

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Citation:

T. B. Rolinger, T. A. Simon and C. D. Krieger, "Parallel Sparse Tensor Decomposition in Chapel," 2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), Vancouver, BC, Canada, 2018, pp. 896-905, doi: 10.1109/IPDPSW.2018.00143.

DOI:

<https://doi.org/10.1109/IPDPSW.2018.00143>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Parallel Sparse Tensor Decomposition in Chapel

Thomas B. Rolinger^{*†}, Tyler A. Simon^{*}, and Christopher D. Krieger^{*}

^{*} Laboratory for Physical Sciences, University of Maryland, College Park, MD USA

[†] Department of Computer Science, University of Maryland, College Park, MD USA
tbrolin@cs.umd.edu, {tasimon,krieger}@lps.umd.edu

Abstract—In big-data analytics, using tensor decomposition to extract patterns from large, sparse multivariate data is a popular technique. Many challenges exist for designing parallel, high performance tensor decomposition algorithms due to irregular data accesses and the growing size of tensors that are processed. There have been many efforts at implementing shared-memory algorithms for tensor decomposition, most of which have focused on the traditional C/C++ with OpenMP framework. However, Chapel is becoming an increasingly popular programming language due to its expressiveness and simplicity for writing scalable parallel programs. In this work, we port a state of the art C/OpenMP parallel sparse tensor decomposition tool, SPLATT, to Chapel. We present a performance study that investigates bottlenecks in our Chapel code and discusses approaches for improving its performance. Also, we discuss features in Chapel that would have been beneficial to our porting effort. We demonstrate that our Chapel code is competitive with the C/OpenMP code for both runtime and scalability, achieving 83%-96% performance of the original code and near linear scalability up to 32 cores.

Keywords—Chapel, OpenMP, sparse, tensor decomposition, performance study

I. INTRODUCTION

With the growing need to process large amounts of sparse, multi-way data, tensor decomposition has become a popular technique within large-scale data analytics. Due to irregular data accesses and the growing scale of tensors, designing parallel, high performance tensor algorithms is difficult. While there exists a wealth of research in implementing tensor algorithms for shared memory systems [1], [2], most implementations are rooted in the traditional C/C++ with OpenMP framework.

Chapel [3] is an emerging open-source parallel programming language that is geared towards enabling high user productivity without sacrificing scalability or code performance. With the expressiveness of high-level language constructs in Chapel, users can focus more on the algorithm they are implementing rather than low-level parallelization details. In recent work, Chapel has been shown to have competitive, and in some cases, higher performance than traditional languages and parallel libraries [4], [5], [6].

As the core algorithms within parallel sparse tensor decomposition are complex, immature, and constantly being refined, programmers would benefit from using a language such as Chapel to quickly prototype algorithms while still writing clean, easy to read code that is maintainable and extensible. However, maintaining competitive performance with implementations written in traditional languages is crucial.

In this work, we port a state of the art parallel sparse tensor decomposition tool, SPLATT [7], from C/OpenMP to Chapel and evaluate its performance. We also describe our experience in porting the code and identify features in Chapel that were most beneficial as well as absent features that would have been useful to us. This paper makes the following contributions:

- 1) Reimplements SPLATT, a sparse tensor decomposition tool, in Chapel. SPLATT is currently at the forefront of performance among tensor decomposition implementations. To the best of our knowledge, this is the first work to implement tensor decomposition in Chapel.
- 2) Presents a performance study that analyzes and compares our Chapel code to the reference implementation of SPLATT written in C. We discuss approaches for improving the performance of the Chapel code and identify bottlenecks in our code related not only to the Chapel language but also the tasking layer used by the Chapel runtime.
- 3) Demonstrates that the Chapel code achieves 83%-96% of the performance of the C/OpenMP code and near linear scalability up to 32 cores.

The rest of this paper is organized as follows. Section II presents a brief overview of the Chapel programming language, focusing on the features most relevant to our work. Background information regarding tensors and tensor decomposition is given in Section III, as well as a description of SPLATT. We discuss our approach and experience porting SPLATT to Chapel in Section IV. Section V presents our performance study and discusses steps taken to improve the Chapel code. Section VI presents work related to designing and implementing tensor decomposition algorithms as well as evaluating and optimizing Chapel programs. Finally, we provide concluding remarks in Section VII.

II. CHAPEL PROGRAMING LANGUAGE OVERVIEW

Chapel is an open-source programming language that contains first-class parallel constructs for both shared- and distributed-memory execution. The goal of Chapel is to provide a productive environment for users to write highly efficient and scalable parallel codes. In this section, we briefly present the features of Chapel that are most relevant to our work. As this work focuses only on single-node execution within Chapel, we will omit details regarding Chapel's multi-node or multi-locale features. For a more extensive description of Chapel and its features, we refer readers to Chamberlain's

overview of Chapel [3] or to the Chapel Documentation¹.

Parallelism within a Chapel program is expressed through the creation and execution of *tasks*, which are defined as units of concurrent computation. A Chapel program begins as a single task in which users can explicitly or implicitly create additional tasks to perform parallel computation. Chapel treats threads as system-level resources that are used to execute tasks and relies on a *tasking layer* to map tasks to threads. From this perspective, thread creation and management are completely abstracted from the user.

