

Tile Tensors: A versatile data structure with descriptive shapes for homomorphic encryption

Ehud Aharoni*, Allon Adir, Moran Baruch, Gilad Ezov, Ariel Farkash,
Lev Greenberg, Ramy Masalha, Dov Murik and Omri Soceanu

IBM Research

*aehud@il.ibm.com

Abstract—Moving from the theoretical promise of Fully Homomorphic Encryption (FHE) to real-world applications with realistic and acceptable time and memory figures is an on-going challenge. After choosing an appropriate FHE scheme, and before implementing privacy-preserving analytics, one needs an efficient packing method that will optimize use of the ciphertext slots, trading-off size, latency, and throughput. We propose a solution to this challenge. We describe a method for efficiently working with tensors (multi-dimensional arrays) in a system that imposes tiles, i.e., fixed-size vectors. The tensors are packed into tiles and then manipulated via operations on those tiles. We further show a novel and concise notation for describing packing details.

Our method reinterprets the tiles as multi-dimensional arrays, and combines them to cover enough space to hold the tensor. An efficient summation algorithm can then sum over any dimension of this construct. We propose a descriptive notation for the shape of this data structure that describes both the original tensor and how it is packed inside the tiles. Our solution can be used to optimize the performance of various algorithms such as consecutive matrix multiplications or neural network inference with varying batch sizes. It can also serve to enhance optimizations done by homomorphic encryption compilers. We describe different applications that take advantage of this data structure through the proposed notation, experiment to evaluate the advantages through different applications, and share our conclusions.

I. INTRODUCTION

Fully Homomorphic Encryption (FHE) schemes allow computation over encrypted data. For example, let d be some data, and let $E(d)$ be the encryption of d using a fully homomorphic encryption scheme. Given $E(d)$ and the scheme’s public key, it is possible to compute $E(f(d))$ for any function f without knowing the decryption key, and without learning anything about d . Thus, FHE can be highly useful for many applications, such as outsourcing computations to an untrusted cloud server even for sensitive, private data. In this paper we’ll use the more general acronym HE (homomorphic encryption), referring to both FHE and to more limited schemes. This distinction will be mostly irrelevant to the methods we show.

In HE schemes, the atomic message that can be encrypted is usually a vector of a fixed size that is determined by the scheme’s configuration. For example, in the HE scheme of CKKS [1], an important configuration parameter is the degree n of the cyclotomic polynomial used in the algebraic constructs underlying the scheme. In CKKS, n is a power of 2 and each ciphertext encrypts a vector of $n/2$ complex

numbers. The length of this vector (i.e., the capacity of a ciphertext) is usually referred to as the number of slots. For example, for a typical value of $n = 16384$, each ciphertext has 8192 slots.

The computation on encrypted data is performed by primitive operations, most notably add and multiply. Both work in Single Instruction Multiple Data (SIMD) fashion on all slots. Rotate is another primitive operator that cyclically rotates the encrypted vector by a specified offset.

Efficiently utilizing these operators is a major challenge when performing computations under encryption. Consider for example matrix-vector multiplication. Suppose that each ciphertext has s slots, and we need to compute the multiplication of an s by s matrix M with a vector V of size s . One possible method is to encrypt each row of M in a separate ciphertext, encrypt V in a ciphertext, and multiply it by all the row ciphertexts (elementwise). Then, compute the inner sums of each of the s resulting ciphertexts by repeated rotations and additions. This results in s ciphertexts, each containing one entry of the matrix multiplication result vector.

Many other techniques for matrix-vector multiplication under these constraints are known in the literature [2]–[5], each with its advantages and disadvantages. Every technique has a different packing algorithm for arranging the data into fixed-size vectors in preparation for encryption. Each packing algorithm requires a different corresponding algorithm to perform the computation itself, respecting the particular pattern by which the data was packed. The results must be extracted from the fixed-sized vectors that this computation outputs—which themselves often have a packing arrangement that is very different from what was initially input.

The challenge gets even harder when handling higher dimensional data and performing more complicated computations. For example, in neural network training or inference, the inputs are typically three or four dimensional arrays, and the computation involves a long sequence of matrix multiplications, additions, convolutions, and various elementwise activation functions.

In this paper, we describe a novel, general purpose data structure we term tile tensors. A tile tensor allows users to store tensors of arbitrary shapes and sizes. The tile tensor automatically packs the tensor data into a collection of vectors of fixed size, as required in HE environments, using a wide

variety of configurable options. Tile tensors also offer a set of operators to manipulate the tensor in its packed form.

The essence of our approach is to reinterpret the fixed-size vectors as multi-dimensional arrays, and then to combine as many of them as needed to cover the space required to hold the given tensor. Hence, these fixed-size multi-dimensional arrays are used as tiles to cover a larger space. An efficient summation algorithm is then able to sum over any dimension of this construct, based on an improved rotate-and-sum variant.

Tile tensors, having a relatively small amount of code, can be easily configured to efficiently implement several widely used techniques, as well as new ones that offer better time and space efficiency. This is done by adjusting the tile shapes used. We demonstrate how the new techniques can achieve a three-fold time performance boost in latency over state-of-the-art neural network inference methods. They also offer better flexibility in adapting to the input batch size. We further show how tile tensors can be used to dramatically reduce the number of bootstrapping operations in a neural network training case study.

Tile tensors have an additional benefit: Similar to regular tensors that are described by their shape (i.e., a vector specifying the sizes of their dimensions), a tile tensor is described by a tile tensor shape. But tile tensor shapes describe both the shape of the tensor that is packed inside, as well as the packing details. We propose a notation to express tile tensor shapes. This notation is easy to read and write, and intuitive to understand. Thus, tile tensor shapes can be used as a language for human-to-human communication such as papers, human-to-machine communication such as configuration files and log messages, and machine-to-machine communication when different components negotiate an agreed upon format.

The rest of the paper is organized as follows: Section II describes related work. Section III describes the notation used in the paper. Section IV introduces the concept of tile tensors and how they are used, and Section V follows with formal definitions. Section VI discusses the summation algorithm underlying tile tensor operators. Section VII demonstrates the use of tile tensors for matrix multiplication, and Section VIII shows applications for neural networks. We summarize our conclusions in Section IX.

II. RELATED WORK

Privacy preserving machine learning has been a long-standing research topic, with results for decision trees [6], [7], linear regression [8], logistic regression [9], support vector machine (SVM) classification [10], and more. Initially most of these works based their core technology on multi-party computation (MPC). Then in 2009, Craig Gentry in his PhD thesis [11] introduced fully homomorphic encryption (FHE). He used a bootstrapping mechanism to reduce the accumulated noise, and thus turned existing leveled homomorphic encryption schemes, able to perform operations up to a bounded depth, into fully homomorphic schemes with unbounded computation depth. This triggered a surge in various additional HE schemes such as the BGV [12], GSW [13], BV [14], and TFHE [15],

which performed integer arithmetic calculations on homomorphically encrypted data, and the CKKS scheme [1] for calculations on real numbers (in fact, complex).

A core component of various machine learning techniques is a mathematical library that supports tensor-operations, such as vector and matrix multiplication. Therefore an efficient implementation of these operators for encrypted data is key to machine learning over FHE.

One method that attempts to speed up vector-matrix multiplications is diagonalization [5], [16]. A similar method was used by GAZELLE [2] for interactive inference over an encrypted network.

Another method described as "packing across the batch dimension", "packing the same dimension of multiple input samples in the same ciphertext", or "SIMD representation" can be shown to offer high throughput for non-interactive neural network inference [17], [18] and was also used for training [19].

Dathathri et al. [20] present a different approach for non-interactive inference, based on a concept termed CipherTensor. CipherTensors, like tile tensors, are aimed at abstracting away the packing details of a tensor. We believe tile tensors have several advantages over CipherTensors, and in Subsection VIII-B we compare the two in detail.

III. NOTATION

We use the term tensor as synonymous with multi-dimensional array; this is common in AI.

We denote tensors using capital letters. A tensor R has shape $[n_1, n_2, \dots, n_n]$, where n_i is the size of its i 'th dimension. For example, a 5×6 matrix M has shape $[5, 6]$. Throughout this paper we sometimes refer to tensor objects by their name and shape together, $M[5, 6]$, and sometimes just refer to them by name, M .

Individual elements of a tensor R are specified with round brackets, $R(j_1, j_2, \dots, j_n)$, where $0 \leq j_i < n_i$ are zero-based indices. Zero-based indices are more convenient since we sometimes perform mod operations on them. However, when specifying dimensions, we use one-based indices (e.g., j_1 is the size of the first dimension). This simplifies lists of dimension sizes. We reiterate this statement where needed to avoid confusion.

Tile tensors are compound objects that contain a tensor packed within tiles; these tiles are vectors of fixed sizes (see explanations and definitions in subsequent sections). As a naming convention, a tile tensor containing the tensor R is named T_R .

Tile tensors also have shapes, with more complicated syntax, e.g., $[\frac{*}{8}, \frac{6}{16}]$. As with regular tensors, we sometimes refer to tile tensor objects by their name and shape together, e.g., $T_R[\frac{*}{8}, \frac{6}{16}]$.

The fraction line that appears as part of a tile tensor shape has a special meaning, as defined in Section V. A fraction line that does not appear as part of a tile tensor shape will have the usual meaning of division.

The symbol $*$ plays two roles:

- As part of a tile tensor shape it has special meaning as defined in Section V.
- Outside a tile tensor shape it specifies elementwise multiplication between tensors or tile tensors, e.g., $M * R$ or $T_M * T_R$. We define elementwise operations in an extended way, as explained in Subsection IV-D.

Matrix multiplication is denoted without a multiplication symbol: the multiplication of matrix M and matrix R is denoted MR .

Tags are used to denote different objects. For example, R , R' , and R'' denote different tensors. The transpose of matrix M is denoted M^T .

IV. INTRODUCTION TO TILE TENSORS

Tile tensors are versatile data structures that can store an arbitrary tensor within a set of tiles in a variety of ways. Their versatility and ease-of-use offer many benefits when performing computations in an environment that enforces tiles, e.g., HE environments.

Similar to a tensor, a tile tensor has a shape but its shape contains more information. It describes both the shape of the regular tensor packed inside the tile tensor and the packing details. Tile tensor shapes make it easy to communicate packing details, for example between researchers or between developers and HE systems.

In this section, we introduce the tile tensor data structure using examples and diagrams. Formal definitions are provided in Section V.

A. Basics

Tile tensors are tensors made of tiles.

A tile is a one dimensional vector of numbers, which can be operated on in a Single Instruction Multiple Data (SIMD) fashion. For example, the result of adding two tiles is a tile containing the elementwise addition of the two vectors they contain. We assume an environment that works with tiles of a fixed length, and supplies SIMD operations such as addition, multiplication, and rotation.

This is exactly the case for HE cryptosystems. Although this paper is primarily oriented towards such systems, tile tensors can be used with any system that imposes tiles, such as GPUs. As is common in HE systems, we use the term 'slots' for the individual cells of a tile. So, a tile of length 8 has 8 slots, each of which can hold a number.

