

On Optimizing Distributed Tucker Decomposition for Dense Tensors

Venkatesan T. Chakaravarthy, Jee W. Choi, Douglas J. Joseph, Xing Liu, Prakash Murali, Yogish Sabharwal, and Dheeraj Sreedhar

IBM Research

{*vechakra,prakmura,ysabharwal,dhsreedh*}@in.ibm.com
{*jwchoi,djoseph,xliu*}@us.ibm.com

Abstract

The Tucker decomposition expresses a given tensor as the product of a small core tensor and a set of factor matrices. Apart from providing data compression, the construction is useful in performing analysis such as principal component analysis (PCA) and finds applications in diverse domains such as signal processing, computer vision and text analytics. Our objective is to develop an efficient distributed implementation for the case of dense tensors. The implementation is based on the HOOI (Higher Order Orthogonal Iterator) procedure, wherein the tensor-times-matrix product forms the core routine. Prior work have proposed heuristics for reducing the computational load and communication volume incurred by the routine. We study the two metrics in a formal and systematic manner, and design strategies that are optimal under the two fundamental metrics. Our experimental evaluation on a large benchmark of tensors shows that the optimal strategies provide significant reduction in load and volume compared to prior heuristics, and provide up to 7x speed-up in the overall running time.

1 Introduction

Tensors are the higher dimensional analogues of matrices. While matrices represent two-dimensional data, tensors are useful in representing data in three or higher dimensions. Tensors have been studied extensively via generalizing concepts pertaining to matrices. The Tucker decomposition [1] is a prominent construction that extends the singular value decomposition (SVD) to the setting of tensors. Given an N -dimensional tensor \mathcal{T} , the decomposition approximately expresses the tensor as the product of a small N -dimensional core tensor \mathcal{G} and a set of N factor matrices, one along each dimension (or *mode*); see Figure 1 for an illustration. The core is much smaller than the original tensor leading to data compression. Prior work [2] has reported compression rates to the tune of 5000 on large real tensors. Apart from data compression, the decomposition is also useful in analysis such as PCA, and finds applications in diverse domains such as computer vision [3] and signal




Figure 1: Illustration for Tucker decomposition on a 3-dimensional tensor: \mathcal{T} - input tensor, \mathcal{G} - core tensor, \mathbf{F}_n - factor matrices.

processing [4]. A detailed discussion on the topic can be found in the excellent survey, by Kolda and Bader [5].

Tucker decomposition has been well-studied in sequential, shared memory and distributed settings for both dense and sparse tensors (e.g., [2, 6, 7]). Our objective is to develop an optimized implementation for dense tensors on distributed memory systems. The implementation is based on the popular STHOSVD/HOOI procedures. The STHOSVD (Sequentially Truncated Higher Order SVD) [8] is used to produce an initial decomposition. The HOOI (Higher Order Orthogonal Iterator) [9] procedure transforms any given decomposition to a new decomposition with the same core size, but with reduced error. The procedure is applied iteratively so as to reduce the error monotonically across the iterations. We focus on the latter HOOI procedure which is invoked multiple times. The tensor-times-matrix product (TTM) component forms the core module of the procedure. Prior work has proposed heuristic schemes for reducing the computational load and communication volume of the component. Our objective is to enhance the performance by constructing schemes which are *optimal* in these two fundamental metrics.

Prior Work

Heuristics for computational load: The TTM component comprises of a set of tensor-times-matrix multiplication operations. Based on the observation that the operations can be rearranged and reused in multiple ways, prior work has proposed heuristics for reducing the computational load. A naive scheme for implementing the component performs $N(N-1)$ TTM operations. Baskaran et al. [10] focused on reducing the number of TTM operations, and proposed a scheme with (approximately) $N^2/2$ operations, which was further improved to $N \log N$ by Kaya and Uçar [11]. However, minimizing the number of TTM operations is not sufficient and it is crucial to consider the cost of the operations, especially in the context of dense tensors. Austin et al. [2] measure the cost in terms of the number of floating point operations (FLOP). They empirically showed that the performance of the naive scheme can be improved by permuting (ordering) the modes of the input tensor and proposed a greedy heuristic for mode ordering. A similar heuristic is given by Vannieuwenhoven et al. [8].

Heuristics for communication volume: Austin et al. [2] presented the first distributed implementation of HOOI. They distribute the tensors among the processors using a Cartesian

parallel distribution which generalizes the block distribution technique used in the context of matrices. The processors are arranged in the form of an N -dimensional grid and a tensor is partitioned into blocks by imposing the grid on the tensor; the blocks are then assigned to the processors. They showed that the communication volume is determined by the choice of the grid, and presented an empirical evaluation of the effect of the grid on the communication volume.

Our Contributions

Our objective is to enhance the performance of distributed Tucker decomposition for dense tensors by designing optimal schemes for the two metrics and we make the following contributions.

- *Optimal TTM-trees*: As observed in prior work [11], the different TTM schemes can be conveniently represented in the form of trees, called TTM-trees. We present an efficient algorithm for constructing the optimal TTM-tree, the one having the least computational load, measured in terms of number of floating operations (FLOP).
- *Dynamic Gridding*: Prior work uses a static gridding scheme, wherein the same grid is used for distributing the tensors arising in the different TTM operations. We propose the concept of *dynamic gridding* that uses different grids tailored for the different operations, leading to significant reduction in communication volume, even when compared to the optimal static grids.
- *Optimal Dynamic Gridding*: We present an efficient algorithm for finding the dynamic gridding scheme achieving the optimal communication volume.

Our distributed implementation builds on the work of Austin et al. [2] and incorporates the optimal schemes described above. We setup a large benchmark consisting of about 1700 tensors whose metadata are derived from real-life tensors. Furthermore, we also include a set of tensors with metadata derived from simulations in combustion science. Our experimental evaluation on the above benchmark demonstrates that the combination of optimal trees and the dynamic gridding scheme offers significant reduction in computational load and communication volume, resulting in up to 7-factor improvement in overall execution time, compared to prior heuristics. To the best of our knowledge, our study is the first to consider optimal algorithms for the Tucker decomposition.

We note that prior work [8, 2] has provided evidence that STHOSVD may be sufficient for particular application domains. They present experimental evaluations on a sample of tensors arising in image processing and combustion science showing that for these tensors, STHOSVD is sufficient and HOOI does not provide significant error reduction. Our optimizations on HOOI would be useful for other tensors/domains where HOOI provides error reduction over STHOSVD. Furthermore, the ideas developed in this paper can be recast and used for improving STHOSVD as well.

Related Work

Tucker decomposition has been studied in sequential and parallel settings for dense and sparse tensors. For dense tensors, the MATLAB Tensor Toolbox provides a sequential im-

plementation [12]. Zhou et al. [13] proposed a randomized algorithm for the case where the tensor fits in the physical memory of a single machine. Li et al. [14] proposed performance enhancements for a single TTM operation and their techniques can be incorporated within our framework. Austin et al. [2] described the first implementation for distributed memory systems, wherein they proposed heuristics for mode ordering and experimentally demonstrated the effect of grid selection on communication time. For sparse tensors, sequential [7], shared memory [10] and distributed implementations [6] are known. Other tensor decompositions have also been considered (see [5]). In particular, CP decomposition which generalizes the concept of rank factorization has been well studied (e.g. [15]). We refer to survey by Kolda and Bader [5] for a detailed treatment of tensor decompositions.

2 Tucker Decomposition

In this section, we briefly discuss tensor concepts pertinent to our problem, and then, describe the Tucker decomposition and the HOOI procedure.

2.1 Preliminaries

Fibers: Consider an N -dimensional tensor \mathcal{T} of size $L_1 \times L_2 \times \dots \times L_N$. Let $|\mathcal{T}|$ denote the cardinality (number of elements) of the tensor. The elements of \mathcal{T} can be canonically indexed by a coordinate vector of the form $\langle l_1, l_2, \dots, l_N \rangle$, where each index l_n belongs to $[1, L_n]$, for all modes $1 \leq n \leq N$. A *mode- n fiber* \vec{x} is a vector of length L_n , containing all the elements that differ on the n th coordinate, but agree on all the other coordinates, i.e., $\langle l_1, \dots, l_{n-1}, *, l_{n+1}, \dots, l_N \rangle$. The number of mode- n fibers is $|\mathcal{T}|/L_n$. In the analogous case of matrices, two types of fibers can be found: row vectors and column vectors.

