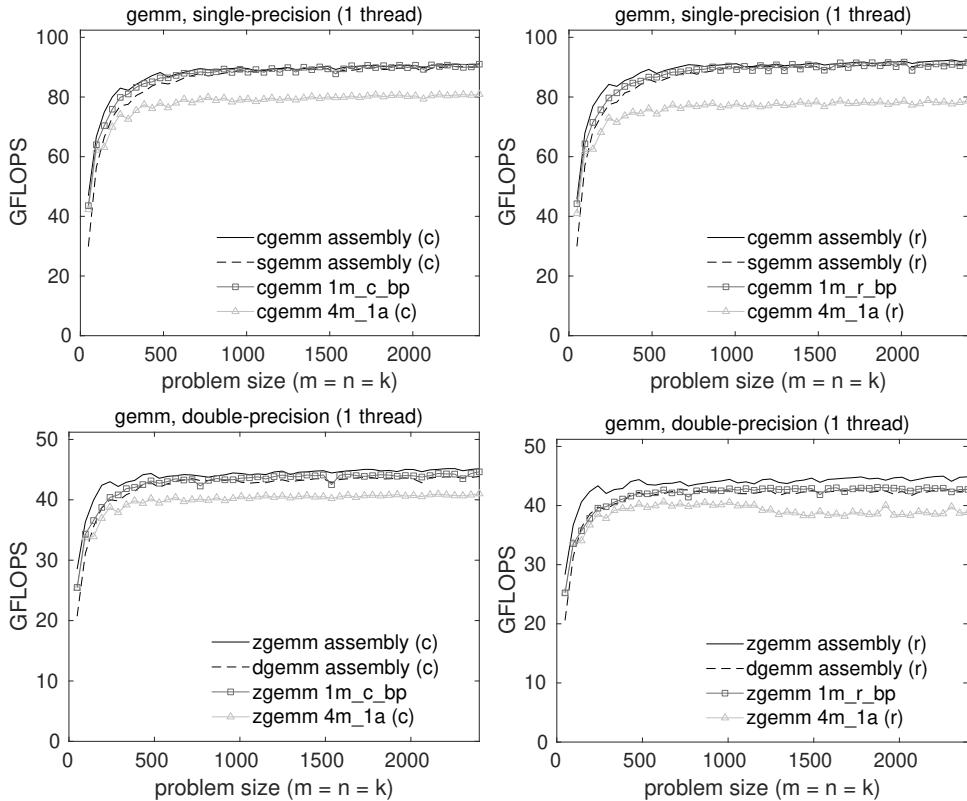


FIG. 4.1. Single-threaded performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core of an Intel Xeon E5-2690 v3 “Haswell” processor. The left and right graphs differ in which 1M implementation they report, with the left graphs reporting 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.

real matrix multiplications, its performance would track nearly identically with that of the real domain GEMM.

4.3. **Multithreaded results.** Figure 4.2 shows single- and double-precision performance using 24 threads, with one thread bound to each physical core of the processor. Performance is presented in units of gigaflops per core to facilitate visual assessment of scalability. For all BLIS implementations, we employed 4-way parallelism within the 5th loop, 3-way parallelism within the 3rd loop, and 2-way parallelism in the 2nd loop for a total of 24 threads. This parallelization scheme was chosen in a manner consistent with that of the previous article using a strategy set forth in [18].

563 Compared to the single-threaded case, we find a more noticeable difference in
 564 multithreaded performance between the 1M algorithms. Specifically, the 1M_R_BP im-
 565 plementation (based on a row-preferring microkernel) outperforms that of 1M_C_BP
 566 (based on a column-preferring microkernel), with the difference more pronounced