Within Chapel, parallelism can be expressed in a task or data parallel manner. The `coforallall` construct provides task-parallelism and allows users to explicitly create a specified number of concurrent tasks. The order in which the tasks are executed is non-deterministic as they are run in parallel but the statements within a given task are executed serially. The OpenMP analogue of a `coforallall` is the `omp parallel` directive. Listing 1 provides an example of a `coforallall` and Listing 2 shows the equivalent code in C/OpenMP, where `numTasks` and `numThreads` are equivalent.

```
1 coforall tid in 0..numTasks-1 {
2   writeln("Hello from Task ", tid);
3   if tid == 0 {
4     writeln("Extra hello from master: ", tid);
5   }
6 }
```

Listing 1. Chapel coforall task-parallel construct

```
1 #pragma omp parallel num_threads(numThreads)
2 {
3   int tid = omp_get_thread_num();
4   printf("Hello from Task %d\n", tid);
5   #pragma omp master
6   {
7     printf("Extra hello from master: %d\n", tid);
8   }
9 }
```

Listing 2. OpenMP equivalent of coforall construct shown in Listing 1

The `forall` loop construct in Chapel allows users to express data-parallelism, where the iterations of the loop are blocked and assigned to different tasks to execute in parallel. The number of tasks assigned to execute a given `forall` loop is a function of the `iterand` expression. The OpenMP equivalent of the `forall` loop is the `omp parallel for` directive. Listing 3 shows an example of the `forall` loop and Listing 4 provides the C/OpenMP equivalent. In both cases, each element of `myArray` is incremented by 1 in parallel. Furthermore, Chapel supports whole-array operations that can be performed implicitly in parallel, so the `forall` loop can be equivalently expressed with the single statement shown on line 4 in Listing 3.

```
1 forall elem in myArray {
2   elem += 1;
3 }
4 myArray += 1; // equivalent to above forall loop
```

Listing 3. Chapel forall data-parallel loop

```
1 #pragma omp parallel for
2 for (int i = 0; i < arraySize; i++) {
3   myArray[i] += 1;
4 }
```

Listing 4. OpenMP equivalent of forall loop

Chapel also provides a mechanism for read/write protection of variables shared among tasks. A variable that is declared with the `sync` qualifier is associated with a state that is either full or empty. The state must be empty before a value can be written to the variable and the state must be full before the variable can be read. Attempts to access a `sync` variable in the incorrect state cause the task to block until the state changes. The `atomic` type qualifier may also be applied to integers, reals and boolean variables. Chapel supports the atomic operations commonly provided by other languages, such as test and set, compare, add and subtract. While Chapel does not have a built-in lock/mutex construct, users can achieve the same effect through the use of `sync` or `atomic` variables.

While Chapel encourages the use of high-level language features, it does provide several mechanism for lower-level C interoperability. Some of these are illustrated in Listing 5. Of particular relevance to our work is the ability to retrieve a C-pointer to a Chapel array or matrix via the `c_ptrTo` procedure, shown on line 4. Once a C-pointer is retrieved, operations such as pointer arithmetic and pointer aliasing can be used, as shown on line 6. Upon execution of this code, the original Chapel matrix will be modified (all elements set to 1).

```
1 var rows, cols = 3;
2 var myDomain : domain(2) = {0..rows-1, 0..cols-1};
3 var myMatrix : [myDomain] int = 0;
4 var myPtr = c_ptrTo(myMatrix);
5 for row in 0..rows-1 {
6   var myRowPtr = myPtr + (row * cols);
7   for col in 0..cols-1 {
8     myRowPtr[col] = 1;
9   }
10 }
```

Listing 5. Use of `c_ptrTo` in Chapel

III. TENSOR DECOMPOSITION

The need to process large amounts of sparse, high dimensional data is common in fields such as signal processing and data mining. Analysts in these fields often need to identify previously unknown relationships among elements in the data. *Tensors*, which are matrices extended to three or more dimensions, are a natural way to model such data and *tensor decomposition* is a popular technique for extracting patterns from large, multi-way data. Tensor decomposition is often presented as the higher-order analogue of matrix singular value decomposition (SVD) [8]. For the remainder of this paper, we refer to the dimensions of a tensor as its *modes* and the number of modes in a tensor as its *order*. We refer the reader to the extensive surveys by Kolda and Bader [8] and Sidiropoulos et al. [9] for more detailed information about tensors and tensor decomposition.

¹<https://chapel-lang.org/docs/latest/>

The Canonical Decomposition/Parallel Factorization (CP) algorithm is the most commonly used approach to extend SVD to tensors. CP computes a set of components whose sum approximates the original tensor. The number of components, R , is referred to as the *rank* of the decomposition. The most popular approach to computing these components is alternating least squares (ALS). A basic sketch of CP-ALS is shown in Algorithm 1 for a 3^{rd} order tensor \mathcal{X} . While CP-ALS is general and can operate on tensors of arbitrary order, we restrict our discussion to 3^{rd} order tensors for simplicity. For each mode n , I_n denotes its length and the dense matrix $\mathbf{A}^{(n)}$ represents the *factor matrix* for that mode and consists of I_n rows and R columns. The columns of the factor matrices represent the rank-one components of the decomposition.

It has been shown that most of the computational and storage complexity of CP-ALS stems from lines 5, 8 and 11 in Algorithm 1, which multiply the mode n unfolded, or *matricized*, tensor $\mathcal{X}^{(n)}$ with the Khatri-Rao product (written as \odot) of the factor matrices [10], [7]. This operation is referred to as the matricized tensor times Khatri-Rao product (MTTKRP). Unless the factors are very sparse, this product is totally dense due to high fill-in and can easily require many times more memory than the original tensor. Due to the performance characteristics of the MTTKRP, researchers have focused on designing efficient implementations with respect to both execution time and memory [11], [2].

