

HPC Formulations of Optimization Algorithms for Tensor Completion

Shaden Smith^{a,*}, Jongsoo Park^b, George Karypis^a

^a*Department of Computer Science & Engineering, University of Minnesota*
^b*Facebook*

Abstract

Tensor completion is a powerful tool used to estimate or recover missing values in multi-way data. It has seen great success in domains such as product recommendation and healthcare. Tensor completion is most often accomplished via low-rank sparse tensor factorization, a computationally expensive non-convex optimization problem which has only recently been studied in the context of parallel computing. In this work, we study three optimization algorithms that have been successfully applied to tensor completion: alternating least squares (ALS), stochastic gradient descent (SGD), and coordinate descent (CCD++). We explore opportunities for parallelism on shared- and distributed-memory systems and address challenges such as memory- and operation-efficiency, load balance, cache locality, and communication. Among our advancements are a communication-efficient CCD++ algorithm, an ALS algorithm rich in level-3 BLAS routines, and an SGD algorithm which combines stratification with asynchronous communication. Furthermore, we show that introducing randomization during ALS and CCD++ can accelerate convergence. We evaluate our parallel formulations on a variety of real datasets on a modern supercomputer and demonstrate speedups through 16384 cores. These improvements reduce time-to-solution from hours to seconds on real-world datasets. We show that after our optimizations, ALS is advantageous on parallel systems of small-to-moderate scale, while both ALS and CCD++ provide the lowest time-to-solution on large-scale distributed systems.

1. Introduction

Many domains rely on *multi-way* data, which are variables that interact in three or more dimensions, or *modes*. An electronic health record is an interaction between variables such as a patient, symptoms, diagnosis, medical

*Corresponding author

Email addresses: shaden@cs.umn.edu (Shaden Smith), jongsoo@fb.com (Jongsoo Park), karypis@cs.umn.edu (George Karypis)

Jongsoo Park was located at and supported by Intel Parallel Computing Lab at the time of submission.

procedures, and outcome. Similarly, how much a customer will like a product is an interaction between the customer, product, and the context in which the purchase occurred (e.g., date of purchase or location). Analyzing multi-way data can provide valuable insights about the underlying relationships of the different variables that are involved. Utilizing these insights, a doctor would be more equipped to provide a successful treatment and a retailer would be able to better recommend products that meet the customer’s needs and preferences. *Tensors* are a natural way of representing multi-way data. Tensors are the generalization of matrices to more than two modes. *Tensor completion* is the problem of estimating or recovering missing values of a tensor. For example, discovering phenotypes in electronic health records is improved by tensor completion due to missing and noisy data [35]. Similarly, predicting how a customer will rate a product under some context can be thought of as estimating a missing value in a tensor [26].

Multi-way data analysis is based on the assumption that the data of interest follows a low-rank model that can be discovered. Tensor factorization is a technique that reduces a tensor to a low-rank representation, which can then be used by applications or domain experts. The task of tensor completion is often accomplished by first finding a low-rank tensor factorization for the observed data, and if a low-rank model exists, then using it to predict the unobserved data. A subtle, but important constraint is that the factorization must only capture the non-zero (or *observed*) entries of the tensor. The remaining entries are treated as missing values, not actual zeros as is often the case in other sparse tensor and matrix operations. This is because a factorization that treats missing entries as zero-valued would simply learn to estimate unobserved entries as zero.

Tensor completion is challenging on modern processors for two reasons. First, floating point processing rates are much higher than memory bandwidth. Insufficient memory bandwidth is detrimental because only a small amount of work is performed for each non-zero, and accessing each non-zero involves many indices. Furthermore, the few computations that are performed depend on the unstructured sparsity pattern of the observed data, leading to poor cache locality and increased pressure on memory bandwidth. Second, tensors often have a combination of long, *sparse* modes (e.g., patients or customers) and short, *dense* modes (e.g., medical procedures or temporal information). Scalability in the presence of highly varied mode lengths requires attention to load balance, degree of parallelism, and communication. This challenge is largely absent in matrix completion, as the number of rows and columns is usually large.

The high performance computing community has addressed scalability challenges with respect to traditional tensor factorizations, for both shared-memory [1, 33, 30, 18, 14] and distributed-memory [6, 15, 31] systems. However, the techniques and optimizations that underlie these methods cannot be directly applied to the problem of tensor completion due to the treatment of missing entries.

In this work, we explore high performance tensor completion with three popular optimization algorithms: alternating least squares (ALS), stochastic gradient descent (SGD), and coordinate descent (CCD++). We explore issues on shared- and distributed-memory systems such as memory and operation-efficient

algorithms, cache locality, load balance, and communication. Our contributions include:

1. Tensor completion algorithms that scale to thousands of cores despite the presence of highly varied mode lengths.
2. An SGD algorithm that uses a hybrid of stratification and asynchronous updates to maintain convergence at scale.
3. A study of applying randomization to ALS and CCD++, resulting in accelerated convergence and higher quality solutions.
4. An experimental evaluation with several real-world datasets on up to 16384 cores. Our ALS and CCD++ algorithms are $153\times$ and $21.4\times$ faster than state-of-the-art parallel methods, respectively. This effectively reduces solution time from hours to seconds.
5. Publicly available MPI+OpenMP implementations that utilize compressed tensor representations in order to improve cache locality and reduce memory consumption and the number of FLOPs performed.

In addition, our work shows that depending on the underlying parallel architecture and the characteristics of the desired solution, the best performing optimization algorithm varies.

The rest of this paper is organized as follows. Section 2 introduces tensor notation and provides a brief background on approaches for tensor completion. Section 3 reviews optimization methods for tensor completion and existing work on their parallelization. Section 4 details our HPC formulations of the ALS, SGD, and CCD++ algorithms. Section 5 accelerates the convergence rates of ALS and CCD++ by including randomization. Section 6 includes experimental methodology and results. Finally, we provide concluding remarks in Section 7.

2. Preliminaries

2.1. Tensor Notation & Factorization

We denote matrices using bold capital letters (\mathbf{A}) and tensors using bold capital calligraphic letters (\mathcal{R}). A tensor occupies three or more dimensions, or *modes*. In future discussions, we focus on three-mode tensors to reduce notational clutter. However, all of the presented algorithms are general to any number of modes. A tensor has $\text{nnz}(\mathcal{R})$ non-zero entries and is of dimension $I \times J \times K$. The entry in position (i, j, k) is written $\mathcal{R}(i, j, k)$. A colon in the place of an index represents all members of that mode. For example, $\mathbf{A}(:, f)$ is column f of the matrix \mathbf{A} . A *fiber* is the result of holding all but one index constant (e.g., $\mathcal{X}(i, j, :)$) and is the generalization of a row or column of a matrix. A *slice* is the result of holding all but two indices constant (e.g., $\mathcal{R}(i, :, :)$) and the result is a matrix. A common operation in tensor algebra is the *Hadamard* product, denoted $\mathbf{A} * \mathbf{B}$, which performs element-wise multiplication.

The *canonical polyadic decomposition* (CPD), also known as CANDECOMP or PARAFAC [16], is a widely-used model for tensor factorization [4, 21, 26, 37, 12, 35]. The CPD is popular in domains dealing with large-scale data (e.g.,

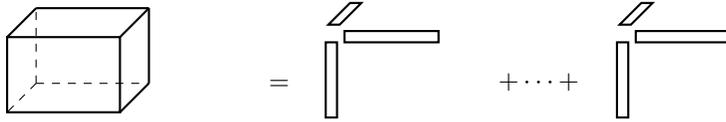


Figure 1: The CPD as a sum of outer products.

machine learning) due to its efficient computation and because it is unique under mild assumptions [34]. The CPD models \mathcal{R} as a rank- F matrix for each mode: $\mathbf{A} \in \mathbb{R}^{I \times F}$, $\mathbf{B} \in \mathbb{R}^{J \times F}$, and $\mathbf{C} \in \mathbb{R}^{K \times F}$. Using the CPD, a tensor \mathcal{R} can be constructed as a summation of F rank-one tensors (Figure 1). Domain experts are almost always interested in a *low-rank* CPD, with a typical value for F being 10 or 50. The CPD can also be written element-wise:

$$\mathcal{R}(i, j, k) = \sum_{f=1}^F \mathbf{A}(i, f) \mathbf{B}(j, f) \mathbf{C}(k, f). \quad (1)$$

More information on tensor factorization can be found in the survey by Kolda and Bader [17].

2.2. Tensor Completion with the CPD

Tensor completion under the CPD model is written as the following non-convex optimization problem:

$$\min_{\mathbf{A}, \mathbf{B}, \mathbf{C}} \frac{1}{2} \sum_{\mathcal{R}(\cdot, \cdot, \cdot)} \mathcal{L}(i, j, k)^2 + \frac{\lambda}{2} (\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2 + \|\mathbf{C}\|_F^2), \quad (2)$$

where λ is a parameter for regularization and $\mathcal{L}(\cdot)$ is the loss function defined as

$$\mathcal{L}(i, j, k) = \mathcal{R}(i, j, k) - \sum_{f=1}^F \mathbf{A}(i, f) \mathbf{B}(j, f) \mathbf{C}(k, f). \quad (3)$$

Note that Equation (2) is only defined over the non-zero (or *observed*) entries of \mathcal{R} and Equation (3) is derived from the element-wise formulation of the CPD in Equation (1).

3. Related Work

The non-convexity of Equation (2) has inspired a substantial amount of research on optimization algorithms that effectively minimize the objective while being operation- and memory-efficient enough to be used in practice. Three optimization algorithms have seen particular success due to their efficiency, opportunities for parallelism, and fast convergence. These methods are summarized in Table 1 and described below.

We follow the convention of the matrix completion community and refer to an *epoch* as the work performed to update \mathbf{A} , \mathbf{B} , and \mathbf{C} one time using the training data. We avoid the term *iteration* in order to emphasize the varying amounts of work performed and progress made.

Table 1: Summary of optimization algorithms for matrix and tensor completion.