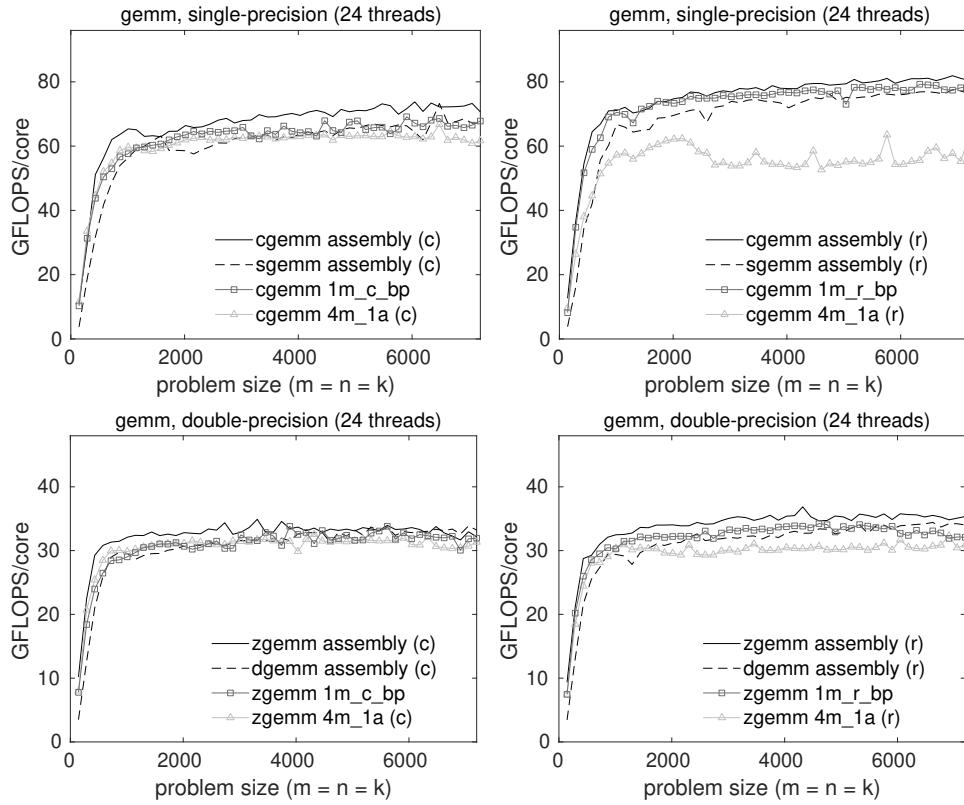


FIG. 4.2. Multithreaded performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on two Intel Xeon E5-2690 v3 “Haswell” processors, each with 12 cores. All data points reflect the use of 24 threads. The left and right graphs differ in which 1M implementation they report, with the left graphs reporting 1M_C_BP (which employs a column-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.

567 in single-precision. We suspect this is rooted not in the algorithms *per se* but in
 568 the differing microkernel implementations used by each 1M algorithm. The 1M_R_BP
 569 algorithm uses a real microkernel that is 6×16 and 6×8 in the single- and double-
 570 precision cases, respectively, while 1M_C_BP uses 16×6 and 8×6 microkernels for
 571 single- and double-precision implementations, respectively. The observed difference
 572 in performance between the 1M algorithms is likely attributable to the fact that the
 573 microkernels’ different values for m_R and n_R place different latency and bandwidth
 574 requirements when reading F.E. from the caches (primarily L1 and L2). More specif-
 575 ically, larger values of m_R place a heavier burden on loading elements from the L2
 576 cache, which is usually disadvantageous since that cache may exhibit higher latency
 577 and/or lower bandwidth. By contrast, a microkernel with larger n_R loads more ele-
 578 ments (per $m_R \times n_R$ rank-1 update) from the L1 cache, which resides closer to the

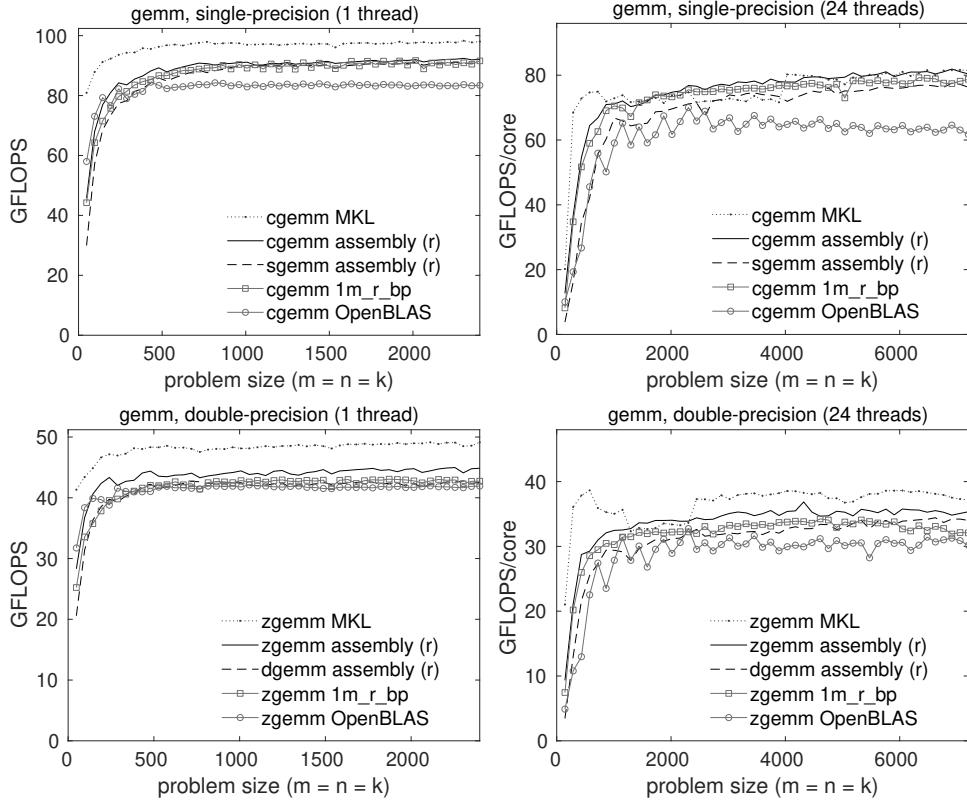


FIG. 4.3. Single-threaded (left) and multithreaded (right) performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core (left) or 12 cores (right) of an Intel Xeon E5-2690 v3 ‘Haswell’ processor. All multithreaded data points reflect the 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.

