



**HAL**  
open science

# High performance tensor-vector multiples on shared memory systems

Filip Pawlowski, Bora Uçar, Albert-Jan Yzelman

► **To cite this version:**

Filip Pawlowski, Bora Uçar, Albert-Jan Yzelman. High performance tensor-vector multiples on shared memory systems. [Research Report] RR-9274, Inria - Research Centre Grenoble – Rhône-Alpes. 2019, pp.1-20. hal-02123526v1

**HAL Id: hal-02123526**

**<https://inria.hal.science/hal-02123526v1>**

Submitted on 8 May 2019 (v1), last revised 24 Oct 2019 (v2)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



# High performance tensor–vector multiples on shared memory systems

Filip Pawłowski, Bora Uçar, Albert-Jan Yzelman

**RESEARCH  
REPORT**

**N° 9274**

May 2019

Project-Team ROMA





## High performance tensor–vector multiples on shared memory systems

Filip Pawłowski\*, Bora Uçar†, Albert-Jan Yzelman‡

Project-Team ROMA

Research Report n° 9274 — May 2019 — 20 pages

**Abstract:** Tensor–vector multiplication is one of the core components in tensor computations. We have recently investigated high performance, single core implementation of this bandwidth-bound operation. In this work, we investigate efficient, shared memory algorithms to carry out this operation. Upon carefully analyzing the design space, we implement a number of alternatives using OpenMP and compare them experimentally. Experimental results on up to 8 socket systems show near peak performance for the proposed algorithms.

**Key-words:** tensors, tensor–vector multiplication, shared-memory parallel machines

---

\* Huawei Technologies France and ENS Lyon

† CNRS and LIP (UMR5668 Université de Lyon - CNRS - ENS Lyon - Inria - UCBL 1),  
46, allée d’Italie, ENS Lyon, Lyon F-69364, France.

‡ Huawei Technologies France

**RESEARCH CENTRE  
GRENOBLE – RHÔNE-ALPES**

Inovallée  
655 avenue de l’Europe Montbonnot  
38334 Saint Ismier Cedex

## Multiplication tenseur–vecteur haute performance sur des machines à mémoire partagée

**Résumé :** La multiplication tenseur–vecteur est l’un des composants essentiels des calculs de tenseurs. Nous avons récemment étudié cette opération, qui consomme la bande passante, sur une plateforme séquentielle. Dans ce travail, nous étudions des algorithmes efficaces pour effectuer cette opération sur des machines à mémoire partagée. Après avoir soigneusement analysé les différentes alternatives, nous mettons en œuvre plusieurs d’entre elles en utilisant OpenMP, et nous les comparons expérimentalement. Les résultats expérimentaux sur un à huit systèmes de sockets montrent une performance quasi maximale pour les algorithmes proposés.

**Mots-clés :** tenseur, multiplication tenseur-vecteur, machines parallèles à mémoire partagée

## 1 Introduction

Tensor–vector multiply (*TVM*) operation, along with its higher level analogues tensor–matrix (*TMM*) and tensor–tensor multiplies (*TTM*), are the building blocks of many algorithms [1]. All these three operations are applied to a given mode (or dimension) in *TVM* and *TMM* or modes in *TTM*. Among these three operations, *TVM* is the most bandwidth-bound. Recently, we have investigated this operation on single core systems and proposed data structures and algorithms to achieve high performance and mode-oblivious behavior [9]. While high performance is a common term in the close by area of matrix computations, mode-obliviousness is mostly related to tensor computations. It requires that a given algorithm for a core operation (*TVM* in our case) should have more or less the same performance no matter which mode it is applied to. In matrix terms, this corresponds to having the same performance in computing matrix–vector and matrix–transpose–vector multiplies. Our aim in this work is to develop high performance and mode oblivious parallel algorithms on shared memory systems.

Let  $\mathcal{A}$  be a tensor with  $d$  modes, or for our purposes in this paper, a  $d$ -dimensional array. The  $k$ -mode tensor–vector multiplication produces another tensor whose  $k$ th mode is of size one. More formally, for  $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$  and  $\mathbf{x} \in \mathbb{R}^{n_k}$ , the  $k$ -mode *TVM* operation  $\mathcal{Y} = \mathcal{A} \times_k \mathbf{x}$  is defined as

$$y_{i_1, \dots, i_{k-1}, 1, i_{k+1}, \dots, i_d} = \sum_{i_k=1}^{n_k} a_{i_1, \dots, i_{k-1}, i_k, i_{k+1}, \dots, i_d} x_{i_k},$$

for all  $i_j \in \{1, \dots, n_j\}$  with  $j \in \{1, \dots, d\}$ , where  $y_{i_1, \dots, i_{k-1}, 1, i_{k+1}, \dots, i_d}$  is an element of  $\mathcal{Y}$ , and  $a_{i_1, \dots, i_{k-1}, i_k, i_{k+1}, \dots, i_d}$  is an element of  $\mathcal{A}$ . The output tensor  $\mathcal{Y} \in \mathbb{R}^{n_1 \times \dots \times n_{k-1} \times 1 \times n_{k+1} \times \dots \times n_d}$  is  $d - 1$  dimensional. That is why one can also state that the  $k$ -mode *TVM* contracts a  $d$ -dimensional tensor along the mode  $k$  and forms a  $d - 1$  dimensional tensor. Note that the operation  $\mathcal{Y} = \mathcal{A} \times_k \mathbf{x}$  is special from the computational point of view. The size of one of its inputs,  $\mathcal{A}$ , is much greater than the other input,  $\mathbf{x}$ .

We have proposed [9] a blocking approach for obtaining efficient, mode-oblivious tensor computations by investigating the case of tensor–vector multiplication. Ballard et al. [2] investigate the communication requirements of a well-known operation called MTTKRP and discuss a blocking approach. MTTKRP is usually formulated by matrix–matrix multiplication using BLAS libraries. Earlier approaches to this and related operations unfold the tensor (reorganizes the whole tensor in the memory) and carry out the overall operation using a single matrix–matrix multiplication [5]. Li et al. [6] instead propose a parallel loop-based algorithm: a loop of the BLAS3 kernels which operate in-place on parts of the tensor such that no unfolding is required. Kjolstad et al. [4] propose The Tensor Algebra Compiler (taco) for tensor computations. Given a tensor algebraic expression, mixing tensors of different dimensions and format (sparse or dense), Taco generates a code, mixing dense and sparse storages for different modes of a tensor according to the operands of the expression. Tensor–tensor multiplication, or contraction, has received considerable attention. This operation is the most general form of the multiplication operation in (multi)linear algebra. CTF [10], TBLIS [7], and GETT [11] are recent libraries carrying out this operation based on principles and lessons learned for high performance matrix–matrix multiplication. Apart from not explicitly considering *TVM*, these do not adapt the tensor layout. As a consequence, they all require transpositions, one way or another. Our *TVM* routines address a special case of *TVM* which is a special case of *TTM*, based on our earlier work [9].

After listing the notation in Section 2, we provide a background on blocking algorithms we proposed earlier for high performance in single core execution. Section 3 contains *TVM* algorithms whose analyses are presented in Section 4. Section 5 contains experiments on up to 8-socket 120 core systems.