Suppose we need to perform calculations with a matrix $M[5, 6]$, using a system that has tiles of length 8. The manner in which the matrix elements are packed into tiles will affect the way we can manipulate it. One arrangement may be better suited for one type of calculation, whereas another type of calculation may favor a different arrangement. This is where tile tensors come in handy, as they can easily store our matrix in different arrangements.

Let's start with a simple example. Let $T_M[5, \frac{6}{8}]$ be a tile tensor that contains M . Its tile tensor shape $[5, \frac{6}{8}]$, which can also be written $[5, 6/8]$, is a generalization of a tensor shape with a richer syntax and semantics. It neatly describes both

the shape of the tensor packed within and the details of the packing arrangements.

The shape $[5, \frac{6}{8}]$ indicates that our matrix is stored with each row in its own separate tile, as shown in Figure 1a. The $6/8$ part indicates that the tile size is 8, while only 6 slots of each each tile are filled with data and the rest are zero. Since the tile size is specified in the second dimension, moving inside the tile means moving along the second dimension of M , whereas moving along M 's first dimension means moving between tiles.

The tile tensor T_M can be created using a *pack* operation that receives the tensor to be packed and the desired tile tensor shape: $T_M = \text{pack}(M, [5, \frac{6}{8}])$. The tile tensor shape is stored as part of T_M . Later, T_M can be *unpacked* to retrieve the original tensor: $M = \text{unpack}(T_M)$. We sometimes describe this relationship as M being stored in T_M or M being the original tensor of T_M , or we may simply say that T_M contains M .

Figure 1b depicts a different way to pack M into a tile tensor: $T'_M = \text{pack}(M, [\frac{5}{8}, 6])$. Here each column is placed in a separate tile.

Figure 1c shows yet another way to pack M : $T''_M = \text{pack}(M, [\frac{5}{2}, \frac{6}{4}])$. We explain first how it's packed and then how to read the tile tensor shape. Here the tiles are themselves matrices of shape $[2, 4]$, and we combined as many of them as needed to cover a space large enough to accommodate our matrix. As usual, unused slots are set to zero. Recall that we assumed our tiles are one dimensional vectors of length 8. However, when used inside tile tensors, they can be reinterpreted as tensors of a specified shape, which in our case is $[2, 4]$. Usually, we require this reinterpretation to have the same total number of slots, although we will mention some exceptions to this rule.

All this information can be read off the tile tensor shape $[\frac{5}{2}, \frac{6}{4}]$ as follows. The numbers above the fraction line define the original tensor's shape, $[5, 6]$. The numbers below the fraction line define the shape of the tile, $[2, 4]$. The order is important: the shape specifies that the tile size along the first dimension is 2, and 4 along the second dimension.

The 6 tiles in Figure 1c themselves form a two dimensional array of size 3×2 . We call this the external tensor. It is a tensor where each element is a tile. The external tensor's shape is easily deduced from the tile tensor shape as having $[\frac{5}{2}] = 3$ rows and $[\frac{6}{4}] = 2$ columns.

Generally, any tensor can be packed into a tile tensor with tiles of any size, interpreted as tensors of arbitrary shape. E.g., say we have tiles of size 1024, and we need to pack a tensor $A[50, 20, 255]$. We can pack A for example into the tile tensor $T_A[\frac{50}{16}, \frac{20}{2}, \frac{255}{32}]$. This means we interpret our tiles as cubes of shape $[16, 2, 32]$. These cubes are combined in an external tensor of shape $[4, 10, 8]$, creating a three-dimensional structure large enough to accommodate A .

Revisiting the first example, $T_M[5, \frac{6}{8}]$, it can also be written as $T_M[\frac{5}{1}, \frac{6}{8}]$. A tile dimension of size 1 may be omitted, hence the original shorter version. Viewed this way, even the original matrix $M[5, 6]$ can be viewed as a tile tensor $M[\frac{5}{1}, \frac{6}{1}]$, using

M(0,0)	M(0,1)	M(0,2)	M(0,3)	M(0,4)	M(0,5)	0	0
M(1,0)	M(1,1)	M(1,2)	M(1,3)	M(1,4)	M(1,5)	0	0
M(2,0)	M(2,1)	M(2,2)	M(2,3)	M(2,4)	M(2,5)	0	0
M(3,0)	M(3,1)	M(3,2)	M(3,3)	M(3,4)	M(3,5)	0	0
M(4,0)	M(4,1)	M(4,2)	M(4,3)	M(4,4)	M(4,5)	0	0

(a) $M[5, 6]$ packed inside $T_M[5, 6/8]$

M(0,0)	M(0,1)	M(0,2)	M(0,3)	M(0,4)	M(0,5)
M(1,0)	M(1,1)	M(1,2)	M(1,3)	M(1,4)	M(1,5)
M(2,0)	M(2,1)	M(2,2)	M(2,3)	M(2,4)	M(2,5)
M(3,0)	M(3,1)	M(3,2)	M(3,3)	M(3,4)	M(3,5)
M(4,0)	M(4,1)	M(4,2)	M(4,3)	M(4,4)	M(4,5)
0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0

(b) $M[5, 6]$ packed inside $T'_M[5/8, 6]$

M(0,0)	M(0,1)	M(0,2)	M(0,3)	M(0,4)	M(0,5)	0	0
M(1,0)	M(1,1)	M(1,2)	M(1,3)	M(1,4)	M(1,5)	0	0
M(2,0)	M(2,1)	M(2,2)	M(2,3)	M(2,4)	M(2,5)	0	0
M(3,0)	M(3,1)	M(3,2)	M(3,3)	M(3,4)	M(3,5)	0	0
M(4,0)	M(4,1)	M(4,2)	M(4,3)	M(4,4)	M(4,5)	0	0
0	0	0	0	0	0	0	0

(c) $M[5, 6]$ packed inside $T''_M[5/2, 6/4]$

Fig. 1: $M[5, 6]$ packed into three different tile tensors with different tile tensor shapes. The rectangles represent the tiles. For each tile tensor we show how M 's elements are placed inside the tiles.

tiles of size $[1, 1]$. Consequently, we regard tile tensors as a generalization of tensors.

B. Replication

For our next example we show various ways to pack a one dimensional vector $V[5]$. The low dimensionality will simplify the diagrams, but everything is applicable to tensors of arbitrary shape.

An obvious way to pack V using tiles of size 8 is $T_V = \text{pack}(V, [\frac{5}{8}])$; this takes a single tile and places the whole vector inside. T_V is shown in Figure 2a.

Another alternative is to first reshape $V[5]$ into a matrix $V[5, 1]$, and then pack it with 2-dimensional tiles. Choosing again tiles of shape $[2, 4]$, we can pack V as $T'_V = \text{pack}(V, [\frac{5}{2}, \frac{1}{4}])$, as shown in Figure 2b.

V(0)	V(1)	V(2)	V(3)	V(4)	0	0	0
------	------	------	------	------	---	---	---

(a) $V[5]$ packed inside $T_V[5/8]$

V(0,0)	0	0	0
V(1,0)	0	0	0
V(2,0)	0	0	0
V(3,0)	0	0	0
V(4,0)	0	0	0
0	0	0	0

(b) $V[5, 1]$ packed inside $T'_V[5/2, 1/4]$

V(0,0)	V(0,0)	V(0,0)	V(0,0)
V(1,0)	V(1,0)	V(1,0)	V(1,0)
V(2,0)	V(2,0)	V(2,0)	V(2,0)
V(3,0)	V(3,0)	V(3,0)	V(3,0)
V(4,0)	V(4,0)	V(4,0)	V(4,0)
0	0	0	0

(c) $V[5, 1]$ packed inside $T''_V[5/2, *3/4]$

V(0,0)	V(0,0)	V(0,0)	0
V(1,0)	V(1,0)	V(1,0)	0
V(2,0)	V(2,0)	V(2,0)	0
V(3,0)	V(3,0)	V(3,0)	0
V(4,0)	V(4,0)	V(4,0)	0
0	0	0	0

(d) $V[5, 1]$ packed inside $T'''_V[5/2, *3/4]$

Fig. 2: $V[5]$ packed into different tile tensors. The rectangles represents the tiles. For each tile tensor, we show how V 's elements are placed inside the tiles.

Yet another way to pack V uses a new bit of notation: $T''_V = \text{pack}(V, [\frac{5}{2}, \frac{*}{4}])$. The $*$ indicates that this dimension in the original tensor V is 1. However, instead of occupying just one slot in the tile, it is replicated in all 4 slots, as shown in Figure 2c. Packing with replication is useful, as explained in Section VII.

We can further denote in the tile tensor shape that we only partially replicate the values along the tile dimension. For example, the tile tensor $T'''_V = \text{pack}(V, [\frac{5}{2}, \frac{*3}{8}])$ has the values replicated only in 3 slots out of the available 4 along this dimension. This is illustrated in Figure 2d.

C. Unknown Values

When tensors are packed into tile tensors, unused slots are filled with zeroes, as shown in Figures 1 and 2.

However, after tile tensors are manipulated, the unused slots might get filled with arbitrary values, as explained in the next

subsection. Although these unused slots are ignored when the tile tensor is unpacked, the presence of arbitrary values in them can still impact additional manipulation. To reflect this state, the tile tensor shape syntax has an additional symbol $?$, used as described below.

Returning to the example $V[5]$, let's say it is stored in the tile tensor $T_V[\frac{5}{2}, \frac{1?}{4}]$. The shape $[\frac{5}{2}, \frac{1?}{4}]$ is similar to the shape $[\frac{5}{2}, \frac{1}{4}]$ that was depicted in 2b, except the $\frac{1?}{4}$ part indicates that here the unused slots along the second tile dimension are filled with unknown values, as shown in Figure 3. As usual, unused slots are ignored when unpacking the tile tensor, thus $V = \text{unpack}(T_V)$.

D. Operators

Tile tensors are meant to be manipulated via a set of operators similar to the ones used to manipulate tensors. In fact, there is a homomorphism between operations on the tile tensors and the corresponding operations on the tensors they contain.

Consider an elementwise operator such as addition. The tile tensors T_A and T_B can be added together if they have equal shapes. (We'll soon relax this requirement.) The homomorphism property means that the result is a tile tensor containing the addition of the two original tensors, such that $\text{unpack}(T_A + T_B) = \text{unpack}(T_A) + \text{unpack}(T_B)$.

Elementwise operators are implemented by executing the operators on corresponding tiles. Thus, tile tensors support the same set of elementwise operators that the tiles support. In HE environments, these include addition and multiplication, and some unary operators such as negate and conjugate. These operators work with yet another homomorphism between the tile ciphertexts and the plaintexts they represent.

Applying binary elementwise operators on two tensors usually requires strict equality between their shapes. It will be useful for us to relax this requirement a little. Say, for example, we compute $M[5, 4] + V[1, 4]$. Strictly speaking, the shapes mismatch. However, whenever a dimension's size of one tensor is the degenerate size 1, we'll allow the other tensor's size along this dimension to be of any size, and interpret the operation as if the smaller tensor is replicated. More specifically, in $M[5, 4] + V[1, 4]$ the vector V will be added to each row of M , as shown in Figure 4.