Tensor Unfolding: The tensor \mathcal{T} is stored as a matrix and there are N different matrix layouts are possible, called the unfoldings. The *mode- n unfolding* of \mathcal{T} refers to the matrix whose columns are the mode- n fibers of the tensor. The columns are arranged in a lexicographic order (the details are not crucial for our discussion). This matrix is of size $L_n \times (|\mathcal{T}|/L_n)$, and we denote it as $T_{(n)}$.

Tensor-Times-Matrix Multiplication (TTM): For any mode n , the tensor \mathcal{T} can be multiplied by a matrix \mathbf{A} along mode n , provided \mathbf{A} has size $K \times L_n$, for some K ; the operation is denoted $\mathcal{Z} = \mathcal{T} \times_n \mathbf{A}$. Conceptually, the operation applies the linear transformation \mathbf{A} to all the mode- n fibers. It is realized via the matrix-matrix multiplication $\mathbf{A} \times T_{(n)}$, and taking the output matrix to be the mode- n unfolding of \mathcal{Z} . While the length along mode n changes from L_n to K , the number of fibers and the lengths along other modes remains the same. Thus, \mathcal{Z} has cardinality $K \cdot (|\mathcal{T}|/L_n)$ and size $L_1 \times \dots \times L_{n-1} \times K \times L_{n+1} \times L_N$.

TTM-Chain: The *TTM-chain* operation refers to multiplying \mathcal{T} along multiple distinct modes. For two modes n_1 and n_2 and matrices \mathbf{A}_1 and \mathbf{A}_2 , we first multiply \mathcal{T} by \mathbf{A}_1 along mode n_1 , and then multiply the output tensor along mode n_2 by \mathbf{A}_2 . An important property of the operation is commutativity [9], namely the two TTM operations can be performed in any order: $(\mathcal{T} \times_{n_1} \mathbf{A}_1) \times_{n_2} \mathbf{A}_2 = (\mathcal{T} \times_{n_2} \mathbf{A}_2) \times_{n_1} \mathbf{A}_1$. In general, for a

subset of distinct modes $S = \{n_1, n_2, \dots, n_r\}$, and matrices $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_r$, where \mathbf{A}_j has size $K_j \times L_{n_j}$, the output is a tensor $\mathcal{Z} = \mathcal{T} \times_{n_1} \mathbf{A}_1 \times \dots \times_{n_r} \mathbf{A}_{n_r}$. The length of \mathcal{Z} remains the same as \mathcal{T} along modes not belonging to S , and changes to K_j , for all $n_j \in S$. Commutativity implies that the multiplications can be performed in any order.

2.2 Tucker Decomposition and HOOI

The Tucker decomposition of \mathcal{T} approximates the tensor as the product of a *core tensor* \mathcal{G} of size $K_1 \times K_2 \times \dots \times K_N$, with each $K_n \leq L_n$, and a set of *factor matrices* $\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_N$: $\mathcal{T} \approx \mathcal{Z} = \mathcal{G} \times_1 \mathbf{F}_1 \times_2 \mathbf{F}_2 \times \dots \times_N \mathbf{F}_N$. The factor matrix \mathbf{F}_n has size $L_n \times K_n$. The decomposition compresses the length of \mathcal{T} along each mode from L_n to K_n . We write the decomposition as $\{\mathcal{G}; \mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_N\}$. The error of the decomposition is measured by comparing the recovered tensor $\tilde{\mathcal{Z}}$ and the input tensor \mathcal{T} under the normalized root mean square metric.

The HOOI procedure [9] transforms a given decomposition into a new decomposition having the same core size, but with reduced the error. Given an initial decomposition, the procedure can be invoked repeatedly to reduce the error monotonically, until a desired convergence is achieved. An initial decomposition can be found using methods such as STHOSVD [8].

The HOOI procedure (a single invocation), shown in Figure 2, takes as input the tensor \mathcal{T} , and a decomposition $\{\mathcal{G}, \mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_N\}$ with core size $K_1 \times K_2 \times \dots \times K_N$. It produces a new decomposition $\{\tilde{\mathcal{G}}, \tilde{\mathbf{F}}_1, \tilde{\mathbf{F}}_2, \dots, \tilde{\mathbf{F}}_N\}$ with lesser error, but having the same core and factor matrix sizes. For computing each new factor matrix $\tilde{\mathbf{F}}_n$, the procedure utilizes the alternating least squares paradigm and works in two steps. First, it performs a TTM-chain operation by skipping mode n and multiplying \mathcal{T} by the transposes of all the other factor matrices \mathbf{F}_j (with $j \neq n$) and obtains a tensor \mathcal{Z} . The tensor \mathcal{Z} has length compressed from L_j to K_j along all modes $j \neq n$. In the next step, it performs an SVD on $Z_{(n)}$, the mode- n unfolding of the \mathcal{Z} . The new factor matrix $\tilde{\mathbf{F}}_n$ is obtained by arranging the leading K_n singular vectors as columns. Once all the new factor matrices are computed, the new core tensor is computed.

Figure 3 (a) depicts the process in the form of a tree. The root represents the input tensor \mathcal{T} , each node with label n represents multiplication along mode n , and each leaf represents a new factor matrix. For dense tensors, the SVD operations tend to be inexpensive (see [2]). Therefore, we focus on optimizing the TTM component comprising of the N TTM-chains, from the perspectives of computational load and communication volume.

3 Computational Load

The TTM component performs N TTM-chains, each involving $(N - 1)$ TTM operations. Commutativity allows us to rearrange and reuse the operations in multiple ways, all of which can be represented in the form of TTM-trees, as observed in prior work [11]. We measure the computational load of a tree by the number of floating point operations incurred. Our objective is to design an efficient algorithm for finding the optimal TTM-trees. Below, we first formalize the above model and rephrase prior schemes, and then describe

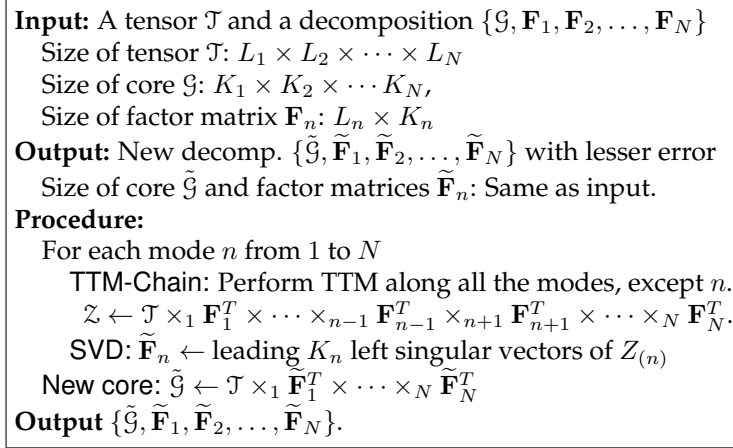


Figure 2: HOOI Procedure

the optimal algorithm.

3.1 TTM-trees and Cost

In a TTM-tree, the root represents the input tensor \mathcal{T} , each leaf node represents a unique new factor matrix and each internal node (nodes other than the root and the leaves) represents TTM along a particular mode. The root-to-leaf path leading to a new factor matrix $\tilde{\mathbf{F}}_n$ realizes the TTM-chain required for computing $\tilde{\mathbf{F}}_n$.

TTM-Trees

Formally, a *TTM-tree* H is a rooted tree with a function $\text{lbl}(\cdot)$ that assigns a label $\text{lbl}(u)$ to each node u such that the following properties are satisfied: (i) the label of the root node is $\text{lbl}(\text{root}) = \mathcal{T}$; (ii) there are exactly N leaves, with each leaf u being labeled with a unique new factor matrix $\text{lbl}(u) = \tilde{\mathbf{F}}_n$; (iii) each internal node u is labeled with a mode $\text{lbl}(u) \in [1, N]$; (iv) for each leaf u with label $\text{lbl}(u) = \tilde{\mathbf{F}}_n$, the path from the root to u has exactly $(N - 1)$ internal nodes and all the modes except n appear on the path.