Algorithm	Complexity	Storage	Traversals	Ref.
ALS	$\mathcal{O}(M(F^2 \text{nnz}(\mathcal{R}) + IF^3))$	$\mathcal{O}(F^2)$	M	[25, 13]
SGD	$\mathcal{O}(MF \text{nnz}(\mathcal{R}))$	$\mathcal{O}(F)$	1	[25, 27]
CCD++	$\mathcal{O}(MF \text{nnz}(\mathcal{R}))$	$\mathcal{O}(I)$	MF	[27, 13]

Complexity is the number of floating-point operations performed in one epoch. **Storage** is the amount of memory required to perform the factorization, excluding matrix and tensor storage. **Traversals** is the number of times the sparsity structure must be traversed in one epoch, excluding checks for convergence. **Ref.** provides references for the *tensor* variant of the algorithm. References for the matrix variants are found in Section 3. F is the rank of the factorization. M is the number of modes in the tensor. I is the length of the longest mode.

3.1. Alternating Least Squares

ALS is an alternating optimization algorithm that cyclically updates one matrix factor at a time while holding all others constant. Factor updates are based on the observation that if \mathbf{B} and \mathbf{C} are treated as constant, solving for a row of \mathbf{A} is a convex optimization problem with a least squares solution.

Illustrated in Figure 2, computing $\mathbf{A}(i, :)$ accesses all non-zeros in $\mathcal{R}(i, :, :)$ and also the rows $\mathbf{B}(j, :)$ and $\mathbf{C}(k, :)$ for each non-zero $\mathcal{R}(i, j, k)$. The rows of \mathbf{B} and \mathbf{C} are used to compute \mathbf{H}_i , a $|\mathcal{R}(i, :, :)| \times F$ matrix. If the l th non-zero in $\mathcal{R}(i, :, :)$ has coordinate (i, j, k) , then $\mathbf{H}_i(l, :) = [\mathbf{B}(j, :) * \mathbf{C}(k, :)]$, where $*$ denotes the Hadamard product. Given \mathbf{H}_i , we can compute $\mathbf{A}(i, :)$ via

$$\mathbf{A}(i, :) \leftarrow \left(\mathbf{H}_i^T \mathbf{H}_i + \lambda \mathbf{I} \right)^{-1} \mathbf{H}_i^T \text{vec}(\mathcal{R}(i, :, :)), \quad (4)$$

where $\text{vec}(\cdot)$ rearranges the non-zero entries of its argument into a dense vector. The matrix $\left(\mathbf{H}_i^T \mathbf{H}_i + \lambda \mathbf{I} \right)$ is symmetric positive-definite and so the inversion is accomplished via a Cholesky factorization and forward/backward substitutions. ALS requires $\mathcal{O}(F^2 \text{nnz}(\mathcal{R}))$ operations to form all of the \mathbf{H}_i , and $\mathcal{O}(F^3)$ operations per row for the matrix inversions. In total, $\mathcal{O}(F^2 \text{nnz}(\mathcal{R}) + IF^3)$ operations are performed to update a factor. After computing all rows of \mathbf{A} , the other factors are computed in the same manner.

Parallel ALS algorithms exploit the independence of the I least squares problems and solve them in parallel. ALS was one of the first optimization algorithms applied to large-scale matrix completion [38]. Recently, a high performance ALS algorithm for matrix completion on CPUs and GPUs was developed [9]. The algorithm exploits level-3 BLAS opportunities during the construction of $\mathbf{H}_i^T \mathbf{H}_i$.

ALS was first extended to tensor completion on shared-memory systems [25]. Shin and Kang presented a distributed-memory implementation based on the MapReduce paradigm [27]. They use a *coarse-grained* tensor decomposition which partitions each mode separately, assigning to each process a set of complete $\mathcal{R}(i, :, :)$, $\mathcal{R}(:, j, :)$, and $\mathcal{R}(:, :, k)$ slices. After a process updates $\mathbf{A}(i, :)$, the new values are broadcasted to all other processes. Karlsson et al. developed a distributed-memory algorithm for MPI [13]. The distributed-memory algorithm assigns non-zeros to processes without restriction, allowing for $\text{nnz}(\mathcal{R})$

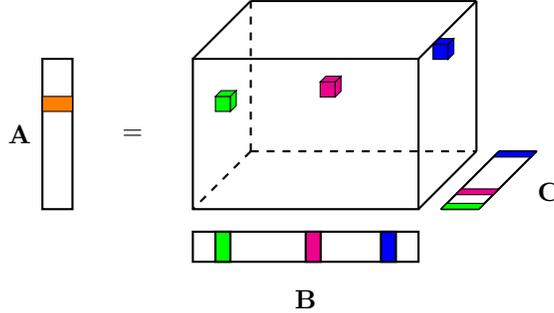


Figure 2: Memory accesses during ALS. Computing $\mathbf{A}(i, :)$ (in orange) requires $\mathcal{R}(i, :, :)$ and the corresponding rows of \mathbf{B} and \mathbf{C} .

parallelism. The added parallelism comes with the cost of communicating partial computations of $\mathbf{H}_i^T \mathbf{H}_i$ and $\mathbf{H}_i^T \text{vec}(\mathcal{R}(i, :, :))$ with an all-reduce, requiring $\mathcal{O}(IF^2)$ words communicated per process.

3.2. Stochastic Gradient Descent

The strategy of SGD is to take many small steps per epoch, each based on the gradient at a single non-zero. At each step, SGD selects one non-zero at random and updates the factorization based on the gradient at $\mathcal{R}(i, j, k)$. Updates are of the form

$$\begin{aligned} \mathbf{A}(i, :) &\leftarrow \mathbf{A}(i, :) + \eta [\mathcal{L}(i, j, k) (\mathbf{B}(j, :) * \mathbf{C}(k, :)) - \lambda \mathbf{A}(i, :)], \\ \mathbf{B}(j, :) &\leftarrow \mathbf{B}(j, :) + \eta [\mathcal{L}(i, j, k) (\mathbf{A}(i, :) * \mathbf{C}(k, :)) - \lambda \mathbf{B}(j, :)], \\ \mathbf{C}(k, :) &\leftarrow \mathbf{C}(k, :) + \eta [\mathcal{L}(i, j, k) (\mathbf{A}(i, :) * \mathbf{B}(j, :)) - \lambda \mathbf{C}(k, :)], \end{aligned}$$

where η is a step size parameter. Each update requires $\mathcal{O}(F)$ operations, resulting in a complexity of $\mathcal{O}(F \text{nnz}(\mathcal{R}))$ per epoch.

SGD is parallelized by exploiting the independence of non-zeros whose coordinates are disjoint. The popular *stratification*-based SGD method [10] partitions an M -mode tensor into a P^M grids for P processes. For example, Figure 3 shows a case with $M=3$ and $P=3$. The grid has P^{M-1} strata, each corresponding to P blocks in a diagonal line (i.e., $\{(1, t_2, \dots, t_M), (2, (t_2+1) \bmod P, \dots, (t_M+1) \bmod P), \dots, (P, (t_2+P-1) \bmod P, \dots, (t_M+P-1) \bmod P)\}$ for all $1 \leq t_2, \dots, t_M \leq P$). Each epoch comprises processing P^{M-1} strata, which covers all the non-zeros of the tensor. Since no two non-zeros in different blocks of a given stratum share the same index in any mode, P processes can work on P blocks of a stratum in parallel.

Instead of stratification, some parallel SGD methods allow non-zeros with overlapping coordinates to be processed in parallel. Hogwild [24], a parallel algorithm for shared-memory systems, exploits the stochastic nature of SGD to have lock-free parallelism. The concept is simple: process the shuffled non-zeros in parallel without stratification or synchronization constructs. Due to the

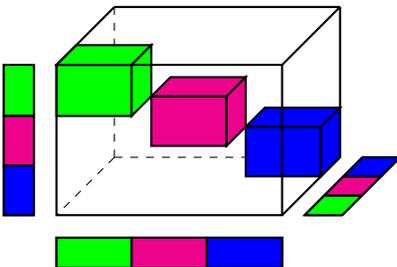


Figure 3: Stratified SGD. Colored blocks of non-zeros can be processed in parallel without conflict.

sparse nature of the input, race conditions are expected to be rare. When they do occur, the stochastic nature of the algorithm will naturally fix any errors and continue to converge. A similar idea for distributed-memory systems is *asynchronous* SGD (ASGD). All non-zeros are processed in parallel, and a few times per epoch processes combine local updates to overlapped rows via weighted averages. The communication and averaging of local updates are performed asynchronously [22].

3.3. Coordinate Descent

In contrast to ALS and SGD which update entire factor rows at a time, coordinate descent methods optimize only one variable at a time. CCD++ is a coordinate descent method originally developed for matrix factorization [36] and later extended to tensors [27, 13]. CCD++ updates columns of \mathbf{A} , \mathbf{B} , and \mathbf{C} in sequence, in effect optimizing the rank-one components of the factorization. Updates take the form:

$$\mathbf{A}(i, f) \leftarrow \frac{\alpha_i}{\lambda + \beta_i}, \quad (5)$$

where

$$\alpha_i = \sum_{\mathcal{R}(i, :, :)} \mathcal{L}(i, j, k) \mathbf{B}(j, f) \mathbf{C}(k, f),$$

and

$$\beta_i = \sum_{\mathcal{R}(i, :, :)} (\mathbf{B}(j, f) \mathbf{C}(k, f))^2.$$

After updating all $\mathbf{A}(:, f)$, the columns $\mathbf{B}(:, f)$ and $\mathbf{C}(:, f)$ are updated similarly. An important optimization for CCD++ is to compute $\mathcal{L}(\cdot)$ only once each epoch and reuse it for each of the F columns [36]. This can be accomplished without additional storage by directly updating \mathcal{R} each iteration with the current residual. The resulting complexity is $\mathcal{O}(F(\text{nnz}(\mathcal{R}) + I + J + K))$, which for most datasets is $\mathcal{O}(F \text{nnz}(\mathcal{R}))$, matching SGD.

CCD++ is parallelized in the same manner as ALS. All of the α_i 's and β_i 's for a mode are computed independently, again leading to a coarse-grained

decomposition of \mathcal{R} [36, 27]. After updating a factor column in parallel, the new column factors are broadcasted to all processes. Karlsson et al. use a non-restrictive decomposition of the tensor non-zeros as in ALS [13]. Partial products are again aggregated with an all-reduce and new columns are broadcasted. Only the components of the current column must be communicated, and so in one epoch there are $2F$ messages per factor matrix, each of size $\mathcal{O}(I)$.