## 2 Notation and background

### 2.1 Notation

We use mostly standard notation [5].  $\mathcal{A}$  is an order- $d$  or a  $d$ -dimensional tensor.  $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$  has size  $n_k$  in mode  $k \in \{1, \dots, d\}$ .  $\mathcal{Y}$  is a  $(d-1)$ -dimensional tensor obtained by multiplying  $\mathcal{A}$  along a given mode  $k$  with a suitably sized vector  $\mathbf{x}$ . Matrices are represented using boldface capital letters; vectors using boldface lowercase letters; and elements in them are represented by lowercase letters with subscripts for each dimension. When a subtensor, matrix, vector, or an element of a higher order object is referred, we retain the name of the parent object. For example  $a_{i,j,k}$  is an element of the tensor  $\mathcal{A}$ . We use Matlab column notation for denoting all indices in a mode. For  $k \in \{1, \dots, d\}$ , we use  $I_k = \{1, \dots, n_k\}$  to denote the index set for the mode  $k$ . We also use  $n = \prod_{i=1}^d n_i$  to refer to the total number of elements in  $\mathcal{A}$ . Likewise,  $I = I_1 \times I_1 \times \dots \times I_d$  is the Cartesian product of all index sets. The elements of  $I$  are marked with boldface letters  $\mathbf{i}$  and  $\mathbf{j}$ . For example,  $a_{\mathbf{i}}$  is an element of  $\mathcal{A}$  whose indices are  $\mathbf{i} = i_1, \dots, i_d$ . A mode- $k$  fiber  $\mathbf{a}_{i_1, i_2, \dots, i_d}$  in a tensor is obtained by fixing the indices in all modes except the mode- $k$ . A hyperslice is obtained by fixing one of the indices and varying all others. In third order tensors, a hyperslice become a slice and is therefore a matrix. For example  $\mathbf{A}_{i, :, :}$  is the  $i$ th mode-1 slice of  $\mathcal{A}$ .

The total number of threads an algorithm employs is  $p$ , which need not be equal to the number of cores a given machine holds; it can be less when considering strong scalability, and it can be more when exploring the use of hyperthreads. Let  $P = \{1, \dots, p\}$  be the set of all possible thread IDs.

All symbols are summarized in Table 1.

### 2.2 Sequential *TVM* and dense tensor memory layouts

We focus on shared-memory parallel algorithms for computing a  $k$ -mode *TVM*. We parallelize the *TVM* by distributing the input tensor between the physical cores of a shared-memory machine while relying on the sequential tensor data layouts and the associated *TVM* kernels as discussed in our earlier work [9], which we summarize below.

A layout  $\rho$  maps tensor elements onto an array of size  $n = \prod_{i=1}^d n_i$ . Most commonly, dense tensors are stored as multidimensional arrays. For instance, a matrix can be stored in two different ways (row-major and column-major), while  $d$ -dimensional tensors can be stored in  $d!$  different ways, giving  $d!$  unfoldings [5]. Let  $\rho_\pi(\mathcal{A})$  be a layout based on some ordering (permutation) of modes  $\pi$  of  $(1, \dots, d)$  such that

$$\rho_\pi(\mathcal{A}) : (i_1, \dots, i_d) \mapsto \sum_{k=1}^d \left( i_{\pi_k} \prod_{j=k+1}^d n_{\pi_j} \right),$$

with the convention that  $\prod_{j=k+1}^d = 1$  for  $k = d$ . The regularity of this layout allows such tensors be processed using BLAS in a loop without explicit tensor unfoldings. Let  $\rho_Z(\mathcal{A})$  be a **Morton**

$\mathcal{A}, \mathcal{Y}$	An input and output tensor, respectively
$\mathbf{x}$	An input vector
$d$	The order of $\mathcal{A}$ and one plus the order of $\mathcal{Y}$
$n_i$	The size of $\mathcal{A}$ in the $i$ -th dimension
$n$	The number of elements in $\mathcal{A}$
$I_i$	The index set corresponding to $n_i$
$I$	The Cartesian product of all $I_i$
$\mathbf{i}$ and $\mathbf{j}$	Members of $I$
$k$	The mode of a <i>TVM</i> computation
$b$	Individual block size of tensors blocked using hypercubes
$p_s$	The number of sockets
$p_t$	The number of threads per socket
$p$	The total number of threads $p_s p_t$
$s$	The ID of a given thread
$P$	The set of all possible thread IDs
$\pi$	Any distribution of $\mathcal{A}$
$\pi_{1D}$	A 1D block distribution
$b_{1D}(j)$	The block size of a load-balanced 1D distribution of $j$ elements
$\rho_\pi$	A unfolded layout for storing a tensor
$\rho_Z$	A Morton-order for storing a tensor
$\rho_Z \rho_\pi$	Blocked tensor layout with a Morton block order
$m_s$	The number of fibers in each slice under a 1D distribution
$\mathcal{A}_s, \mathcal{Y}_s$	Thread-local versions of $\mathcal{A}, \mathcal{Y}$
$M_s$	Memory requirement at thread $s$
$S$	Number of barriers required
$U_{s,i}$	Intra-socket data movement for thread $s$ in phase $i$
$V_{s,i}$	Inter-socket data movement for thread $s$ in phase $i$
$W_{s,i}$	Work (in flops) local to thread $s$ in phase $i$
$g, h$	Intra- and inter-socket throughput (seconds per word)
$r$	Thread compute speed (seconds per flop)
$L$	The time for a barrier to complete (in seconds)
$T_{\text{seq}}$	Best sequential time
$T(n, p)$	Time taken for parallel execution on $n$ elements and $p$ threads
$O(n, p)$	The overhead of the parallel algorithm
$E(n, p)$	The efficiency of the parallel algorithm

Table 1 – Notation used throughout the paper

**layout** defined by the space-filling Morton order [8]. The Morton order is defined recursively, where at every step the covered space is subdivided into two within every dimension; for 2D planar areas this creates four cells, while for 3D it creates eight cells. In every two dimensions the order between cells is given by a (possibly rotated) Z-shape. Let  $w$  be the number of bits used to represent a single coordinate, and let  $i_k = (l_1^k \dots l_w^k)_2$  for  $k \in \{1, \dots, d\}$  be the bit representation of each coordinate. The Morton order in  $d$  dimensions  $\rho_Z(\mathcal{A})$  can then be defined as

$$\rho_Z(\mathcal{A}) : (i_1, \dots, i_d) \mapsto (l_1^1 l_1^2 \dots l_1^d l_2^1 l_2^2 \dots l_2^d \dots l_w^1 l_w^2 \dots l_w^d)_2.$$

Such layout improves performance on systems with multi-level caches due to the locality preserving properties of the Morton order. However, as such  $\rho_Z(\mathcal{A})$  is an irregular layout, and thus unsuitable for processing with BLAS routines.

Blocking is a well-known technique improving the locality of elements. A blocked tensor consists of a total of  $\prod_{i=1}^d a_i$  blocks  $\mathcal{A}_j \in \mathbb{R}^{b_1 \times \dots \times b_d}$  where  $n_k = a_k b_k$  for all  $k \in \{1, \dots, d\}$ . We introduced [9] a  $\rho_Z \rho_\pi$  blocked layout, which organizes elements into blocks, and stores blocks consecutively using the two layouts:  $\rho_\pi$  to store the elements of individual blocks, and  $\rho_Z$  to order the blocks in memory. By using the regular layout at the lower level, we are still able to employ BLAS routines for processing the individual blocks, while benefiting from the properties of the Morton order (increased data reused between blocks, and mode-obliviousness performance). The detailed description of sequential algorithms using these layouts can be found in our earlier work [9].

### 3 Shared-memory parallel TVM algorithms

#### 3.1 Shared-memory parallelization

We assume a shared-memory architecture consisting of  $p_s$  connected processors. Each processor supports running  $p_t$  threads for a total of  $p = p_s p_t$  threads. Each processor has local memory to which access is faster than accessing remote memory areas. We assume threads taking part in a parallel TVM computation are *pinned* to a specific core, meaning that threads will not move from one core to another while a TVM is executed.

A pinned thread has a notion of local memory: namely, all addresses that are mapped to the memory controlled by the processor the thread is pinned to. This gives rise to two distinct modes of use for shared memory areas: the *explicit* versus *interleaved* modes. If a thread allocates, initialises, and remains the only thread using this memory area, we dub its use explicit. In contrast, if the memory pages associated with an area cycle through all available memories, then the use is called interleaved. If a memory area is accessed by all threads in a uniformly random fashion, then it is advisable to interleave to achieve high throughput.

We will consider both parallelizations of for-loops as well as parallelizations following the Single Program, Multiple Data (SPMD) paradigm. In the former, we identify a for-loop where each iterant can be processed concurrently without causing race conditions. Such a for-loop can either be cut *statically* or *dynamically*; the former cuts a loop of size  $n$  in exactly  $p$  parts and has each thread execute a unique part of the loop, while the latter typically employs a form of work stealing to assign parts of the loop to threads. In both modes, we assume that one does not explicitly control which thread will execute which part of the loop.