V(0,0)	?	?	?
V(1,0)	?	?	?
V(2,0)	?	?	?
V(3,0)	?	?	?
V(4,0)	?	?	?
0	?	?	?

Fig. 3: $V[5, 1]$ stored in $T_M[5/2, 1?/4]$. Unused space along the second dimension has unknown values, marked as cells containing $?$.

Similarly, tile tensors can be added together if their shapes are compatible, but not necessarily identical. To be compatible, the two tile tensor shapes must have the same number of dimensions, and for each i the i 'th dimension specification in each shape must be compatible. Two dimension specifications are compatible if the tile sizes (below the fraction line) are equal, and the sizes of the original tensors along this dimension (above the fraction line) are either equal, or one of them is fully replicated. For example, $T_M[\frac{18}{8}, \frac{4}{16}]$ can be added with $T_V[\frac{*}{8}, \frac{4}{16}]$.

Another operator commonly used with tensors is summation. The operator $\text{sum}(A, i)$ sums over the i 's dimension of A . E.g., for $M[5, 4]$, summing all the rows together is $R = \text{sum}(M, 1)$, resulting in $R[1, 4]$. Recall that for dimension indices we use one-based indices; hence, $\text{sum}(M, 1)$ indicates summing over the first dimension. After summation, its size is reduced to 1.

Tile tensors can also be efficiently summed. Summing over a tile tensor is defined to be homomorphic to summing the tensor it contains, thus $\text{unpack}(\text{sum}(T_A, i)) = \text{sum}(\text{unpack}(T_A), i)$. For regular tensors, the operator $\text{sum}(A, i)$ reduces the i 'th dimension size to 1. For tile tensors, the operator $\text{sum}(T_A, i)$ has a similar effect, which follows simple rules we refer to as the summation rules.

We assume T_A 's shape along the i 'th dimension is specified as $\frac{n_i}{t_i}$. Summation changes the shape according to the following rules:

- If $t_i = 1$, then the resulting shape along the i 'th dimension is $\frac{1}{1}$, or simply 1.
- If i is the lowest non-trivial tile dimension (i.e., the smallest i such that $t_i > 1$), the resulting shape along the i 'th dimension is $\frac{*}{t_i}$.
- Otherwise, the resulting shape along the i 'th dimension is $\frac{1?}{t_i}$.

If before summation the dimension was $\frac{n_i}{t_i}$, then after summation it will always be $\frac{1?}{t_i}$.

The reason for these rules lies in the way tile tensors are summed, as described in Subsection VI-B and in Appendix A.

As an example, let T_A be a tile tensor with the shape $[4, \frac{3}{8}, \frac{5}{16}]$. Table I depicts the resulting shape after summing over each of the three dimensions.

M(0,0)	M(0,1)	M(0,2)	M(0,3)
M(1,0)	M(1,1)	M(1,2)	M(1,3)
M(2,0)	M(2,1)	M(2,2)	M(2,3)
M(3,0)	M(3,1)	M(3,2)	M(3,3)
M(4,0)	M(4,1)	M(4,2)	M(4,3)

+

V(0,0)	V(0,1)	V(0,2)	V(0,3)
V(0,0)	V(0,1)	V(0,2)	V(0,3)
V(0,0)	V(0,1)	V(0,2)	V(0,3)
V(0,0)	V(0,1)	V(0,2)	V(0,3)
V(0,0)	V(0,1)	V(0,2)	V(0,3)

Fig. 4: Interpretation of the elementwise operation: $M[5, 4] + V[1, 4]$ operation. V is first replicated to reach the appropriate shape $[5, 4]$ and then the addition is performed. This effectively adds V to every row of M .

TABLE I: Summation rules example. The resulting shape after summing over any of the three dimensions of $T_A[4, \frac{3}{8}, \frac{5}{16}]$.

Operator	Resulting shape
$sum(T_A, 1)$	$[1, \frac{3}{8}, \frac{5}{16}]$
$sum(T_A, 2)$	$[4, \frac{*}{8}, \frac{5}{16}]$
$sum(T_A, 3)$	$[4, \frac{3}{8}, \frac{17}{16}]$

E. Benefits of Tile Tensor Shape Notation

We carefully chose the syntax of tile tensor shapes to be intuitive and readable. After getting used to it, one can work with tile tensors without the aid of diagrams and without trying to figure out how the tensor data is laid out inside the tiles. Instead, both the elementwise operators and the summation operator have syntactic rules that allow putting tile tensors to use with complicated tasks. In Subsection VII-A, for example, we demonstrate this point for matrix multiplication.

Tile tensor shapes can also be used as a language for communicating information about packing details. These shapes can be used by researchers when describing their packing methods in papers and reports. For example, a known packing method sometimes referred to as stacked vector [4], and sometimes as input packing [2], can be concisely and precisely described with a tile tensor shape, replacing lengthy descriptions and additional terminology. The same as true for other known packing methods as well (see Subsection VII-B).

In addition, tile tensor shapes can be used for human-machine communication between developers and HE platforms. The tile tensor shape syntax was chosen to contain only ordinary symbols, easily typed in standard keyboards and easily displayed on simple consoles. This allows developers to specify the desired shapes to the system and then track how they change in log file printouts. If an error occurs, it allows the system to produce a clear error message explaining what went wrong. Thus, tile tensor shapes can play the same role that regular tensor shapes have in systems such as TensorFlow [21].

Tile tensor shapes can also be used for machine-machine communication. For example, HE components can negotiate agreed-upon tile tensor shapes when communicating among themselves.

Finally, tile tensor shapes can be used by HE compilers such as [18], [20]. These compilers transform a computation specification in some high-level format to a detailed HE execution plan and optimize various parameter choices. The Chet compiler [20] already uses a data structure similar to tile tensors, called CipherTensor. As explained in Subsection VIII-B, we believe tile tensors are more powerful, and their shapes offer an easy way for compilers to scan a wide variety of layout options.

V. TILE TENSORS DEFINITION

In this section we formally define the concept of tile tensors.

We start in the first subsection with a simple definition, without replication or unknown values. Subsequent subsections extend it with these two additional features.

Formal definitions of how operators act on tile tensors are given in Appendix A.

A. Basic Definition

A tile tensor $T_A[\frac{n_1}{t_1}, \frac{n_2}{t_2}, \dots, \frac{n_k}{t_k}]$ is a data structure that holds the values of a tensor of shape $[n_1, n_2, \dots, n_k]$ packed into tiles.

Each tile can hold a vector of length $s = \prod_i t_i$. We refer to its cells as *slots*. We usually assume each tile has exactly s slots. A tile tensor can also be constructed with a relaxed assumption using tiles with more than s slots, offering reduced functionality as explained in Subsection VI-B.

The tiles are arranged in a multi-dimensional array we refer to as the external tensor, denoted by $et(T_A)$. The shape of $et(T_A)$ is $[e_1, e_2, \dots, e_k]$, where $e_i = \lceil \frac{n_i}{t_i} \rceil$. Each element in this tensor is a tile. We denote by $et(T_A)(l_1, l_2, \dots, l_k)$ the tile at indices (l_1, l_2, \dots, l_k) , and $et(T_A)(l_1, l_2, \dots, l_k)(h)$ is the slot at index h within this tile.

While being flat, we interpret a tile as a tensor of shape $[t_1, t_2, \dots, t_k]$ by mapping the tensor elements to the flat vector in a row-major order. The tensor indices of (m_1, m_2, \dots, m_k) will be mapped to the vector element $h = \sum_i m_i \prod_{x=i+1}^k t_x$. The inverse mapping is: $m_i = \lfloor h / \prod_{x=i+1}^k t_x \rfloor \bmod t_i$.

A specific slot in the tile tensor can be identified by the tile indices (l_1, l_2, \dots, l_k) , and its index h within the tile. We further associate each slot with a set of indices (j_1, j_2, \dots, j_k) ; we call these its *logical indices*. We reinterpret each tile as a tensor of shape $[t_1, t_2, \dots, t_k]$ and weld them together in a grid of size $[e_1, e_2, \dots, e_k]$, forming one big tensor of shape $[t_1 e_1, t_2 e_2, \dots, t_k e_k]$. The indices of a slot in this tensor are its logical indices. Formally, the slot at index h of tile (l_1, l_2, \dots, l_k) has logical indices (j_1, j_2, \dots, j_k) defined as follows:

$$j_i = l_i t_i + \lfloor h / \prod_{x=i+1}^k t_x \rfloor \bmod t_i \quad (1)$$

Let A be a tensor of shape $[n_1, n_2, \dots, n_k]$. We say that T_A contains A if the elements of A are mapped to slots according to their logical indices. Formally, if a slot has logical indices (j_1, j_2, \dots, j_k) such that $\forall_i j_i < n_i$, then this slot contains the value $A(j_1, j_2, \dots, j_k)$. Otherwise, it is an unused slot and it contains 0.

For brevity, whenever $t_i = 1$, it can be omitted from the shape definition. For example, $[\frac{2}{1}, \frac{5}{8}]$ can be written $[2, \frac{5}{8}]$.

The operator $pack(A, [\frac{n_1}{t_1}, \frac{n_2}{t_2}, \dots, \frac{n_k}{t_k}])$ is used to denote the construction of a tile tensor according to the definition above, and $unpack(T_A)$ the reconstruction of the original tensor A from the tile tensor.

In a tile tensor shape notation, the fraction line can be equivalently replaced with a slash $/$.

B. Replication

The symbol $*$ within a tile tensor shape denotes replication as follows.

A tile tensor $[\frac{n_1*d_1}{t_1}, \frac{n_1*d_2}{t_2}, \dots, \frac{n_k*d_k}{t_k}]$, such that $\forall_i (d_i = 1 \vee n_i = 1) \wedge d_i \leq t_i$, is a data structure holding the values of a tensor of shape $[n_1, n_2, \dots, n_k]$ packed into tiles.

The definition is the same as before, except for the following.

The shape of the external tensor is now $[e_1, e_2, \dots, e_k]$, where $e_i = \lceil \frac{n_i d_i}{t_i} \rceil$.

Let A be a tensor of shape $[n_1, n_2, \dots, n_k]$. We say that T_A contains A if the following holds: If a slot has logical indices (j_1, j_2, \dots, j_k) such that $\forall_i j_i < n_i d_i$, then this slot contains the value $A(j_1 \bmod n_1, j_2 \bmod n_2, \dots, j_k \bmod n_k)$. Otherwise, it is an unused slot and it contains 0.

Note that we required $\forall_i (d_i = 1 \vee n_i = 1)$. If $d_i = 1$ then the i 'th dimension behaves just as in the previous subsection. If $d_i > 1$ and $n_i = 1$ then the operator $\bmod n_i$ simply reduces the index to 0, causing this value to be replicated d_i times. The replication could also work for $n_i > 1$, but this can also be achieved by adding another dimension and replicating along it. Thus, for simplicity, we require that any replicated dimension has $n_i = 1$.

We further required $\forall_i d_i \leq t_i$. Replicating values more times than the tile size can usually be replaced with logical replication, hence we restricted physical replication to tile size. Replicating values within a single tile can be useful, as will be evident when we discuss matrix multiplication in Section VII.