Figure 3 (a) - (c) provides example TTM-trees for the case of $N = 4$. Although the trees differ in the order in which the modes are processed and the total number of TTMs performed, they all realize the necessary TTM chains.

Given a tree H , the HOOI procedure can be executed via a natural top-down process by associating each node with an input tensor $\mathcal{I}_n(u)$ and an output tensor $\mathcal{O}_n(u)$. For the root node, $\mathcal{I}_n(\text{root}) = \mathcal{O}_n(\text{root}) = \mathcal{T}$. Each internal node u with $\text{lbl}(u) = n$ takes as input the tensor output by its parent v , multiplies it along mode n by the factor matrix \mathbf{F}_n^T , and outputs the resultant tensor, i.e., $\mathcal{I}_n(u) = \mathcal{O}_n(v)$ and $\mathcal{O}_n(u) = \mathcal{I}_n(u) \times_n \mathbf{F}_n^T$. Each leaf node u with $\text{lbl}(u) = \tilde{\mathbf{F}}_n$ constructs the new factor matrix $\tilde{\mathbf{F}}_n$ by performing an SVD on the tensor output by its parent. The correctness of the procedure follows from the commutativity property of the TTM-chain operation. In the above procedure, we reuse the tensor output by a node for processing all its children. By executing the process via

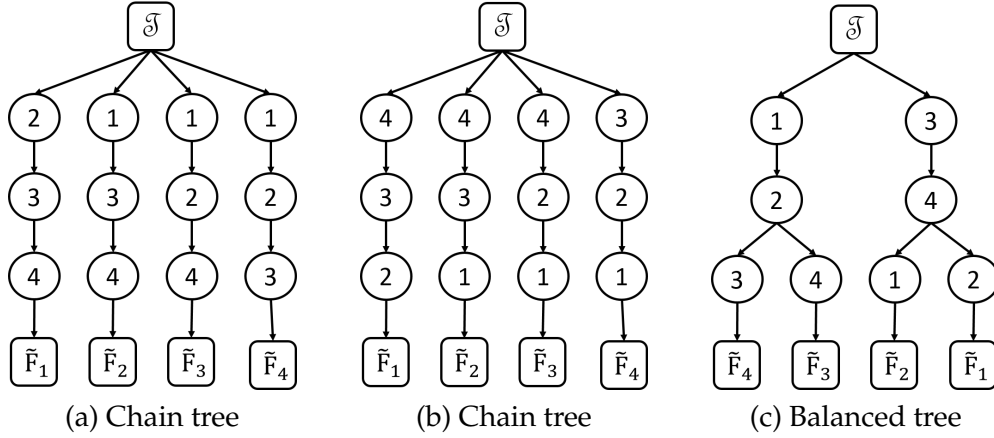


Figure 3: Example TTM-trees. Tree (a) and (b) are both chain trees, but use different orderings, $\langle 1, 2, 3, 4 \rangle$ and $\langle 4, 3, 2, 1 \rangle$, respectively.

an in-order traversal, we can ensure that the maximum number of intermediate tensors stored at any point is bounded by the depth of the tree.

Computational Load

We define the *cost* (or *computational load*) of a TTM-tree H to be the number of floating point operations (FLOP) performed. Each internal node u with label $\text{lbl}(u) = n$ executes the TTM $\text{Out}(u) = \mathcal{J}n(u) \times_n \mathbf{F}_n^T$. Recall that the operation involves the matrix-matrix multiplication, wherein the matrix \mathbf{F}_n^T is multiplied by the mode- n unfolding of $\mathcal{J}n(u)$. The matrix has size $K_n \times L_n$ and the unfolded tensor has size $L_n \times (|\mathcal{J}n(u)|/L_n)$ and so, the cost of the TTM is $K_n \cdot |\mathcal{J}n(u)|$. The cardinality of the output tensor is $|\text{Out}(u)| = (K_n/L_n)|\mathcal{J}n(u)|$; namely, the node compresses the tensor by a factor (K_n/L_n) . We can compute the cost incurred at all the nodes and the cardinality of their output tensors by performing the above calculations in a top-down manner. Then, the cost of the tree H is given by the sum of costs of its internal nodes. We can see that each mode n is associated with two parameters: a *cost factor* K_n and a *compression factor* (K_n/L_n) , which we denote as h_n . At each node, the cost incurred and the cardinality of the output tensor can be expressed in terms of these two parameters.

Figure 4 provides an illustration. The cost incurred and the cardinality of the output tensor are shown at each node. For the ease of exposition, we have normalized all the quantities by $|\mathcal{J}|$. The root node has cost 0 and its cardinality of its output is $|\mathcal{J}|$, which is 1 after normalization. Each node u with label n incurs a cost of K_n times the cardinality of the tensor output by its parent; it outputs a tensor having cardinality compressed by a factor h_n .

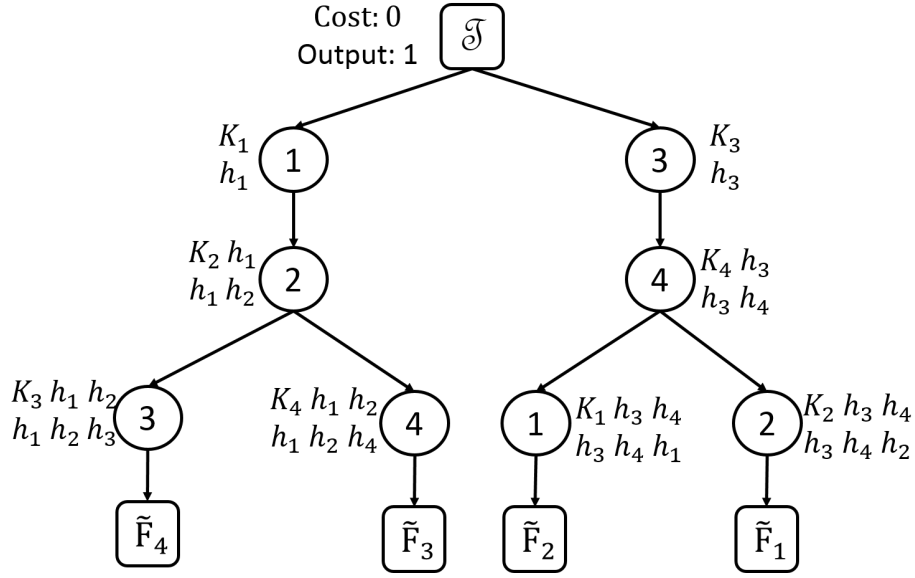


Figure 4: Cost analysis

3.2 Prior Schemes

We rephrase the prior schemes in terms of TTM-trees.

Chain trees: These trees encode the naive scheme, with N independent chains, each comprising of $(N - 1)$ nodes (see figure 3 (a) and (b)).

Balanced trees: Chain trees perform $N(N - 1)$ TTMs. Kaya and Uçar [11] improved the count to approximately $N \log N$, via a divide-and-conquer strategy. The idea is to divide the modes into two groups $\{1, 2, \dots, m\}$ and $\{m + 1, m + 2, \dots, N\}$, where $m = \lfloor N/2 \rfloor$. We create a chain of nodes of length m with labels from the first group and attach it to the root. Then, we recursively construct a subtree for the second group and attach it at the bottom of the chain. We then repeat the process by reversing the roles of the two groups. Figure 3 (c) shows an example for $N = 4$. The number of internal nodes is approximately $N \log N$.

Mode Ordering: Since the TTM-chain operation is commutative, the TTM products within a chain can be performed in any order. Based on this fact, Austin et al. [2] propose the concept of mode ordering, wherein the modes of the input tensor are rearranged according to some permutation. For example, Figure 3 (a) and (b) are both chain trees, but have different mode orderings. They proposed two greedy heuristic for mode ordering. The first heuristic arranges the modes in increasing order of cost factor K_n , placing lower cost modes at the top of the tree where large tensors are encountered. The second heuristic arranges the modes in increasing order of compression factor h_n , aiming at higher compression at the top layers of the tree. We are not aware of any prior work on mode ordering with respect to balanced trees.