At the finest level, a distribution of an order- $d$  tensor of size  $n_1 \times \dots \times n_d$  over  $p$  processors is given by a map  $\pi : I \rightarrow \{1, \dots, p\}$ . A one-dimensional distribution has  $\pi(i_1, \dots, i_d) = \pi_{1D}(i_1)$ , where

$\pi_{1D} : I_1 \rightarrow \{1, \dots, p\}$  distributes a tensor across the mode-1 hyperslices. That is all elements in  $\mathcal{A}_{i_1, \dots, i_d}$  are mapped to the same processor. Let  $m_s = |\pi_{1D}^{-1}(s)|$  count the number of  $d-1$ -dimensional hyperslices of  $\mathcal{A}$  each part contains. We demand that a 1D distribution be *load-balanced*:

$$\max_{s \in P} m_s - \min_{s \in P} m_s \leq 1.$$

The regular 1D block distribution,

$$\pi_{1D}(i) = \lfloor (i-1)/b_{1D}(n_1) \rfloor \text{ with } \textit{block size } b_{1D}(j) = \lceil j/p \rceil, \quad (1)$$

attains this. The choices to distribute over the first mode and to use a block distribution are without loss of generality. The dimensions of  $\mathcal{A}$  and their fibers could be permuted to fit any other load-balanced 1D distribution. In the remainder of the text we furthermore assume  $n_1 \geq p$ ; for smaller sizes the dimensions could be reordered or fewer threads be used.

### 3.2 Baseline: loopedBLAS

The *TVM* operation could naively be written using  $d$  nested for-loops, where the outermost loop that does not equal the mode  $k$  of the *TVM* is executed concurrently using a dynamic schedule. For a better performing parallel baseline, however, we instead assume  $\mathcal{A}$  and  $\mathcal{Y}$  have the default unfolded layout for which the  $d-k$  inner for-loops correspond to a dense matrix–vector multiplication if  $k < d$ ; we can thus write the parallel *TVM* as a loop over BLAS-2 calls, and use highly optimized libraries to execute those. For  $k = d$  the naively nested for-loops actually correspond to a dense transposed-matrix–vector multiplication, which is a standard BLAS-2 call as well. Note that the matrices involved with these BLAS-2 calls generally are rectangular.

We execute the loop over the BLAS-2 calls in parallel using a dynamic schedule. For  $k = d$  (and for smaller tensors) this may not expose enough parallelism to make use of all available threads. We use any such left-over threads to parallelize the BLAS-2 calls themselves, while taking care that threads collaborating on the same BLAS-2 call are pinned close to each other to exploit shared caches as much as possible. Since all threads access both the input tensor and input vector, and since it cannot be predicted which thread accesses which part of the output tensor, all memory areas corresponding to  $\mathcal{A}$ ,  $\mathcal{Y}$ , and  $\mathbf{x}$  must be interleaved. We dub the thus resulting algorithm *loopedBLAS*, which is for  $p = 1$  equivalent to *tvLooped* in our earlier work [9].

### 3.3 1D *TVM* algorithms

We present a family of algorithms assuming a 1D distribution of the input and output tensors, thus resulting in  $p$  disjoint input tensors  $\mathcal{A}_s$  and  $p$  disjoint output tensors  $\mathcal{Y}_s$  where each of their unions correspond to  $\mathcal{A}$  and  $\mathcal{Y}$ , respectively. For all but  $k = 1$ , a parallel *TVM* in this scheme amounts to a thread-local call to a sequential *TVM* computing  $\mathcal{Y}_s = \mathcal{A}_s \times_k \mathbf{x}$ ; each thread reads from its own part of  $\mathcal{A}$  while writing to its own part of  $\mathcal{Y}$ . We may thus employ the  $\rho_Z \rho_\pi$  layout for  $\mathcal{A}_s$  and  $\mathcal{Y}_s$  and use its high-performance sequential mode-oblivious kernel; here,  $\mathbf{x}$  is allocated interleaved while  $\mathcal{A}_s$  and  $\mathcal{Y}_s$  are explicit. The global tensors  $\mathcal{A}$  and  $\mathcal{Y}$  are never materialized in shared-memory—only their distributed variants are required in memory. We expect the explicit allocation of these two largest data entities involved with the *TVM* computation to induce much better parallel efficiency compared to the *loopedBLAS* baseline where all data is interleaved.

For  $k = 1$ , the output tensor  $\mathcal{Y}$  cannot be distributed, because the size of its first dimension is 1. We define that the  $\mathcal{Y}$  are then instead subject to a 1D block distribution over mode 2, and will assume that  $n_2 \geq p$ ; this is again without loss of generality via permutation of the dimensions. Since the distributions of  $\mathcal{A}$  and  $\mathcal{Y}$  do not match in this case, communication ensues. We suggest three variants that minimise data movement, which are characterized chiefly by the number of synchronisation barriers they require: zero, one, or  $p - 1$ . Before describing these variants, we first motivate why it is sufficient to only consider one-dimensional partitionings of  $\mathcal{A}$ .

Consider a succession of *TVMs* over all  $d$  modes. Assume  $p > 1$  and any load-balanced distribution  $\pi$  of  $\mathcal{A}$  and  $\mathcal{Y}$  such that thread  $s$  has at most  $2d \left( \frac{n}{p} + 1 \right)$  work. For any  $\mathbf{i} \in I$  and corresponding  $\mathbf{j} = (i_1, \dots, i_{k-1}, 1, i_{k+1}, \dots, i_d)$ , the distribution  $\pi$  defines which thread is multiplying the input tensor element  $a_{\mathbf{i}}$  with its corresponding input vector element  $x_{i_k}$ , performs local reductions of multiplicands that should reduce to the same output tensor element  $y_{\mathbf{j}}$ , while final reductions of these are done by one or more threads in  $\pi(i_1, i_2, \dots, i_{k-1}, I_k, i_{k+1}, \dots, i_d)$ ; any thread in this set is said to *contribute* to the reduction of  $y_{\mathbf{j}}$ . We do not assume a specific reduction algorithm over threads and instead ignore the cost of final reductions altogether—we only count the minimal work involved. We assume that  $n_1 \geq p$  and  $n_i \leq n_{i+1}$  for  $i = 1, \dots, d - 1$ .

For any  $\mathbf{i} \in I$ , let  $X_{\mathbf{i}} = \pi(\{\mathbf{j} \in I \mid \bigvee_{k=1}^d i_k = j_k\})$ . If there is an  $\mathbf{i}$  for which  $|X_{\mathbf{i}}| > 1$  there must be at least one mode  $k$  for which more than one thread contribute to  $y_{\mathbf{j}}$ ,  $\mathbf{j} = (i_1, \dots, i_{k-1}, 1, i_{k+1}, \dots, i_d)$ . Note that if  $|X_{\mathbf{i}}| > 1$  for all  $\mathbf{i} \in I$ , then all elements of  $\mathcal{Y}$  will be involved with at least one reduction during one of the  $d$  *TVMs*. This amounts to either  $n/n_1$  or  $n/n_2$  reductions depending on whether they occur for  $k > 1$  or  $k = 1$ , respectively. Note that the latter is exactly the number of reductions a 1D distribution incurs for  $k = 1$  (and  $k = 1$  only), and note that  $n/n_1 \geq n/n_2$ .

Now suppose there exist  $n/n_1/p + 1$  coordinates  $\mathbf{i} \in I$  such that  $X_{\mathbf{i}} = \{s\}$ . For each such  $\mathbf{i}$  there exist  $d$  fibers that are assigned to  $s$  so that for any  $k$ -mode *TVM* thread  $s$  requires only  $2n_k$  flops to compute  $y_{i_1, \dots, i_{k-1}, 1, i_{k+1}, \dots, i_d}$  locally. Summing over all possible  $d$  modes, this corresponds to a total of  $\sum_{j=1}^d 2n_j(n/n_1/p + 1) \geq 2(dn/p + \sum_{j=1}^d n_j)$  flops executed by thread  $s$  (with any remainder work distributed over the  $p - 1$  other processes), which breaks the assumption of load balance since  $2d(p - 1) < 2dp \leq 2 \sum_{j=1}^d n_j$ . Any  $\pi$  must therefore have at most  $n/n_1/p$  coordinates  $\mathbf{i}$  with  $|X_{\mathbf{i}}| = 1$  assigned to a single thread.