579 processor and offers lower latency and/or higher bandwidth than the L2 cache.

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

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

591 **4.4. Comparing to other implementations.** While our primary goal is not
592 to compare the performance of the newly developed 1M implementations with that
593 of other established BLAS solutions, some basic comparison is merited and thus we
594 have included Figure 4.3 (left). These graphs are similar to those in Figure 4.1,
595 except that: we show only implementations based on row-preferential microkernels;

596 we omit 4M_1A; and we include results for complex GEMM implementations provided
 597 by OpenBLAS 0.3.6 [16] and Intel MKL 2019 Update 4 [12].

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

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

603 **4.5. Additional results.** Additional performance results were gathered on a
 604 Marvell ThunderX2 compute server as well as an AMD EPYC (Zen) system. For
 605 brevity, we present and discuss that data in the appendix available online as sup-
 606 plementary materials. Those results reinforce the narrative provided here, lending
 607 even more evidence that the 1M method is capable of yielding high-performance im-
 608 plementations of complex matrix multiplication that are competitive with (and often
 609 outperform) other leading library solutions.

610 **5. Observations.**

611 **5.1. 4m limitations circumvented.** The previous article concluded by iden-
 612 tifying a number of limitations inherent in the 4M method. We now revisit this list
 613 and briefly discuss whether, to what degree, and how those limitations are overcome
 614 by algorithms based on the 1M method.

615 **Number of calls to primitive.** The most versatile 4M algorithm, 4M_1A, incurs
 616 up to a four-fold increase in function call overhead over a comparable assembly-based
 617 implementation. By comparison, 1M algorithms require at most a doubling of micro-
 618 kernel function call overhead, and in certain common cases (e.g., when $\beta \in \mathbb{R}$ and C
 619 is row- or column-stored), this overhead can be avoided completely. The 1M method
 620 is a clear improvement over 4M due to its one-to-one substitution of the matrix mul-
 621 tiplication primitive.

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

627 **Non-contiguous output to C .** Algorithms based on the 4M method must up-
 628 date only the real and then only the imaginary parts of the output matrix, twice
 629 each. When C is stored (by rows or columns) in the standard format, with real and
 630 imaginary F.E. interleaved, this piecemeal approach prevents the real microkernel
 631 from using vector load and store instructions on C during those four updates. The
 632 1M method avoids this issue altogether by packing A and B to formats that allow the
 633 real microkernel to update contiguous real and imaginary F.E. of C simultaneously.

634 **Reduction of k_C .** Algorithm 4M_1A requires that the real microkernel’s pre-
 635 ferred k_C blocksize be halved in the complex algorithm in order to maintain proper
 636 cache footprints of \tilde{A}_i and \tilde{B}_p as well the footprints of their constituent micropans.²¹
 637 Using these sub-optimally sized micropans can noticeably hobble the performance
 638 of 4M_1A. Looking back at Table 3.1, it may seem like 1M suffers a similar handicap;

²¹ Recall that the halving of k_C for 4M_1A was motivated by the desire to keep not just two, but four real micropans in the L1 cache simultaneously. These correspond to the real and imaginary parts of the current micropans of \tilde{A}_i and \tilde{B}_p .

639 however, the reason for halving k_C and its effect are both completely different. In the
 640 case of 1M, the use of $k_C^z = \frac{1}{2}k_C$ is simply a conversion of units (complex elements
 641 to real F.E.) for the purposes of identifying the size of the complex submatrices to
 642 be packed that will induce the optimal k_C value from the perspective of the real mi-
 643 crokernel, *not* a reduction in the F.E. footprint of the micropanels operated upon by
 644 that real microkernel. The ability of 1M to achieve high performance when $k = \frac{1}{2}k_C$
 645 is actually a strength for certain higher-level applications, such as Cholesky, LU, and
 646 QR factorizations based on rank- k update. Those operations tend to perform better
 647 when the algorithmic blocksize (corresponding to k_C) is as narrow as possible in order
 648 to limit the amount of computation in the lower-performing unblocked subproblem.

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

653 **Interference with multithreading.** Because the 1M algorithms are imple-
 654 mented entirely within the packing routines and virtual microkernel, they parallelize
 655 just as easily as the most thread-friendly of the 4M algorithms, 4M_1A, and entirely
 656 avoid the threading difficulties of higher-level 4M algorithms.²²

657 **Non-applicability to two-operand operations.** Certain higher-level appli-
 658 cations of 4M are inherently incompatible with two-operand operations because they
 659 would overwrite the original contents of the input/output operand even though subse-
 660 quent stages of computation depend on that original input. 1M avoids this limitation
 661 entirely. Like 4M_1A, 1M can easily be applied to two-operand level-3 operations such
 662 as TRMM and TRSM.²³

663 **5.2. Summary.** The analysis above suggests that the 1M method solves or
 664 avoids most of the performance-degrading weaknesses of 4M and in the remaining
 665 cases is no worse off than the best 4M algorithm.