Figure 5: Illustration of binary tree transformation

3.3 Constructing Optimal Trees

In this section, we present our algorithm for constructing the optimal TTM-tree, the tree with the minimum cost. The algorithm is based on dynamic programming and runs in time $O(4^N)$. In practice, the algorithm takes negligible time, since the number of dimensions of dense tensors is fairly small (typically, $N \leq 10$).

Towards developing the dynamic programming algorithm, we first claim that the optimal TTM-tree is binary, namely every node has at most two children. The proof is based on the observation that if a node u has three children, then the children can be rearranged so that only two of the nodes remain as children of u . We then identify a set of subproblems and derive a recurrence relation relating them. This is followed by a description of the algorithm and an analysis of the running time.

Lemma 3.1. *There exists an optimal binary tree.*

Proof. Let H^* be an optimal tree. Suppose a node u has three children v_1 , v_2 and v_3 . The properties of TTM-trees imply that none of the three nodes can be a leaf node. Let H_1 , H_2 and H_3 be the subtrees rooted at the three nodes and let n be the label of v_1 . Without loss of generality, assume that the leaf node bearing the label \tilde{F}_n appears in the subtree H_3 . Below, we argue that H^* can be transformed into a new tree H' , without increasing the cost, such that v_2 is no longer a child of u . Thus, the transformation reduces the number of children of u by one. By repeating the process, we can get a binary tree having cost not more than H^* . The transformation is discussed next.

If the label of v_2 is also n , then we can merge v_1 and v_2 . Otherwise, node n must appear on all the paths from v_2 to the leaves in H_2 . We perform two operations: (i) for any node z in H_2 with label n , we delete the node (by making its children the children of its parent); (ii) make v_2 as a child of v_1 . Let H' be the new tree created by the process. See Figure 5 for an illustration.

We can verify that H' is a valid TTM-tree. Furthermore, the cost of H' cannot be more than H^* , as argued next. The cost of the nodes outside H_2 does not change. Let z' be any node in H_2 and consider three cases: (i) if z' is one of the deleted nodes, then we save its

cost; (ii) if z' is the descendant of a deleted node, its cost does not change; (iii) if z' is an ancestor of a deleted node, the cost cannot increase, since under H' , the tensor input to z' is further shrunk by TTM along mode n . \square

Subproblems

Consider any binary tree H and let u be an internal node in it. With respect to u , the modes n can be partitioned into three groups: (i) *pre-multiplied*: n is found along the path from the root to u , including u ; (ii) *computed under u* : the leaf bearing label $\tilde{\mathbf{F}}_n$ is found under the sub-tree rooted at u ; (iii) n does not belong to either category. Let P , Q and R denote the set of modes belonging to the three categories. For an illustration, consider the tree in Figure 3 (c) and let u denote the right child of the root labeled 3; with respect this node, $P = \{3\}$, $Q = \{1, 2\}$ and $R = \{4\}$. Notice that the triple (P, Q, R) forms a partitioning of $[1, N]$. We can characterize any node u in a TTM-tree via the above 3-partition.

We next make an observation regarding the set R . Consider the stage in the HOOI execution, wherein we have completed the processing of the node u . At this stage, we have already completed multiplication along all modes in P . For any mode $n \in Q$, the corresponding TTM-chain involves multiplication along all modes, except n . Of these modes, we are yet to perform multiplication along the modes in R and $Q \setminus \{n\}$. The multiplications along modes in R are common to the TTM-chains corresponding to all the modes in Q . Therefore, at this stage, we can potentially select any mode from R , perform multiplication along the mode and reuse the output tensor. Hence, we call the modes in R as *reusable*. For instance, mode 4 is reusable in the example discussed earlier (Figure 3).

The idea behind the dynamic programming algorithm is to consider a subproblem for each possible triple (P, Q, R) as follows: construct the optimal subtree given that the modes in P have been multiplied already, the modes in Q needs to be computed and R are the reusable modes. We formalize the concept using the notion of partial TTM-trees. These trees are similar to the usual TTM-trees, except that the root represents a partially processed tensor and we only need compute a partial set of factor matrices.

Partial TTM-tree: Consider a triple (P, Q, R) with $|Q| \geq 1$. Let $\mathcal{T}[P]$ denote the tensor obtained by multiplying \mathcal{T} by the factor matrices along all the modes found in P . A *partial TTM-tree* for (P, Q, R) is a rooted tree with labels on its nodes such that the following properties are satisfied: (i) the root is labeled $\mathcal{X} = \mathcal{T}[P]$; (ii) there are exactly $|Q|$ leaves, with each leaf u being labeled with a unique factor matrix $\tilde{\mathbf{F}}_n$, for $n \in Q$; (iii) each internal node u is labeled with a mode from $[1, N] \setminus P$; (iv) for each leaf node u with label $\tilde{\mathbf{F}}_n$, the path from the root to u has exactly $N - |P| - 1$ internal nodes and all the modes except $P \cup \{n\}$ appear on them. Figure 6 shows two example partial-TTM trees for the triple $P = \{3\}$, $Q = \{1, 2\}$ and $R = \{4\}$ discussed earlier.

The cost of a partial-TTM tree is defined analogous to the usual TTM-trees. Let $H^*(P, Q, R)$ denote the optimal partial TTM-tree for the triple (P, Q, R) and let $\text{cost}^*(P, Q, R)$ be the cost of the optimal tree. The optimal tree for the original problem is given by $H^*(P, Q, R)$ with $P = \emptyset$, $Q = [1, N]$ and $R = \emptyset$.

Recurrence Relation

We discuss the subproblem structure and derive a recurrence relation. Consider a triple (P, Q, R) . Since optimal trees are binary, the root of $H^*(P, Q, R)$ can have either one or two children. The recurrence relation considers both the possibilities, which we refer to as *reuse* and *splitting*.

Reuse: This option is available, if $R \neq \emptyset$. In this case, we select a mode $n \in R$ and multiply $\mathcal{X} = \mathcal{T}[P]$ along mode n . The result is then reused for computing the new factor matrices of all the modes in Q . In terms of TTM-trees, the operation corresponds to adding a single child with label n to the root of the partial TTM-tree. Once the above TTM operation is performed, we are left with solving the subproblem corresponding to the triple $(P \cup \{n\}, Q, R \setminus \{n\})$. The cost is given by sum of the cost of the TTM operation $\mathcal{X} \times_n \mathbf{F}_n^T$ and the cost of recursively solving the subproblem. Recall that the former cost is $K_n \cdot |\mathcal{X}|$. The latter cost is $\text{cost}^*(P \cup \{n\}, Q, R \setminus \{n\})$. In the above process, any mode from R can be reused and we can find the best option by considering all the choices.

Splitting: The second possibility is to split (or partition) Q into sets Q_1 and Q_2 and independently solve the triples (P, Q_1, R) and (P, Q_2, R) . The total cost is given by the sum of costs of optimal subtrees of the two subproblems, i.e., $\text{cost}^*(P, Q_1, R) + \text{cost}^*(P, Q_2, R)$. Any (non-trivial) partition (Q_1, Q_2) of Q with $Q_1, Q_2 \neq \emptyset$ can be used in the above process and the best choice can be found by an exhaustive search.

The above discussion yields the following recurrence relation for computing the optimal cost of a triple (P, Q, R) :

$$\text{cost}^*(P, Q, R) = \min\{\text{cost}_{\text{reuse}}^*, \text{cost}_{\text{split}}^*\}, \text{ where}$$

$$\begin{aligned} \text{cost}_{\text{reuse}}^* &= \min_{n \in R} K_n \cdot |\mathcal{T}[P]| + \text{cost}^*(P \cup \{n\}, Q, R \setminus \{n\}) \\ \text{cost}_{\text{split}}^* &= \min_{\langle Q_1, Q_2 \rangle \subseteq Q} \text{cost}^*(P, Q_1, R) + \text{cost}^*(P, Q_2, R). \end{aligned}$$

Algorithm and Running Time Analysis

The algorithm constructs a dynamic programming table having at most 3^N entries, one for each triple (P, Q, R) with $|Q| \geq 1$. The base cases for the recurrence relation are triples with $|P| = N - 1$, $|Q| = 1$ and $|R| = 0$, and the cost is 0 in these cases. The other entries get computed by looking up previously computed entries as per the recurrence relation. The entries can be considered according to the following partial ordering: (P', Q', R') precedes (P, Q, R) , if either $P \subset P'$ or $P = P'$ and $Q' \subset Q$. The optimal partial trees can be constructed in a similar manner. The optimal tree for the original problem corresponds to the triple $P = \emptyset$, $Q = [1, N]$ and $R = \emptyset$.