Algorithm 1: CP-ALS for 3^{rd} order tensors

```

1: Procedure CP-ALS( $\mathcal{X}$ ,  $R$ )
2: initialize  $\mathbf{A}^{(n)} \in \mathbb{R}^{I_n \times R}$  for  $n = 1, 2, 3$ 
3: repeat
4:    $\mathbf{V} \leftarrow \mathbf{A}^{(2)T} \mathbf{A}^{(2)} * \mathbf{A}^{(3)T} \mathbf{A}^{(3)}$ 
5:    $\mathbf{A}^{(1)} \leftarrow \mathcal{X}^{(1)} (\mathbf{A}^{(3)} \odot \mathbf{A}^{(2)}) \mathbf{V}^\dagger$ 
6:   normalize columns of  $\mathbf{A}^{(1)}$  (storing norms as  $\lambda$ )
7:    $\mathbf{V} \leftarrow \mathbf{A}^{(1)T} \mathbf{A}^{(1)} * \mathbf{A}^{(3)T} \mathbf{A}^{(3)}$ 
8:    $\mathbf{A}^{(2)} \leftarrow \mathcal{X}^{(2)} (\mathbf{A}^{(3)} \odot \mathbf{A}^{(1)}) \mathbf{V}^\dagger$ 
9:   normalize columns of  $\mathbf{A}^{(2)}$  (storing norms as  $\lambda$ )
10:   $\mathbf{V} \leftarrow \mathbf{A}^{(1)T} \mathbf{A}^{(1)} * \mathbf{A}^{(2)T} \mathbf{A}^{(2)}$ 
11:   $\mathbf{A}^{(3)} \leftarrow \mathcal{X}^{(3)} (\mathbf{A}^{(2)} \odot \mathbf{A}^{(1)}) \mathbf{V}^\dagger$ 
12:  normalize columns of  $\mathbf{A}^{(3)}$  (storing norms as  $\lambda$ )
13: until fit ceases to improve or maximum iterations
    reached
14: return  $\lambda$ ,  $\mathbf{A}^{(1)}$ ,  $\mathbf{A}^{(2)}$ ,  $\mathbf{A}^{(3)}$ 

```

SPLATT² is an open source software toolbox for sparse tensor factorization and related kernels [11], [12], [13]. SPLATT includes routines for computing least-squares CP, as well as constrained CP and CP with missing values (i.e., tensor completion). SPLATT is written in C and uses OpenMP+MPI for a hybrid of shared- and distributed-memory parallelism. In this paper, we focus on SPLATT’s shared-memory implementation of CP-ALS and leave the distributed implementation for future work.

²SPLATT source code is available from <https://github.com/ShadenSmith/splatt>

SPLATT uses a novel data structure for the storage of sparse tensors, a *compressed sparse fiber* (CSF), that addresses the memory/computation trade-off of computing tensor-matrix products [11]. To efficiently perform MTTKRP on tensors stored in CSF, SPLATT uses a custom shared-memory parallelized algorithm. Furthermore, CSF allows SPLATT to store and operate on tensors with an arbitrary number of modes. We refer readers to the original publication of CSF for more details on the design of its data structure and parallelized algorithms [11].

SPLATT’s approach to storing sparse tensors and its shared-memory parallel algorithm for MTTKRP has been shown to be generally the fastest with respect to execution time as well as the most efficient with regard to memory usage [10].

IV. PORTING SPLATT TO CHAPEL

Our goal when implementing the SPLATT code in Chapel was to simplify the code while preserving the original implementation’s design and approach. Therefore, we made no significant algorithmic or data structure changes. However, where possible, we employed Chapel’s high-level language features to simplify the code. Much of the SPLATT C code was ported to Chapel in a straightforward manner, but a few portions required a significantly different approach. In this section, we describe the challenges we encountered during porting and outline features lacking in Chapel that would have been beneficial to our work.

A. Mutex Pool

In the C/OpenMP implementation of SPLATT, a mutex pool is used in some of the parallel MTTKRP routines. As Chapel does not have a built-in mutex or lock, we used an array of `sync bool` variables. When the pool is initialized, the `sync` variables are set to true, putting them in the “full” state, as discussed in Section II. Acquiring a lock is equivalent to reading the `sync` variable and releasing a lock is equivalent to writing to the `sync` variable. Similar approaches have been taken elsewhere to create arrays of OpenMP locks in Chapel [5]. This technique was functionally correct, but resulted in a significant loss of performance for our application, as discussed in Section V-D. Therefore, we modified our approach to use `atomic` variables instead of `sync` variables. Listing 6 shows how acquiring and releasing a lock is implemented with `atomic` variables.

```

1 proc set(pool : [] atomic bool, lockID : int) {
2   while pool[lockID].testAndSet() {
3     chpl_task_yield();
4   }
5 }
6 proc unset(pool : [] atomic bool, lockID : int) {
7   pool[lockID].clear();
8 }

```

Listing 6. Acquiring/Releasing Locks via `atomic` variables

B. Work Sharing Constructs

One pattern of parallel computation in the C SPLATT implementation involves each thread operating on its own

private buffer but also iterating over a designated slice of a matrix. This is accomplished with a `omp for` nested within a `omp parallel` section, as illustrated in Listing 7. In this code, each thread updates its own copy of `myVals` but accesses only a certain slice of rows of the matrix `vals`, as the outermost for-loop is parallelized and split up among the team of threads spawned by the `omp parallel` directive. While a `coforall` construct can be used in place of the `omp parallel` directive, we cannot use a `forall` loop in place of the `omp for` directive. Doing so would result in each task created by the `coforall` construct executing line 8 for all values of `i` rather than only a subset. To the best of our knowledge, there is no direct translation of Listing 7 to Chapel using only its built-in constructs. Instead, we must use a normal for-loop inside of a `coforall` construct and manually compute the loop bounds for each task.

```

1 #pragma omp parallel
2 {
3     int tid = omp_get_thread_num();
4     double *myVals = thdData[tid];
5     #pragma omp for
6     for (int i = 0; i < rows; i++) {
7         for (int j = 0; j < cols; j++) {
8             myVals[j] += vals[i][j] * 2;
9         }
10    }
11    // do reduction on myVals
12 }

```

Listing 7. Nesting `omp for` within `omp parallel`

C. Array of Arrays

Some of the data structures within the reference SPLATT implementation use an array of arrays, where the sub-arrays are of varying sizes. This is accomplished by declaring an array of N pointers, where N is the number of sub-arrays, and then allocating each sub-array to be of the desired length. While Chapel does support the concept of an array of arrays, it currently does not allow the sub-arrays to be of different sizes. To overcome this limitation, we resorted to using an array of objects, where each object contains its own array. This approach was effective, but made the Chapel code less readable.