4. Algorithms for High Performance Tensor Completion

4.1. Efficient Loss Computation with a Compressed Sparse Tensor

The choice of data structure for representing a sparse tensor affects performance in ways such as memory bandwidth, number of FLOPs performed, and opportunities for parallelism. The HPC formulations of the various optimization algorithms that we developed use the *compressed sparse fiber* (CSF) data structure [30], illustrated in Figure 4. CSF can be thought of as a generalization of the compressed sparse row data structure for matrices. The CSF data structure represents a tensor as a forest of I trees, with each mode recursively compressed one after the other and stored in the next level of the tree.

The evaluation of $\mathcal{L}(\cdot)$ in Equation (3) benefits from exploiting the CSF tensor representation. Consider the computation of $\mathcal{L}(\cdot)$ associated with two successive non-zeros:

$$\begin{aligned}\mathcal{L}(i, j, k) &= \mathcal{R}(i, j, k) - \sum_{f=1}^F \mathbf{A}(i, f) \mathbf{B}(j, f) \mathbf{C}(k, f), \\ \mathcal{L}(i, j, k') &= \mathcal{R}(i, j, k') - \sum_{f=1}^F \mathbf{A}(i, f) \mathbf{B}(j, f) \mathbf{C}(k', f).\end{aligned}$$

We reuse partial results by storing the Hadamard (element-wise) product of $\mathbf{A}(i, :)$ and $\mathbf{B}(j, :)$ in a row vector \mathbf{v} :

$$\begin{aligned}\mathbf{v} &\leftarrow \mathbf{A}(i, :) * \mathbf{B}(j, :) \\ \mathcal{L}(i, j, k) &= \mathcal{R}(i, j, k) - \sum_{f=1}^F \mathbf{v}(f) \mathbf{C}(k, f), \\ \mathcal{L}(i, j, k') &= \mathcal{R}(i, j, k') - \sum_{f=1}^F \mathbf{v}(f) \mathbf{C}(k', f).\end{aligned}$$

This reduces the computation from $2F \text{nnz}(\mathcal{R})$ multiplications to $F \text{nnz}(\mathcal{R}) + FP$, where P is the number of unique $\mathcal{R}(i, j, :)$ fibers. The *matricized tensor times Khatri-Rao product* (MTTKRP) operation present in ALS uses the same technique for operation reduction [33].

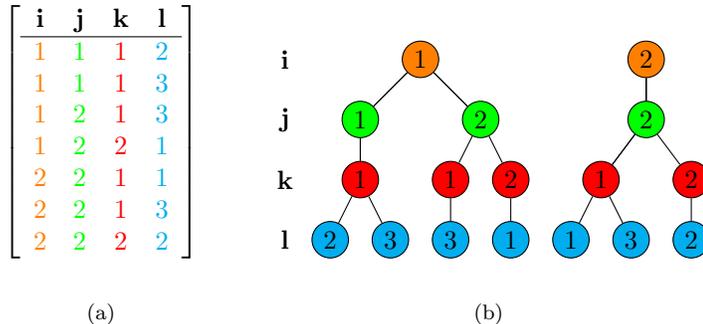


Figure 4: Data structures for the sparsity pattern of a four-mode tensor. (a) Coordinate format: non-zeros are represented as an uncompressed list. (b) Compressed sparse fiber (CSF) format: the sparsity structure is recursively compressed and non-zeros are stored as paths from roots to leaves.

4.2. Parallel ALS

We follow the strategy of parallelizing over the rows of \mathbf{A} for shared-memory parallel systems. Recall from Equation (4) that \mathbf{H}_i has $|\mathcal{R}(i, :, :)|$ rows and F columns. A major challenge when designing ALS algorithms is that if multiple rows of \mathbf{A} are computed at once, the various \mathbf{H}_i matrices must somehow be represented in memory. However, the collective storage for all \mathbf{H}_i prohibitively requires $F \text{nnz}(\mathcal{R})$ storage. Another option, and the one used by existing work [13], is to aggregate rank-1 updates. For each non-zero $\mathcal{R}(i, j, k)$, a rank-1 update with the row vector $(\mathbf{B}(j, :) * \mathbf{C}(k, :))$ is accumulated into $\mathbf{H}_i^T \mathbf{H}_i$. A naive implementation must aggregate the rank-1 updates for all possible i , requiring IF^2 storage. It was observed that this storage overhead can be reduced by sorting the tensor non-zeros before processing long modes [13].

We instead store M CSF representations of \mathcal{R} , in which the m th CSF places the m th mode at the top level of the tree and sorts the remaining ones by length. Since the mode ordering forces all non-zeros in $\mathcal{R}(i, :, :)$ to be grouped into the same tree, we can process all of the non-zeros in $\mathcal{R}(i, :, :)$ sequentially and thus only the $\mathcal{O}(F^2)$ memory associated with a single row is allocated. For multicore systems, we parallelize over the rows of \mathbf{A} with a memory overhead of only $\mathcal{O}(F^2)$ per thread. The total storage overhead with P threads is $(M-1) \text{nnz}(\mathcal{R}) + PF^2 = \mathcal{O}(M \text{nnz}(\mathcal{R}))$, which is smaller than $F \text{nnz}(\mathcal{R})$ and IF^2 for most problems. We note that this multi-CSF strategy was used in prior work [33], but was motivated by exposing parallelism in MTTKRP instead of reducing the size of intermediate data.

The construction of $\mathbf{H}_i^T \mathbf{H}_i$ and $\mathbf{H}_i^T \text{vec}(\mathcal{R}(i, :, :))$ are performed together during a single pass over the sparsity structure of \mathcal{R} using similar strategies as presented in Section 4.1. Interestingly, the expression $\mathbf{H}_i^T \text{vec}(\mathcal{R}(i, :, :))$ is equivalent to computing one transposed row of the MTTKRP operation, for which operation-efficient CSF algorithms exist [30].

Accumulating rank-1 updates into $\mathbf{H}_i^T \mathbf{H}_i$ causes $\mathcal{O}(F^2)$ data to be accessed for each non-zero and leads to memory-bound computation. We collect the

intermediate Hadamard products formed during MTTKRP into a thread-local matrix of fixed size. When that matrix fills, or all non-zeros in $\mathcal{R}(i, :, :)$ are processed, the thread performs one rank- k update, where k is the number of rows in the thread-local matrix. We empirically found that a matrix of size $2048 \times F$ is sufficient to see significant benefits from the BLAS-3 performance. For typical values of F , this equates to storage overheads of up to a few megabytes per thread.

We follow existing work and use a coarse-grained decomposition in the distributed-memory setting [27]. Coarse-grained decompositions impose separate 1D decompositions on each tensor mode, eliminating the need to communicate and aggregate partial computations. As a result, the worst-case communication volume is reduced from $\mathcal{O}(IF^2)$ to $\mathcal{O}(IF)$ per process, which is the cost of exchanging updated rows of \mathbf{A} . Our coarse-grained decomposition uses chains-on-chains partitioning to optimally load-balance the slices of each mode [23]. After partitioning and distributing non-zeros, we use our shared-memory ALS kernels without modification. After a factor matrix is updated, we use an all-to-all collective communication to exchange the new rows.

4.3. Parallel SGD

Processing non-zeros in a totally random fashion requires $\mathcal{O}(\text{nnz}(\mathcal{R}))$ work per epoch to shuffle the non-zeros and results in random access patterns to the matrix factors. We instead use a more coarse-grained approach and randomize only one mode of the tensor, i.e., we randomize the processing order of the trees in the CSF structure (Figure 4) but sequentially access non-zeros within each tree. This approach reduces the cost of shuffling to $\mathcal{O}(I)$ and retains the cache-friendly access pattern to the matrix factors that is provided by the sorted indices in CSF. Our shared-memory SGD algorithm uses a Hogwild approach and processes the trees in parallel without synchronization constructs [24]. Choosing which mode to randomize requires careful consideration, because SGD may not converge if updates are not sufficiently stochastic. This is demonstrated in Section 6.3.

In a distributed-memory setting, we refer to a group of blocks in the P^M stratification grid (Figure 3) that share a coordinate as a *stratum layer*. For example, blocks of coordinate $(i, :, :)$ are in the i th stratum layer along the first mode. We partition the grid by first selecting the longest mode and decomposing it in a 1D fashion. Each process is assigned to a unique stratum layer in the longest mode. This privatizes the largest matrix factor, meaning no processes can produce updates which overlap with the another process' local portion. Thus, there will not be any communication associated with the largest mode during the factorization.

The number of strata increases exponentially with the number of modes. Assuming the number of non-zeros is constant, and therefore the work per epoch is constant, the average work per stratum that can be processed in parallel decreases exponentially. Many real-world tensors have modes with skewed dimensions. For example, a tensor of health records will have significantly more unique patients than unique medical procedures or doctors. The lengths of the

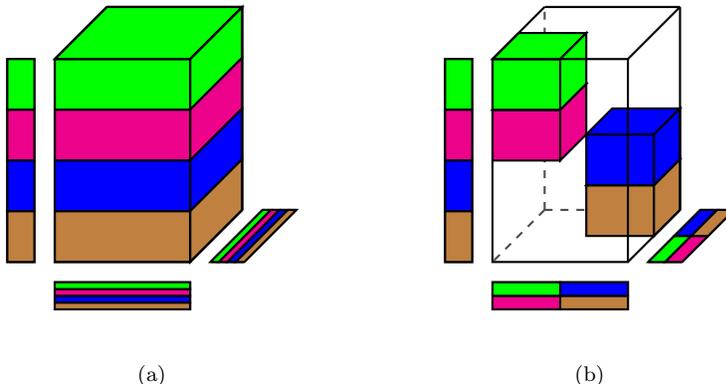


Figure 5: Asynchronous SGD strategies for $P=4$ processes. (a) $S=1$ stratum layers. The longest mode is partitioned and other modes are updated asynchronously. (b) $S=2$ stratum layers. The longest mode is partitioned and S teams of P/S asynchronously update at the end of each stratum.

dense modes can be much smaller than the number of available cores. In this case, the amount of parallelism is limited by the shortest mode (the number of blocks in a stratum equals $\min(I_1, \dots, I_M)$, where I_m is the length of m th mode).

In order to address these parallelism challenges, we extend ASGD to tensors, which allows multiple processes to update the same row of a factor matrix with their local copies. ASGD can trade-off the staleness of factor matrices for increased parallelism by adjusting the number of copies and the frequency of synchronizing the copies. Our implementation is parameterized with the number of stratum layers, S , which determines the number of strata as S^{M-1} . We can set $S = P$ (which reduces ASGD to the usual stratified SGD) for tensors with a few modes or set $S < P$ for tensors with more modes. When $S < P$, since each stratum has only S independent blocks, P/S processes need to update the same range of factor matrices simultaneously, resulting in up to P/S copies of a factor matrix row. Specifically, we partition an M -mode tensor into a $P \times S^{M-1}$ grid and assign P mode-1 layers to each process. Then, we group every P/S processes as a *team* with total S teams. This process is shown in Figure 5.