The running time of the algorithm can be analyzed by counting the number of dynamic programming table lookups performed. Each lookup can be specified by a configuration of the form $\langle P \cup \{n\}, Q, R - \{n\}, n \rangle$ in the reuse scenario, and by $\langle P, Q_1, Q_2, R \rangle$ in the splitting scenario. In either case, there are at most 4^N possible configurations. The algorithm does not perform lookup on the same configuration twice and hence, the total number of lookups is at most $2 \cdot 4^N$. Thus, algorithm runs in time $O(4^N)$.




Figure 6: Example partial TTM-trees with $N = 4$, $P = \{3\}$, $Q = \{1, 2\}$, and $R = \{4\}$. $\mathcal{X} = \mathcal{T}[P] = \mathcal{T} \times_3 \mathbf{F}_3^T$

Remarks: In the recurrence relation, we may intuitively think that whenever $R \neq \emptyset$, we should always reuse some mode from R ; see tree H_2 in Figure 6 for an illustration not reusing even though $R \neq \emptyset$. However, the strategy is incorrect. We can construct examples, wherein the optimal tree sacrifices the reuse option on modes having high cost factor so as to postpone multiplication along these modes till the tensor shrinks sufficiently

Given that N is small, we may consider constructing the optimal TTM-trees via an exhaustive search. A naive search over all TTM-trees is prohibitively expensive. The TTM operation corresponding to a mode n involves multiplication along all the other $(N - 1)$ modes, which can be performed in any of the $((N - 1)!)^N$ orderings. Over all the nodes, the number of combinations is $((N - 1)!)^N$, all which can be realized as chain trees. We can expedite the search by considering only the binary TTM-trees. We are not aware of any closed form expression for the number of binary TTM-trees. We note that our algorithm can be modified to enumerate all these trees. Instead of enumeration, the algorithm incorporates memoization and computes the optimal tree efficiently in time $O(4^N)$.

4 Communication Volume

Our strategy is to fix a TTM-tree H (based on the heuristic or the optimal tree) and devise schemes for minimizing the volume. Our distributed implementation uses the same strategy as that of Austin et al. [2] for distributing the tensors and performing TTM in a distributed manner. We propose a dynamic gridding scheme that offers significant reduction in volume and design an efficient algorithm for finding the optimal scheme.

4.1 Distributed Setup

Tensor Distribution: Fix a TTM-tree H and let P be the number of processors. We arrange the processors in an N -dimensional grid $g = q_1 \times q_2 \times \dots \times q_N$ such that $P = \prod_j q_j$. To distribute a tensor, we impose the grid on the tensor and partition it into P blocks, and assign each block to a processor; see Figure 7 for an illustration. The input tensor \mathcal{T} and all

	$N = 5$	6	7	8	9	10
$P = 2^5$	126	252	562	792	1287	2002
2^{10}	1001	3003	8008	19448	43758	92378
2^{20}	10626	53130	230K	880K	3.1M	10M

Table 1: Number of grids for differnt values of P and N

the intermediate tensors gets partitioned using the same grid.

Distributed TTM and Volume: Each node u with label n and parent v performs the TTM operation $\mathcal{O}ut(u) = \mathcal{I}n(u) \times_n \mathbf{F}_n^T$. For the grid g , we denote the communication volume incurred by the operation as $\text{vol}(u, g)$. As observed in the prior work $\text{vol}(u, g) = (q_n - 1)|\mathcal{O}ut(u)|$; a brief outline of the argument in the following paragraph. The total communication volume of g , denoted $\text{vol}(H, g)$, is defined to be the sum of volumes incurred at all the internal nodes.

Recall that the TTM operation $\mathcal{O}ut(u) = \mathcal{I}n(u) \times_n \mathbf{F}_n^T$ can be viewed as applying the linear transformation \mathbf{F}_n^T to every mode- n fiber \vec{x} of $\mathcal{I}n(u)$. That is, we need to perform the matrix-vector product $\vec{y} = \mathbf{F}_n^T \cdot \vec{x}$. Since the factor matrices are small in size, we can afford to keep a copy of them at every processor. However, each mode- n fiber \vec{x} gets distributed equally among some q_n processors and so, computing the product requires a reduce operation. Similarly, the output fiber \vec{y} must be distributed among the same processors using a scatter operation. See Figure 8 for an illustration. The reduce-scatter operation is performed over the output fiber \vec{y} of K_n , for which we incur $(q_n - 1)K_n$ units of communication. Summed up over all the fibers, the total communication volume for the TTM is $(q_n - 1)|\mathcal{O}ut(u)|$.

In the above distribution method, if $q_n > L_n$ for some mode n , then some processor would receive an empty block while partitioning \mathcal{T} . Similarly, if $q_n > K_n$ then same scenario would arise on some intermediate tensor. We avoid the load imbalance by considering only grids with $q_n \leq K_n$, for all n ; we call these *valid grids*. In the rest of the discussion, unless explicitly mentioned, we shall only consider valid grids.

4.2 Finding the Optimal Static Grid

We observe that the optimal static grid, the one achieving the minimum communication volume, can be found via an exhaustive search in negligible time. The number of grids, including the invalid ones, is the same as number of ways in which the integer P can be expressed as the product of N factors, which we denote $\psi(P, N)$. If the prime factorization of P is $P = p_1^{e_1} \cdot p_2^{e_2} \cdots p_s^{e_s}$, then we have that

$$\psi(P, N) = \prod_{i=1}^s \binom{e_i + N - 1}{N - 1}$$

Table 1 shows the quantity for example values of P and N . When the quantity becomes large, the search can be parallelized in a straightforward manner. Even for the extreme case of $P = 2^{20}$ and $N = 10$, the number of grids to be scanned per processor is approximately 10.




Figure 7: Example grids: the two figures use the grids $\langle 4, 2, 1 \rangle$ and $\langle 2, 2, 2 \rangle$, respectively.




Figure 8: Matrix-fiber multiplication

4.3 Dynamic Gridding Scheme

The idea of dynamic gridding is as follows. Consider a node u , and let its parent be v and label be n . The node u performs the TTM operation $Out(u) = \mathcal{I}n(u) \times_n \mathbf{F}_n^T$. If the tensor $\mathcal{I}n(u)$ is represented in a grid $g = \langle q_1, q_2, \dots, q_N \rangle$ then we incur a volume of $(q_n - 1)|Out(u)|$. Thus, it is beneficial to represent $\mathcal{I}n(u)$ under a grid with a small assignment q_n , and in fact, the operation can be made communication-free by assigning $q_n = 1$. The static gridding scheme selects a single grid by considering the cumulative effect of the above communication volume over all the nodes. The idea of dynamic gridding is to select different grids for representing the intermediate tensors, as appropriate for each node. However, we need to pay a price for dynamic gridding: if the tensor output by the parent v is represented in a grid g and we have selected a different grid g' for representing it at u , then the tensor must be regridded (redistributed) among the processors. The process incurs a volume of $|\mathcal{I}n(u)|$. Thus, a dynamic grid scheme must decide whether or not to regrid at each node, and furthermore, if it decides to regrid, the new grid must be selected in a manner beneficial for the TTM operations performed later in the subtree, so that the overall communication is minimized.