If a dimension is replicated, $d_i > 1$ and $n_i = 1$, then for brevity, n_i can be omitted in the shape definition e.g., $[\frac{1*3}{4}, \frac{5}{8}]$ can be written $[\frac{*3}{4}, \frac{5}{8}]$. If we replicate across the entire tile size, $d_i = t_i$, then d_i can be omitted as well e.g., $[\frac{1*4}{4}, \frac{5}{8}]$ can be written simply $[\frac{*}{4}, \frac{5}{8}]$. In the degenerate case where $d_i = 1$, both the value of d_i and the $*$ symbol can be omitted altogether, returning to the basic definition from the previous subsection.

C. Unknown Values in Unused Slots

We now further extend the tile tensor shape with an additional annotation of $?$ in dimension i , to indicate that the unused space along this dimension is not necessarily filled with zeroes.

A tile tensor $T_A[\frac{n_1*d_1 I_1(?) }{t_1}, \frac{n_1*d_2 I_2(?) }{t_2}, \dots, \frac{n_k*d_k I_k(?) }{t_k}]$, such that $I_i(?)$ is either $?$ or empty, is a data structure holding the values of a tensor of shape $[n_1, n_2, \dots, n_k]$ packed into tiles.

The definition is the same as before, except for the following. Say a slot has logical indices (j_1, j_2, \dots, j_k) . As before, it is an unused slot if $\exists_i j_i \geq n_i d_i$. If $\exists_i j_i \geq n_i d_i \wedge I_i(?) = ?$, then it contains an unknown value; otherwise, it contains 0 as usual. I.e., it contains an unknown value if its logical indices exceed the size of the used range along a dimension annotated with $?$.

```

10  sum_tile_a(L, n) :
20      e=1
30      S=L
40      for j=numBits(n)-2 downto 0
50          S=S+rot(S, e)
60          e=e*2
70          if (bit(n, j)==1) then
80              S=L+rot(S, 1)
90              e=e+1
100     return S

```

(a) Sum the first n elements of L , adapted from left-to-right repeated squaring

```

10  sum_tile_b(L, n) :
20      e=1
30      X=L
40      Y=null
50      while true
60          if (n mod 2==1) then
70              if (Y==null) then
80                  Y=X
90              else
100                 Y=X+rot(Y, e)
110                 n=(n-1)/2
120            else
130                 n=n/2
140            if (n==0) then
150                return Y
160            X=X+rot(X, e)
170            e=e*2

```

(b) Sum the first n elements of L , adapted from right-to-left repeated squaring.

Fig. 5: Pseudocode of two rotate-and-sum algorithms adapted from the evaluation of powers algorithms. $numBits(n)$ denotes the number of bits in the integer expansion of n , and $bit(n, j)$ the j 'th bit in this expansion.

VI. SUMMATION ALGORITHMS

The ability to sum over a tile tensor depends primarily on the ability to sum over a tile. As usual, we focus on HE environments in which summing requires rotations.

We start with summing over a flat tile and demonstrate an improved algorithm for this purpose. Next, we extend the algorithm for a tile interpreted as a multi-dimensional tensor.

A rather simple further extension to summing over a tile tensor is deferred to Appendix A. The time performance of this algorithm depends on the tile tensor shape, which will later allow us to perform optimizations.

A. Summation Over a Single Flat Tile

In HE environments, a known technique for summing involves repeated rotations and additions, as described in [5] and shown here in Figure 5a.

When summing n elements where n is a power of 2, the algorithm collapses to a simpler algorithm. This is because the "if" condition in line 70 is always false. In this subsection, we discuss the optimality of this algorithm when n is not a power of 2. This may arise in situations where the tile size is not a power of 2, or we need to sum only the first n elements of the tile with n that is not a power of 2.

We show here a simple reduction between this problem and the problem of efficient evaluation of powers. The immediate benefit of this reduction is a new algorithm for summing over a tile that uses rotations by offsets that are powers of 2, which in some environments is more efficient. Additional benefits are some further insights regarding optimality.

Let L be a tile of length s . Let $L(j)$ be its element at index j for $0 \leq j < s$. For convenience, we define a notation for cyclic indices, $L\{j\} = L(j \bmod s)$.

We use two basic operations: rotate left $L' = \text{rot}(L, i)$, where L' is a vector of length s such that $L'\{j\} = L\{j + i\}$, and elementwise addition.

Let L^x denote a vector of length s such that $L^x\{j\} = \sum_{i=j}^{j+x-1} L\{i\}$. Note that $L^x(0)$ contains the sum of the first x elements of L . Note further that L^s is a vector where each of its elements is equal to $\text{sum}(L)$, hence it contains the result replicated in all slots.

To compute L^x for any x , we define an operator $L^x \otimes L^y = L^x + \text{rot}(L^y, x)$. It's easy to show that $L^x \otimes L^y = L^{x+y}$. Let $L' = L^x \otimes L^y$. Thus, $L'\{j\} = L^x\{j\} + L^y\{j + x\} = \sum_{i=j}^{j+x-1} L\{i\} + \sum_{i=j+x}^{j+x+y-1} L\{i\} = \sum_{i=j}^{j+x+y-1} L\{i\} = L^{x+y}$.

Now the problem of summing the first x elements of L can be reduced to efficiently computing L^x using the operator \otimes , obeying the law of exponents $L^x \otimes L^y = L^{x+y}$. This allows using algorithms for efficient evaluation of powers [22], [23].

Adapting the left-to-right repeated squaring algorithm to our operators results in the algorithm shown in Figure 5a, which is the same as shown in [5]. The second algorithm in Figure 5b is a novel variant adapted from the right-to-left repeated squaring algorithm.

In Figure 5a the rotation offset e can take any value, whereas in Figure 5b it can only be a power of 2 (e only changes by doubling in line 170). This is advantageous in some HE systems since support for efficient rotations is usually prepared for selected offsets, and powers of 2 is a reasonable general purpose choice, as is done in some major HE libraries.

Table II shows some performance measures. The technical details of these runs are given in Appendix F. Each tile has 4096 slots. We filled these slots with 1190 random numbers, and the rest with zeroes. The number 1190 was chosen because it is a non-power of 2 with about half of its bits set to 1. We compared three summation methods. Two of them are the two variants of Figure 5. The third sums over the first 2048 elements, in which both variants collapse to a much simpler algorithm as mentioned above. This third variant is only applicable when unused slots are zeroes. We measured time performance, number of rotations executed, and the mean square error (MSE) of noise in the result (the noise is due to

using the CKKS scheme). The results are the mean of 10,000 repetitions.

TABLE II: Comparison of rotate-and-sum algorithms for summing 1190 out of 4096 elements. The first two algorithms are depicted in Figure 5, and executed with $n = 1190$. The third is either algorithm with $n = 2048$.

Algorithm	Time (ms)	Noise (MSE)	Rotations
Left-to-right	151	2.5e-13	14
Right-to-left	74	3.8e-13	14
Sum over 2048	58	3.4e-13	11

Table II shows that the right-to-left variant we propose is about twice as fast as the left-to-right version, though a bit more noisy. The simple summation over the closest power of 2 is the fastest of the three, and not the least noisy. As mentioned above, it is not always applicable.

The reduction to power evaluation also shows that neither the left-to-right nor the right-to-left algorithms are generally optimal in the number of rotations. However, achieving an optimal number of rotations generally is complicated, and probably not worthwhile since the rotation offsets also have a large impact on performance, as shown in Table II.

B. Summation Over a Multi-dimensional Tile

Let $T[t_1, t_2, \dots, t_k]$ be a tensor mapped into a flat tile $L[s]$ in row-major order, i.e., $T(j_1, j_2, \dots, j_k) = L(\sum_i j_i \prod_{x=i+1}^k t_x)$. Thus, moving along the i 'th dimension of T means moving in strides of $d = \prod_{x=i+1}^n t_x$ inside the flat tile L . We also assume $\prod_x t_x = s$.

To sum over the i 'th dimension, we can therefore apply the summation algorithms of the previous subsection, by replacing the rotation $\text{rot}(L, x)$ with $\text{rot}(L, xd)$, but with a subtle difference. A $\text{rot}(L, xd)$ operation moves each element x steps backwards along the i 'th dimension. If it falls off the lower end, it doesn't rotate back to the other end, but actually its index of the previous dimension $i - 1$ is decreased. Therefore, $\text{rot}(L, xd)$ is equivalent to an operator $\text{shift}(L, i, x)$ that shifts elements x steps down along the i 'th dimension, introducing unknown values at the vacant slots.

The algorithms in Subsection VI-A can still be applied but with one exception. Previously we observed that summing over the entire size of the tile results in an outcome that is replicated in all slots. This observation doesn't hold when we replace rotate with shift.

For the special case of the first dimension $i = 1$, however, $\text{rot}(L, xd)$ does serve as full rotation along the first dimension, since there is no previous dimension. Hence, for this case, the result will still be replicated.

This is the underlying root cause for the summation rules of Subsection IV-D. Note further that if $\prod_x t_x < s$, replication will not occur. Since replication after summation of the first dimension is a useful property, the general tile tensor definition requires $\prod_x t_x = s$.

VII. MATRIX-VECTOR MULTIPLICATION

Let's assume we have two tensors: a matrix M of shape $[a, b]$ and a vector V of shape $[b]$. We would like to compute their product $R[a] = M[a, b]V[b]$. We now describe a general method of performing this using tile tensors.

Our new approach has several benefits. First, it captures four different known techniques for matrix vector multiplication in HE settings as special cases. These four techniques were previously described separately, and a completely different code was required to implement them. Using tile tensors we can implement all of them in a short piece of code and switch between them with easy parameter settings.

Second, the general case offers new and more efficient techniques for some common use cases, e.g., a series of consecutive matrix-vector multiplications.

Finally, the new formulation extends naturally to matrix-matrix multiplication.

A. Matrix-vector Multiplication General Formulas

Let $T_M[\frac{a}{t_1}, \frac{b}{t_2}]$ and $T_V[\frac{1}{t_1}, \frac{b}{t_2}]$ be tile tensors that contain M and V respectively, for some chosen tile shape $[t_1, t_2]$.

The result can be computed using the following formula:

$$T_R[\frac{a}{t_1}, \frac{1}{t_2}] = \text{sum}(T_M[\frac{a}{t_1}, \frac{b}{t_2}] * T_V[\frac{*}{t_1}, \frac{b}{t_2}], 2). \quad (2)$$

This formula works for any value of a, b, t_1, t_2 . To see why, note first that the shapes of T_M and T_V allow elementwise multiplication as explained in Subsection IV-D. Using the homomorphism of both elementwise operators and the summation operator, we can conclude that the result T_R contains a tensor R satisfying:

$$R[a, 1] = \text{sum}(M[a, b] * V[1, b], 2). \quad (3)$$

According to our extended definition of elementwise operators, V is multiplied elementwise with each row of M . We then sum over the second dimension. Thus, $R(i)$ is the inner product between the i 'th row of M and V , as required.