In the best case, all these  $n/n_1/p$  elements of  $\mathcal{A}$  over all  $p$  threads lie in a single  $d - 1$  dimensional hyperslice perpendicular to the output tensor, thus saving exactly one  $d - 2$ -dimensional hyperslice on  $\mathcal{Y}$  from being reduced. In the best case these reductions only occur for a single mode  $j$ , thus saving  $n/(n_1 n_{\max\{j, 2\}}) \leq \frac{n}{n_1 n_2}$  reductions; a significantly lower-order factor versus the original number of reductions ( $n/n_1$  or  $n/n_2$ ). The data movement a 1D distribution attains on  $\mathbf{x}$  is  $(p - 1)n_1$  for all modes. Since this is similarly significantly less than the cost of reduction  $n/n_1 \geq pn_k$  for  $d > 2$  (or if  $n_1 = n_2$ ), we thus conclude the use of  $\pi_{1D}$  is asymptotically optimal under the above assumptions.

Note this assumed all modes of the *TVM* are equally important.

### 3.3.1 0-sync

We avoid performing a reduction on  $\mathcal{Y}$  for  $k = 1$  by storing  $\mathcal{A}$  twice; once with a 1D distribution over mode 1, another time using a 1D distribution over mode  $d$ . Although the storage requirement is doubled, data movement remains minimal while explicit reduction for  $k = 1$  is completely eliminated, since the copy with the 1D distribution over mode  $d$  can then be used without penalty.

In either case, the parallel *TVM* computation completes after a sequential thread-local *TVM*; this variant requires no barriers to resolve data dependencies.

### 3.3.2 1-sync

This variant performs an explicit reduction of the  $\mathcal{Y}_s$  for  $k = 1$  and behaves as the 0-sync variant otherwise. It requires a larger buffer for the  $\mathcal{Y}_s$  to cope with  $k = 1$ , since each thread computes a full output tensor that contains partial results only; i.e.,  $\mathcal{Y}_t = \sum_{s=1}^p (\mathcal{A}_s \times_1 \mathbf{x})_t$ , for all  $t \in P$ . Element  $\mathcal{Y}_i$ ,  $i \in I$  is reduced by thread  $\pi_{1D}(i_2)$ , and by load balance each thread has at most  $\frac{b_{1D}(n_2)n}{n_1 n_2}$  such elements. A barrier must separate the local *TVM* from the reduction phase to ensure no incomplete  $\mathcal{Y}_s$  are reduced.

### 3.3.3 Interleaved $p - 1$ -sync

This variant assumes an interleaved  $\mathcal{Y}$  in lieu of thread-local  $\mathcal{Y}_s$ . Each  $\mathcal{A}_s$  is further split into  $p$  parts, each stored thread-locally, giving rise to the input tensors  $\mathcal{A}_{s,t}$ ; this split is equal across all threads and is used to synchronise writing into the output tensor. Each thread  $s$  executes:

```

1:  $\mathcal{Y} = \mathcal{A}_{s,s} \times_k \mathbf{x}$ 
2: for  $t = 1$  to  $p$  do
3:   if  $k = 1$  then
4:     barrier
5:      $\mathcal{Y} += \mathcal{A}_{s,((s+t-1) \bmod p)+1} \times_k \mathbf{x}$ 

```

This algorithm avoids doubling the storage requirement yet still eliminates explicit reductions for  $k = 1$ , replacing reduction with synchronization. For  $k > 1$  no barriers are required since each thread writes into disjoint areas of  $\mathcal{Y}$  due to the 1D block distribution over mode 1.

### 3.3.4 Explicit $p - 1$ -sync

We split each  $\mathcal{A}_s$  into  $p$  subtensors  $\mathcal{A}_{s,t}$  according to  $\pi_{1D}$  on mode 1. The explicit  $p - 1$ -sync variant allocates all  $\mathcal{A}_{s,t}$  and  $\mathcal{Y}_s$  explicitly local to thread  $s$  and keeps  $\mathbf{x}$  interleaved. Each thread executes:

```

1: if  $k = 1$  then
2:    $\mathcal{Y}_s = \mathcal{A}_{s,t} \times_k \mathbf{x}$ 
3:   for  $t = 2$  to  $p$  do
4:     barrier
5:      $\mathcal{Y}_{(s+t-1) \bmod p+1} += \mathcal{A}_{s,t} \times_k \mathbf{x}$ 
6: else
7:   for  $t = 1$  to  $p$  do
8:      $\mathcal{Y}_s = \mathcal{A}_{s,t} \times_k \mathbf{x}$ 

```

Note that on line 5 inter-socket data movement occurs on the output tensor. To cope with cases where the output tensors act as input on successive *TVM* calls,  $\mathcal{A}$  must be split in  $p$  parts *across each dimension* while  $\mathcal{Y}$  should likewise be split in  $p^{d-1}$  parts. We emphasize a careful implementation will never allocate  $p^d$  separate subtensors.

### 3.3.5 Hybrid p-sync

This variant stores two versions of the output tensor: one interleaved  $\mathcal{Y}$ , and one thread-local  $\mathcal{Y}_s$ . It combines the best features of the interleaved and explicit  $p - 1$ -sync variants. Both  $\mathcal{A}_s$  and  $\mathcal{Y}_s$  are split in  $q = \prod_{i=2}^d p_i \geq p$  parts, splitting each object in  $p_i$  parts across mode  $i$  using a  $\pi_{1D}$  distribution. We index the resulting objects as  $\mathcal{A}_{s,t}$  and  $\mathcal{Y}_{s,t}$  which remain explicitly allocated to thread  $s$ . The input vector  $\mathbf{x}$  remains interleaved.

```

1: if  $k = 1$  then
2:    $\mathcal{Y} = \mathcal{A}_{s,s} \times_k \mathbf{x}$ 
3:   for  $t = 2$  to  $q$  do
4:     barrier
5:      $\mathcal{Y} += \mathcal{A}_{s,(t+s-1) \bmod p+1} \times_k \mathbf{x}$ 
6:   else
7:     for  $t = 1$  to  $q$  do
8:        $\mathcal{Y}_{s,t} = \mathcal{A}_{s,t} \times_k \mathbf{x}$ 

```

If this algorithm is to re-use output of a mode-0 *TVM*, then, similarly to the 0-sync variant, each thread must re-synchronize its local  $\mathcal{Y}_{s,t}$  with  $\mathcal{Y}$ . Thus unless the need explicitly arises, implementations need not distribute  $\mathcal{Y}$  over  $n_2$  as part of a mode-1 *TVM* (at the cost of interleaved data movement on  $\mathcal{Y}$ ).

### 3.4 Explicit $p - 1$ -sync

In the text we left in the middle how the explicit  $p - 1$  sync should be modified to cope with successive *TVMs* of the (worst-case) form  $(\mathcal{A} \times_0 x_0) \times_1 x_1 \dots \times_d k_d$ . This variant assumes  $\mathcal{A}$  and  $\mathcal{Y}$  are split into  $p^d$  parts  $\mathcal{A}_{s_0, s_1, \dots, s_{d-1}}$  and  $p^{d-1}$  parts  $\mathcal{Y}_{s_0, s_1, \dots, s_{d-1}}$ , respectively. All of these are explicitly used and allocated by thread  $s_0$ , in-line with  $\mathcal{A}_s$  and  $\mathcal{Y}_s$ , with each dimension distributed according to  $\pi_{1D}$ . We assume all sizes are sufficiently large; i.e.,  $n_i \geq p$ . To ease the algorithmic description, we virtually assume  $x$  is split into  $p$  parts  $x_s$  also according to  $\pi_{1D}$ , but with each part remaining interleaved. Each thread  $s_1 = s$  executes:

```

1:  $\mathcal{Y}_{s_1, s_1, \dots, s_d} = \mathcal{A}_{s_1, s_1, \dots, s_d} \times_k x_{s_k}$ 
2: for  $t = 2$  to  $p$  do
3:    $s_2 = (s_1 + t - 1) \bmod p$ 
4:   if  $k = 1$  then
5:     barrier
6:     for  $(s_3, s_4, \dots, s_d) = (1, \dots, 1)$  to  $(p, \dots, p)$  do
7:        $\mathcal{Y}_{s_1, \dots, s_{k-1}, 1, s_{k+1}, s_d} += \mathcal{A}_{s_1, \dots, s_d} \times_k x_{s_k}$ 

```

The advantage of this algorithm is that all  $\mathcal{Y}_{s,t}$  remain explicitly allocated. We require  $p$  splits across each dimension of  $\mathcal{A}$  to ensure a valid 1D distribution can always exist in case of successive tensor–vector multiplications which repeatedly contract mode 0.