In Figure 9, we have shown an example, carefully constructed so as to highlight the

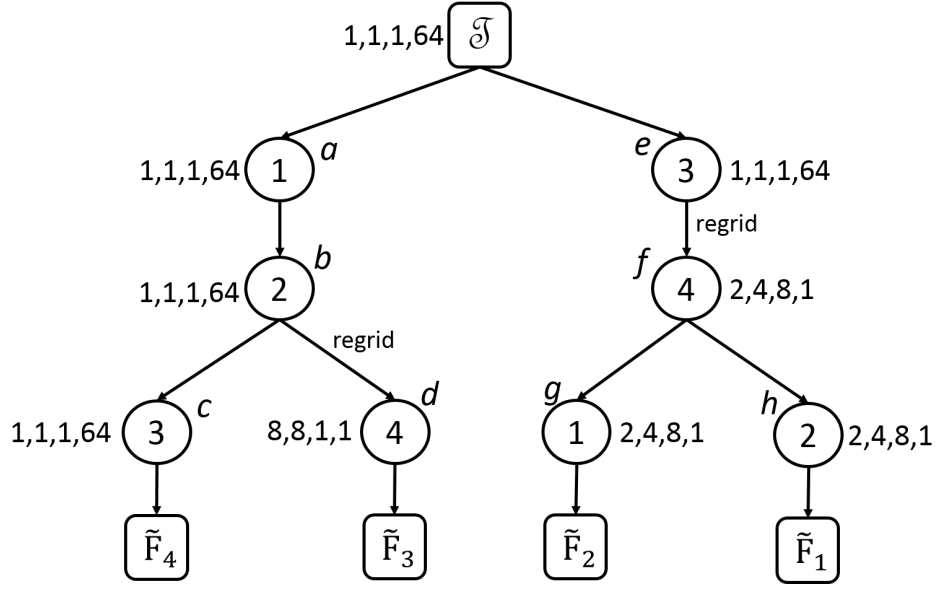


Figure 9: Example dynamic grid scheme

different aspects of dynamic gridding. Assume that the number of processors is $P = 64$ and the core is of size $8 \times 8 \times 8 \times 64$. The choice of the initial grid $\langle 1, 1, 1, 64 \rangle$ makes the TTM operations at nodes a, b, c and e are communication-free. However, the grid is not suitable for the TTM at node d , since the volume incurred is $63 \times |\text{Out}(d)|$. Instead, we switch to a new grid $\langle 8, 8, 1, 1 \rangle$, making the operation communication-free. We perform another regrid operation at node f by selecting the new grid $\langle 2, 4, 8, 1 \rangle$. The choice of the new grid is motivated by the following considerations. The subtree beneath f does not involve any TTM along mode 3 and so, it is prudent to assign a high value along the mode. However, we must select a valid grid, and the constraint implies that the maximum possible value is 8 (since the core length along mode 3 is $K_3 = 8$). We next assign a value of 1 to mode 4, thereby making the TTM at node d communication-free. The remaining of value of 8 is assigned to the modes 1 and 2 in a balanced manner.

Dynamic Grid Scheme: Formally, a *dynamic grid scheme* is a mapping π that associates a grid $\pi(u)$ with each node u . The volume incurred by the scheme, denoted $\text{dvol}(H, \pi)$ is defined as follows. For each node u with label n and parent v , we compute the volume incurred at the node as the sum of two components: (i) TTM operation volume: $(q_n - 1)|\text{Out}(u)|$, where q_n is the assignment to mode n under $\pi(u)$; (ii) regridding volume: if $\pi(u)$ is the same as the parent grid $\pi(v)$, then the volume is zero, and otherwise, it is $|\mathcal{J}n(u)|$. The volume of the scheme π , denoted $\text{dvol}(H, \pi)$, is defined to be the sum of communication incurred over all the nodes u . At the root node, we represent the input tensor \mathcal{T} under the grid $\pi(\text{root})$ and we do not have the regrid option. Let $\text{dvol}^*(H)$ denote the optimal communication volume achievable among all dynamic grid schemes, and let $\text{Opt}(H)$ denote an optimal scheme.

4.4 Optimal Dynamic Gridding Scheme

In this section, we develop an efficient dynamic programming algorithm for computing the optimal dynamic grid scheme for a given tree H . For a node u , let $H(u)$ denote the subtree rooted at u . A *partial grid scheme* for $H(u)$ refers to a mapping π that specifies a grid for each node in $H(u)$. For each node u and each grid g_{par} , we shall define a subproblem with the following connotation: assuming that the tensor output by the parent is represented under the grid g_{par} , find the optimal partial grid scheme for the subtree $H(u)$. We solve these subproblems via a bottom-up traversal of the tree, wherein the optimal solution at u is computed from the optimal solutions of its children.

Subproblems

Consider a pair (u, g_{par}) , where u is a node and g_{par} is a grid. For a partial grid scheme π for the subtree, let $\text{dvol}(H(u), \pi|g_{\text{par}})$ denote the volume incurred by π given that the tensor output by the parent of u is represented in the grid g_{par} . Formally, it is computed as follows. For each node $z \in H(u)$, define a parent grid $\text{pg}(z)$: for the node u , $\text{pg}(u) = g_{\text{par}}$, and for the other nodes, $\text{pg}(z) = \pi(z')$, where z' is the parent of z . For any node $z \in H(u)$, associate the volume given by the sum of the following two components: (i) TTM operation volume: $(q_n - 1)|\text{Out}(z)|$, where n is the mode label of z and q_n is the assignment to mode n under $\pi(z)$; (ii) regridding volume: if $\pi(u)$ is the same as $\text{pg}(z)$, then the volume is zero, and otherwise, it is $|\mathcal{J}n(z)|$. Then, the volume $\text{dvol}(H_u, \pi|g_{\text{par}})$ is defined to be the sum of volumes associated with all the nodes $z \in H(u)$. Let $\text{dvol}^*(H(u)|g_{\text{par}})$ denote the minimum volume possible among all partial grid schemes π . We do not regrid at root and so, define $\text{dvol}^*(H|g_{\text{par}})$ to be the minimum volume given that \mathcal{T} is represented under g_{par} .

Recurrence Relation

We derive a recurrence relation for computing $\text{dvol}^*(H(u)|g_{\text{par}})$. Let v_1, v_2, \dots, v_s be the children of u . In determining the optimal partial scheme, we have two options: (i) regrid: select a new grid $\text{rg}^*(u)$ for representing $\mathcal{J}n(u)$; (ii) do no regrid: represent $\mathcal{J}n(u)$ under the given grid g_{par} . In the first case, we select $\text{rg}^*(u)$ to be the grid yielding the minimum volume for the child subtrees:

$$\text{rg}^*(u) = \underset{g}{\text{argmin}} \sum_{j=1}^s \text{dvol}^*(H(v_j)|g).$$

We can now write the recurrence for $\text{dvol}^*(H(u)|g_{\text{par}})$. Let n be the label of u and v be the parent of u . Let $g_{\text{par}} = \langle p_1, p_2, \dots, p_N \rangle$ and let $\text{rg}^*(u) = \langle q_1, q_2, \dots, q_N \rangle$. Then:

$$\text{dvol}^*(H(u)|g_{\text{par}}) = \min\{\text{vol}_1^*, \text{vol}_2^*\}, \text{ where}$$

$$\text{vol}_1^* = |\mathcal{J}n(u)| + (q_n - 1)|\text{Out}(u)| + \sum_{j=1}^s \text{dvol}^*(H(v_j)|\text{rg}^*(u))$$

$$\text{vol}_2^* = (p_n - 1)|\text{Out}(u)| + \sum_{j=1}^s \text{dvol}^*(H(v_j)|g_{\text{par}})$$

The two quantities correspond to the optimal solutions for the two choices of regridding and not regridding. In both the cases, we incur communication for the TTM operation and communication in the subtrees. In addition, the first case incurs a regrid volume of $|\mathcal{I}n(u)|$. Under the two choices, the tensors $\mathcal{I}n(u)$ and $\mathcal{O}ut(u)$ get represented under the grids $rg^*(u)$ and g_{par} , respectively. Consequently, the recursive calls for the two choices are made with the corresponding grids. At the root node, we represent \mathcal{T} under g_{par} and do not regrid and so, we consider only the first choice at the root. The optimal volume for the whole tree $dvol^*(H)$ is given by minimum of $dvol^*(H|g_{par})$, over all the choices of g_{par} and can be computed via enumerating the choices.

Algorithm and Running Time Analysis

The algorithm constructs a dynamic programming table containing an entry for each pair (u, g_{par}) . Thus, the number of entries is $|H| \cdot \psi(P, N)$, where $|H|$ is number of nodes in the tree. The entries are computed via a bottom-up traversal of the tree. For each entry (u, g_{par}) , we need to compute $rg^*(u)$, which requires a search over all the grids. However, the selection of the grid $rg^*(u)$ is independent of the parameter g_{par} and so, it is sufficient to compute it once per node. The recurrence relation involves a table lookup for each child and an entry for a node is looked up only by its parent, and so the total number of table lookups is $O(|H| \cdot \psi(P, N))$. Similar to the case of static grids (Section 4.2), the exhaustive search involved in computing $rg^*(u)$ can be parallelized in a straightforward manner, if P is large. Thus, the algorithm executes in negligible time in practice.