D. Function Pointers

In the original C/OpenMP SPLATT implementation, a pointer to a non-trivial function was passed as an argument to another function. To the best of our knowledge, this is not currently supported in Chapel. To work around this, we created a class that consists of the function we wish to pass as an argument. We then create an instance of the class and pass that object as an argument to the desired procedure.

E. Commentary on Chapel Features

From a programmer productivity perspective, we found porting the C/OpenMP code to Chapel to be relatively straightforward, even for inexperienced Chapel users such as ourselves. We particularly found features such as built-in reductions, whole array assignments and operations, array

resizing and the BLAS/LAPACK interface to be of significant value for reducing code complexity. The lack of a mutex or lock capability, having limited control over parallelization of nested loops, and the absence of native support for function pointers and jagged arrays complicated the porting effort and increased the complexity of the final source code.

V. PERFORMANCE EVALUATION

Having completed our Chapel implementation of SPLATT, we turned our attention to its performance, particularly in comparison to the C implementation. Our initial Chapel code was considerably slower than the C version. We conducted a performance analysis to locate the most significant performance issues and identified the sorting routine, MTTKRP kernel and LAPACK-based matrix inverse routine as the code segments with the largest performance differences. To improve the performance of these routines, we not only modified the Chapel code but also investigated the runtime environment used by Chapel and discovered significant performance issues related to the tasking layer. Sections V-C through V-E provide details on these issues and their resolutions. We present the performance of our final, optimized code in Section V-F.

A. System Configuration and Data Sets

Our Chapel code is based on SPLATT v2.0.0 (repository branch `b4bbad4-master`). Our code currently only includes SPLATT’s shared-memory parallel implementation. SPLATT’s optional feature to tile the modes of a tensor was omitted from our port, as it is not commonly used, and is not evaluated in our experiments. Also, we limited our port to operate on 3^{rd} order tensors. This was done to simplify our port and because most tensor data sets available are 3^{rd} order tensors.

We conducted performance tests on 5 different 3^{rd} order tensors. Across these different input tensors, we found the performance trends between the C reference implementation and our Chapel code to be similar. Therefore, for the sake of brevity, in this paper we only present the results from the YELP and NELL-2 data sets, as described in Table I. The Yelp Phoenix Academic Data set, from the Yelp Dataset Challenge³, contains reviews of businesses. The Never Ending Language Learner (NELL) data [14] contains subject-verb-object relationships between words. The YELP data set is representative of small sparse tensors while the NELL-2 data set is representative of moderate to large sparse tensors. Furthermore, the YELP data set showcases a particular feature that resulted in significant performance differences between the Chapel code and the reference C code, as discussed in Section V-D.

The details of our evaluation hardware and software are summarized in Table II. Both the original C/OpenMP code and our Chapel code make use of the `syrk`, `potrf` and `potrs` BLAS/LAPACK routines via OpenBLAS. Unless otherwise stated, the Chapel environment described in Table II is the configuration we used.

³https://www.yelp.com/dataset_challenge/

TABLE I
PROPERTIES OF DATA SETS

Name	Dimensions	Non-Zeros	Density	Size on Disk
YELP	41k x 11k x 75k	8M	1.97E-7	240 MB
RATE-BEER	27k x 105k x 262k	62M	8.3E-8	1.85 GB
BEER-ADVOCATE	31k x 61k x 182k	63M	1.84E-7	1.88 GB
NELL-2	12k x 9k x 29k	77M	2.4E-5	2.3 GB
NETFLIX	480k x 18k x 2k	100M	5.4E-6	3 GB

TABLE II
ENVIRONMENT AND SYSTEM PROPERTIES

Hardware	Software	Chapel
CPU: 2x E5-2697v4 Xeon Broadwell Cores: 36 Frequency: 2.3 GHz Last-level cache: 45 MB Memory: 512 GB DDR4	CentOS Linux 7.4.1708 gcc 4.8.5 OpenMP 3.1 OpenBLAS 0.2.20 SPLATT v2.0.0	Version: 1.16 Tasks: Qthreads Memory: jemalloc flags: -fast OMP_NUM_THREADS=1

TABLE III
RUNTIME IN SECONDS FOR CP-ALS ROUTINES - INITIAL RESULTS

Data set	Threads/Tasks	Code	MTTKRP	Sort	Mat A ^{TA}	Mat norm	CPD fit	Inverse
YELP	1	C	13.31	0.82	0.34	0.14	0.04	0.94
		Chapel-initial	225.11	7.21	0.36	0.14	0.04	0.98
	32	C	0.73	0.07	0.41	0.01	0.01	0.05
		Chapel-initial	118.93	0.47	0.56	0.06	0.01	0.98
NELL-2	1	C	109.25	7.90	0.13	0.06	0.01	0.37
		Chapel-initial	1999	69.04	0.14	0.06	0.01	0.39
	32	C	5.81	0.63	0.24	0.01	0.01	0.04
		Chapel-initial	88.3	5.01	0.19	0.02	0.01	0.39