666 **5.3. Limitations of 1m.** Although the 1M method avoids most of the weak-
 667 nesses inherent to the 4M method, a few notable caveats remain.

668 **Non-real values of beta.** In the most common cases where $\beta^i = 0$, the 1M
 669 implementation may employ the optimization described in Section 3.6.5. However,
 670 when $\beta^i \neq 0$, the virtual microkernel must be called. In such cases, 1M yields slightly
 671 lower performance due to extra memops.²⁴

672 **Algorithmic dependence on I/O preference.** If the real domain microkernel
 673 is row-preferential (and thus performs row-oriented I/O on C), then the 1M implemen-
 674 tation must choose an algorithm based on the 1M_R variant. But (in this scenario),
 675 if 1M_C is instead preferred for some reason, then either the underlying microkernel
 676 needs to be updated to handle both row- and column-oriented I/O, or a new column-
 677 preferential microkernel must be written. A similar caveat holds if the real domain
 678 microkernel is column-preferential and the 1M_R variant is preferred.

²² This thread-friendly property holds even when the virtual microkernel is bypassed altogether as discussed in Section 3.6.5

²³ 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.

²⁴ The 4M method suffers lower performance when $\beta^i \neq 0$ for similar reasons.

679 **Higher bandwidth on \tilde{A}_i and \tilde{B}_p .** Compared to a conventional, assembly-
 680 based GEMM, implementations based on the 1M method require twice as much mem-
 681 ory bandwidth when reading packed matrices \tilde{A}_i and \tilde{B}_p . Microkernels that encode
 682 complex arithmetic at the assembly level are able to load real and imaginary F.E.
 683 and then reuse those F.E. from registers, thus increasing the microkernel's arithmetic
 684 intensity. By contrast, the 1M method's reliance on real domain microkernels means
 685 that it must reuse real and imaginary F.E. from some level of cache and thus incur
 686 additional memory traffic.²⁵ The relative benefit of the conventional approach is likely
 687 to be most visible when parallelizing GEMM across all cores of a many-core system
 688 since that situation tends to saturate memory bandwidth.

689 **5.4. Further discussion.** Before concluding, we offer some final thoughts on
 690 the 1M method and its place in the larger spectrum of approaches to implementing
 691 complex matrix multiplication.

692 **5.4.1. Geometric interpretation.** Matrix multiplication is sometimes thought
 693 of as a three-dimensional operation with a contraction (accumulation) over the k di-
 694 mension. This interpretation carries into the complex domain as well. However, when
 695 each complex element is viewed in terms of its real and imaginary components, we
 696 find that a fourth pseudo-dimension of computation (of fixed size 2) emerges, one
 697 which also involves a contraction. The 1M method reorders and duplicates elements
 698 of A and B in such a way that exposes and “flattens” this extra dimension of com-
 699 putation. This, combined with the exposed treatment of real and imaginary F.E.,
 700 causes the resulting floating-point operations to appear indistinguishable from a real
 701 domain matrix multiplication with m and k dimensions (for column-stored C) or k
 702 and n dimensions (for row-stored C) that are twice as large.

703 **5.4.2. Data reuse: efficiency vs. programmability.** Both the conventional
 704 approach and 1M move data efficiently through the memory hierarchy.²⁶ However,
 705 once in registers, a conventional complex microkernel reuses those loaded values to
 706 perform twice as many flops as 1M. The previous article observes that all 4M algo-
 707 rithms make different variations of the same tradeoff: by forgoing the reuse of F.E.
 708 from registers and instead reusing those data from some level of cache, the algorithms
 709 avoid the need to explicitly encode complex arithmetic at the assembly level. As it
 710 turns out, 1M makes a similar tradeoff, but gives up less while gaining more: it is
 711 able to effectively reuse F.E. from two of the three matrix operands from registers
 712 while still avoiding the need for a complex microkernel, and it manages to replace
 713 that kernel operation with a single real matrix multiplication. And we would argue
 714 that increasing programmability and productivity by forfeiting a modest performance
 715 advantage is a good trade to make under almost any circumstance.

716 **5.4.3. Storage.** The supremacy of the 1M method is closely tied to the inter-
 717 leaved storage of real and imaginary values—specifically, of the output matrix C . If
 718 applications instead store complex matrices with their real and imaginary compo-
 719 nents split into two separate real matrices, the 4M approach (for numerically sensitive
 720 settings) as well as low-level applications of 3M (for numerically insensitive settings)
 721 may become more appropriate [19, 21].

²⁵ The 4M method suffers the same “bandwidth penalty” as 1M for the same reason.

²⁶ 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.