5 Distributed Implementation

The distributed implementation consists of two modules, a *planner* and an *engine*. The planner constructs a TTM-tree, either based on the heuristics or the optimal tree, and selects grids, either the optimal static or dynamic gridding scheme. The module only requires the meta-data as input: the dimension lengths of the input tensor \mathcal{T} and the core tensor. It needs to be executed only once and the output can be used across multiple invocations of the HOOI procedure. All the processors use the same TTM-tree and there is synchronization at each tree node.

The second module, called the *engine* maintains tensors in a distributed manner, and implements the TTM and SVD routines. Tensors are distributed according to the block distribution method (Section 4). To change the grid under which a tensor is represented, the engine implements an element redistribution procedure via the `MPI_Alltoallv` collective. The TTM operation is implemented in a distributed manner using the algorithm proposed by Austin et al. [2]. The naive method for computing the TTM product along a specified mode first requires an unfolding of the tensor along the mode. Their algorithm cleverly avoids the unfolding operation by employing a blocking strategy which breaks down the TTM product into a series of matrix multiplication calls. The matrix multiplication calls are performed using `dgemm`. We implement the SVD component using distributed Gram matrix computation (AA^T) followed by eigen value decomposition (EVD). The Gram matrix computation is performed using `dysrk` calls which exploits the symmetry in the product. The EVD is computed sequentially by invoking the `dsyevx` routine;

this is acceptable since it operates on small square matrices of size $L_n \times L_n$, and $L_n \leq 2000$ in our setting.

6 Experimental Evaluation

In this section, we present an experimental evaluation of the algorithms described in the paper.

6.1 Setup

System

The experiments were conducted on an IBM BG/Q system. Each BG/Q node has 16 cores and 16 GB memory. Our implementation is based on MPI and OpenMP, with `gcc 4.4.6` and `ESSL 5.1`. Each MPI rank was mapped to a single node and spawns 16 threads which are mapped to the cores. All the experiments use 32 nodes.

Tensors

As discussed in the introduction, the execution time of the HOOI algorithm is crucially dependent on the metadata (dimension lengths of the input tensor and the core tensor), and independent of the elements in the tensor. We exploit this property to construct a large benchmark of tensors with metadata derived from real world tensors considered in prior work.

We also include a set of tensors with metadata derived from simulations in combustion science [2]. The metadata of these tensors is shown in Table 2. Due to memory limitations, we curtailed the length along certain dimensions; while the length along all the spatial dimensions were retained as such, we reduced the length along the axes of variables/timesteps and proportionately reduced the length of the core along these axes. We fill these tensors with randomly generated data.

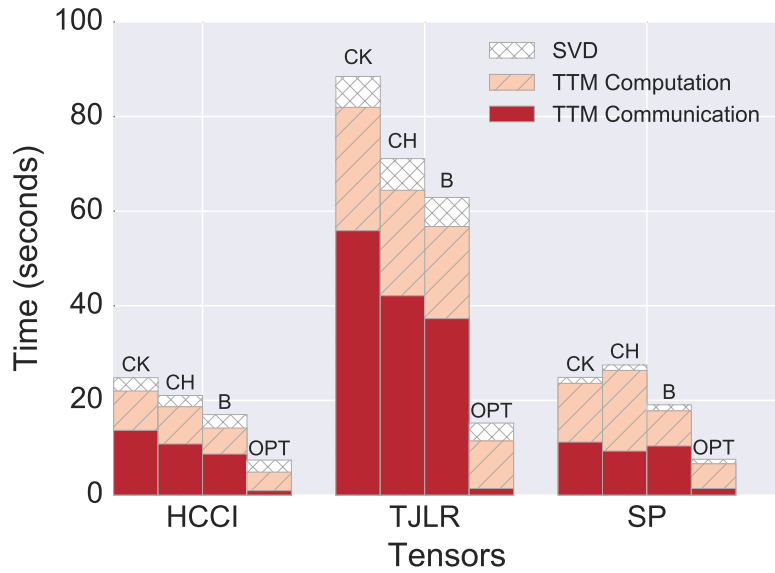
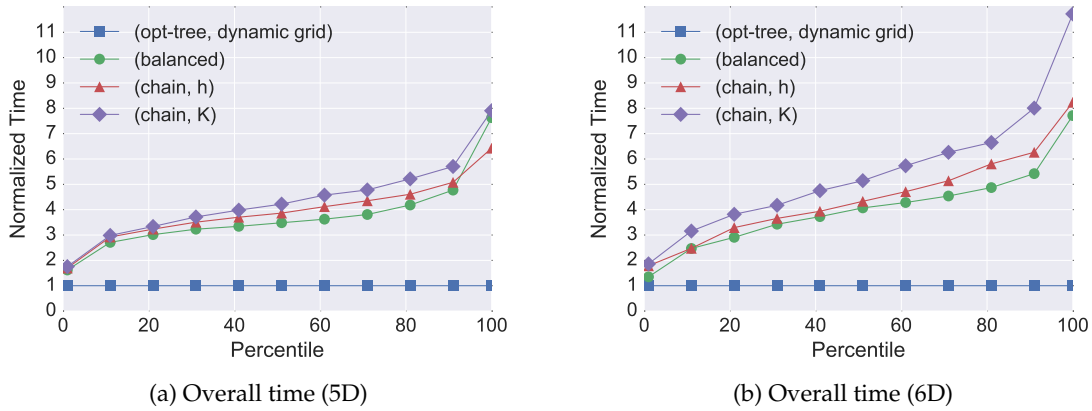
The benchmark is constructed as follows. We constructed 5 and 6-dimensional tensors with dimension lengths L_n drawn from the set $\{20, 50, 100, 400\}$. We selected the core dimension lengths K_n by fixing the compression ratio $h_n = K_n/L_n$. The value for h_n was drawn from the set $\{1.25, 2, 5, 10\}$. Given the above two sets of choices, an input for the HOOI procedure can be generated as follows: for each dimension $n \in [1, N]$, we select L_n from the first set of choices, and select h_n from the second set of choices, and set $K_n = h_n \cdot L_n$. We placed an upper limit of $8 \cdot 10^9$ on the cardinality of \mathcal{T} . We enumerated all possible HOOI inputs in the above manner and obtained a benchmark consisting of 1134 5-dimensional and 642 6-dimensional tensors.

6.2 Evaluation

The experiments involved comparing our algorithm and prior heuristics (Section 3.2 and 4.1). The heuristics are obtained by fixing the tree class to be chain and balanced trees, and the mode ordering to be K -ordering and h -ordering. In the case of balanced trees,

	Tensor Dimensions	Core Tensor Dimensions
HCCI	(672, 672, 627, 16)	(279, 279, 153, 14)
TJLR	(460, 700, 360, 16, 4)	(306, 232, 239, 16, 4)
SP	(500, 500, 500, 11, 10)	(81, 129, 127, 7, 6)


Table 2: Real tensors used in our study



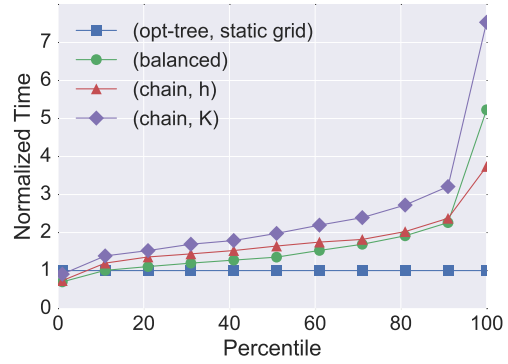
(c) Real Tensors. CK:(chain, K), CH: (chain, h), B: (balanced), OPT:(opt-tree, dynamic grid)

Figure 10: Overall Execution Time


we observed that K -ordering and h -ordering do not impact the execution time and so, we use the input (naive) mode ordering. For all these heuristics, we use the optimal static grids. We compare the heuristics with our algorithms: the optimal tree algorithm with static grids and the same algorithm with dynamic grids.