There is no need to actually allocate  $p^d$  separate objects for  $\mathcal{A}$ , which would be prohibitive for larger  $p$  or  $d$ . Instead, one can maintain a two-level block layout for all but the first mode, first imposing a block size of  $\lceil n_i/p \rceil$  and applying the default block size  $b$  of the underlying  $\rho_Z \rho_\pi$  layout within each such superblock. If the superblock order is a natural layout ( $\rho_\pi$ ) all of the  $p^{d-1}$  splits of  $\mathcal{A}_{s_0, \dots}$  will be contiguous in memory local to thread  $s_0$ , causing line 7 to translate to index offsets for  $s_i$ ,  $i > 0$ . The same arguments hold for  $\mathcal{Y}$  which ensures this variant can be efficiently

implemented.

## 4 Analysis of the algorithms

We analyse the parallel *TVM* algorithms from the previous section, restricting ourselves not only to the amount of data moved during a *TVM* computation, but also consider mode-obliviousness, memory, and work. We furthermore subdivide data movement in intra-socket data movement (where cores contend for resources) as well as inter-socket data movement (where data is moved over a communication bus, instead of only to and from local memory). For quantifying data movement we assume perfect caching, meaning that all required data elements are touched exactly once. Once we quantify algorithm properties in each of these five dimensions, we consider their *iso-efficiencies* [3].

Consider the memory requirement  $M_s$  in number of words to store, and the number of barriers  $S$ . Each thread  $s$  thus executes  $S + 1$  different phases, which are numbered using integer  $i$ ,  $0 \leq i \leq S$ . For each thread and phase, let intra-socket data movement  $U_{s,i}$  and inter-socket data movement  $V_{s,i}$  be in number of words, while work  $W_{s,i}$  is in number of floating point operations (flops). These quantities fully quantify an algorithm, while the following quantifies a machine: intra-socket throughput  $g$  and the inter-socket throughput  $h$ , both in seconds per word (per socket); thread compute speed by  $r$  seconds per flop, and the time  $L$  in which a barrier completes in seconds.

Since barriers require active participation by processor cores while they also make use of communication, the time in which a given algorithm completes a *TVM* computation is given by

$$T(n, p) = \sum_{i=0}^S \left[ \max_{k \in \{1, \dots, p\}} (\max\{U_{k,i}g + V_{k,i}h, W_{k,i}r\} + L) \right]. \quad (2)$$

A perfect sequential *TVM* algorithm would complete in  $T_{\text{seq}}(n) = \max\{2nr, (n + n/n_k + n_k)g\}$  time. The *parallel overhead*  $O(n, p)$  then is  $pT(n, p) - T_{\text{seq}}(n)$  while the *parallel efficiency*  $E(n, p)$  is the *speedup* of an algorithm divided by  $p$ :

$$E(n, p) = T_{\text{seq}}(n)/pT(n, p) = 1 / \left( \frac{O(n, p)}{T_{\text{seq}}(n)} + 1 \right). \quad (3)$$

The above formula allows us to compute how fast  $n$  should grow to retain the same parallel efficiency as  $p$  increases and vice versa, giving rise to the concept of iso-efficiency. Strongly scaling algorithms have that  $O(n, p)$  is independent of  $p$  (which is unrealistic), while weakly scalable algorithms have that the ratio  $O(pn, p)/T_{\text{seq}}(pn)$  is constant; iso-efficiency instead tells us a much wider range of conditions under which the algorithm scales.

The *TVM* computation is a heavily bandwidth-bound operation. It performs  $2n$  flops on  $n + n/n_k + n_k$  data elements, and thus has arithmetic intensity equal to  $1 < \frac{2n}{n + n/n_k + n_k} < 2$  flop per word. This amounts to a heavily bandwidth-bound computation even when considering a *sequential TVM* [9]. The multi-threaded case is even more challenging, as cores on a single socket compete for the same local memory bandwidth. In our subsequent analyses we will thus ignore the computational part of the equations for  $T$ ,  $O$ , and  $E$ . We also consider memory overhead and efficiency versus the sequential memory requirement  $M_{\text{seq}} = n + \max_k (n/n_k + n_k)$  words.

## 4.1 loopedBLAS

The *loopedBLAS* variant interleaves  $\mathcal{A}$ ,  $\mathcal{Y}$ , and  $\mathbf{x}$ , storing them once while it performs  $2n$  flops to complete the *TVM*; it is thus both memory- and work-optimal. It does not include any cache-oblivious nor mode-oblivious optimizations, and requires no barrier synchronisations ( $S = 0$ ). Since all memory used is interleaved we assume their effective bandwidth is spread over  $g$  and  $h$  proportional to the number of CPU sockets  $p_s$ . Assuming a uniform work balance is achieved at run time, all  $p$  threads read  $n/p$  data from  $\mathcal{A}$ , write  $n/(pn_k)$  to  $\mathcal{Y}$ , and read the full  $n_k$  elements of  $\mathbf{x}$ . Thus  $U(s, 0) = \frac{1}{p_s}v$  and  $V(s, 0) = \frac{p_s-1}{p_s}v$  with  $v = \left(\frac{n+n/n_k}{p} + n_k\right)$  and the parallel overhead becomes

$$O(n, p) = \frac{p_s - 1}{p_s}(n + n/n_k)(g - h) + n_k((p_t - 1)g + p_t(p_s - 1)h). \quad (4)$$

We see the overhead is dominated by  $\Theta(n(g - h))$  as  $p_s$  increases, while for  $p_s = 1$  the overhead simplifies to  $\Theta(p_t n_k g)$ .

## 4.2 0-sync

This algorithm incurs  $n$  words of extra storage and thus is not memory optimal. In both cases of  $k = 1$  and  $k > 1$  the amount of work executed remains optimal at  $2n$  flops. The cache- and mode-oblivious optimisations from our earlier work are fully exploited by this algorithm, while, like *loopedBLAS*,  $S$  remains zero. Here,  $\mathcal{A}$  and  $\mathcal{Y}$  are allocated explicitly while  $\mathbf{x}$  remains interleaved; hence  $U(s, 0) = \frac{1}{p}(n + n/n_k) + \frac{1}{p_s}n_k$  and  $V(s, 0) = \frac{p_s-1}{p_s}n_k$ , resulting in

$$O(n, p) = (p_t - 1)n_k g + p_t(p_s - 1)n_k h. \quad (5)$$

This overhead is bounded by  $\Theta(pn_k h)$  for  $p_s > 1$ , a significant improvement over *loopedBLAS*.

If the output tensor  $\mathcal{Y}$  must assume a similar datastructure to  $\mathcal{A}$  for further processing, the output must also be stored twice. The minimal cost for this incurs extra overhead at  $\Theta(ng/n_k)$ ; i.e., a thread-local data copy, which remains asymptotically smaller than  $T_{\text{seq}} = \Theta(ng)$ . We do note that applications which require repeated *TVMs* such as the higher-order power method actually do *not* require this extra step; multilinearity can be exploited to perform the HOPM block-by-block [9, Section 5.6].