We partition each factor matrix among P processes, aligning with the grid used for the tensor partitioning. At the beginning of a stratum, each process sends the rows of the factor matrices that it owns to the other processors that need them. By our construction, a factor matrix row will be sent to one team, thus limiting the number of copies to P/S . Then, each process goes through the non-zeros of the current stratum it owns, updating the corresponding rows of factor matrices. After the update, processes send the updated rows back to their owners. Finally, processes compute weighted sums of the received updated rows, where the weights are the number of non-zeros which updated the particular row. For example, for a 3-mode tensor \mathcal{R} and a given stratum, suppose process p_1 processes 2 non-zeros in a mode-2 slice $\mathcal{R}(:, i, :)$ and p_2 processes 1 non-zeros in the same slice. At the end of the stratum, the owner of i th row computes

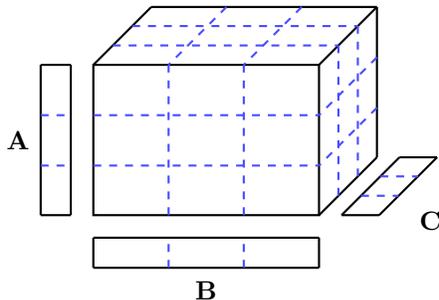


Figure 6: The P^M -way tiling scheme used in CCD++ for $P=3$ threads.

a weighted sum of the local copies of i th row of mode-2 factor matrix from p_1 and p_2 , using $2/3$ and $1/3$ as their weights, respectively. Since we synchronize the factor matrices every stratum, the number of synchronizations per epoch is S^{M-1} .

ASGD allows us to alleviate the limited amount of parallelism and frequent communication, the primary challenges of SGD, especially for high-mode tensors. Still, compared to ALS and CCD++, SGD has higher communication volume, which can be analyzed as follows. For each stratum, a process receives the rows of factor matrices that correspond to the non-zeros it needs to process, which equals the sum of number of non-empty slices of each mode except for the first mode (which is completely privatized). In a worst case (and not an uncommon case for highly sparse tensors), we have only a few non-zeros per slice, leading to receiving $\mathcal{O}(F \text{nnz}(\mathcal{R}_{p,s}))$ floating point numbers for each mode for the sub-tensor $\mathcal{R}_{p,s}$ processed by process p at stratum s . Summing over all processes and strata results in $\mathcal{O}(MF \text{nnz}(\mathcal{R}))$ total communication volume. It is important to receive only the required factor matrix rows corresponding to non-empty slices to be processed. Otherwise, the total communication volume will be $\mathcal{O}\left(S^{M-2}F \sum_{m=2}^M I_m\right)$ as analyzed by Shin and Kang [27].

4.4. Parallel CCD++

As discussed in Section 3.3, CCD++ implementations follow a similar parallelization strategy as ALS on shared-memory systems. Following Equation (5), all α_i 's and β_i 's are independent subproblems and can be computed in parallel. However, unlike ALS, it is not advantageous to use separate CSF representations for each mode in order to extract coarse-grained parallelism. This is due to the added $(M-1)\text{nnz}(\mathcal{R})$ operations that would be required for updating multiple residual tensors. Therefore, we restrict ourselves to a single tensor and turn to other decomposition strategies.

We leverage the P^M -way tiling strategy developed for parallelizing MT-TKRP with a single CSF [30]. Shown in Figure 6, an M -dimensional grid is imposed on \mathcal{R} , with each dimension having P chunks. This allows P threads to

partition any mode of \mathcal{R} into P independent chunks, each consisting of P^{M-1} tiles. Each mode of \mathcal{R} can thus be updated without parallel overheads such as reductions or synchronization.

Conveniently, because threads access whole layers of tiles at a time, we do not need each of the P^M tiles to have a balanced number of non-zeros. Instead, only the layers themselves need to be balanced. We use chains-on-chains partitioning to determine the layer boundaries in each mode, resulting in load-balanced parallel execution.

In a distributed-memory setting, the communication requirements of CCD++ closely follow those of MTTKRP. A minor variation comes from CCD++ being a column-major method, and thus we must exchange partial results and updated columns F times per mode instead of exchanging full rows of size F once per mode. Unlike ALS, exchanging partial results does not cause a prohibitive amount of communication. We can therefore choose from the recently proposed fine-grained [15] and medium-grained [31] decompositions for MTTKRP. We opt for the medium-grained decomposition, which imposes an M dimensional grid over \mathcal{R} . The medium-grained decomposition varies from our tiling strategy in that there are only as many cells in the grid as processes. After distributing \mathcal{R} over a grid, our shared-memory parallelization strategy is applied by each process independently.

4.5. Dense Mode Replication

ALS and CCD++ parallelize over the dimensions of \mathcal{R} . As discussed in Section 4.3, real-world tensors have skewed mode lengths. Simply parallelizing over the short, dense modes is insufficient because the number of threads can easily outnumber the slices. Additionally, the dense modes often have non-zeros that are not uniformly distributed, leading to further load imbalance.

The issue of dense modes was first addressed by Shao [25] for shared-memory ALS when $I < P$. Non-zeros are instead divided among threads and each thread computes a local set of $\mathbf{H}_i^T \mathbf{H}_i$ and $\mathbf{H}_i^T \text{vec}(\mathcal{R}(i, :, :))$. A parallel reduction is then used to combine all partial results before the Cholesky factorization. We adopt this solution in our own ALS implementation and parallelize directly over non-zeros when the mode is dense. We use the tiling mechanism employed by CCD++ to load balance non-zeros. The tensor is still stored using CSF, and thus the optimizations to achieve BLAS-3 performance can still be employed.

CCD++ uses a similar mechanism for handling dense modes. CCD++ only constructs one representation of \mathcal{R} which is already tiled for parallelism. There is no advantage to tiling the dense modes because they will not be partitioned, and so we use a P^{M-d} -way tiling on a tensor with d dense modes. Each thread then computes local α_i and β_i and we aggregate them with a parallel reduction when the mode is dense.

The same decomposition strategies apply to distributed-memory computation. The factors representing dense modes are replicated on all processes. Instead of doing a coarse- or medium-grained decomposition to establish communication patterns, a simple all-to-all reduction is used to aggregate partial computations.

5. Improving Convergence of ALS and CCD++ via Randomization

Non-convex optimization problems such as (2) are susceptible to converging to local minima. Algorithms such as SGD and randomized block coordinate descent [20, 8] have been used with great success because they have the ability to escape local minima when a deterministic descent algorithm would instead converge. While randomized algorithms typically provide no guarantees on converging to a globally optimal solution, in practice they often converge faster and arrive at higher quality local minima than their deterministic counterparts.

The success of these randomized methods motivates the addition of randomization to ALS and CCD++. Careful attention must be paid in order to avoid introducing additional computation or communication and to maintain the degree of available parallelism. In the following discussion we refer to the factor matrices as $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(M)}$ for notational convenience.

5.1. Randomized ALS

An epoch of ALS traditionally updates $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(M)}$ in a cyclic fashion. A cyclic scheme means that $\mathbf{A}^{(m)}$ will always be updated using the latest $\mathbf{A}^{(m-1)}$, which in turn was updated using the latest $\mathbf{A}^{(m-2)}$, and so on. Thus, after the factor matrices are initialized, the optimization process is entirely deterministic. Instead, we propose to introduce *mode-level* randomization and apply a random permutation $\pi : \{1, \dots, M\} \rightarrow \{1, \dots, M\}$ to the tensor modes each epoch. This randomization scheme allows $\mathbf{A}^{(\pi(m))}$ to be updated using the latest state of a different mode each epoch.

Mode-level randomization can be implemented with negligible overhead. If each process uses the same reproducible pseudo-random number generator, then π can be constructed locally by each process and does not need to be communicated. Thus, the only cost of randomization is the construction of π , which is negligible.

5.2. Randomized CCD++

CCD++ presents two opportunities for low-overhead randomization: *rank-level* and *mode-level*. Rank-level randomization updates the F rank-one tensors in a random order instead of cyclically. Within each rank-one optimization, CCD++ can leverage mode-level randomization by altering the order of the updates to columns $\mathbf{A}^{(1)}(:, f), \dots, \mathbf{A}^{(M)}(:, f)$. Like ALS, the only overheads associated with randomization are the construction of two permutations of size F and M , both of which are negligible.

Additional randomized behavior could conceivably be incorporated into ALS and CCD++ by making multiple passes over the factors and updating random subsets of the rows each time. However, we do not explore this option because it reduces the degree of available parallelism from the length of the mode to the number of rows in the current block. Moreover, the already-unstructured communication pattern becomes dynamic due to different sets of rows being updated and communicated each pass.

6. Experimental Methodology and Results

6.1. Experimental Setup

We use the Cori supercomputer at NERSC. Each compute node has 128 GB of memory and is equipped with two sockets of 16-core Intel Xeon E5-2698 v3 that has 40 MB last-level cache. The compute nodes are interconnected via Cray Aries with Dragonfly topology. Our ALS, SGD, and CCD++ implementations are made part of the open source tensor factorization library, SPLATT [29]. We use double-precision floating-point numbers and 64-bit integers. We use the Intel compiler version 16.0.0 with the `-xCORE-AVX2` option for instructions available in the Haswell generation of Xeon processors, Cray MPI version 7.3.1, and Intel MKL version 11.3.0 for LAPACK routines used in ALS. Compute jobs are scheduled with Slurm version 16.05.03. We run one MPI rank per socket (two ranks per node) and one OpenMP thread per core for SGD and CCD++, and one MPI rank per node for ALS. We use the *bold driver* heuristic [2] for a dynamic step size parameter in SGD with an initial value of 10^{-3} .

We follow the recommender systems community and use *root-mean-square error* (RMSE) as a measure of factorization quality. RMSE was the metric used by the Netflix Prize [3], where the first algorithm to improve the baseline RMSE by 10% was awarded one million dollars. RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{\sum_{\mathcal{R}(i,j,k)} \mathcal{L}(i,j,k)^2}{\text{nnz}(\mathcal{R})}}.$$

Datasets are split into 80% *training*, 10% *validation*, and 10% *test* sets. The training set is used to compute the factorization. RMSE is computed each epoch using the validation set, and convergence is detected when the RMSE does not improve for twenty epochs. The final factorization quality is determined by the test set. All algorithms are given the same random initialization for fairness.