722 **6. Conclusions.** We began the article by reviewing the general motivations for
 723 induced methods for complex matrix multiplication as well as the specific methods,
 724 3M and 4M, studied in the previous article. Then, we recast complex scalar multipli-
 725 cation (and accumulation) in such a way that revealed a template that could be used
 726 to fashion a new induced method, one that casts complex matrix multiplication in
 727 terms of a single real matrix product. The key is the application of two new packing
 728 formats on the left- and right-hand matrix product operands that allows us to dis-
 729 guise the complex matrix multiplication as a real matrix multiplication with slightly
 730 modified input parameters. This 1M method is shown to have two variants, one each
 731 favoring row-stored and column-stored output matrices. When implemented in the
 732 BLIS framework, competitive performance was observed for 1M algorithms on three
 733 modern microarchitectures. Finally, we reviewed the limitations of the 4M method
 734 that are overcome by 1M and concluded by discussing a few high-level observations.

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

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

754

REFERENCES

755 [1] J. J. DONGARRA, J. DU CROZ, S. HAMMARLING, AND I. DUFF, *A set of level 3 basic linear*
 756 *algebra subprograms*, ACM Trans. Math. Soft., 16 (1990), pp. 1–17.

757 [2] G. FRISON, D. KOUZOUPI, T. SARTOR, A. ZANELLI, AND M. DIEHL, *BLASFEO: Basic lin-*
 758 *ear algebra subroutines for embedded optimization*, ACM Trans. Math. Soft., 44 (2018),
 759 pp. 42:1–42:30, <https://doi.org/10.1145/3210754>, <http://doi.acm.org/10.1145/3210754>.

760 [3] K. GOTO AND R. A. VAN DE GEIJN, *Anatomy of high-performance matrix multiplication*, ACM
 761 Trans. Math. Soft., 34 (2008), pp. 12:1–12:25, <https://doi.org/10.1145/1356052.1356053>,
 762 <http://doi.acm.org/10.1145/1356052.1356053>.

763 [4] K. GOTO AND R. A. VAN DE GEIJN, *High-performance implementation of the level-3 BLAS*,
 764 ACM Trans. Math. Soft., 35 (2008), pp. 4:1–4:14, <https://doi.org/10.1145/1377603>,
 765 <http://doi.acm.org/10.1145/1377603.1377607>.

766 [5] J. A. GUNNELS, G. M. HENRY, AND R. A. VAN DE GEIJN, *A family of high-performance matrix*
 767 *multiplication algorithms*, in Proceedings of the International Conference on Computational
 768 Sciences-Part I, ICCS ’01, Berlin, Heidelberg, 2001, Springer-Verlag, pp. 51–60, <http://dl.acm.org/citation.cfm?id=645455.653765>.

770 [6] N. J. HIGHAM, *Stability of a method for multiplying complex matrices with three real matrix*
 771 *multiplications*, SIAM J. Matrix Anal. App., 13 (1992), pp. 681–687, <https://doi.org/10.1137/0613043>,
 772 <https://doi.org/10.1137/0613043>.

773 [7] J. HUANG, *Practical fast matrix multiplication algorithms*, (2018). PhD thesis, The University
774 of Texas at Austin.

775 [8] J. HUANG, D. A. MATTHEWS, AND R. A. VAN DE GEIJN, *Strassen's algorithm for tensor contrac-*
776 *tion*, SIAM Journal on Scientific Computing, 40 (2018), pp. C305–C326, <https://doi.org/10.1137/17M1135578>, <https://doi.org/10.1137/17M1135578>, <https://arxiv.org/abs/https://doi.org/10.1137/17M1135578>.

777 [9] J. HUANG, L. RICE, D. A. MATTHEWS, AND R. A. VAN DE GEIJN, *Generating families of*
778 *practical fast matrix multiplication algorithms*, in 31th IEEE International Parallel and
779 Distributed Processing Symposium (IPDPS 2017), May 2017, pp. 656–667, <https://doi.org/10.1109/IPDPS.2017.56>.

780 [10] J. HUANG, T. M. SMITH, G. M. HENRY, AND R. A. VAN DE GEIJN, *Strassen's algorithm reloaded*,
781 in Proceedings of the International Conference for High Performance Computing, Network-
782 ing, Storage and Analysis, SC '16, Piscataway, NJ, USA, 2016, IEEE Press, pp. 59:1–59:12,
783 <http://dl.acm.org/citation.cfm?id=3014904.3014983>.

784 [11] J. HUANG, C. D. YU, AND R. A. VAN DE GEIJN, *Implementing Strassen's algorithm with CUT-*
785 *LASS on NVIDIA Volta GPUs*, FLAME Working Note #88, TR-18-08, The University
786 of Texas at Austin, Department of Computer Science, 2018, https://apps.cs.utexas.edu/apps/sites/default/files/tech_reports/GPUStrassen.pdf.

787 [12] INTEL, *Math Kernel Library*. <https://software.intel.com/en-us/mkl>, 2019.

788 [13] INTEL CORPORATION, *Intel® 64 and IA-32 Architectures Optimization Reference Manual*,
789 no. 248966-033, June 2016.