A second general approach is to initially transpose M , as follows: Let $T'_M[\frac{b}{t_1}, \frac{a}{t_2}]$ and $T'_V[\frac{b}{t_1}, \frac{1}{t_2}]$ be tile tensors that contain M^T and V respectively. Formula 4 shows how to multiply them.

$$T_R[\frac{*}{t_1}, \frac{a}{t_2}] = \text{sum}(T'_M[\frac{b}{t_1}, \frac{a}{t_2}] * T'_V[\frac{b}{t_1}, \frac{*}{t_2}], 2). \quad (4)$$

This computes the correct result using the same reasoning as before.

The benefit of this second approach, is that the result $T_R[\frac{*}{t_1}, \frac{a}{t_2}]$ is replicated along the first dimension due to the summation rules of Subsection IV-D. Thus, it is ready to play the role of V in Formula 2, and we can perform two matrix-vector multiplications consecutively without any processing in between.

Figure 6 shows a short Python function that implements both Formulas 2 and 4. All it assumes is that T_M and T_V are already correctly packed according to either formula, and the dimension to sum over is specified accordingly.

```
def matVecMult (T_M, T_V, sumDim) :
    res = T_M.elementwiseMult (T_V)
    res.sum(sumDim)
    return res
```

Fig. 6: Pseudo-code for matrix multiplication using tile tensors. We first multiply the two input tile tensors T_M and T_V elementwise, then sum over the specified dimension sumDim .

B. Known Special Cases

Specifying specific values for t_1 and t_2 in Formula 2, we can obtain three different known methods for matrix-vector multiplication in HE settings. A fourth known method is obtained with a small further generalization of Formula 2.

The first known method is referred to as matrix in row order [5], or the naïve approach [2]. In this approach each row of M is stored in a separate tile, and V is stored in a tile. This method is obtained by simply setting $t_1 = 1$ and $t_2 = s$ in the general case of Formula 2, where s is the total length of the tile.

A second method termed matrix in column order [5] or output packing [2] packs the matrix with each column in a separate tile. Here, of course, we can simply set $t_1 = s$ and $t_2 = 1$ in the general case of Formula 2.

Although the second method seems a trivial modification of the first method, each method was previously described as a different algorithm. In the first method, V 's tile is multiplied with all tiles of M , and each tile is summed using rotate-and-sum. In the second method, V is stored in multiple tiles and each is multiplied with one of M 's tiles. The results are then added together to produce just one tile. Using tile tensors, the same code snippet shown in Figure 6 covers both methods. The different behaviors described above emerge automatically from the definitions of elementwise multiplication and summation operators. Tile tensors can also easily handle the case where the matrix rows or columns exceed the tile sizes. This is an issue that would previously complicate each algorithm in a different way.

A third method is termed input packing [2] or stacked-vector - row major [4]. This method assumes $b \ll s$, allowing us to pack multiple rows of the matrix in a single tile, as many as can fit. Similarly, we replicate V the same number of times in a tile. After multiplication, the computation proceeds by summing inside the tile using rotate-and-sum, such that each replicate of V is summed separately.

Using tile tensors, this method is described as follows. We set $t_2 = b$ in Formula 2, forcing $t_1 = \lfloor \frac{s}{b} \rfloor$. This may violate the requirement we usually demand that $t_1 t_2 = s$. However, tile tensors can still operate without it just as well, except results after summation will never be replicated.

A better way though, would be to use Formula 4, i.e., transpose M prior to packing. We'll keep the requirement $t_1 t_2 = s$. E.g., we can require that t_1 be the smallest integer that divides s and satisfies $t_1 \geq b$. M is now packed in column-major order, with the addition of padding after each column

to make sure strides along the row divide s , and the vector V is packed with each element replicated consecutively several times. The benefit of this somewhat counter-intuitive approach is that as mentioned previously, the output of Formula 4 is ready to act as input for Formula 2. Hence, we can perform a second matrix-vector multiplication without processing in between.

A fourth method is usually described in the context of non-interactive neural network training or inference [17]–[19], where multiple matrices and vectors are processed in a batch. This too is a special case of our formulation, and is in fact obtainable by feeding the function in Figure 6 with appropriately packed inputs. A discussion of this fourth method appears in the Section VIII.

C. Advantages of the General Case

In the previous subsection we demonstrated how to capture four known techniques in a simple unified formulation. We also shared new insights on the input-packing (or vector stacking) method, and offered an improved version. In this subsection, we show that tile tensors offer new and sometimes better ways to pack matrices and handle matrix-vector multiplication.

Say our tile size is $s = 1024$, and our matrix is $M[768, 768]$. All of the methods described above pack M in tiles with at least one quarter of the slots unused, thus wasting both memory and time during the computation. Specifically, the input-packing method has no inherent mechanism for packing multiple rows together when two rows can't fit within a single tile. Using tile tensors, we can pack M to $T_M[\frac{768}{4}, \frac{768}{256}]$, which uses all the slots of each tile, while still supporting efficient elementwise multiplication and summation. Note that each tile in T_M contains one third of 4 different rows of M ; this is a packing unlike any of the ones described above.

Now say we are required to perform a series of consecutive matrix vector multiplications. I.e., given matrices M_1, M_2, \dots, M_k and V , we need to compute $(M_k \dots (M_2(M_1 V)))$. The shapes of the matrices are not necessarily the same, but are valid with respect to the rules of matrix multiplication.

One way to do it is to use the simple row-order and column-order methods described above in an alternating fashion. Matrices with even indices will be stored in row-order and the odd ones in column-order. Thus, the output of each stage will fit as input to the next stage. But if some of these matrices are smaller than the tile size, this will be highly wasteful. Using input-packing won't work, since matrices that have different dimensions will each require a different arrangement of data within the tiles, requiring costly processing of the output of one matrix to fit the next.

Using tile tensors, we can choose some tile size $[t_1, t_2]$, and pack all matrices consistently with it. We alternate the use of Formula 2 and 4, since the output of the latter is fit as input to the former.

For the opposite direction, some processing is required, but this can be done efficiently. The output of Formula 2 is

$T_R[\frac{a}{t_1}, \frac{1?}{t_2}]$. We first multiply it by a mask to clean the unknown values, resulting in $T_R[\frac{a}{t_1}, \frac{1}{t_2}]$. Then, we replicate the second dimension, resulting in $T_R[\frac{a}{t_1}, \frac{*}{t_2}]$, which is ready as input for Formula 4. The operator for performing this replication is not explicitly described in this paper, but it is very similar to the summation operator, performing the replication using similar rotate-and-sum technique.

We demonstrate the efficiency of this method in Section VIII, where we use it for neural network inference.

D. Extension to Matrix-Matrix Multiplication

Let's say we would like to multiply two matrices: $M_1[a, b]$ and $M_2[b, c]$. Using tile tensors, this can be done as follows:

$$T_R[\frac{a}{t_1}, \frac{1?, c}{t_2, t_3}] = \text{sum}(T_{M_1}[\frac{a}{t_1}, \frac{b}{t_2}, \frac{*}{t_3}] * T_{M_2}[\frac{*}{t_1}, \frac{b}{t_2}, \frac{c}{t_3}], 2). \quad (5)$$

Correctness follows in the same way as for the matrix-vector case: Formula 5 homomorphically computes the following formula.

$$R[a, 1, c] = \text{sum}(M_1[a, b, 1] * M_2[1, b, c], 2). \quad (6)$$

Where $M_1[a, b, 1]$ and $M_2[1, b, c]$ indicate trivial reshaping of $M_1[a, b]$ and $M_2[b, c]$. It's easy to verify that Formula 6 indeed matches matrix-matrix multiplication.

As before, we can transpose M_1 first to achieve a different formula, whose output can be used as input to the former one:

$$T_R[\frac{*}{t_1}, \frac{a, c}{t_2, t_3}] = \text{sum}(T_{M_1}[\frac{b}{t_1}, \frac{a}{t_2}, \frac{*}{t_3}] * T_{M_2}[\frac{b}{t_1}, \frac{*}{t_2}, \frac{c}{t_3}], 1). \quad (7)$$

The code in Figure 6 still works for this case as well, assuming the two inputs were properly packed.

E. Diagonalization Techniques

Some known techniques for matrix multiplication involve diagonalization, i.e., building a new matrix containing in some way diagonals from the original one [3], [5]. While being highly useful in some settings, they are usually inefficient in long non-interactive computation under encryption. See a short discussion in Appendix E.

VIII. NON-INTERACTIVE INFERENCE AND TRAINING OF NEURAL NETWORKS UNDER ENCRYPTION

In this section we show how tile tensors can improve both inference and the training of neural networks under encryption, especially when working in non-interactive mode.

In interactive mode, the client who has the decryption key is involved in the process and reducing communication overhead plays an important part in the optimization. For example, in the interactive inference algorithm proposed by Juvekar et al. [2] they opted to use diagonalization techniques for matrix multiplication, to reduce the size of the inputs and outputs. During the interaction, the client can repack the data in the manner suitable for the next step of the computation.

In non-interactive mode, a series of consecutive operations are performed under encryption. There’s a simple packing technique that is often used in such a setting; it is usually described as ”packing across the batch dimension”, ”packing the same dimension of multiple input samples in the same ciphertext”, or ”SIMD representation” [4], [17]–[19]. In this paper we refer to it as batch-packing.

The batch-packing method is yet another special case of the versatile range of tile tensors. It can be eloquently described using tile tensor shapes as follows. Let $B[x, n]$ be a batch of n input samples, each a vector of shape $[x]$. In batch-packing, it is packed into a tile tensor as $T_B[x, \frac{n}{s}]$, using tiles with s slots. To enable running this batch through the network, the network parameters must be stored replicated across all slots. I.e., if $W[a, b]$ is a weight matrix of one of the network layers, it is packed as $T_W[a, b, \frac{*}{s}]$.

Using batch-packing computation can be done simply and efficiently in SIMD (Single Instruction Multiple Data) fashion on all slots, usually with minimum or no rotation operations. Batch-packing is also stable in the sense that it remains the same throughout the computation. I.e., if a batch-packing of a matrix M is multiplied by a batch-packing of a vector V , the result is a batch-packing of the result R . Hence, there’s usually no overhead for preparing the output of one step to be correctly packed for the next.

However, batch-packing has some disadvantages for both inference and training, as explained in the following subsections.

A. Inference

Batch-packing for inference can reach a very high, possibly maximal, throughput [18]. However, this throughput is achieved when the batch size equals the number of slots, $n = s$. If the inference server is measured for latency on a single sample $n = 1$, or for batches much smaller than the number of slots $n \ll s$, batch-packing is highly inefficient.

Also, when using batch-packing the network parameters are replicated s times. For large values of s , they can burden the machine’s memory or may not fit in memory at all, requiring expensive disk I/O instead.

Finally, to avoid bootstrapping operations for inference on a deep neural network, the HE system should be configured for a multiplication depth that is high enough to support all the layers. In schemes commonly used in AI, such as CKKS [1], this entails high values of s , e.g., 8192 or 16384. This makes the above two problems more acute.

To solve these problems using our general technique for matrix multiplication, we further extend it using three-dimensional tiles; the third dimension is used for the batch dimension. E.g., a batch of n inputs vectors $V[a, n]$ we packed as $T_V[\frac{a}{t_1}, \frac{*}{t_2}, \frac{n}{t_3}]$, and a matrix $W[a, b]$ representing the weights of a fully connected layer we packed as $T_W[\frac{a}{t_1}, \frac{b}{t_2}, \frac{*}{t_3}]$. They can be multiplied by $sum(T_W * T_V, 1)$.