6.2. Datasets

Table 2 summarizes the tensors that we use for evaluation. The reported non-zeros and memory requirements reflect only that of the training data, because that is the portion of the computation that we focus on in this work. We work with real-world datasets coming from a variety of domains. MovieLens, Netflix and Yahoo! are $(user, item, month)$ product rating tuples. Values in MovieLens and Netflix range from 1-5 and Yahoo! values range from 1-100. Amazon is formed from $(user, item, word)$ tuples taken from product reviews. Patents is formed from $(year, word, word)$ pairwise co-occurrences taken from United States utility patents from years 1969 through 2014. Non-zero $\mathcal{R}(i, j, k)$ is equal to $\log(f_{jk})$, where f_{jk} is the number of times terms j and k appeared in a seven-word window during year i . Outpatient is a six-mode tensor of $(patient, institution, physician, diagnoses, procedure, day)$ tuples formed from synthetic outpatient Medicare claims. We selected non-zeros from the original data in order to have three-, four-, five-, and six-mode versions of this dataset with the

Table 2: Summary of training datasets.

Dataset	NNZ	Dimensions	Mem. (GB)
Movielens [11]	20M	138K, 27K, 234	0.5
Netflix [3]	80M	480K, 17K, 73	2.4
Outpatient [5]	87M	1.6M, 6K, 13K, 6K, 1K, 192K	4.5
Yahoo! [7]	210M	1M, 625K, 133	6.3
Amazon [19]	1.4B	4.8M, 1.7M, 1.8M	41.5
Patents [28]	2.9B	46, 240K, 240K	85.7

K, **M**, and **B** stand for thousand, million, and billion, respectively. **NNZ** is the number of non-zeros in the training dataset. **Mem.** is the memory required to store the tensor as a list of (*coordinate*, *value*) tuples, measured in gigabytes.

same number of non-zeros. The Amazon, Patents, and Outpatient datasets are publicly available in the FROSTT tensor collection [28].

Throughout the remaining discussion, we will use appropriate subsets of our datasets based on the subject of the experiment. For example, when discussing convergence properties of the various optimization algorithms, we will focus on the ratings tensors (i.e., Movielens, Netflix, and Yahoo!) due to them requiring many iterations to converge and their prevalence in the matrix and tensor completion community. As another example, the synthetic Outpatient tensor converges too quickly for a convergence study, but has a large number of modes and allows for an evaluation on the scalability for higher order tensors.

6.3. Intra-Method Evaluation

6.3.1. ALS

Figure 7 shows the effects of BLAS-3 routines and dense mode replication while factoring the Yahoo! tensor with $F=10$. With neither optimization, each ALS epoch takes 359 seconds on average and we achieve a $22.3\times$ speedup on 32 cores. BLAS-3 routines improve the runtime by $9\times$, but speedup is reduced to $20.5\times$. Finally, by replicating the dense mode across cores and using a parallel reduction on partial products, we achieve $24.2\times$ speedup with 32 cores and a resulting $219\times$ improvement over the original serial implementation.

6.3.2. CCD++

Figure 8 shows the effects of dense mode replication on CCD++. Dense mode replication will not affect serial runtime, and so we show speedup as we scale the number of threads. Speedup is improved from $4.6\times$ to $16.2\times$ due to improved load balance from mode-replication and also from temporal locality. Note that super-linear speedup is observed, reaching $2.3\times$ with two cores. The $p\times p$ tiling of \mathcal{R} for p threads results in a smaller portion of the factor matrices accessed at a time, improving temporal locality. Interestingly, super-linear speedup was also observed in the original CCD++ evaluation for matrices [36].

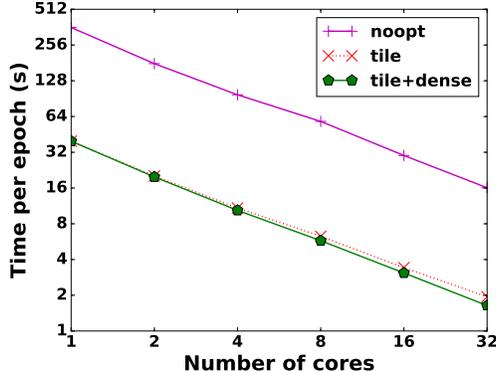


Figure 7: Effects of ALS optimizations during a rank-10 factorization of the Yahoo! tensor. **noopt** is a baseline ALS implementation with a CSF data structure. **tile** delays rank-1 updates to use the BLAS-3 *dsyrk* routine. **tile+dense** includes **tile** and also dense mode replication.

Despite the improved speedup, CCD++ sees little improvement after 16 cores (one full socket) due to NUMA effects and the memory-bound nature of the algorithm. Compared to ALS, CCD++ performs a factor of F fewer FLOPs on every non-zero that is accessed. Additionally, CCD++ being a column-oriented method requires MF passes per epoch over the sparsity structure of \mathcal{R} , compared to M times for ALS and one time for SGD. We show in Section 6.8 that this limitation can be solved by simply using one MPI process per socket.

6.3.3. SGD

Figure 9 shows the effects of coarse-grained randomization on SGD convergence. We use a full random traversal of the tensor as a baseline and compare against two CSF configurations. The first configuration, CSF-S, sorts the modes in non-decreasing order with the smallest mode placed at the top of the data structure. CSF-S is the default mode ordering used by SPLATT due to it typically resulting in the highest level of compression [33]. CSF-L sorts the modes in non-increasing order, with the longest mode placed at the top of the CSF structure. CSF-L effectively trades additional storage for increased randomization and higher degrees of parallelism. CSF-S has the fastest per-epoch runtime but fails to converge. CSF-L exhibits similar per-epoch convergence compared to the baseline, but the computational savings afforded by the CSF data structure results in a faster time-to-solution. We therefore use CSF-L as the default in the remaining experiments.

Figure 10 shows the effects of stratification on convergence. We evaluate three algorithms across two factorizations: fully asynchronous, fully stratified, and a hybrid algorithm with 16 stratum layers, totaling 256 strata. The hybrid configuration outperforms both baselines for the rank-10 factorization and converges to a higher quality solution in less time. All three algorithms are competitive for rank-40, but the hybrid ultimately reaches a lower RMSE.

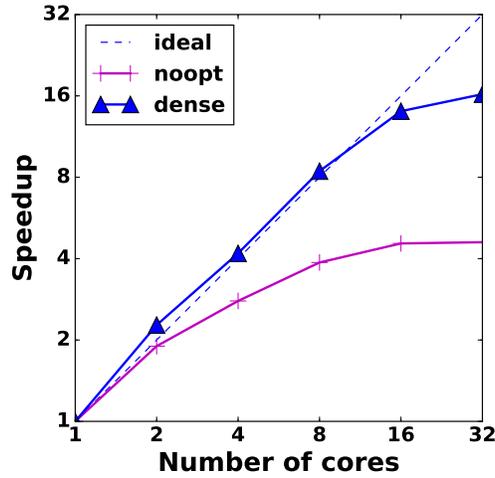


Figure 8: Effects of CCD++ optimizations during a rank-10 factorization of the Yahoo! tensor. **noopt** is a baseline CCD++ implementation with the CSF data structure. **dense** uses dense mode replication.

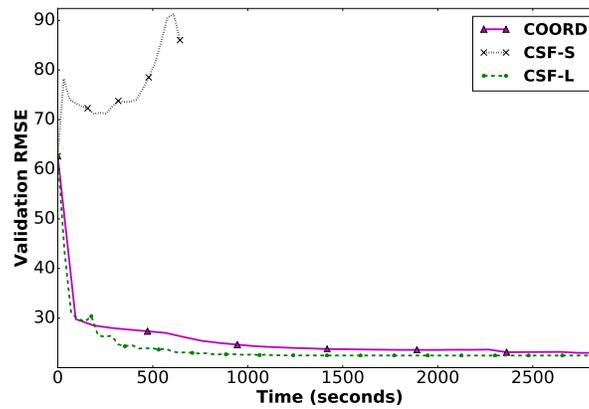


Figure 9: Effects of randomization strategy on SGD convergence rate during a serial rank-40 factorization of the Yahoo! tensor. **COORD** uses complete randomization on a coordinate form tensor. **CSF-S** randomizes over the shortest mode. **CSF-L** randomizes over the longest mode.

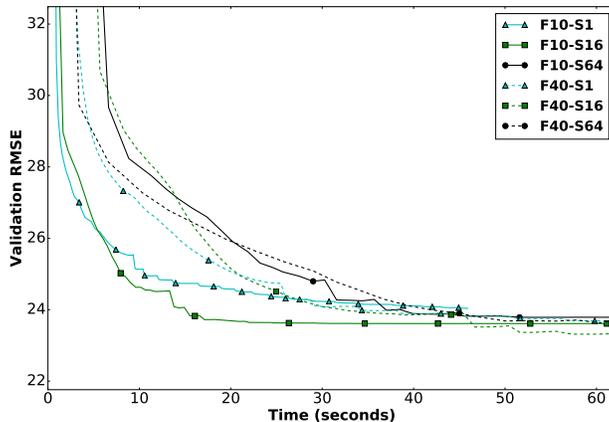


Figure 10: Convergence rates for SGD parallelization strategies using 32 compute nodes on the Yahoo! dataset. \mathbf{F} denotes the rank of the factorization. \mathbf{S} denotes the number of stratum layers, scaling from fully asynchronous (S1) to fully stratified (S64).

6.4. Communication Volume

Figure 11 shows the average communication volume per epoch. CCD++ consistently has a lower communication volume than SGD and ALS due to its medium-grained decomposition. The communication volume for SGD increases until four nodes (eight ranks) are used, and then sharply decreases. The communication volume of SGD scales with the number of strata used, which we limit to 64. Recall that SGD uses S^{M-1} strata for an M -mode tensor. Therefore, we see volume increase until we reach the maximum number of strata. From that point communication is limited to within strata, and we see communication volume decrease. The communication volume of ALS increases until eight nodes and then stays near constant. ALS uses a coarse-grained decomposition which in the worst case requires each communication of entire factors per epoch.