For our experiments, the reference C/OpenMP code and our Chapel code performed 20 iterations of CP-ALS on each data set with a decomposition rank of 35. The reported per-routine runtimes represent the total time for those routines at the end of the 20 iterations. Furthermore, we performed 10 trials for both codes and report the average of those runs. For the C/OpenMP code, we varied the number of OpenMP threads from 1 to 32 via the `OMP_NUM_THREADS` environment variable. While the Chapel code also utilizes OpenBLAS and can leverage OpenMP, we opted to fix the number of OpenMP threads to 1 for our Chapel experiments. This is further explained in Section V-E. We control the number of threads used by Qthreads within our Chapel code by setting the `CHPL_RT_NUM_THREADS_PER_LOCALE` environment variable. This restricts both the C/OpenMP and Chapel code to use the same number of threads during each run. Otherwise, Qthreads defaults to a thread count that is equal to the number of cores on the system. To control the number of tasks used for the `coforall` constructs, we varied a user-level configuration variable that is used in the loop iterator.

B. Initial Results

Table III shows the execution times of the main routines for the C and Chapel code on the YELP and NELL-2 data sets when using 1 and 32 threads/tasks. Inverse refers to computing the Moore-Penrose inverse (denoted as \mathbf{V}^\dagger in Algorithm 1), Mat A^{TA} refers to performing lines 4, 7, and 10 in Algorithm 1, Mat norm refers to normalizing the columns of the factor matrices (lines 6, 9, and 12), CPD fit refers to computing the

fit of the decomposition (line 13) and Sort refers to the pre-processing step that SPLATT performs to sort the non-zeros in the tensor prior to performing CP-ALS.

The routines that exhibit the most severe performance issues for our Chapel code are sorting, MTTKRP, and inverse (shown in bold-face in Table III). MTTKRP and sorting in Chapel are about 18x and 8.7x slower than the C code, respectively. It is also clear that there is a scalability issue on the YELP data set for the Chapel MTTKRP, as it only achieves a 1.9x speed up from 1 to 32 tasks. Furthermore, as mentioned in Section V-A, we executed our Chapel experiments with `OMP_NUM_THREADS` set to 1, preventing any parallelization of the inverse procedure that is implemented via LAPACK routines from OpenBLAS. Our reasoning for doing this is explained in Section V-E. In the following sections, we investigate these routines and present modifications to the Chapel code to improve performance.

C. Sorting Optimizations

As a pre-processing step, SPLATT employs a parallel counting sort, which consists of a recursive quicksort routine that is modified to specifically sort non-zeros from 3^{rd} order tensors. We profiled our Chapel code to identify what portions of this sorting routine were consuming the most time. We identified one line of code in the quicksort routine that declared a local auxiliary array of two integers. In the case of the NELL-2 data set when using one thread or task, a total of 46 million recursive calls are made to the quicksort routine, and thus 46 million array allocations. While creating an array

of that size by itself exhibits very little overhead, it is still a more expensive task than the equivalent C code since Chapel arrays are high-level constructs. We found that this reoccurring array creation can account for as much as 10% of the sorting runtime. To mitigate this, we simply declared two separate integer variables rather than an array of two integers, requiring only minimal changes to the code.

The other portion of the code, which contributed the most to the sorting runtime, was a small loop executed before the quicksort routine was performed. For two array of arrays, A and B , that are of length N where N is the order of the tensor, the loop reassigns the sub-arrays of A to the corresponding sub-arrays of B . The length of each sub-array is equal to the number of non-zeros in the tensor, which can be significant in the case of large tensors. In C, the sub-arrays are pointers, making this a simple assignment operation. In our initial Chapel code, we implemented each array of arrays as a 2D matrix and performed the sub-array assignment using slicing. As slicing can create excessive overhead (see Section V-D), we opted to use an array of arrays rather than a 2D matrix. However, in reassigning each sub-array, Chapel performs a copy of one sub-array to another while the C code is just reassigning pointers. To match this behavior, we used the `c_ptrTo` procedure in Chapel to obtain pointers to the main arrays and then used the equivalent C syntax to reassign the sub-arrays. This change improved the entire sorting routine by roughly 4x on both the YELP and NELL-2 data set.

Combining both the array creation and slicing modifications mentioned above, we improved the sorting routine by as much as 8x from our original version. Figure 1 shows how the different modifications affected the overall sorting runtime for the NELL-2 data set as the number of threads/tasks are varied from 1 to 32. We observed similar trends for the YELP data set.

D. MTTKRP Optimizations

As the MTTKRP is the critical routine of CP-ALS, we focused heavily on improving its performance within our Chapel code. Through our performance profiling, we identified two areas that contribute the most to the performance of our implementation: matrix row accessing and mutexes/locks.

1) *Matrix Row Accesses*: A frequent pattern in the MTTKRP code is to obtain a pointer to a particular row in a factor matrix and then perform a multiply/accumulate across the elements in the row, similar to the code shown in Listing 5. In the C implementation of SPLATT, the factor matrices are stored as 1D arrays in row-major order, so accessing any given row can be done simply through pointer arithmetic. Our Chapel code represents the factor matrices as 2D matrices, with the goal of leveraging Chapel’s built-in features such as slicing.