Our method captures batch-packing as a special case by setting $t_3 = s$ and $t_1 = t_2 = 1$. It still allows us the

freedom to chose many other settings, which will be useful as demonstrated by the results below.

For our experiments’ benchmarks we chose the CryptoNets network [17], described in Appendix B-A. The network was trained for classifying the MNIST dataset [24] and reaches an accuracy of 98.95%.

This network starts with a convolutional layer followed by two fully connected layers. When using batch-packing, we computed the convolution in the usual manner. But when we chose $t_3 < s$, we reduced computing the convolution to matrix-matrix multiplication by preprocessing the input as follows: for each filter F of the convolutional layer we identified all possible window locations on the input image. We extracted each such window, flattened it into a row, and created a matrix M_1 with all these rows. A second matrix M_2 was populated with the corresponding flattened filter for each row. Computing $M_1 M_2^T$ ends up with computing the convolution. We assumed M_1 was created on the client side prior to encryption. Thus, for the case $t_3 < s$, we ended up with 3 consecutive matrix multiplications, which we performed using our alternating scheme described in Subsection VII-C.

Our experiments use CKKS, configured for 8192 slots. The technical details are given in Appendix F. Since the input batch is packed as $T_V[\frac{a}{t_1}, \frac{*}{t_2}, \frac{n}{t_3}]$, it is most efficient when the batch size n equals t_3 . This prevents unused slots, and minimizes latency. Thus our experiments assume $n = t_3$. For each $t_3 = 1, 2, 4, \dots, 8192$ we tested all possible alternatives for t_1 and t_2 . All the results are the average of 10 runs.

In Table III we show some of the results for $t_3 = 1$. The most efficient tile size is $[32, 256, 1]$, which is globally the most efficient in terms of latency. The tile shapes $[1, 8192, 1]$ and $[8192, 1, 1]$ are equivalent to using matrix in column order and matrix in row order methods.

TABLE III: Time performance using different tile dimensions with $t_3 = 1$ and batch size $n = 1$, for inference time over the CryptoNets network. Latency and throughput measure the time performance of the inference calculation. The enc/dec column reports the client-side time spent on encryption and decryption. The memory column shows the server side’s memory usage.

Tile shape	Latency (sec)	Throughput (samples/min)	Enc/Dec (sec)	Memory (GB)
[1, 8192, 1]	2.62	22.9	0.12	1.66
[8, 1024, 1]	1.03	58.5	0.05	0.84
[32, 256, 1]	0.71	84.5	0.05	0.81
[64, 128, 1]	0.87	69.0	0.08	0.84
[128, 64, 1]	0.94	63.6	0.09	1.03
[256, 32, 1]	1.53	39.2	0.15	1.37
[1024, 8, 1]	5.59	10.7	0.45	3.13
[8192, 1, 1]	37.2	1.6	2.56	15.46

In Table IV we show some of the results for $t_3 = 1, 16, 64, 256, 1024, 4096, 8192$ and batch size $n = t_3$. For each value of t_3 we show the optimal value of t_1, t_2 . This demonstrates how latency degrades and throughput increases as t_3 grows. For $t_3 = 8192$ we get the ordinary batch-

packing, and therefore switched to the usual way of computing convolution, which reduces some overhead in input packing.

TABLE IV: Time performance using different tile dimensions. For $t_3 = 1, 4, 16, 64, 256, 1024, 4096, 8192$, the optimal choice for t_1 and t_2 is shown. We set batch size $n = t_3$. Latency and throughput measure the time performance of the inference calculation. The enc/dec column reports the client-side time spent on encryption and decryption. The memory column shows the server side’s memory usage.

Tile shape	Latency (sec)	Throughput (samples/min)	Enc/Dec (sec)	Memory (GB)
[32, 256, 1]	0.71	84.5	0.05	0.81
[16, 128, 4]	0.86	278.4	0.09	1.18
[8, 64, 16]	1.29	741.9	0.27	2.38
[4, 32, 64]	2.75	1399.9	0.78	6.39
[1, 32, 256]	5.45	2811.1	2.31	16.62
[1, 8, 1024]	16.42	3741.6	8.93	62.54
[1, 2, 4096]	60.32	4073.79	18.33	246.99
[1, 1, 8192]	111.71	4399.8	2.7	361.85

Our results demonstrate several advantages when using tile tensors. First, they offer a knob, namely t_3 , allowing us to adapt to batch size or to control the trade off between throughput and latency. In addition, t_3 can be used for controlling memory usage. Recall that all this is done without any code changes.

Second, the latency of 0.71 seconds is to our knowledge the lowest latency reported for an MNIST neural network classifier under encryption that achieves close to 99% accuracy - see a more detailed comparison in the next subsection. This is achievable thanks to our ability to optimize t_1 and t_2 once the value of t_3 was set. It is also possible thanks to the alternating scheme of Subsection VII-C, which reduces latency further by eliminating the processing required after the first layer.

B. Comparison with Advanced Neural Network Inference Techniques

The disadvantages of batch-packing was noticed and addressed in previous works as well [4], [20], [25]. An exact performance comparison under equal conditions is difficult due to various constraints, but in this subsection we offer an in-depth comparison, aimed at pointing out the advantages of tile tensors for neural network inference.

The LoLa network [4] performs predictions on the same CryptoNets architecture (among others), using a mixture of known matrix-vector multiplication methods (described in Subsection VII-B). Switching between different methods within a single inference computation requires a processing stage between layers, resulting in extra additions and rotations. They achieve a latency of 2.2 seconds using 8 threads. Our lowest latency is 0.71 seconds using 20 threads; reducing to 8 threads changes this result slightly to 0.72. The LoLa network uses 150 ciphertext-ciphertext multiplications, 279 rotations, and 399 additions for a single prediction. (We deduced these numbers from LoLa’s detailed description.) Our approach requires 32 multiplications, 89 rotations, and 113 additions.

This roughly three-fold reduction matches the observed latency results. The reduction is due to the advantages of tile tensors, as described in the previous subsection. Using tile tensors has many more advantages over a specific tailored computation, as they allow us to easily switch between different methods, adapt to different batch sizes, simplify logging and debugging, and more.

Another advanced tool, the Chet compiler [20], performs inference of encrypted data in a non-encrypted network. For this easier problem, they report 2.5 seconds latency on a similarly sized, though less accurate, MNIST neural network classifier using 16 threads. This result was later improved by the EVA [25] compiler to 0.6 seconds using 56 threads and various optimizations unrelated to packing, of a kind outside the scope of this paper. Our best result on the more accurate CryptoNets architecture, when the network is not encrypted, reduces to 0.55 seconds using 16 threads. Our best latency result relies heavily on rotations of the data flowing through the network, thus inference on a non-encrypted network is only slightly better than when it is encrypted.

A direct comparison with EVA is difficult here because of multiple optimizations in EVA beyond the scope of this work (e.g., eliminating rescale operations so as to reduce the overall prime chain length) and the different architecture with multiple convolution layers, (for which tile tensor support will be added soon). However, regarding the packing techniques, we can argue for the usual benefits of tile tensors.

Both EVA and the Chet compiler use a data structure termed CipherTensor. Like tile tensors, CipherTensor abstracts away the packing details using some user-defined meta data. However, we believe tile tensors have several advantages. EVA and the Chet compiler also optimize the CipherTensor’s meta-data given a specific computation. As per our understanding, this optimization is restricted to a small set of implemented layouts, whereas tile tensors shapes offer a wider variety of options. Moreover, CipherTensors don’t have the summation algorithm that can result with replication, hence replication is needed before each layer. Using the tile tensor summation algorithm, it can be skipped for every other layer. Finally, it wasn’t demonstrated that CipherTensors offer an easy method to trade latency for throughput, and control memory consumption, as demonstrated for tile tensors in the previous subsection.

C. Training

Our goals in this subsection are more modest. We will not show a fully working non-interactive training system, but will point to a new type of reasoning that leads to a drastic reduction in the number of bootstrapping operations. Tile tensors play an important role in this optimization.

A previous attempt to train neural networks non-interactively under encryption by Nandakumar et al. [19] used the BGV scheme in HELib [5], where bootstrapping operations are done automatically whenever needed. As mentioned above, they used batch-packing.

We will use CKKS instead of BGV. Other than natively handling real numbers (in fact, complex), it has another trait interesting for our purposes: it allows tracking the state of a ciphertext in the modulus chain through a property called chain-index. Once the chain-index reaches 0, the ciphertext needs to be bootstrapped before additional operations are performed. This allows us to lay out a plan for the entire computation that will minimize the number of bootstrapping operations.

Bootstrapping is currently not implemented for CKKS in any of the major libraries, however a known method exists [26], and we expect it will be implemented soon. Thus, we conducted experiments in which we implemented bootstrapping by simply decrypting and re-encrypting the ciphertext, but counting the times this operation was required.

For our case study, we used a network based on the one described as *CNN-static* by Yoon Kim [27] designed for classifying newsgroup texts. It was identified as interesting by a customer requiring training under encryption.

We first adapted it to be HE friendly. Its adapted architecture is detailed in Appendix B-B; it contains a convolutional layer with square activation, mean-pooling, dropout, and a fully connected layer. Using Euclidean loss and stochastic gradient descent, we trained it under encryption on a subset of 6 news groups from the original dataset.

Batch-packing seems especially well-suited for training a neural network under encryption. The training process is complicated and circular: the output gradients of one iteration are added to the network weights, which are then used as an input in the next iteration. The stability property of batch-packing therefore greatly simplifies the implementation using HE.

However, our implementation using batch-packing resulted in 1286 bootstrapping operations required per iteration. Since a single bootstrapping operation is likely to cost a few seconds at least, this number renders training under encryption unfeasible. As mentioned above, the number of bootstrapping operations can be reduced by carefully planning the chain indices at every step of the computation. However, it can be shown that using pure batch-packing for this architecture entails a lower bound of 1088 operations, still an unfeasible count. Proof of this claim appears in Appendix C.

For this reason, we switched to using tile tensors. We manually crafted a plan where we designated a specific tile tensor shape for each component in this computation. The details of this plan are in Appendix D.

Our new method turned out to have several disadvantages over batch-packing: it was slower and more noisy. The additional noise required us to change the HE configuration to allow more precision and higher multiplication depth (4 instead of 3). However, it drastically reduced the number of bootstrapping operations per iteration, from 1286 to 14, making the entire computation feasible.

The improvement in bootstrapping count stems from a more efficient representation of the filter weights, whereas using batch-packing they are replicated multiple times. This allows

bootstrapping the filters instead of the convolution output, reducing their overall count.

Our hand-crafted plan exploits the properties of tile tensors. The versatility and wide range of packing options allowed us to implement each part of the network with a different shape. The data flowing through it continuously shape-shifts as well.

Table V shows the results for the batch-packing method versus our new specific method we label "tile tensors". Technical details of the runs are in Appendix F.

TABLE V: Performance results of training with batch-packing vs. our new handcrafted method specific for tile tensor shapes.