## 4.3 Mode-obliviousness

The *loopedBLAS* algorithm is highly sensitive to the mode  $k$  in which a *TVM* is executed, while those algorithms based on the  $\rho_Z \rho_\pi$  tensor layout are, by design, not sensitive to  $k$  [9]. The 0-sync and 1-sync variants exploit the  $\rho_Z \rho_\pi$  maximally; the thread-local tensors use a single such layout, and each thread thus behaves fully mode-oblivious.

For the  $p - 1$ -sync variants, however, each locally stored input tensor is split in several slices. Suppose mode  $k$  has  $\mathcal{A}_s$  split in  $p_k$  parts, and each part is stored using a  $\rho_Z \rho_\pi$  layout; then  $\mathcal{Y}$  incurs a *TVM*  $p_k$  times, and thus has all its elements touched  $p_k - 1$  times more than for a 0- or 1-sync variant. This hampers both cache efficiency and mode-obliviousness; hence if  $p = \prod_{i=1}^d p_i$  slices are required, we should make sure to minimize  $\max_i p_i$ —i.e.,  $p - 1$ -sync behaves optimally if the  $p_i$  are the result of an integer factorization of  $p$ .

Method	Work	Memory	Movement	Barrier	Oblivious	Implicit	Explicit	$k$
<i>loopedBLAS</i>	<b>0</b>	<b>0</b>	$p_s n(g - h)$	<b>0</b>	none	$\mathcal{A}, \mathcal{Y}, x$	-	all
0-sync	<b>0</b>	$pn$	<b><math>pn_k h</math></b>	<b>0</b>	<b>full</b>	<b>x</b>	$\mathcal{A}, \mathcal{Y}$	all
1-sync	$pn/n_k r$	<b>0</b>	$pn/n_k h$	$pL$	<b>full</b>	<b>x</b>	$\mathcal{A}, \mathcal{Y}$	1
i-sync	<b>0</b>	<b>0</b>	$p_s n/n_k(g - h)$	$p^2 L$	good	<b>x, Y</b>	$\mathcal{A}$	all
e-sync	<b>0</b>	<b>0</b>	$p_s n/n_k(g - h)$	$p^2 L$	fair	<b>x</b>	$\mathcal{A}, \mathcal{Y}$	1
h-sync	<b>0</b>	$pn/n_k$	$p_s n/n_k(g - h)$	$p^2 L$	good	<b>x, Y</b>	$\mathcal{A}, \mathcal{Y}$	1

Table 2 – Summary of overheads for each parallel shared-memory TVM algorithm, plus the allocation mode of  $\mathcal{A}, \mathcal{Y}$ , and **x**. We display the worst-case asymptotics; i.e., assuming  $p_s > 1$  and the worst-case  $k$  for non mode-oblivious algorithms. Optimal overheads are in bold.

Method	Work	Memory	Movement	Barrier	Oblivious	Implicit	Explicit	$k$
<i>loopedBLAS</i>	<b>0</b>	<b>0</b>	<b><math>pn_k g</math></b>	<b>0</b>	none	$\mathcal{A}, \mathcal{Y}, x$	-	all
0-sync	<b>0</b>	$pn$	<b><math>pn_k g</math></b>	<b>0</b>	<b>full</b>	<b>x</b>	$\mathcal{A}, \mathcal{Y}$	all
1-sync	$pn/n_k r$	<b>0</b>	$pn/n_k g$	$pL$	<b>full</b>	<b>x</b>	$\mathcal{A}, \mathcal{Y}$	1
i-sync	<b>0</b>	<b>0</b>	<b><math>pn_k g</math></b>	$p^2 L$	good	<b>x, Y</b>	$\mathcal{A}$	all
e-sync	<b>0</b>	<b>0</b>	<b><math>pn_k g</math></b>	$p^2 L$	fair	<b>x</b>	$\mathcal{A}, \mathcal{Y}$	1
h-sync	<b>0</b>	$pn/n_k$	<b><math>pn_k g</math></b>	$p^2 L$	good	<b>x, Y</b>	$\mathcal{A}, \mathcal{Y}$	1

Table 3 – Like Table 2, but assuming  $p_s = 1$ .

#### 4.4 Iso-efficiencies

Table 2 summarizes the results from this section. Note that the parallel efficiency depends solely on the ratio  $O(n, p)$  versus  $T_{\text{seq}}(n)$ ; when considering the work- and communication-optimal 1-sync algorithm, efficiency thus is proportional to  $\frac{pn_k h}{ng}$ ; i.e., the algorithm scales as long as  $p$  grows proportionally with  $n/n_k$ . For *loopedBLAS*, efficiency is proportional to  $p_s(h/g - 1)$ . Since  $h < g$  this drops as the number of sockets increases; *loopedBLAS* does not scale at all.

For 1-sync, iso-efficiency is attained whenever  $\frac{p}{n_k}(r + h) + \frac{p}{n}L$  is constant; i.e.,  $p$  should grow linearly with  $n$  when computation is latency-bound, and with linearly with  $n_k$  otherwise. Both are unfavorable. The i-sync, e-sync, and h-sync algorithms attain iso-efficiency when  $p_s$  grows linearly with  $n_k$  and  $p^2$  grows linearly with  $n$ , also unfavorable.

Table 3 summarizes the overheads of each algorithms assuming  $p_s = 1$ .

#### 4.5 1-sync

This variant requires the  $\mathcal{Y}_s$  are all of size  $v = n/n_k$  instead of  $v/p$  elements, which constitutes a memory overhead of  $(p - 1)v$ . It benefits from the same cache- and mode-oblivious properties as the 0-sync variant, and achieves the same overhead (Eq. 5) when  $k > 1$ . For  $k = 1$ , however, we must account for the reduction phase on the  $\mathcal{Y}_s$  and for the barrier that precedes it. Reduction proceeds with minimal cost by having each thread reduce  $v/p$  elements corresponding to those elements it should locally store, resulting in an overhead for a 0-mode TVM of

$$\begin{aligned}
 O(n, p) = & (p_t - 1)n_k g + p_t(p_s - 1)n_k h + \\
 & p_t v g + p_t(p_s - 1)v h + (p - 1)vr + pL.
 \end{aligned} \tag{6}$$

This extra overhead is proportional to  $pv$  for both memory movement and flops.

#### 4.6 Interleaved $p - 1$ -sync

This variant remains both memory and work optimal, while it incurs  $p - 1$  barriers for  $k = 1$ . Accesses to  $\mathcal{A}_{s,p}$  remain explicit while all others are interleaved. Let  $v = n/(pn_k) + n_k$ . Then  $\sum_i U(s, i) = n/p + \frac{1}{p_s}v$  and  $\sum_i \frac{p_s-1}{p_s}v$  leading to

$$O(n, p) = \frac{p_s - 1}{p_s} n/n_k(g - h) + (p_t - 1)n_k g + p_t(p_s - 1)n_k h, \quad (7)$$

which increases with  $p(p - 1)L$  if  $k = 1$ . This parallel overhead is a significant increase over that of the 0-sync variant (Eq. 5) if its output is not doubly stored, and equivalent otherwise. It still improves significantly over *loopedBLAS* (compare Eq. 4).

Mode-obliviousness is affected adversely by splitting  $\mathcal{A}_s$  in  $p$  parts since reading from  $\mathbf{x}$  and writing to  $\mathcal{Y}$  now only partially follow a Morton order.

#### 4.7 Explicit $p - 1$ -sync

The interleaved  $p - 1$ -sync variant requires interleaved storage of  $\mathcal{Y}$  which causes significant overhead since writing output requires inter-process data movement. This explicit variant remains work- and memory-optimal but reduces the parallel overhead to

$$O(n, p) = (p_t - 1)n_k g + p_t(p_s - 1)n_k h + \begin{cases} \frac{p_s-1}{p_s} n/n_k(g - h) + p(p - 1)L, & \text{if } k = 1 \\ 0, & \text{otherwise.} \end{cases}$$

The effect of Morton ordering for cache-obliviousness is affected more strongly than for interleaved  $p - 1$ -sync, but all modes are affected equally here. This explicit variant improves on the interleaved one in that inter-socket communication related to  $\mathcal{Y}$  is now only incurred for  $k = 1$ . Compared to the 0-sync variant, this  $p - 1$ -sync variant trades synchronisation and inter-socket communication for enhanced memory efficiency.