In our initial code, we obtained a row reference by slicing the factor matrix. An issue with using slicing in this manner in our code is that the amount of computation that is performed on each slice is negligible when compared to the overhead of obtaining the slice. Each slice only consists of R elements,




Fig. 1. Chapel sorting runtime on the NELL-2 data set when using 1 to 32 threads/tasks. Initial refers to our unoptimized Chapel code, Array-opt refers to the optimization that eliminates the creation of a small integer array in a frequently called function. Slices-opt refers to the optimization that eliminates the use of slicing in array reassignment operations and All-opts refers to the Chapel code when using both Array-opt and Slices-opt. All runtimes are shown in seconds.

where R is the rank of the decomposition and typically small (35 in our case), and the computation is fairly light. As described in a Chapel GitHub issue⁴, array slicing can be expensive due to computing and creating the domain of the resulting array view and creating and setting up the array descriptor for the view.

Our first approach was to eliminate slicing by using direct 2D indexing for matrices, even though it deviated from the reference implementation of SPLATT. We found that this modification led to as much as a 12x and 17x speed-up for the YELP and NELL-2 MTTKRP runtimes, respectively. Next, we adopted a more direct translation of the C code into our Chapel code by obtaining a pointer to the factor matrices and then using pointer arithmetic to access the particular rows. This change produced about a 1.26x speed-up over the 2D indexing approach. Figures 2 and 3 present these optimizations for the YELP and NELL-2 data sets, respectively, when varying the number of threads/tasks from 1 to 32.

2) *Mutexes/Locks*: It is obvious from Figure 2 that the MTTKRP runtime for the YELP data set is exhibiting poor scalability. This behavior is not seen with the reference C/OpenMP implementation. However, the results for the NELL-2 data set in Figure 3 do not show the same poor scalability, rather showing near linear speed-up. We noticed that the main difference between the YELP and NELL-2 data sets with respect to the MTTKRP is that for all thread/task counts beyond two for the YELP data set, the SPLATT algorithm will require the use of locks during the MTTKRP, as briefly mentioned in Section IV, while the NELL-2 data set will perform “no-lock” versions of the MTTKRP for all

⁴<https://github.com/chapel-lang/chapel/issues/8203>




Fig. 2. Chapel MTTKRP runtime on the YELP data set when using 1 to 32 threads/tasks. Initial refers to our unoptimized Chapel code using slicing, 2D Index refers to replacing slicing with more direct 2D indexing and Pointer refers to retrieving C-pointers and using pointer arithmetic. All runtimes are shown in seconds and the vertical axis is logarithmic.




Fig. 3. Chapel MTTKRP runtime on the NELL-2 data set when using 1 to 32 threads/tasks. Initial refers to our unoptimized Chapel code using slicing within the MTTKRP, 2D Index refers to replacing slicing with more direct 2D indexing and Pointer refers to retrieving C-pointers and using pointer arithmetic. All runtimes are shown in seconds and the vertical axis is logarithmic.

thread/task counts. The decision of whether or not to use locks is highly dependent on the tensor and the number of threads being used.

In our initial code, locks were implemented with `sync` variables. We discussed our implementation with the Chapel team and they suggested using `atomic` variables instead, as `sync` variables can be very expensive per operation, especially within the Qthreads tasking layer. Within Qthreads, `sync` variables are better suited for heavily contended locks held for long periods of time. This is because if the variable is




Fig. 4. Chapel MTTKRP runtime on the YELP data set when using 1 to 32 threads/tasks. Sync refers to using `sync` variables for our mutex pool implementation, Atomic refers to using `atomic` variables for our mutex pool implementation and FIFO-sync refers to using `sync` variables but with the `fifo` tasking layer. All runtimes are shown in seconds and the vertical axis is logarithmic. All implementations shown use the `Pointer` optimization for accessing the factor matrices within the MTTKRP.

locked when requested, the requesting task will be put to sleep rather than execute a spin-wait. However, the critical sections in the MTTKRP that are protected by locks consist of very short, non-intensive computation. Using this information, we modified our mutex implementation to use `atomic` variables, as described in Section IV, which employ a spin-wait behavior that is better suited for our application. By doing so, we significantly improved both the scalability and overall MTTKRP runtime for the YELP data set, as shown in Figure 4. As the NELL-2 data set does not require the use of locks, the performance of `sync` vs `atomic` variables is irrelevant. The particular case of the YELP data set speaks to the high irregularity in the sparse tensor decomposition application and the issues that can arise due to over-optimization for a particular data set.

As part of our investigation into the MTTKRP runtime and the mutex implementation, we tried using an alternative tasking layer to Qthreads, namely `fifo` (i.e., POSIX threads). Figure 4 includes the results for using `fifo` and `sync` variables, which is competitive with the Qthreads and `atomic` implementation. This is because under the `fifo` tasking layer, `sync` variables are implemented with a spin-wait like behavior rather than putting the task to sleep. This observation underscores the importance of the tasking layer used and its effects on Chapel’s features.

E. Conflicts Between Qthreads and OpenMP

In both the C and Chapel codes, the matrix inverse procedure is computed using two LAPACK routines provided by OpenBLAS, which are parallelized via OpenMP. When running our Chapel code with multiple OpenMP threads, we encountered performance issues for not only the inverse procedure but other routines using Qthreads. We have isolated




Fig. 5. Runtimes for CP-ALS routines on the YELP data set when using 1 thread/task. Times are shown in seconds. Lower bars represent better performance. C refers to the original C/OpenMP code and Chapel-optimize refers to our fully optimized Chapel code (i.e., all sorting optimizations, C-pointers within MTTKRP and `atomic` variables for mutexes).




Fig. 7. Runtimes for CP-ALS routines on the YELP data set when using 32 threads/tasks. Times are shown in seconds. Lower bars represent better performance. C refers to the original C/OpenMP code and Chapel-optimize refers to our fully optimized Chapel code (i.e., all sorting optimizations, C-pointers within MTTKRP and `atomic` variables for mutexes).