Performance measure	Batch-packing	Tile tensors
Runtime per iteration (excluding bootstrapping)	135 seconds	633 seconds
Bootstrapping operations count	1286	14
Noise after 10 iterations	0.000008	0.00024

IX. CONCLUSIONS

We presented tile tensors, an efficient data structure for working with tensors in a system that imposes tiles.

We demonstrated improved latency for neural network inference under encryption using a new technique for matrix multiplication, and some automated optimizations. For neural network training, we managed to dramatically reduce the number of bootstrapping operations using a handcrafted plan. We believe that an automated optimizer can use tile tensor shapes to find an even better plan.

Tile tensor shapes can further be used as a new language for describing packing details.

REFERENCES

- [1] J. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *Proceedings of Advances in Cryptology - ASIACRYPT 2017*. Springer Cham, 11 2017, pp. 409–437.
- [2] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, "GAZELLE: A low latency framework for secure neural network inference," in *27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, Aug. 2018, pp. 1651–1669. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity18/presentation/juvekar>
- [3] X. Jiang, M. Kim, K. Lauter, and Y. Song, "Secure outsourced matrix computation and application to neural networks," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1209–1222. [Online]. Available: <https://doi.org/10.1145/3243734.3243837>
- [4] A. Brutzkus, R. Gilad-Bachrach, and O. Elisha, "Low latency privacy preserving inference," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. Long Beach, California, USA: PMLR, 09–15 Jun 2019, pp. 812–821. [Online]. Available: <http://proceedings.mlr.press/v97/brutzkus19a.html>
- [5] S. Halevi and V. Shoup, "Algorithms in helib," in *Advances in Cryptology - CRYPTO 2014*, J. A. Garay and R. Gennaro, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 554–571.
- [6] R. Agrawal and R. Srikant, "Privacy-preserving data mining," in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '00. New York, NY, USA: Association for Computing Machinery, 2000, p. 439–450. [Online]. Available: <https://doi.org/10.1145/342009.335438>

- [7] Y. Lindell and B. Pinkas, “Privacy preserving data mining,” in *Advances in Cryptology — CRYPTO 2000*, M. Bellare, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 36–54.
- [8] M. J. A. W. Du, “Privacy-preserving cooperative scientific computations,” in *Proceedings of the 14th IEEE Computer Security Foundations Workshop*. Cape Breton, Nova Scotia, Canada: IEEE Computer Society, 2001, pp. 273–282.
- [9] A. B. Slavkovic, Y. Nardi, and M. M. Tibbits, “Secure logistic regression of horizontally and vertically partitioned distributed databases,” in *Proceedings of the Seventh IEEE International Conference on Data Mining Workshops*, ser. ICDMW ’07. USA: IEEE Computer Society, 2007, p. 723–728. [Online]. Available: <https://doi.org/10.1109/ICDMW.2007.84>
- [10] J. Vaidya, H. Yu, and X. Jiang, “Privacy-preserving svm classification,” in *Proceedings of Knowl Inf Syst 14*. Springer-Verlag, 2008, p. 161–178. [Online]. Available: <https://doi.org/>
- [11] C. Gentry, “A fully homomorphic encryption scheme,” Ph.D. dissertation, Stanford University, Palo Alto, CA, 2009. [Online]. Available: <https://crypto.stanford.edu/craig/craig-thesis.pdf>
- [12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, “(leveled) fully homomorphic encryption without bootstrapping,” in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ser. ITCS ’12. New York, NY, USA: Association for Computing Machinery, 2012, p. 309–325. [Online]. Available: <https://doi.org/10.1145/2090236.2090262>
- [13] C. Gentry, A. Sahai, and B. Waters, “Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based,” in *Advances in Cryptology – CRYPTO 2013*, R. Canetti and J. A. Garay, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 75–92.
- [14] Z. Brakerski and V. Vaikuntanathan, “Lattice-based fhe as secure as pke,” in *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ser. ITCS 14. New York, NY, USA: Association for Computing Machinery, 2014, p. 1–12. [Online]. Available: <https://doi.org/10.1145/2554797.2554799>
- [15] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, “Tfhe: Fast fully homomorphic encryption over the torus,” *Journal of Cryptology*, 04 2019.
- [16] e. a. Jiang, Xiaoqian, “Secure outsourced matrix computation and application to neural networks,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 2018, 2018. [Online]. Available: <https://dl.acm.org/doi/pdf/10.1145/3243734.3243837>
- [17] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, “Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy,” in *International Conference on Machine Learning*, 2016, pp. 201–210.
- [18] F. Boemer, Y. Lao, R. Cammarota, and C. Wierzynski, “Ngraph-he: A graph compiler for deep learning on homomorphically encrypted data,” in *Proceedings of the 16th ACM International Conference on Computing Frontiers*, ser. CF ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 3–13. [Online]. Available: <https://doi.org/10.1145/3310273.3323047>
- [19] K. Nandakumar, N. Ratha, S. Pankanti, and S. Halevi, “Towards deep neural network training on encrypted data,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [20] R. Dathathri, O. Saarikivi, H. Chen, K. Laine, K. Lauter, S. Maleki, M. Musuvathi, and T. Mytkowicz, “Chet: An optimizing compiler for fully-homomorphic neural-network inferencing,” in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 142–156. [Online]. Available: <https://doi.org/10.1145/3314221.3314628>
- [21] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [22] D. E. Knuth, *The Art of Computer Programming, vol. 2, Seminumerical Algorithms*, 2nd ed. Addison-Wesley Pub (Sd), 1 1981, vol. 2.
- [23] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. PHI Learning Pvt. Ltd. (Originally MIT Press), 2010.
- [24] Y. LeCun, C. Cortes, and C. J. Burges, “The mnist database of handwritten digits, 1998,” *URL http://yann.lecun.com/exdb/mnist*, vol. 10, p. 34, 1998.
- [25] R. Dathathri, B. Kostova, O. Saarikivi, W. Dai, K. Laine, and M. Musuvathi, “Eva: An encrypted vector arithmetic language and compiler for efficient homomorphic computation,” in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2020. New York, NY, USA: Association for Computing Machinery, 2020, p. 546–561. [Online]. Available: <https://doi.org/10.1145/3385412.3386023>
- [26] H. Chen, I. Chillotti, and Y. Song, “Improved bootstrapping for approximate homomorphic encryption,” in *Advances in Cryptology – EUROCRYPT 2019*, Y. Ishai and V. Rijmen, Eds. Cham: Springer International Publishing, 2019, pp. 34–54.
- [27] Y. Kim, “Convolutional neural networks for sentence classification,” 2014.
- [28] K. Lang. (2000) Cmu text learning group data archives - 20_newsgroup. [Online]. Available: <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-20/www/data/news20.html>
- [29] “Microsoft SEAL (release 3.5),” <https://github.com/Microsoft/SEAL>, Apr. 2020, microsoft Research, Redmond, WA.

APPENDIX A TILE TENSOR OPERATORS

Tile tensor can support a wide variety of operators. In this paper we focus on elementwise operators, and a summation operator. We show that both can be implemented efficiently, and together they make tile tensors a powerful tool.

Both are defined homomorphically, i.e., applying them on a tile tensor is equivalent to applying them on the tensor it contains. Hence we’ll start this section with a definition of these operators on regular tensors.

A. Operators on regular tensors

We’ll start with defining elementwise operators on regular tensors in an extended way. For notation simplicity and w.l.o.g., we’ll use the addition operator.

Let the shapes of the tensors A and A' be $[n_1, n_2, \dots, n_k]$ and $[n'_1, n'_2, \dots, n'_k]$ respectively. We’ll say that their shapes are compatible if $n_i = n'_i \vee n_i = 1 \vee n'_i = 1$. We’ll define the result of the elementwise operation to be equivalent to first replicating the mismatching dimensions of A and A' and then executing the elementwise operation.

Formally, the result of $A + A'$ is a tensor A'' of shape $[m_1, m_2, \dots, m_k]$ such that $m_i = \max(n_i, n'_i)$, and $A''(j_1, j_2, \dots, j_k) = A(j_1 \bmod n_1, j_2 \bmod n_2, \dots, j_k \bmod n_k) + A'(j_1 \bmod n'_1, j_2 \bmod n'_2, \dots, j_k \bmod n'_k)$.

The summation operator is defined as follows. Let A be a tensor of shape $[n_1, n_2, \dots, n_k]$. We’ll denote by $sum(A, i)$ the result of summing of A along dimension i . Let’s look at the result for $i = 1$ for example (recall our dimension indices are one-based, hence $i = 1$ is the first dimension. For element indices however we use zero-based indices). The result, $A' = sum(A, 1)$, is a tensor of shape $[1, n_2, \dots, n_k]$, such that $A'(0, j_2, \dots, j_k) = \sum_{j_1=0}^{n_1-1} A(j_1, j_2, \dots, j_k)$. Similarly, summing over any dimension $1 \leq i \leq k$ reduces the corresponding dimension size to 1.

B. Elementwise operators on tile tensors

Again, for notation simplicity and w.l.o.g., we'll use the addition operator.

Let T_A and T'_A be two tile tensors with shapes in the general form $[\frac{n_1*d_1}{t_1}, \frac{n_2*d_2}{t_2}, \dots, \frac{n_k*d_k}{t_k}]$ and $[\frac{n'_1*d'_1}{t'_1}, \frac{n'_2*d'_2}{t'_2}, \dots, \frac{n'_k*d'_k}{t'_k}]$ respectively. We assume for simplicity no dimension is annotated by the unknown ? marking.

The result of addition: $T''_A = T_A + T'_A$ is defined as follows. The external tensor of T''_A is computed by an elementwise operator on the external tensors of T_A and T'_A , i.e., $et(T''_A) = et(T_A) + et(T'_A)$. We treat here the external tensors as a tensor of tiles, i.e., the operator is ultimately applied on corresponding tiles. According to the definition in Subsection A-A, some dimensions may be logically replicated.

The above computation is well defined whenever $et(T_A)$ and $et(T'_A)$ have compatible shapes in the sense of Subsection A-A, but to achieve the homomorphic property, $unpack(T_A + T'_A) = unpack(T_A) + unpack(T'_A)$, additional constraints are required. The following conditions are sufficient to guarantee both compatibility of the external tensors and the homomorphism:

- $t_i = t'_i$, and
- $n_i = n'_i \vee (n_i = 1 \wedge d_i = t_i) \vee (n'_i = 1 \wedge d'_i = t'_i)$

In words, the tile dimensions must match exactly. Also, the dimensions of the tensors packed inside must either match or one of them is fully replicated.

The result tile tensor has the shape: $[\frac{n''_1*d''_1 I_1(?) }{t_1}, \frac{n''_2*d''_2 I_2(?) }{t_2}, \dots, \frac{n''_k*d''_k I_k(?) }{t_k}]$, where:

- $n''_i = \max(n_i, n'_i)$
- $d''_i = \min(d_i, d'_i)$
- $I_i(?) = ?$ if $n_i d_i \neq n'_i d'_i$, and empty otherwise.