790 [14] INTEL CORPORATION, *Intel® Xeon® Processor E5 v3 Product Family: Processor Specification*
791 *Update*, no. 330785-010US, September 2016.

792 [15] T. M. LOW, F. D. IGUAL, T. M. SMITH, AND E. S. QUINTANA-ORTÍ, *Analytical modeling is*
793 *enough for high-performance BLIS*, ACM Trans. Math. Soft., 43 (2016), pp. 12:1–12:18,
794 <https://doi.org/10.1145/2925987>, <http://doi.acm.org/10.1145/2925987>.

795 [16] *OpenBLAS*. <http://xianyi.github.com/OpenBLAS/>, 2019.

796 [17] D. T. POPOVICI, F. FRANCHETTI, AND T. M. LOW, *Mixed data layout kernels for vector-*
797 *ized complex arithmetic*, in 2017 IEEE High Performance Extreme Computing Conference
798 (HPEC), Sep. 2017, pp. 1–7, <https://doi.org/10.1109/HPEC.2017.8091024>.

799 [18] T. M. SMITH, R. A. VAN DE GEIJN, M. SMELYANSKIY, J. R. HAMMOND, AND F. G. VAN
800 ZEE, *Anatomy of high-performance many-threaded matrix multiplication*, in Proceedings
801 of the 28th IEEE International Parallel & Distributed Processing Symposium (IPDPS),
802 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>.

803 [19] F. G. VAN ZEE, *Inducing complex matrix multiplication via the 1m method*, FLAME Working
804 Note #85 TR-17-03, The University of Texas at Austin, Department of Computer Sciences,
805 February 2017.

806 [20] F. G. VAN ZEE, T. SMITH, F. D. IGUAL, M. SMELYANSKIY, X. ZHANG, M. KISTLER, V. AUSTEL,
807 J. GUNNELS, T. M. LOW, B. MARKER, L. KILLOUGH, AND R. A. VAN DE GEIJN, *The BLIS*
808 *framework: Experiments in portability*, ACM Trans. Math. Soft., 42 (2016), pp. 12:1–12:19,
809 <http://doi.acm.org/10.1145/2755561>.

810 [21] F. G. VAN ZEE AND T. M. SMITH, *Implementing high-performance complex matrix multiplica-*
811 *tion via the 3m and 4m methods*, ACM Trans. Math. Soft., 44 (2017), pp. 7:1–7:36.

812 [22] F. G. VAN ZEE AND R. A. VAN DE GEIJN, *BLIS: A framework for rapidly instantiating BLAS*
813 *functionality*, ACM Trans. Math. Soft., 41 (2015), pp. 14:1–14:33, <http://doi.acm.org/10.1145/2764454>.

814 [23] R. C. WHALEY, A. PETITET, AND J. J. DONGARRA, *Automated empirical optimization of soft-*
815 *ware and the ATLAS project*, Parallel Computing, 27 (2001), pp. 3–35, [https://doi.org/10.1016/S0167-8191\(00\)00087-9](https://doi.org/10.1016/S0167-8191(00)00087-9), <http://www.sciencedirect.com/science/article/pii/S0167819100000879>. New Trends in High Performance Computing.

816 [24] C. D. YU, J. HUANG, W. AUSTIN, B. XIAO, AND G. BIROS, *Performance optimization for*
817 *the k-nearest neighbors kernel on x86 architectures*, in Proceedings of the International
818 Conference for High Performance Computing, Networking, Storage and Analysis, SC '15,
819 New York, NY, USA, 2015, ACM, pp. 7:1–7:12, <https://doi.org/10.1145/2807591.2807601>,
820 <http://doi.acm.org/10.1145/2807591.2807601>.

829 **Appendix A. Additional Performance Results.**

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

834 **A.1. Marvell ThunderX2.** In this section, we report the performance of the
 835 1M method on the Marvell ThunderX2, a high-performance ARMv8 microarchitec-
 836 ture.

837 **A.1.1. Platform and implementation details.** Results presented in this sec-
 838 tion were gathered on a single compute node consisting of two 28-core Marvell Thun-
 839 derX2 CN9975 processors.²⁷ Each core, running at a clock rate of 2.2 GHz, provides
 840 a single-core peak performance of 17.6 gflops (GFLOPS) in double precision and
 841 35.2 GFLOPS in single precision. Each socket has a 32MB L3 cache that is shared
 842 among cores, and each core has a private 256KB L2 cache and 32KB L1 (data) cache.
 843 Performance experiments were gathered under the Ubuntu 16.04 operating system
 844 running the Linux 4.15.0 kernel. Source code was compiled by the GNU C compiler
 845 (`gcc`) version 7.3.0.²⁸ The version of BLIS used in these tests was version 0.5.0-1.²⁹

846 In this section, we show 1M results for only Algorithm 1M_C_BP. Unlike the
 847 results shown in the main article, we did not develop conventional assembly-based
 848 microkernels and thus cannot compare against a complex domain solution based on
 849 those kernels. For further comparison, we measured performance for the complex
 850 GEMM implementations found in OpenBLAS³⁰ and ARMPL 18.4.0.