Fig. 6. Runtimes for CP-ALS routines on the NELL-2 data set when using 1 thread/task. Times are shown in seconds. Lower bars represent better performance. C refers to the original C/OpenMP code and Chapel-optimize refers to our fully optimized Chapel code (i.e., all sorting optimizations, C-pointers within MTTKRP and `atomic` variables for mutexes).




Fig. 8. Runtimes for CP-ALS routines on the NELL-2 data set when using 32 threads/tasks. Times are shown in seconds. Lower bars represent better performance. C refers to the original C/OpenMP code and Chapel-optimize refers to our fully optimized Chapel code (i.e., all sorting optimizations, C-pointers within MTTKRP and `atomic` variables for mutexes).

these issues to conflicts between the Qthreads tasking layer and the OpenMP threads used by OpenBLAS.

As we increased the number of OpenMP threads from 1 to 32, we observed that the inverse routine's runtime on the YELP data set became 15x slower at 32 threads than the serial case. By default, Qthreads pins workers to cores for optimal affinity. This can negatively impact OpenMP threads by using resources on cores for overhead tasks such as checking for more work. We tried setting `QT_AFFINITY=no`, which allows spin-waiting threads to migrate to different cores that are out of the way of the OpenMP threads. We found that

this drastically improved the performance of the matrix inverse routine, achieving a 2x speed-up rather than the initial 15x slow down. However, this runtime is still roughly 10x slower than the C code for the YELP data set.

To further reduce the conflicts between OpenMP and QThreads, we reduced the spin-wait interval. By default, Qthreads workers will spin-wait for 300,000 iterations before suspending. While this spin-waiting is occurring, it can degrade the performance of the OpenMP threads sharing those cores. In our case, the contention is large because the inverse routine has a relatively short runtime and follows a section of

QThreads activity. With a long spin-wait period, a significant spin-wait overlap occurs. Following the suggestion in a Chapel GitHub issue⁵, we shortened the amount of spin-waiting by setting `QT_SPINCOUNT=300`. Doing so further improved the performance of the inverse procedure on 32 threads by 2.3x. However, even this result is still 4x slower than the C code.

Just as QThreads can degrade OpenMP performance, the opposite is also true. Using `QT_AFFINITY=no` resulted in unpredictable performance when using 32 threads for the matrix normalization routine, which is written entirely in Chapel and follows directly after the OpenMP-based inverse procedure. We observed a 7x – 13x slow down in the matrix normalization runtime at 32 threads when compared to our default configuration. We theorize that, without QThreads core pinning, spin-waiting QThreads and OpenMP threads migrate away from one another. Once the program exits the OpenMP region, these QThreads will re-migrate back out to idle cores if necessary. The contention is minor when the system has enough cores to accommodate all QThreads and OpenMP threads. When using 32 threads, a majority of the cores in our 36 core system are being used by the OpenMP threads during the inverse procedure and by the Qthread workers during the matrix normalization routine, leading to significant migration with its accompanying cache rewarming and context switch overhead.

There does not appear to be a clear solution to overcoming these issues. If there were BLAS routines implemented entirely in Chapel, that could avoid the conflicts between Qthreads and OpenMP. Because the inverse routine generally contributes a small amount to the overall CP-ALS runtime and is the only procedure in our code that can benefit from using OpenMP threads, we have opted to set the number of OpenMP threads to 1 for our Chapel experiments.

F. Final Results

Figures 5 and 6 present the per-routine runtimes of the reference C/OpenMP implementation and our fully optimized Chapel implementation when using one thread/task on the YELP and NELL-2 data sets, respectively. We also show the 32 threads/tasks results in Figures 7 and 8. Furthermore, the MTTKRP runtime on the YELP and NELL-2 data sets as the number of threads/tasks is varied from 1 to 32 are presented in Figures 9 and 10, respectively.

We observe drastic improvements with respect to our original Chapel code (Table III) for the MTTKRP and sorting routines. While our final Chapel code’s performance for sorting is still slightly worst than the C/OpenMP code and could be investigated further for more improvement, it only represents a small portion of the overall CP-ALS runtime. It is evident that our Chapel code achieves competitive parallel performance and scalability for the MTTKRP when compared to the reference C/OpenMP implementation. Overall, the Chapel code achieves 83%-93% of the MTTKRP performance




Fig. 9. MTTKRP runtime on the YELP data set when using 1 to 32 threads/tasks. C refers to the original C/OpenMP code, Chapel-initial refers to our original unoptimized version of the Chapel code and Chapel-optimize refers to our fully optimized Chapel code (i.e., all sorting optimizations, C-pointers within MTTKRP and `atomic` variables for mutexes). The vertical axis is logarithmic and shown in seconds.




Fig. 10. MTTKRP runtime on the NELL-2 data set when using 1 to 32 threads/tasks. C refers to the original C/OpenMP code, Chapel-initial refers to our original unoptimized version of the Chapel code and Chapel-optimize refers to our fully optimized Chapel code (i.e., all sorting optimizations, C-pointers within MTTKRP and `atomic` variables for mutexes). The vertical axis is logarithmic and shown in seconds.

of the C/OpenMP code on the YELP data set and 84%-96% of the MTTKRP performance on the NELL-2 data set. It also achieves near linear scalability up to 32 cores on both data sets. However, it is clear that the inverse routine suffers in performance (most notably in Figure 7) due to the parallelization issues mentioned in Section V-E.

VI. RELATED WORK

Designing and implementing sparse parallel tensor decomposition algorithms has been a rich research area in recent

⁵<https://github.com/chapel-lang/chapel/issues/8337>

years. There are several shared-memory based tensor decomposition implementations using novel approaches to performing MTTKRP in a scalable and efficient manner [11], [1], [2]. However, to the best of our knowledge, the work presented in this paper is the first to implement parallel sparse tensor decomposition in a high productivity programming language such as Chapel.