(a) Computational Time (5D)




(b) Computational Time (6D)




(c) Computational Load (5D)



(d) Computational Load (6D)



(e) Communication Time



(f) Communication Volume

Figure 11: Analysis of Benchmark Results

The following metrics were studied: overall execution time, computational load and time, and communication volume and time. The dimensions of the tensors/matrices arising in the computations are identical across different HOOI iterations (only data elements change). Consequently, any two HOOI iterations will incur the same computational load and communication volume. Thus, the running times would be approximately the same across iterations. Hence, we executed each algorithm on all the benchmark tensors and measured these metrics for a single HOOI invocation.

Overall Execution Time

We compared overall execution time of the opt-tree algorithm with dynamic gridding against the prior heuristics. For each tensor, we normalized the execution times w.r.t the execution time of the opt-tree algorithm (which becomes 1 unit). Given that the benchmark is large, we summarize the results using a percentile plot. Figure 10a and 10b shows the plots for 5D and 6D tensors. In these plots, normalized time of t on percentile value k means that for $k\%$ of tensors, the normalized execution time is less than t . For example, in Figure 10a, the 60th percentile value for the (chain, K) is 4.7, meaning that the improvement factor obtained by the opt-tree algorithm is at most 4.7x for 60% of the tensors and at least 4.7x for the remaining 40% of the tensors. These plots reveal the overall performance of the heuristics across the benchmark; a lower curve means that the heuristic performs better.

The curves corresponding to the prior work lie above the opt-tree algorithm, i.e., it outperforms all the prior algorithms on every tensor in the benchmark. The performance gain is dependent on the meta-data. and varies from 1.5x to 7x. The tensors that achieved the minimum and the maximum gains are: Min - $400 \times 400 \times 20 \times 20 \times 20$ compressed to $320 \times 40 \times 10 \times 10 \times 10$; Max - $400 \times 100 \times 100 \times 50 \times 20$ compressed to $80 \times 80 \times 10 \times 40 \times 10$. The median improvement is 3.4x for 5D and 4.0x for 6D tensors. A detailed study is required to characterize the gain in terms of meta-data.

We also studied the performance of the algorithms on the real tensors. Figure 10c shows the actual execution time for one HOOI invocation. For each tensor, we show 4 bars, corresponding to three prior algorithms and the opt-tree algorithm with dynamic grids. For all the tensors, we see that balanced tree outperforms the chain algorithms, because it reuses TTM operations. The opt-tree algorithm offers improvements as high as 4.6x over (chain, h), 5.8x over (chain, K) and 4.1x over (balanced). For these tensors, the superior performance of the opt-tree algorithm is mainly because of drastic reduction in communication time and partial reduction in computation time. Remarkably, the opt-tree algorithm becomes near communication-free under all the three tensors.

Computation Optimization

Here, we study the performance gains from optimal computation tree construction by comparing heuristics and the opt-tree algorithm on computation time and load for the TTM-component. We normalized the quantities with respect to the opt-tree algorithm. The time and load for each algorithm-tensor pair was normalized w.r.t the time and load of the opt-tree algorithm. The comparison of the time for 5D and 6D tensors are reported in

Figure 11a and 11b. The opt-tree algorithm offers 1.5-1.7x median improvement compared to prior algorithms for 5D tensors and 1.4-2.0x median improvement for 6D tensors. The maximum gain is as high as 2.8x and 3.7x for 5D and 6D.

Figure 11c and 11d show the normalized computational load for 5D and 6D. We see that the opt-tree algorithm offers up to 2.8x (5D) and 3.6x (6D) reduction in load over the best prior algorithm, corroborating the improvements seen in time. The improvements are higher for 6D, compared to 5D, because opt-tree has more opportunities for careful placement and reuse of the TTMs.

Communication Optimization

In this experiment, we study the benefits of dynamic gridding. To do so, we compare the opt-tree algorithm with the static and the dynamic gridding schemes under the metrics of communication time and volume. For the latter, we include the time incurred in TTM multiplication, as well as regridding. The quantities are normalized with respect to the dynamic gridding scheme. The results are shown in Figure 11e and 11f. In Figure 11f, we can see that dynamic gridding offers up to 6x factor improvement in communication volume over static gridding, whereas in Figure 11e, we can see improvements up to 17x factor (median 9.4x) in communication time. The reason for higher improvements on communication time is that regridding (based on all-to-all collective) turns out to be faster than TTM multiplication (based on reduce-scatter over group communicators) for the same communication volume. Remarkably, the dynamic grid scheme outperforms static grid scheme on almost all the tensors in the benchmark, with a gain of at least 3-factor on 90% of the tensors. The gain in communication time is a result of improvement in communication volume, a machine independent statistic. Thus, we expect similar gains on other distributed memory systems as well.

7 Conclusions

We studied the Tucker decomposition for dense tensors for the distributed memory setting. We proposed efficient algorithms for computing the optimal trees and dynamic gridding schemes. Our experimental evaluation on a large benchmark demonstrates that the proposed algorithms lead to significant reduction in computational load and communication volume, and offers up to 7x improvement in performance. To further improve the performance of HOOI on dense tensors, a distributed SVD solver could be used instead of the Gram product followed by sequential EVD. Investigating the applicability of the techniques developed in this paper to the case of sparse tensors is a potential avenue of future work.

Acknowledgements: We thank Woody Austin, Grey Ballard and Tamara G. Kolda for sharing their insights with us, and the reviewers for helpful comments.

References

- [1] L. R. Tucker, "Some mathematical notes on three-mode factor analysis," *Psychometrika*, vol. 31, pp. 279–311, 1966.
- [2] W. Austin, G. Ballard, and T. G. Kolda, "Parallel tensor compression for large-scale scientific data," in *IPDPS*, 2016.
- [3] M. A. O. Vasilescu and D. Terzopoulos, "Multilinear analysis of image ensembles: Tensorfaces," in *ECCV*, 2002.
- [4] D. Muti and S. Bourennane, "Multidimensional filtering based on a tensor approach," *Signal Processing*, vol. 85, pp. 2338–2353, 2005.
- [5] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, pp. 455–500, 2009.
- [6] O. Kaya and B. Uçar, "High performance parallel algorithms for the tucker decomposition of sparse tensors," in *ICPP*, 2016.
- [7] T. G. Kolda and J. Sun, "Scalable tensor decompositions for multi-aspect data mining," in *ICDM*, 2008.
- [8] N. Vannieuwenhoven, R. Vandebril, and K. Meerbergen, "A new truncation strategy for the higher-order singular value decomposition," *SIAM J. on Scientific Computing*, vol. 34, no. 2, pp. 1027–1052, 2012.
- [9] L. D. Lathauwer, B. D. Moor, and J. Vandewalle, "On the best rank-1 and rank- (R_1, R_2, \dots, R_N) approximation of higherorder tensors," *SIAM J. Matrix Analysis and Applications*, vol. 21, pp. 1324–1342, 2000.
- [10] M. Baskaran, B. Meister, N. Vasilache, and R. Lethin, "Efficient and scalable computations with sparse tensors," in *HPEC*, 2012.
- [11] O. Kaya and B. Uçar, "High-performance parallel algorithms for the tucker decomposition of higher order sparse tensors," Inria, Tech. Rep. RR-8801, HAL-01219316, 2015.
- [12] B. W. Bader and T. G. Kolda, "Efficient MATLAB computations with sparse and factored tensors," *SIAM J. on Scientific Comp.*, vol. 30, no. 1, pp. 205–231, 2007.
- [13] G. Zhou, A. Cichocki, and S. Xie, "Decomposition of big tensors with low multilinear rank," *CoRR*, *arXiv:1412.1885*, 2015.
- [14] J. Li, C. Battaglino, I. Perros, J. Sun, and R. Vuduc, "An input-adaptive and in-place approach to dense tensor-times-matrix multiply," in *SC*, 2015.
- [15] U. Kang, E. Papalexakis, A. Harpale, and C. Faloutsos, "GigaTensor: Scaling tensor analysis up by 100 times - algorithms and discoveries," in *KDD*, 2012.