If any of the input dimensions is annotated with ?, then so will the corresponding output dimension. For the multiplication operator, since multiplying by zero is zero, some further deduction can be made that reduce the number of ? annotations, which we omit here.

C. Summing over a tile tensor

We'll now show how to efficiently compute $sum(T_A, i)$ for any tile tensor T_A and any dimension i , while maintaining the homomorphic property, $unpack(sum(T_A, i)) = sum(unpack(T_A), i)$.

The computation starts by summing over the external tensor $et(T_A)$. I.e., we reduce the size of $et(T_A)$ along the i 'th dimension to 1 by summing the tiles together. Then we sum along the i 'th dimension of each tile, as explained in Subsection VI-B.

If a dimension of T_A is annotated with ?, then the last tiles along this dimension, those that contain the unused unknown slots, should be handled separately. We execute the summation on all other tiles, then sum inside these last tiles, and add them to the rest. When summing inside the tiles that contain unknown values, we'll sum over just the known values, an operation that is generally more expensive, and prevents replication from occurring. Thus in the summation rules of

subsection IV-D summing over a dimension annotated with ? always ends up with ? in the result.

The complexity of summing a tile tensor of shape $[\frac{n_1}{t_1}, \frac{n_2}{t_2}, \dots, \frac{n_k}{t_k}]$ over dimension j is $O(\prod_i [n_i/t_i] + \log(t_j) \prod_{i \neq j} [n_i/t_i])$ addition operations, and $O(\log(t_j) \prod_{i \neq j} [n_i/t_i])$ rotations. Replication and unknowns do not change this complexity.

APPENDIX B NEURAL NETWORKS ARCHITECTURE

A. Cryptonets

Architecture defined in [17] with activation function $Act(x) = x^2$.

- 1) Conv: [Input: 28×28 , 5 filters of size 5×5 , stride=2, output: 845]+Act..
- 2) FC: [Input: 845, Output: 100]+Act.
- 3) FC: [Input: 100, Output: 10].

B. HE friendly CNN-static

This network is based on the one described as *CNN-static* by Yoon Kim [27], that was also implemented by the Kaggle community¹. The network was trained with the accompanying dataset from the CMU Text Learning Group Data Archives [28].

For our case study we simplified it to be smaller and more HE friendly, while reaching similar accuracy on a subset of 6 news groups from the original dataset. The simplified network is described in figure 7.

The input is 20×50 matrix representing an embedded sentence. It has one convolutional layer with one filter group of 32 filters of size 3×50 , producing output of size 18×32 , and another filter group of 32 filters of size 5×50 , producing output of size 16×32 , both followed by square activation. A mean pooling layer reduces both outputs to 1×32 , then concatenated to 1×64 . This is followed by a dropout layer, a fully connected layer with output size of 6. The loss function is Euclidean.

¹<https://www.kaggle.com/au1206/text-classification-using-cnn>

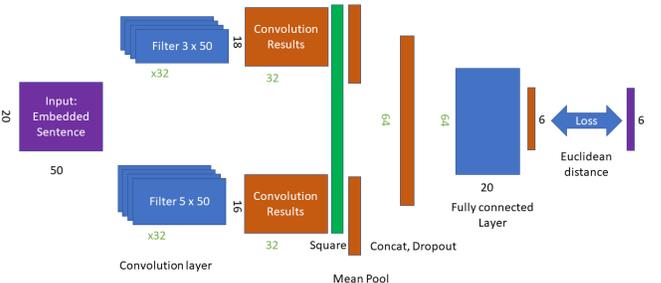


Fig. 7: An HE friendly version of the CNN-Static network.

APPENDIX C
BOOTSTRAPPING OPTIMALITY RESULT

In this appendix we’ll show how to compute a lower bound on the number of bootstrapping operations for training the case study network described in Appendix B-B, assuming pure batch-packing.

Let $T_F[3, 50, \frac{*}{s}]$ be the tile tensor containing a single filter of size 3×50 of the convolutional layer. Let $T_X[20, 50, \frac{n}{s}]$ be its input, and let $T_O[18, 1, \frac{*}{s}]$ be the result of the convolution of T_X and T_F , all using batch-packing.

During the back-propagation stage, the convolutional layer receives as input from the subsequent square activation layer the gradient of T_O , which we’ll denote by $T_{G_O}[18, 1, \frac{n}{s}]$. From this we’ll derive the gradient of T_F , which we’ll denote as $T_{G_F}[3, 50, \frac{*}{s}]$.

Let $C(T_{G_F}), C(T_{G_O}), C(T_O), C(T_F)$ be the corresponding chain indices of each of the tile tensors (i.e., the chain indices of the cipher text tiles they are made of. We assume it kept consistent for all tiles in a tile tensor). T_{G_F} is computed from a multiplication involving T_{G_O} , hence $C(T_{G_F}) < C(T_{G_O})$. Similarly, T_{G_O} is computed from a multiplication involving T_O , and T_O from a multiplication involving T_F , hence without bootstrapping intervention we get: $C(T_{G_F}) < C(T_{G_O}) < C(T_O) < C(T_F)$.

Finally we execute $T_F += \alpha T_{G_F}$ after each iteration (α is the learning rate. This multiplication by a constant can be avoided by various ways). Thus without bootstrapping $C(T_F)$ will drop by 3 points every iteration. Therefore we can conclude that if our configuration offers a multiplication depth of d , then we will have to bootstrap one of $T_{G_F}, T_{G_O}, T_O, T_F$ an amortized count of $3/d$ times.

The smallest of these in number of ciphertexts are either T_O or T_{G_O} having only 18 ciphertexts each, whereas the other two have 150 ciphertexts. Thus we can conclude an amortized cost of at least $54/d$ ciphertexts per filter of size 3×50 per iteration. For filter of size 5×50 a similar computation yields $48/d$, and therefore if 32 such filters of each groups are used and $d = 3$, this amounts to at least 1088 bootstrapping operations amortized per iteration. For $d = 4$ this amounts to at least 816 bootstrapping operations.

APPENDIX D
HANDCRAFTED PLAN FOR TRAINING A CONVOLUTIONAL NETWORK USING TILE TENSORS

In this appendix we’ll show a handcrafted plan for training the convolutional network described in Appendix B-B, using tile tensors.

A filter of size 3×50 was first flattened to a single vector of length 150. Thus the entire group of 32 filters of this size are stored in tensor $F_1[32, 150]$. The input batch $B[20, 50, n]$ was preprocessed to $B_1[18, 150, n]$, such that $B[i, 150, j]$ are the elements of sample j that correspond to the filter when placed over row i of the input matrix. Similarly filters of size 5×50 were flattened and stored in a tensor $F_2[32, 250]$, and the same batch B was preprocessed to $B_2[16, 150, n]$.

For batch size of $n = 255$, tables VI and VII depict the tile tensor shapes used. The plan uses a reshaping operator (for generating T_Q , which is straightforward but not specified in this paper. The improvement in bootstrapping counts stems from the compact representation of the filters, T_{F_1} and T_{F_2} , allowing to bootstrap them instead of the larger convolution output. Using batch-packing, the filters are too large to be bootstrapped, so the convolution output must be bootstrapped.

TABLE VI: Tile tensor packing for different network parts.

Network part	Tile tensor
Filter group 1	$T_{F_1}[1, \frac{32}{32}, \frac{150}{256}]$
Filter group 1 biases	$T_{G_1}[1, \frac{32}{32}, \frac{*}{256}]$
Filter group 2	$T_{F_2}[1, \frac{32}{32}, \frac{250}{256}]$
Filter group 2 biases	$T_{G_2}[1, \frac{32}{32}, \frac{*}{256}]$
Fully connected weights	$T_W[6, \frac{64}{32}, \frac{*}{256}]$
Fully connected biases	$T_X[6, \frac{*}{32}, \frac{*}{256}]$

TABLE VII: Tile tensor packing for input, labels, and intermediate values. The gradients have the same shape as the tile tensors they are the gradient of.

Network part	Tile tensor
B1	$T_{B_1}[18, \frac{*}{32}, \frac{150}{256}, 255]$
B2	$T_{B_2}[16, \frac{*}{32}, \frac{250}{256}, 255]$
Convolution output 1	$T_{O_1}[18, \frac{32}{32}, \frac{17}{256}, 255]$
Convolution output 2	$T_{O_2}[16, \frac{32}{32}, \frac{17}{256}, 255]$
After mean pooling 1	$T_{P_1}[1, \frac{32}{32}, \frac{17}{256}, 255]$
After mean pooling 2	$T_{P_2}[1, \frac{32}{32}, \frac{17}{256}, 255]$
After concatenate and reshape	$T_Q[1, \frac{64}{32}, \frac{255}{256}]$
After fully connected	$T_R[6, \frac{*}{32}, \frac{255}{256}]$
Labels	$T_L[6, \frac{*}{32}, \frac{*}{256}]$

APPENDIX E
DIAGONALIZATION TECHNIQUES

Some known techniques for matrix multiplication involve diagonalization, i.e., building a new matrix containing in some way diagonals from the original one [3], [5].

These methods are highly useful in some settings. While tile tensors can’t help with the diagonalization computation, they can still be used to pack the diagonalized matrix.

However, we argue that diagonalization techniques are inefficient for settings where long non-interactive computations are performed under encryption, e.g., non-interactive neural network inference or training.

The method of Jiang et al. [3] requires two extra multiplications by a constant to prepare the diagonalized matrices. This increases the multiplication depth which is the most costly resource in non-interactive computations under HE. Configuring an HE system for a larger multiplication depth dramatically reduces the computational efficiency of all operations, or alternatively, expensive bootstrapping operations will be incurred.

The method of [5] has an advantage that its input and output are both minimal in size, but it requires $O(s)$ rotations of the input vector of size s , in contrast to $O(\log(s))$ rotations required when using tile tensors. In [2] a variant of this diagonalization method was used in the context of interactive neural network inference to reduce communication overhead. In a non-interactive settings there’s no communication during the computation, hence communication overhead is less important, whereas the cost of the $O(s)$ repeated rotations far exceeds the cost of the multiplication operation itself.

APPENDIX F

EXPERIMENT RESULTS SPECIFICATIONS

All experiments results reported in this paper use the same machine, an Intel(R) Xeon(R) CPU E5-2620 v3 @ 2.40GHz machine with 24 CPUs and 384GB memory. We used the CKKS implementation in SEAL [29].

Summation algorithms experiment was done with poly-degree 8192 and modulus chain $\{50, 40, 40, 50\}$, single thread.

In the neural network inference experiment, we used poly-degree 16384. The modulus chain was $\{45, 35, 35, 35, 35, 35, 45\}$ when $t_3 = s$, and for $t_3 < s$ the modulus chain was $\{45, 35, 35, 35, 35, 35, 45\}$, allowing for a multiplication depth larger by 1, needed for replicating the results after layer 2 as is required by our alternating scheme. We used 20 threads. All results are the average of 10 runs.

In the neural network training experiment, for batch-packing we used poly-degree 8192, and modulus chain $\{49, 40, 40, 40, 49\}$. For the tile tensor handcrafted method, to accommodate large depth and greater noise we used poly-degree 16384, and modulus chain $\{59, 50, 50, 50, 50, 59\}$. We used 24 threads.