There has been a significant effort to evaluate and analyze the performance of Chapel programs for both single- and multi-node environments. Johnson and Hollingsworth ported and optimized several C/OpenMP based benchmarks to single-node Chapel including LULESH, MiniMD, and CLOMP [5]. Haque and Richards implemented an optimized multi-node version of CoMD in Chapel as well as identified key limitations of Chapel in regards to scope-based code locality [6]. Our work, while similar, differs from these efforts in that SPLATT is a sparse application from a different problem domain. It also is a full application with several components, ranging from file I/O and sorting to custom sparse data structures and parallel algorithms, rather than a benchmark or proxy application.

There has also been work on developing techniques to more effectively measure the performance of Chapel programs, where a data-centric view of performance data is studied as opposed to more traditional code-centric views [15]. In our work, we employed code-centric profiling to identify performance bottlenecks via gprof and source-code level timers. In the future, we would be interested in applying such data-centric techniques to improve our code.

VII. CONCLUSION

In this work, we implemented SPLATT, the highest performing parallel sparse tensor decomposition implementation, using Chapel. We identified features missing from Chapel that would have benefited our porting effort, such as a mutex library, varying length sub-arrays within an array of arrays, and more support for functions as first class constructs. Through our performance study, we identified performance bottlenecks in our initial Chapel code, which involved not only Chapel's language features but also its tasking layer. We highlighted three key performance limiters:

- Array slicing overhead: We identified that the excessive use of array slicing was significantly degrading our performance in several parts of the code. We mitigated this by obtaining pointers to the arrays and then using pointer-arithmetic to access the required slices. This improved performance by as much as 18x but resulted in less readable code.
- Sync vs atomic variables: The use of `sync` variables under the Qthreads tasking layer for our mutex pool implementation resulted in poor parallel performance and scalability due to our very short critical sections. We redesigned the mutex pool to use `atomic` variables and achieved a 14.5x speed-up.
- Conflicts between Qthreads and OpenMP: Using the default Qthread settings drastically degraded the performance of LAPACK routines in our code when leveraging

OpenMP threads. While adjusting Qthread's affinity and spin-waiting settings improved the performance, the run-times were still 2.3x slower than the C code, at best. Furthermore, adjusting Qthread's affinity had negative affects on other Qthreads-based routines in our code.

Our Chapel code now achieves up to 96% of the performance of the C language implementation with respect to the critical routine, MTTKRP.

For future work, we intend to extend our Chapel port to include some of the features left out, such as support for tensors of arbitrary order and tensor mode tiling. We also plan to incorporate SPLATT's novel distributed-memory features [16] for tensor decomposition in our code, leveraging Chapel's multi-locales.

REFERENCES

- [1] S. Smith, J. Park, and G. Karypis, "Sparse tensor factorization on many-core processors with high-bandwidth memory," *31st IEEE International Parallel & Distributed Processing Symposium (IPDPS'17)*, 2017.
- [2] M. Baskaran, T. Henretty, B. Pradelle, M. H. Langston, D. Bruns-Smith, J. Ezick, and R. Lethin, "Memory-efficient parallel tensor decompositions," in *2017 IEEE High Performance Extreme Computing Conference, HPEC 2017, Waltham, MA, USA, September 12-14, 2017*, 2017.
- [3] B. L. Chamberlain, "Chapel," in *Programming Models for Parallel Computing (Scientific and Engineering Computation)*. The MIT Press, 2015, pp. 129–159.
- [4] E. Kayraklioglu, W. Chang, and T. El-Ghazawi, "Comparative performance and optimization of chapel in modern manycore architectures," in *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, May 2017, pp. 1105–1114.
- [5] R. B. Johnson and J. Hollingsworth, "Optimizing Chapel for single-node environments," in *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, May 2016, pp. 1558–1567.
- [6] R. Haque and D. Richards, "Optimizing PGAS overhead in a multi-locale chapel implementation of comd," in *2016 PGAS Applications Workshop (PAW)*, Nov 2016, pp. 25–32.
- [7] S. Smith, N. Ravindran, N. D. Sidiropoulos, and G. Karypis, "SPLATT: Efficient and parallel sparse tensor-matrix multiplication," *29th IEEE International Parallel & Distributed Processing Symposium*, 2015.
- [8] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, September 2009.
- [9] N. D. Sidiropoulos, L. D. Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, "Tensor decomposition for signal processing and machine learning," *IEEE Transactions on Signal Processing*, vol. 65, no. 13, pp. 3551–3582, July 2017.
- [10] T. B. Rolinger, T. A. Simon, and C. D. Krieger, "Performance considerations for scalable parallel tensor decomposition," *Journal of Parallel and Distributed Computing*, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0743731517302897>
- [11] S. Smith and G. Karypis, "Tensor-matrix products with a compressed sparse tensor," *Proceedings of the 5th Workshop on Irregular Applications: Architectures and Algorithms*, 2015.
- [12] —, "Accelerating the tucker decomposition with compressed sparse tensors," in *European Conference on Parallel Processing*. Springer, 2017.
- [13] S. Smith, J. Park, and G. Karypis, "HPC formulations of optimization algorithms for tensor completion," *Parallel Computing*, 2017.
- [14] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. H. Jr., and T. M. Mitchell, "Toward an architecture for never-ending language learning," in *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence (AAAI 2010)*, 2010.
- [15] H. Zhang and J. K. Hollingsworth, "Data centric performance measurement techniques for chapel programs," in *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2017, pp. 377–386.
- [16] S. Smith and G. Karypis, "A medium-grained algorithm for distributed sparse tensor factorization," *30th IEEE International Parallel & Distributed Processing Symposium*, 2016.