6.5. Strong Scaling

Figure 12 shows strong scaling results. For SGD, we again limit the number of strata to 64 which provides a good trade-off between convergence rate and parallelism. As discussed in the evaluation of communication volume, SGD primarily introduces overhead until the maximum number of strata is reached, and after that point we begin to overcome the communication overheads. SGD only scales to eight nodes on the Amazon tensor. This is because Amazon is significantly more sparse than Netflix and Yahoo! and as a result SGD performs less work per stratum, resulting in high communication costs.

ALS is unable to process the Amazon and Patents tensors on as few nodes as SGD and CCD++ due to it requiring three copies of \mathcal{R} during factorization. Neither ALS nor CCD++ are consistently faster per epoch at 32 nodes. CCD++ begins slower on all datasets but out-scales ALS on all but Yahoo!, which does not have enough non-zeros to effectively parallelize over 64 nodes (2048 cores).

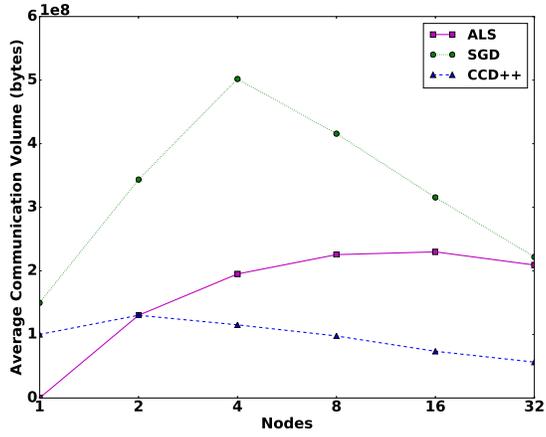


Figure 11: Average communication volume per node on the Yahoo! dataset. CCD++ and SGD use two MPI ranks per node and ALS uses one.

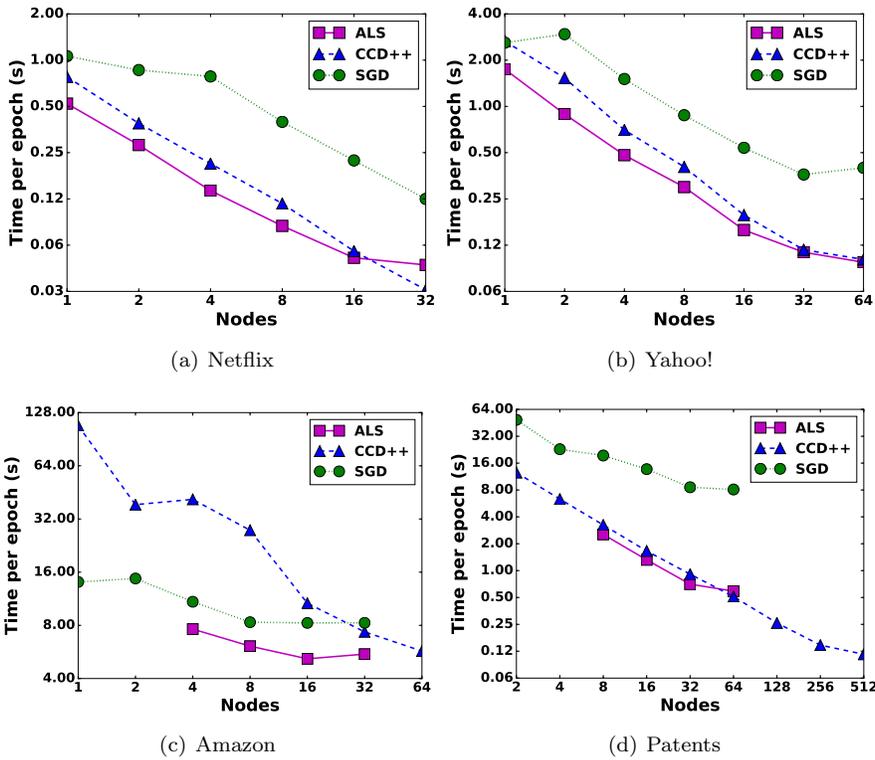


Figure 12: Strong scaling the optimized ALS, SGD, and CCD++ algorithms. Each node has 32 cores.

The disparity between CCD++ and other algorithms on Amazon is due to its large size being more taxing on memory bandwidth. CCD++ is able to scale through 512 nodes (16384 cores) on Patents due to its high density, resulting in a small amount of communication and memory bandwidth relative to the computational load. ALS is unable to scale past 32 nodes on Patents due to its short mode lengths, which result in a load-imbalanced coarse-grained decomposition. The second and third modes are too long to treat as dense (i.e., the storage costs of normal equations are prohibitive). In contrast, the medium-grained decomposition used by CCD++ is able to find a load balanced distribution on Patents.

6.6. Rank Scaling

Figure 13 shows the effects varying the rank of the factorization. We scale from rank 10 to 80 on the Yahoo! dataset. CCD++ and SGD both have $\mathcal{O}(F \text{nnz}(\mathcal{R}))$ complexity, so we expect the runtime to increase by $8\times$ as we scale F . ALS, on the other hand, has complexity $\mathcal{O}(F^2 \text{nnz}(\mathcal{R}) + IF^3)$. The $F^2 \text{nnz}(\mathcal{R})$ term will dominate in most scenarios because users are interested in low rank factorizations and because $I \ll \text{nnz}(\mathcal{R})$ for most tensors. Under this assumption, we expect the runtime of ALS to increase by a factor of $80^2/10^2 = 64$.

CCD++ sees the expected linear increase in runtime on both 32 and 1024 cores: $7.9\times$ and $7.6\times$, respectively. SGD scales sub-linearly and only sees $2.4\times$ and $4.2\times$ increases on 32 and 1024 cores, respectively. The sub-linear effects are due to the way SGD accesses the matrix factors. SGD only ever accesses entire rows of the factors, leading to spatial locality and vectorized inner loops. We do not see the same effects for CCD++ because it accesses the factors in a strided manner that is dependent on the sparsity pattern. Additionally, CCD++ must traverse the sparsity pattern of the tensor F times for each mode, compared to once for SGD. Surprisingly, ALS only sees $9.8\times$ and $10.1\times$ increase in runtime at 32 and 1024 cores, respectively. While the work does increase quadratically, all of the quadratic functions are performed by BLAS-3 routines on small, dense matrices. The work that depends on the sparsity pattern of \mathcal{R} is an MTTKRP operation, which has the same spatial locality as SGD and a cheaper complexity of $\mathcal{O}(F \text{nnz}(\mathcal{R}))$. The BLAS-3 routines will eventually out-scale the cost of the MTTKRP operation, but factorizations of such a high rank are unlikely to be useful to a domain expert.

6.7. Mode Scaling

Figure 14 shows the scalability of our algorithms on the Outpatient dataset as we increase the number of tensor modes while keeping the number of non-zeros constant. ALS sees a roughly linear increase in runtime, matching the computational complexity in Table 1. The runtime of the sixth mode increases super-linearly, which we attribute to the sixth mode being longer than most others and ALS having a $\mathcal{O}(IF^3)$ complexity component. CCD++ exhibits severe slowdown as the number of modes is increased. This is due to CCD++ doing

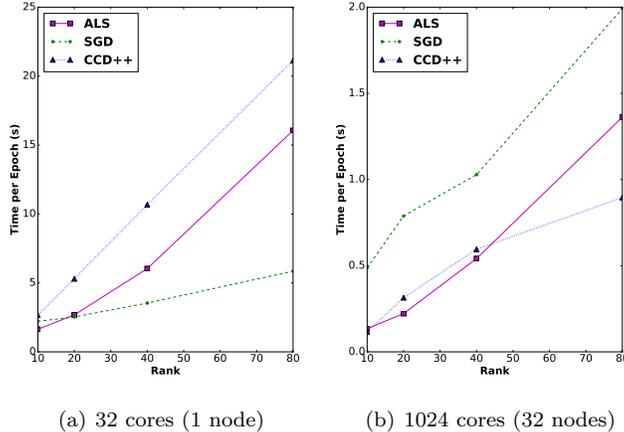


Figure 13: Effects of increasing factorization rank on the Yahoo! dataset.

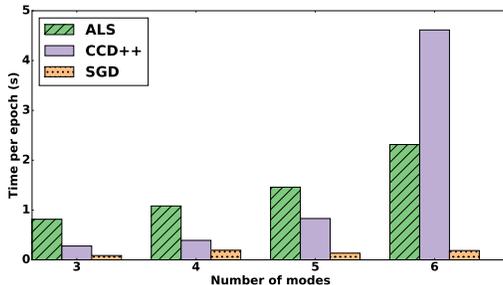


Figure 14: Average time per epoch with 16 nodes while scaling the number of modes using the Outpatient dataset.

MF passes over \mathcal{R} per epoch and performing only $\mathcal{O}(\text{nnz}(\mathcal{R}))$ work per pass. The memory-bound nature of CCD++ is exaggerated as the number of modes increases. SGD has a nearly constant runtime due to it only performing one pass over \mathcal{R} per epoch, regardless of the number of modes. Additionally, higher-order tensors such as Outpatient have several dense modes which will exhibit high temporal locality, leaving the system’s memory bandwidth free for streaming through the single representation of \mathcal{R} . SGD appears to be an attractive choice for higher-order tensors. We cautiously recommend it, however, because situations may arise in which many stratum layers are required to maintain convergence, negatively impacting performance.

6.8. Comparison Against the State-of-the-Art

In Figure 15 we compare our ALS and CCD++ algorithms against the state-of-the-art MPI implementations [13]. We scale from 1 to 1024 cores on the

Yahoo! tensor with $F=10$. We use one MPI rank per node for opt-ALS due to the high speedup it achieves on 32 cores and also due to the high communication volume that comes with a coarse-grained decomposition. Throughout the comparison we refer to the existing implementations as “base-ALS” and “base-CCD++” and our own as “opt-ALS” and “opt-CCD++”.

On one core, opt-ALS is $7\times$ faster than base-ALS due to the BLAS-3 performance. opt-ALS then scales to achieve $353\times$ speedup at 1024 cores, compared to the $13.5\times$ speedup of base-ALS. The improvements in speedup are due to the coarse-grained decomposition used by opt-ALS which reduces the communication volume from $\mathcal{O}(IF^2)$ to $\mathcal{O}(IF)$ words. The difference in communication requirements is observed in the ratio of communication to computation: base-ALS spends 95% of the total runtime communicating, compared to opt-ALS which spends 39% of its runtime communicating. opt-ALS is $185\times$ faster than base-ALS when both use 1024 cores.