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

854 **A.1.2. Analysis.** Figure A.1 contains single-threaded (left) and multithreaded
 855 (right) performance of single-precision (top) and double-precision (bottom) complex
 856 GEMM implementations. In addition to the 1M_C_BP implementation within BLIS,
 857 we also show the corresponding real domain GEMM implementation and the `cgemm` or
 858 `zgemm` found in OpenBLAS and ARMPL. For all BLIS implementations, we employed
 859 4-way parallelism within the 5th loop and 14-way parallelism within the 3rd loop for
 860 a total of 56 threads.

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

²⁷ 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.

²⁸ The following optimization flags were used during compilation of BLIS and its test drivers: `-O3 -ftracer -fvectorize -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`.

²⁹ This version of BLIS may be uniquely identified, with high probability, by the first 10 digits of its `git` “commit” (SHA1 hash) number: e90e7f309b.

³⁰ This version of OpenBLAS may be uniquely identified, with high probability, by the first 10 digits of its `git` commit number: 52d3f7af50.

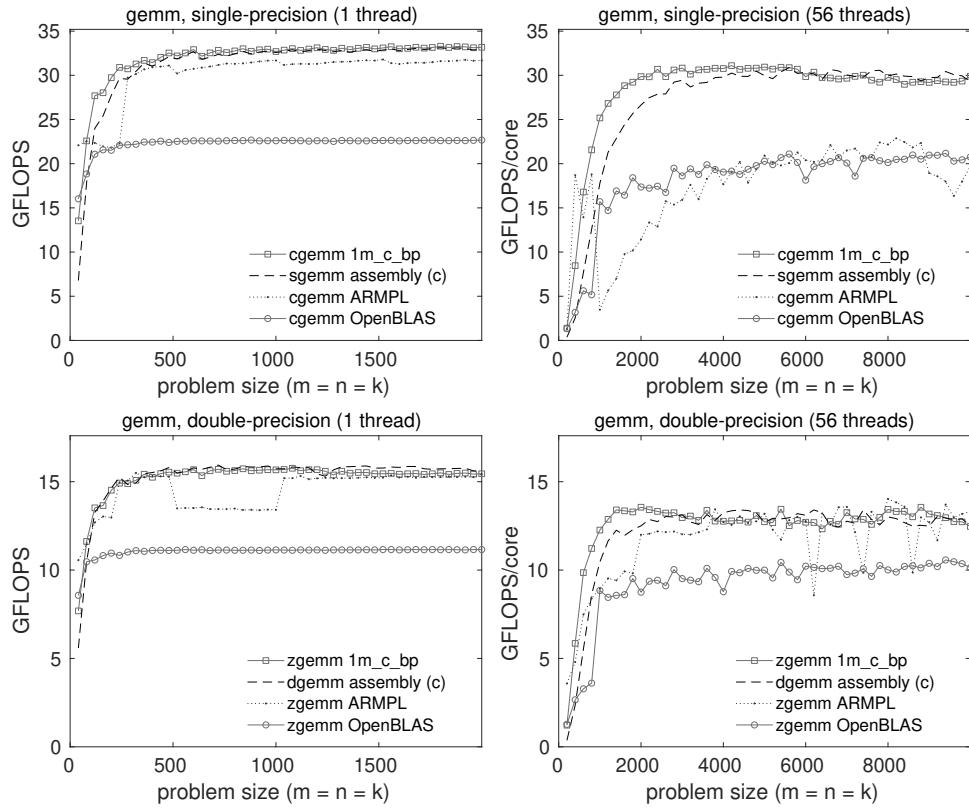


FIG. A.1. Single-threaded (left) and multithreaded (right) performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core (left) or 56 cores (right) of a Marvell ThunderX2 CN9975 processor. All multithreaded data points reflect the use of 56 threads. The real domain GEMM implementation from BLIS uses a column-preferential microkernel, as indicated by 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.

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

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

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

882 **A.2.1. Platform and implementation details.** Results presented in this section
 883 were gathered on a single compute node consisting of two 32-core AMD EPYC
 884 7551 (Zen) processors.³¹ Each core runs at a clock rate of 3.0 GHz when using a
 885 single core and 2.55 GHz when utilizing all cores simultaneously. The former clock
 886 rate yields a single-core peak performance of 24.0 GFLOPS in double precision and
 887 48.0 GFLOPS in single precision, and the latter clock rate yields a multicore peak
 888 performance of 20.4 GFLOPS/core and 40.8 GFLOPS/core for single- and double-
 889 precision computation, respectively. Each socket has a 64MB of L3 cache (distributed
 890 as 8MB for each four-core complex) that is shared among cores, and each core has a
 891 private 512KB L2 cache and 32KB L1 (data) cache. Performance experiments were
 892 gathered under the Ubuntu 18.04 operating system running the Linux 4.15.0 kernel.
 893 Source code was compiled by the GNU C compiler (`gcc`) version 7.4.0.³² The version
 894 of BLIS used in these tests was version 0.6.0-1266.³³

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

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

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

909 In Figure A.2 (top-left), all implementations track closely together except for
 910 MKL.³⁴ We see a similar pattern for single-threaded double precision in Figure A.2
 911 (bottom-left).