## 5 Experiments

We run our experiments on a number of different Intel Ivy Bridge nodes with different specifications summarized in Table 4. As we do not use hyperthreading, we limit the algorithms to use at most  $p/2$  threads equal the number of cores (each core supports 2 hyperthreads). We measure the maximum bandwidth of the systems using several variants of the STREAM benchmark, reporting the maximum measured performance only.

The system uses CentOS 7 with Linux kernel 3.10.0 and software is compiled using GCC version 6.1. We use Intel MKL version 2018.2.199 for *loopedBLAS*. We also run with LIBXSMM version 1.9-864 for algorithms based on blocked layouts (0- and  $i$ -sync), and retain only the result with the library which runs faster of the two. To benchmark a kernel, we conduct 10 experiments for each combination of dimension, mode, and algorithm.

To illustrate and experiment with the various possible trade offs in parallel *TVM*, we implemented the baseline synchronization-optimal *loopedBLAS*, the work- and communication-optimal

Node	CPU (clock speed)	$p_s$	$p_t$	$p$	Memory size (clock speed)	Bandwidth	
						STREAM	Theoretical
1	E5-2690 v2 (3 GHz)	2	20	40	256 GB (1600 MHz)	76.7 GB/s	95.37 GB/s
2	E7-4890 v2 (2.8 GHz)	4	30	120	512 GB (1333 MHz)	133.6 GB/s	158.91 GB/s
3	E7-8890 v2 (2.8 GHz)	8	30	240	2048 GB (1333 MHz)	441.9 GB/s	635.62 GB/s

Table 4 – An overview of machine configurations used. Memory runs in quad channel mode on nodes 1,2, and 3, and in octa-channel on node 4. Each processor has 32 KB of L1 cache memory per core, 256 KB of L2 cache memory per core, and  $1.25p_t$  MB of L3 cache memory amongst the cores.

$d$ / Node	1	2	3
2	$45600 \times 45600$ (15.49)	$68400 \times 68400$ (34.86)	$136800 \times 136800$ (139.43)
3	$1360 \times 1360 \times 1360$ (18.74)	$4080 \times 680 \times 4080$ (84.34)	$4080 \times 680 \times 4080$ (84.34)
4	$440 \times 110 \times 88 \times 440$ (13.96)	$1320 \times 110 \times 132 \times 720$ (102.81)	$1440 \times 110 \times 66 \times 1440$ (112.16)
5	$240 \times 60 \times 36 \times 24 \times 240$ (22.25)	$720 \times 60 \times 36 \times 24 \times 360$ (100.11)	$720 \times 50 \times 36 \times 20 \times 720$ (139.05)

Table 5 – Table of tensor sizes  $n_1 \times \dots \times n_d$  per tensor-order  $d$  and the node as given in Table 4. The exact size in GBs is given in parentheses.

0-sync variant, and the work-optimal  $h$ -sync variant. and investigate their performance and mode-obliviousness. We measure algorithmic performance using the formula for effective bandwidth (GB/s). We benchmark tensors of order-two up to order-5. We choose  $n$  such that the combined input and output memory areas during a single  $TVM$  call have a combined size of at least 10 GBs; The exact array of tensor sizes and block sizes are given in Table 5, and Table 6. respectively. The block sizes selected ensure that computing a  $TVM$  on such a block fits L3 cache. This combination of tensor and block sizes ensures all algorithms run with perfect load balance and without requiring any padding of blocks; to ensure this, we chose block sizes that correspond to 0.5–1 MB of our L3 cache; note that for our parallel  $TVM$  variants, the Morton order ensures the remainder cache remains obviously well-used. We additionally kept the sizes of tensors equal through all pairs of  $(d, p_s)$ , which enables comparison of different algorithms within the same  $d$  and  $p_s$ .

$d$ / Node	1	2	3
2	$570 \times 570$	$570 \times 570$	$570 \times 570$
3	$68 \times 68 \times 68$	$68 \times 68 \times 68$	$34 \times 68 \times 34$
4	$22 \times 22 \times 22 \times 22$	$22 \times 22 \times 22 \times 12$	$12 \times 22 \times 22 \times 12$
5	$12 \times 12 \times 12 \times 12 \times 12$	$12 \times 12 \times 12 \times 12 \times 6$	$6 \times 10 \times 12 \times 10 \times 6$

Table 6 – Table of block sizes  $b_1 \times \dots \times b_d$  per tensor-order  $d$  and the node as given in Table 4. Sizes are chosen such that all elements of a single block can be stored in L3 cache.

$d$	Average performance			Sample stddev.		
	<i>loopedBLAS</i>	0-sync	<i>h</i> -sync	<i>loopedBLAS</i>	0-sync	<i>h</i> -sync
2	40.23	42.28	<b>42.54</b>	0.63	<b>0.55</b>	0.65
3	36.43	39.34	<b>39.87</b>	24.93	2.55	<b>2.50</b>
4	37.63	39.02	<b>39.05</b>	21.29	<b>4.35</b>	4.40
5	34.56	36.53	<b>36.65</b>	22.43	5.14	<b>4.26</b>

Table 7 – Average effective bandwidth (in GB/s) and relative standard deviation (in %, versus the average bandwidth) over all possible  $k \in \{1, \dots, d\}$  of algorithms running on a single processor (Node 1). The highest bandwidth and lowest standard deviation for each  $d$  are stated in **bold**.

algorithm / $p_s$	Sample stddev.			Average performance		
	2	4	8	2	4	8
<i>loopedBLAS</i>	1.74	17.15	3.72	68.52	50.68	9.68
0-sync	<b>0.03</b>	<b>0.10</b>	<b>0.34</b>	<b>84.19</b>	<b>150.82</b>	<b>492.32</b>
<i>h</i> -sync	0.12	0.11	0.74	84.17	150.39	487.38

Table 8 – Average effective bandwidth (in GB/s) and relative standard deviation (in % of average) over all possible  $k \in \{1, \dots, d\}$  of order-2 tensors of algorithms executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3). The highest bandwidth and lowest standard deviation for different  $d$  are stated in **bold**.

## 5.1 Single-socket results

Table 7 shows the experimental results for the single-socket of Node 1. Note that it drops to half of the peak numbers measured in Table 4. Note that as there is no inter-socket communication, all memory regions are exceptionally allocated locally for intra-socket experiments.

As the *loopedBLAS* algorithm relies on the unfold storage, whose structure does require a loop over subtensors for modes 1 and  $d$ ; thus, no for-loop parallelisation is possible for these modes and the algorithm employs the internal MKL parallelization. Thus, the Table shows its performance is highly mode-dependent, and that the algorithms based on blocked  $\rho_Z\rho_\pi$ -storage perform faster than the *loopedBLAS* algorithm. The block Morton-order storage transfers the mode-obliviousness to parallel *TVMs* (the standard deviation oscillates within 1 %), as the Morton-order induces mode-oblivious behavior on each core.

## 5.2 Inter-socket results

Table 8, 9, 10, and 11 show the parallel-*TVM* results on machines with different numbers of sockets for tensors of order-2, 3, 4, and 5, respectively. These runtime results show a lack of scalability of *loopedBLAS*. This is due to the data structures being interleaved 4.1 instead of making use of a 1D distribution. Interleaving or not only matters for multi-socket results, but since Table 7 conclusively shows that approaches based on our  $\rho_Z\rho_\pi$  storage remain superior on single sockets, we may conclude our approach is superior at all scales.