Serial opt-CCD++ is $2.2\times$ faster than base-CCD++ due to the operation reduction and improved cache locality resulting from the CSF data structure. The locality improvements are also present in the ALS results, but due to ALS being a compute-bound algorithm they are not observed except for very small values of F . On 1024 cores, opt-CCD++ and base-CCD++ achieve $685\times$ and $74.2\times$ speedup, respectively. The disparity in speedups is attributed to the large amount of communication performed by base-CCD++, in which 69% of the total runtime is spent in communication routines. In comparison, opt-CCD++ spends 25% of the total runtime communicating. The medium-grained decomposition used by opt-CCD++ results in a smaller communication volume than the arbitrary decomposition used by base-CCD++. Communication volume is further reduced by utilizing *sparse* communication and only sends an updated value to the processes that require it.

We also note that with the configuration using one MPI rank per socket, opt-CCD++ achieves a $30\times$ speedup on 32 cores compared to $16.2\times$ with a pure OpenMP configuration.

6.9. Convergence

We now evaluate the time-to-solution of ALS, CCD++, and SGD. Figure 16 shows convergence of our optimized algorithms using 1, 32, and 1024 cores. SGD is the most successful algorithm in a serial setting. For $F=10$, SGD converges within 1500 seconds and achieves a quality that takes ALS over twice as long to reach. SGD sees a similar advantage when $F=40$, but ALS ultimately reaches a higher quality solution shortly after SGD converges.

On a single node, ALS outperforms SGD and CCD++ for both $F=10$ and $F=40$. Since ALS has the fastest per-epoch times in addition to making the most progress per epoch, it is the recommended algorithm for small-to-moderate node counts. CCD++ is also competitive in the multi-core environment and has the added benefit of having a smaller memory footprint due to only storing a single copy of \mathcal{R} . Trends continue as we move to a large-scale distributed system and ALS narrowly bests CCD++ for $F=40$. While ALS still converges

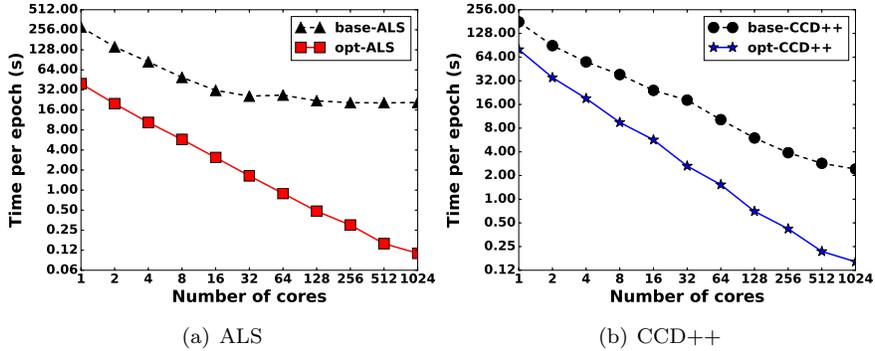


Figure 15: Comparison of the presented ALS and CCD++ algorithms (prefixed **opt**) against the state-of-the-art MPI implementations (prefixed **base**) on a rank-10 factorization of the Yahoo! tensor.

faster than SGD and CCD++, CCD++ is more scalable in distribute-memory environments due to its lower communication volume.

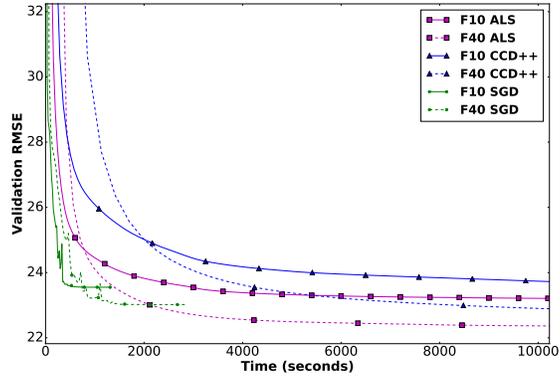
6.10. Effects of Randomization on ALS and CCD++

We now explore the convergence benefits offered by randomization during ALS and CCD++. Figure 17 shows the convergence per epoch when applying mode-level randomization to ALS and combinations of mode- and rank-level randomization to CCD++. In each case, we show both the value of the objective function and the validation RMSE. Observing the objective function offers insights into the effects on the actual optimization algorithm, whereas observing the validation RMSE measures the predictive abilities of the resulting model (i.e., its usefulness to a domain specialist).

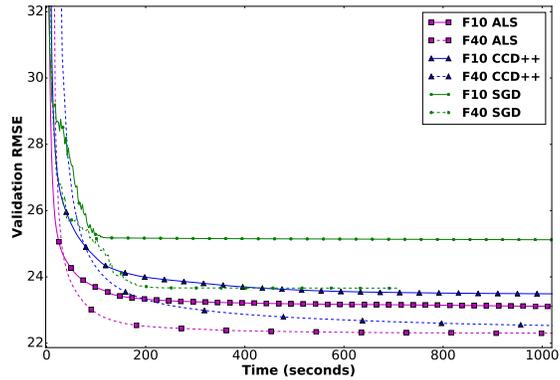
Both algorithms benefit from randomization on the Movielens dataset. Randomized ALS requires 66% fewer epochs to arrive at a validation RMSE which is 0.3% lower than cyclic. Likewise, all variations of randomization improve the validation RMSE of CCD++. The best performance is achieved by rank-level randomization, which converges 16% faster than cyclic and improves the solution again by 0.3%. The objective function is improved in all cases, and therefore the gains are attributed to an optimization algorithm which converges more quickly to a better local minimum.

Netflix likewise benefits from randomization in terms of both convergence rate and validation RMSE. Interestingly, mode-level randomization achieves the best validation RMSE for both ALS and CCD++, but at the same time arrives at the worst objective value. Randomization allows the optimization algorithm to learn a more general, albeit less optimized solution.

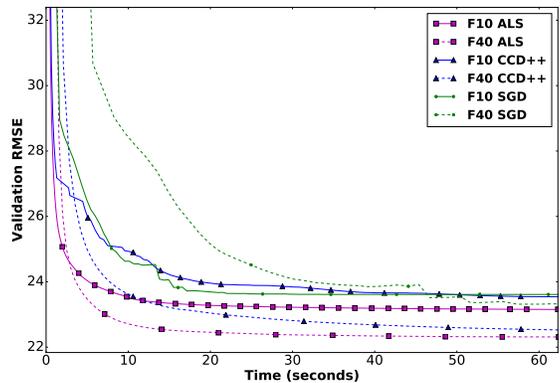
We find that algorithms which use mode-level randomization tend to learn models with empty columns, reducing the rank of the factorization. We explore this phenomenon in Figure 18. We plot the number of non-zero rank-one components, which refer to as the *effective rank* of the factorization. ALS and



(a) Serial



(b) 32 cores (1 node)



(c) 1024 cores (32 nodes)

Figure 16: Convergence rates for parallel methods on the Yahoo! dataset with factorization ranks 10 (F10) and 40 (F40). Ticks are placed every five epochs.

CCD++ learn models with ranks 32 and 25, respectively. The lower-rank models do not fit the training data as well, and thus have a higher objective value. However, due to their lower complexity and sufficiently good model of the data, they are able to better predict unknown entries in the validation set.

Lastly, Yahoo! does not benefit from randomization and the algorithms which use mode-level randomization achieve notably worse objective values and validation RMSEs. The lower-rank factorizations learned via randomization are not able to capture the training data sufficiently well to result in a predictive model.

7. Conclusions

We explored the design and implementation of three optimization algorithms for tensor completion: ALS, SGD, and CCD++. We focused on modern architectures with shared- and distributed-memory parallelism. We solved issues such as memory- and operation-efficiency, cache locality, load balance, and communication. Following our improvements, we achieve speedups up through 16384 cores and are up to $185\times$ faster than state-of-the-art parallel methods. We further improved the convergence rates of ALS and CCD++ by introducing randomization during the optimization procedure.

When comparing algorithms for tensor completion, time-to-solution is the most important detail for end users. We compared convergence rates in three configurations: serial, a multi-core system, and a large-scale distributed system and showed that no algorithm performs best in all three environments. SGD is most competitive in a serial environment, ALS is recommended for shared-memory systems, and both ALS and CCD++ are competitive on distributed systems. Using our developments, the time-to-solution on large datasets is reduced from hours to seconds.