912 In Figure A.2 (top-right) and (bottom-right), we see multithreaded performance
 913 when utilizing all 64 cores of the AMD EPYC system. The relative performance of
 914 1M_R_BP is consistent with the results seen previously on Haswell. That is, the 1M
 915 method facilitates performance that meets or exceeds the performance of an optimized
 916 real domain implementation of GEMM (i.e., one that uses the same microkernels as
 917 1M), but falls slightly short of the performance of a conventional assembly-coded
 918 complex domain GEMM. Once again, MKL performance suffers noticeably on AMD
 919 hardware.³⁵ OpenBLAS lags somewhat behind the BLIS-based implementations, but

³¹ While two-way symmetric multithreading is available on this hardware, a maximum of one logical core per physical core was utilized during our tests.

³² 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`.

³³ This version of BLIS may be uniquely identified, with high probability, by the first 10 digits of its `git` “commit” (SHA1 hash) number: f391b3e2e7.

³⁴ 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.

³⁵ In order to keep the legends in Figure A.2 readable, the curves for MKL in were clipped beyond the first 20 data points. In all four graphs, the omitted data points depict a plateauing of the curve that is consistent with the data that is shown.

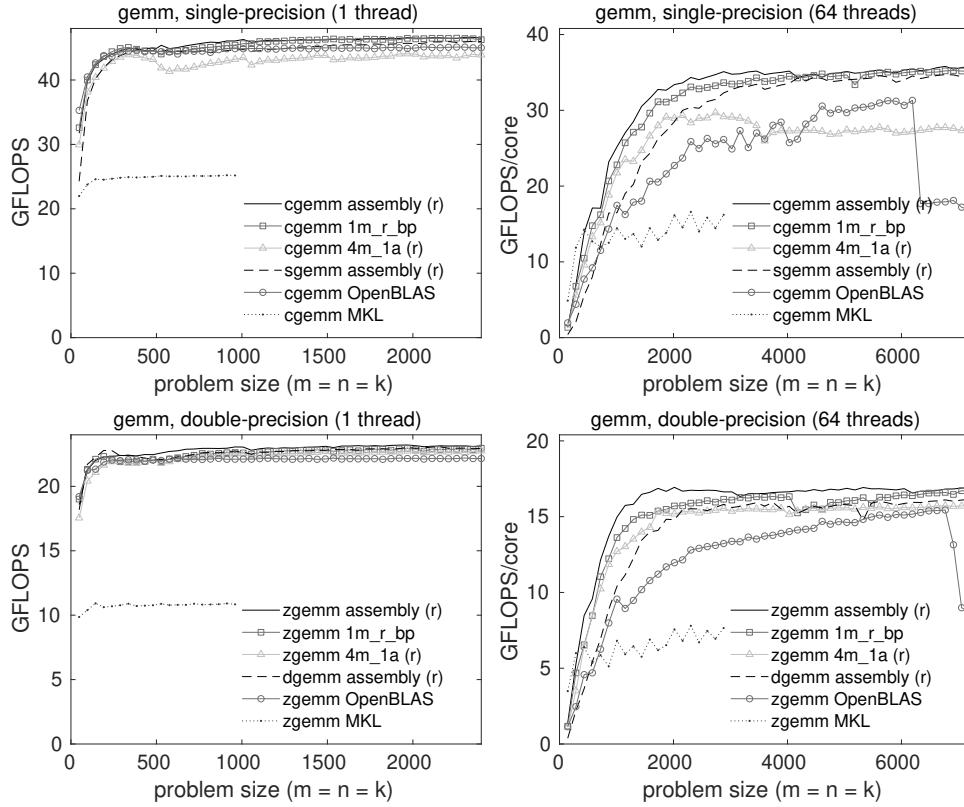


FIG. A.2. Single-threaded (left) and multithreaded (right) performance of various implementations of single-precision (top) and double-precision (bottom) complex GEMM on a single core (left) or 64 cores (right) of an AMD EPYC 7551 (Zen) processor. All multithreaded data points reflect the use of 64 threads. 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.

920 performance unexpectedly drops for very large problem sizes. This behavior was
 921 reproducible, though the exact problem size at which the drop-off occurred shifted
 922 across repeated experiments.

923 In summary, BLIS’s 1M implementation performs very well on the AMD EPYC
 924 7551 when computing in single and double precision, exceeding the performance of
 925 both OpenBLAS and MKL. Scalability (relative to theoretical peak) is also quite
 926 good in both precisions considering the challenges that NUMA-based architectures
 927 sometimes pose to parallelization efforts.