The performance drops slightly when  $d$  increases for all variants. This is inherent to the BLAS libraries handling matrices with a lower row-to-column ratio better than tall-skinny or short-wide

algorithm / $p_s$	Sample stddev.			Average performance		
	2	4	8	2	4	8
<i>loopedBLAS</i>	9.57	16.52	23.05	63.89	55.68	13.66
0-sync	2.80	<b>1.38</b>	<b>3.42</b>	<b>77.06</b>	<b>145.07</b>	<b>467.31</b>
<i>h</i> -sync	<b>1.90</b>	3.86	6.56	76.27	143.17	441.65

Table 9 – Average effective bandwidth (in GB/s) and relative standard deviation (in % of average) over all possible  $k \in \{1, \dots, d\}$  of order-3 tensors of algorithms executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3). The highest bandwidth and lowest standard deviation for different  $d$  are stated in **bold**.

algorithm / $p_s$	Sample stddev.			Average performance		
	2	4	8	2	4	8
<i>loopedBLAS</i>	16.99	23.43	15.45	61.60	47.73	12.59
0-sync	<b>2.84</b>	<b>2.44</b>	<b>1.88</b>	<b>77.12</b>	<b>138.54</b>	<b>446.01</b>
<i>h</i> -sync	3.67	5.47	9.98	76.82	137.79	424.85

Table 10 – Average effective bandwidth (in GB/s) and relative standard deviation (in % of average) over all possible  $k \in \{1, \dots, d\}$  of order-4 tensors of algorithms executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3). The highest bandwidth and lowest standard deviation for different  $d$  are stated in **bold**.

algorithm / $p_s$	Sample stddev.			Average performance		
	2	4	8	2	4	8
<i>loopedBLAS</i>	15.37	19.70	32.03	56.11	54.04	12.43
0-sync	<b>3.47</b>	<b>5.01</b>	<b>5.02</b>	<b>71.71</b>	<b>129.80</b>	<b>421.98</b>
<i>h</i> -sync	4.17	9.37	14.83	71.65	129.60	397.25

Table 11 – Average effective bandwidth (in GB/s) and relative standard deviation (in % of average) over all possible  $k \in \{1, \dots, d\}$  of order-5 tensors of algorithms executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3). The highest bandwidth and lowest standard deviation for different  $d$  are stated in **bold**.

algorithm / $p_s$	2	4	8
<i>loopedBLAS</i>	0.81	0.31	0.02
0-sync	0.99	0.93	0.98
<i>h-sync</i>	0.99	0.93	0.97

Table 12 – Parallel efficiency of algorithms on order-2 tensors executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3), calculated against the single-socket runtime on a given node of *h-sync* algorithm on the same problem size and tensor order.

algorithm / $p_s$	2	4	8
<i>loopedBLAS</i>	0.80	0.34	0.03
0-sync	0.97	0.88	0.96
<i>h-sync</i>	0.96	0.87	0.91

Table 13 – Parallel efficiency of algorithms on order-3 tensors executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3), calculated against the single-socket runtime on a given node of *h-sync* algorithm on the same problem size and tensor order.

matrices [9]—and this ratio increases when processing higher-order tensors.

For first-mode *TVMs*, the 0-sync algorithm slightly outperforms the *h-sync*, while they achieve almost identical performance for all the other modes. The cause lies with the 0-sync not requiring any synchronization for  $k = 1$ , and the effect is the 0-sync achieves the lowest standard deviation. Our measured performances are within the impressive range of 75–88%, 81–95%, and 66–77% of theoretical peak performance for node 1, 2, and 3, respectively.

Tables 12, 13, 14, and 15 display the parallel efficiency versus the performance of the *h-sync* on a single socket. Each node takes its own baseline since the tensor sizes differ between nodes as per Table 5; one can thus only compare parallel efficiencies over the *columns* of these tables, and cannot compare rows; we compare algorithms, and do not investigate inter-socket scalability.

The astute reader will note parallel efficiencies larger than one; these are commonly due to cache-effects, in this case likely output tensors that fit in the combined cache of eight CPUs, but did not fit in cache of a single CPU. These tests conclusively show that both 0- and *h-sync* algorithms scale significantly better than *loopedBLAS* for  $p_s > 1$ , resulting in up to 35x higher efficiencies (for order-4 tensors on node 3).

## 6 Conclusions

We investigate the tensor–vector multiplication operation on shared-memory multicore machines. Building on an earlier work, where we developed blocked and mode-oblivious layouts for tensors, we here explore the design space of parallel shared-memory algorithms based on this same mode-oblivious layout, and propose several candidate algorithms. After analyzing those for work, memory, intra- and inter-socket data movement, the number of barriers, and mode obliviousness, we choose to implement two of them. These algorithms, called 0-sync and *h-sync*, deliver close to peak

algorithm / $p_s$	2	4	8
<i>loopedBLAS</i>	0.79	0.28	0.03
0-sync	0.99	0.83	1.05
<i>h-sync</i>	0.98	0.82	1.00

Table 14 – Parallel efficiency of algorithms on order-4 tensors executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3), calculated against the single-socket runtime on a given node of *h-sync* algorithm on the same problem size and tensor order.

algorithm / $p_s$	2	4	8
<i>loopedBLAS</i>	0.77	0.32	0.05
0-sync	0.98	0.76	1.53
<i>h-sync</i>	0.98	0.76	1.44

Table 15 – Parallel efficiency of algorithms on order-5 tensors executed on different nodes (2 socket node 1, 4 socket node 2, and 8-socket node 3), calculated against the single-socket runtime on a given node of *h-sync* algorithm on the same problem size and tensor order.

performance on up four different systems, with 1, 2, 4, and 8 sockets, and surpass a baseline algorithm based on looped BLAS calls that we optimized.

For future work, we plan to investigate the use of the proposed algorithms in distributed memory systems. All proposed variants should work well, after modifying them to use explicit broadcasts of the input vector and/or explicit reductions on (parts of) the output tensor. While additional buffer spaces may be required, we expect that the other shared-memory cost analyses will naturally transfer to the distributed-memory case, and (thus) favor the 0-sync variant for speed and the *h-sync* variant when memory

## References

- [1] B. W. Bader and T. G. Kolda. Algorithm 862: MATLAB tensor classes for fast algorithm prototyping. *ACM TOMS*, 32(4):635–653, December 2006.
- [2] G. Ballard, N. Knight, and K. Rouse. Communication lower bounds for matricized tensor times Khatri-Rao product. In *2018 IPDPS*, pages 557–567. IEEE, 2018.
- [3] A. Y. Grama, A. Gupta, and V. Kumar. Isoefficiency: Measuring the scalability of parallel algorithms and architectures. *IEEE Parallel & Distributed Technology: Systems & Applications*, 1(3):12–21, 1993.
- [4] F. Kjolstad, S. Kamil, S. Chou, D. Lugato, and S. Amarasinghe. The Tensor Algebra Compiler. *Proc. ACM Program. Lang.*, 1(OOPSLA):77:1–77:29, Oct. 2017. ISSN 2475-1421.
- [5] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, September 2009.

- 
- [6] J. Li, C. Battaglini, I. Perros, J. Sun, and R. Vuduc. An input-adaptive and in-place approach to dense tensor-times-matrix multiply. In *SC'15*, pages 76:1–76:12, 2015.
  - [7] D. Matthews. High-performance tensor contraction without transposition. *SIAM Journal on Scientific Computing*, 40(1):C1–C24, 2018.
  - [8] G. M. Morton. A computer oriented geodetic data base and a new technique in file sequencing. 1966.
  - [9] F. Pawlowski, B. Uçar, and A. N. Yzelman. A multi-dimensional Morton-ordered block storage for mode-oblivious tensor computations. *Journal of Computational Science*, 2019. ISSN 1877-7503. doi: <https://doi.org/10.1016/j.jocs.2019.02.007>.
  - [10] E. Solomonik, D. Matthews, J. R. Hammond, J. F. Stanton, and J. Demmel. A massively parallel tensor contraction framework for coupled-cluster computations. *Journal of Parallel and Distributed Computing*, 74(12):3176–3190, 2014.
  - [11] P. Springer and P. Bientinesi. Design of a high-performance gemm-like tensor–tensor multiplication. *ACM TOMS*, 44(3):28:1–28:29, 2018.



**RESEARCH CENTRE  
GRENOBLE – RHÔNE-ALPES**

Inovallée  
655 avenue de l'Europe Montbonnot  
38334 Saint Ismier Cedex

Publisher  
Inria  
Domaine de Voluceau - Rocquencourt  
BP 105 - 78153 Le Chesnay Cedex  
[inria.fr](http://inria.fr)

ISSN 0249-6399