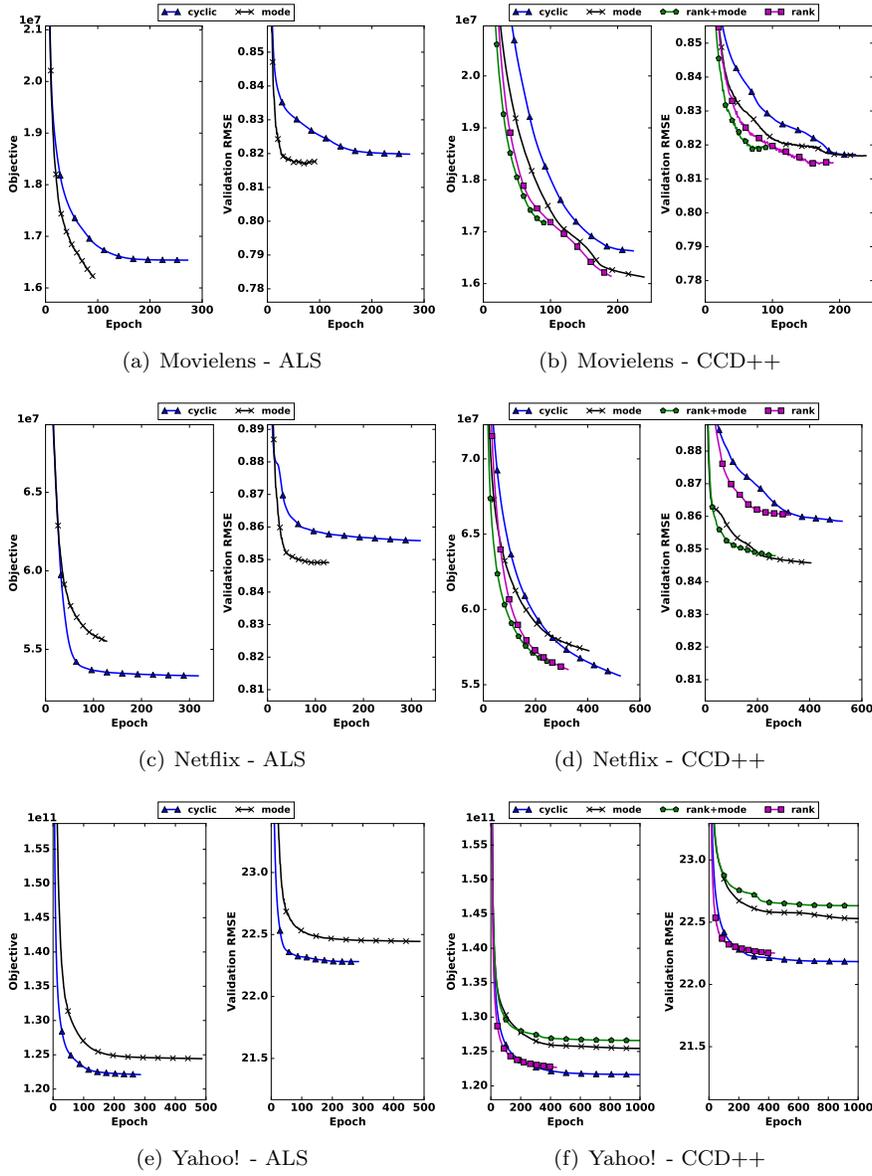


Figure 17: Effects of randomization on ALS and CCD++ convergence during rank-40 factorizations of three datasets. **cyclic** uses no randomization, **mode** uses mode-level, **rank** uses rank-level, and **rank+mode** uses both rank- and mode-level randomization.

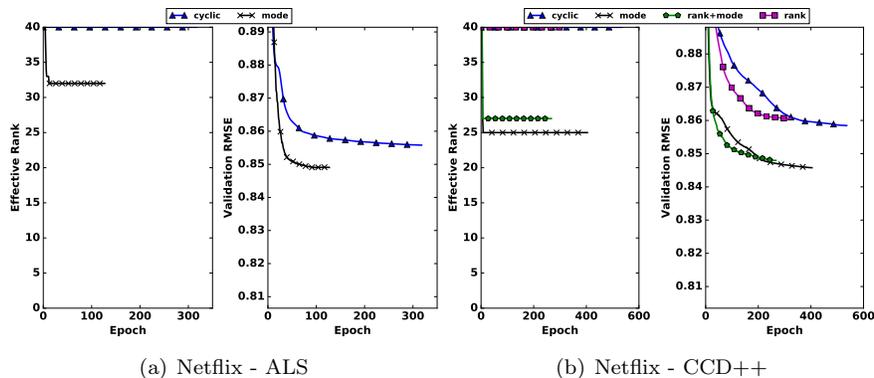


Figure 18: The number of non-zero rank-one components (**effective rank**) during a rank-40 factorization of the Netflix dataset.

Acknowledgments

This work is an extended and revised version of a preliminary conference paper [32]. The authors would like to thank Mikhail Smelyanskiy for valuable discussions, Jeff Hammond for generously donating computing time at NERSC, and Karlsson et al. for sharing source code used for evaluation. This work was supported in part by NSF (IIS-0905220, OCI-1048018, CNS-1162405, IIS-1247632, IIP-1414153, IIS-1447788), Army Research Office (W911NF-14-1-0316), a University of Minnesota Doctoral Dissertation Fellowship, Intel Software and Services Group, and the Digital Technology Center at the University of Minnesota. Access to research and computing facilities was provided by the Digital Technology Center and the Minnesota Supercomputing Institute. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

- [1] Muthu Baskaran, Benoît Meister, Nicolas Vasilache, and Richard Lethin. Efficient and scalable computations with sparse tensors. In *High Performance Extreme Computing (HPEC), 2012 IEEE Conference on*, pages 1–6. IEEE, 2012.
- [2] Roberto Battiti. Accelerated backpropagation learning: Two optimization methods. *Complex systems*, 3(4):331–342, 1989.
- [3] James Bennett and Stan Lanning. The netflix prize. In *Proceedings of KDD cup and workshop*, volume 2007, page 35, 2007.

- [4] J. Douglas Carroll and Jih-Jie Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of “Eckart-Young” decomposition. *Psychometrika*, 35(3):283–319, 1970.
- [5] Center for Medicare and Medicaid Services. CMS data entrepreneurs synthetic public use file (DE-SynPUF), 2010.
- [6] Joon Hee Choi and S. Vishwanathan. DFacTo: Distributed factorization of tensors. In *Advances in Neural Information Processing Systems*, pages 1296–1304, 2014.
- [7] Gideon Dror, Noam Koenigstein, Yehuda Koren, and Markus Weimer. The yahoo! music dataset and kdd-cup’11. In *KDD Cup*, pages 8–18, 2012.
- [8] Olivier Fercoq and Peter Richtárik. Accelerated, parallel, and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- [9] Mark Gates, Hartwig Anzt, Jakub Kurzak, and Jack Dongarra. Accelerating collaborative filtering using concepts from high performance computing. In *Big Data (Big Data), 2015 IEEE International Conference on*, pages 667–676. IEEE, 2015.
- [10] Rainer Gemulla, Erik Nijkamp, Peter J. Haas, and Yannis Sismanis. Large-scale matrix factorization with distributed stochastic gradient descent. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 69–77. ACM, 2011.
- [11] F. Maxwell Harper and Joseph A. Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 5(4):19, 2016.
- [12] Joyce C. Ho, Joydeep Ghosh, and Jimeng Sun. Marble: high-throughput phenotyping from electronic health records via sparse nonnegative tensor factorization. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 115–124. ACM, 2014.
- [13] Lars Karlsson, Daniel Kressner, and André Uschmajew. Parallel algorithms for tensor completion in the cp format. *Parallel Computing*, 57:222–234, 2016.
- [14] Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. PhD thesis, Inria-Research Centre Grenoble–Rhône-Alpes, 2015.
- [15] Oguz Kaya and Bora Uçar. Scalable sparse tensor decompositions in distributed memory systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, page 77. ACM, 2015.

- [16] Henk A. L. Kiers. Towards a standardized notation and terminology in multiway analysis. *Journal of chemometrics*, 14(3):105–122, 2000.
- [17] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- [18] Jiajia Li, Casey Battaglino, Ioakeim Perros, Jimeng Sun, and Richard Vuduc. An input-adaptive and in-place approach to dense tensor-times-matrix multiply. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, page 76. ACM, 2015.
- [19] Julian McAuley and Jure Leskovec. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172. ACM, 2013.
- [20] Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [21] Dimitri Nion and Nicholas D. Sidiropoulos. Tensor algebra and multidimensional harmonic retrieval in signal processing for mimo radar. *Signal Processing, IEEE Transactions on*, 58(11):5693–5705, 2010.
- [22] Fabio Petroni and Leonardo Querzoni. GASGD: stochastic gradient descent for distributed asynchronous matrix completion via graph partitioning. In *Proceedings of the 8th ACM Conference on Recommender systems*, pages 241–248. ACM, 2014.
- [23] Ali Pinar and Cevdet Aykanat. Fast optimal load balancing algorithms for 1d partitioning. *Journal of Parallel and Distributed Computing*, 64(8):974–996, 2004.
- [24] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 693–701, 2011.
- [25] Weijia Shao. Tensor completion. Master’s thesis, Universität des Saarlandes Saarbrücken, 2012.
- [26] Yue Shi, Alexandros Karatzoglou, Linas Baltrunas, Martha Larson, Alan Hanjalic, and Nuria Oliver. Tfmap: optimizing map for top-n context-aware recommendation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 155–164. ACM, 2012.
- [27] Kijung Shin and U. Kang. Distributed methods for high-dimensional and large-scale tensor factorization. In *Data Mining (ICDM), 2014 IEEE International Conference on*, pages 989–994, Dec 2014.

- [28] Shaden Smith, Jee W. Choi, Jiajia Li, Richard Vuduc, Jongsoo Park, Xing Liu, and George Karypis. FROSTT: The formidable repository of open sparse tensors and tools, 2017.
- [29] Shaden Smith and George Karypis. SPLATT: The Surprisingly Parallel sparse Tensor Toolkit. <http://cs.umn.edu/~splatt/>.
- [30] Shaden Smith and George Karypis. Tensor-matrix products with a compressed sparse tensor. In *Proceedings of the 5th Workshop on Irregular Applications: Architectures and Algorithms*, page 7. ACM, 2015.
- [31] Shaden Smith and George Karypis. A medium-grained algorithm for distributed sparse tensor factorization. In *30th IEEE International Parallel & Distributed Processing Symposium (IPDPS'16)*, 2016.
- [32] Shaden Smith, Jongsoo Park, and George Karypis. An exploration of optimization algorithms for high performance tensor completion. In *Proceedings of the 2016 ACM/IEEE conference on Supercomputing*, 2016.
- [33] Shaden Smith, Niranjay Ravindran, Nicholas D. Sidiropoulos, and George Karypis. SPLATT: Efficient and parallel sparse tensor-matrix multiplication. In *International Parallel & Distributed Processing Symposium (IPDPS'15)*, 2015.
- [34] Jos M.F. ten Berge and Nicholas D. Sidiropoulos. On uniqueness in can-decomp/parafac. *Psychometrika*, 67(3):399–409, 2002.
- [35] Yichen Wang, Robert Chen, Joydeep Ghosh, Joshua C. Denny, Abel Kho, You Chen, Bradley A. Malin, and Jimeng Sun. Rubik: Knowledge guided tensor factorization and completion for health data analytics. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1265–1274. ACM, 2015.
- [36] Hsiang-Fu Yu, Cho-Jui Hsieh, Si Si, and Inderjit Dhillon. Scalable coordinate descent approaches to parallel matrix factorization for recommender systems. In *2012 IEEE 12th International Conference on Data Mining*, pages 765–774. IEEE, 2012.
- [37] Qiang Zhang, Michael W. Berry, Brian T. Lamb, and Tabitha Samuel. A parallel nonnegative tensor factorization algorithm for mining global climate data. In *Computational Science–ICCS 2009*, pages 405–415. Springer, 2009.
- [38] Yunhong Zhou, Dennis Wilkinson, Robert Schreiber, and Rong Pan. Large-scale parallel collaborative filtering for the netflix prize. In *Algorithmic Aspects in Information and Management*, pages 337–348. Springer, 2008.